

IB-GAN: A Unified Approach for Multivariate Time Series Classification under Class Imbalance

Grace Deng
Dept. of Statistics & Data Science
Cornell University
gd345@cornell.edu

Cuize Han
Amazon Search
cuize@amazon.com

Tommaso Dreossi
Amazon Search
tommasodreossi@gmail.com

Clarence Lee
Cornell University
clarence.lee@cornell.edu

David S. Matteson
Dept. of Statistics & Data Science
Cornell University
dm484@cornell.edu

Abstract

Classification of large multivariate time series with strong class imbalance is an important task in real-world applications. Standard methods of class weights, oversampling, or parametric data augmentation do not always yield significant improvements for predicting minority classes of interest. Non-parametric data augmentation with Generative Adversarial Networks (GANs) offers a promising solution. We propose Imputation Balanced GAN (IB-GAN), a novel method that joins data augmentation and classification in a one-step process via an imputation-balancing approach. IB-GAN uses imputation and resampling techniques to generate higher quality samples from randomly masked vectors than from white noise, and augments classification through a class-balanced set of real and synthetic samples. Imputation hyperparameter p_{miss} allows for regularization of classifier variability by tuning innovations introduced via generator imputation. IB-GAN is simple to train and model-agnostic, pairing any deep learning classifier with a generator-discriminator duo and resulting in higher accuracy for under-observed classes. Empirical experiments on open-source UCR data and a 90K product dataset show significant performance gains against state-of-the-art parametric and GAN baselines.

Keywords: time series, imbalanced classification, data augmentation, machine learning, generative adversarial networks

1 Background and Motivation

Multivariate time series classification (MTSC) is a growing field where complex features are collected in the

form of time-indexed sequences with varying lengths. Often a single observation unit (product, video, recommendation, etc.) can be described by multiple time series metrics with strong inter-temporal dependencies and a mixture of numeric, categorical, and semantic features. Many recent works using variants of CNNs or RNNs have shown impressive performance in MTSC tasks [1, 2, 3, 4, 5]. However, many real-world datasets consists of strong class imbalance where far fewer samples are observed from minority classes of interest. The three baseline methods for addressing imbalances are class weights, upsampling, and downsampling, each with drawbacks [6, 7]. Downsampling leads to poor classifier performance when too many samples are discarded, while upsampling leads to overfitting by reusing data from the minority class. Data augmentation techniques such as SMOTE [8, 9] propose creating additional synthetic samples from oversampling and linear interpolation of k-nearest neighbors, but is still a parametric model constrained by computation time and does not generalize well to high-dimensional datasets [10]. Techniques for image augmentation [11] include transformations, cropping, noise injection, and random erasing [12]; however, far fewer methods have been developed for time series or more general data types.

Large-scale non-parametric data augmentation [13, 11, 14] with Generative Adversarial Networks (GANs) [15] poses a promising solution. Our paper presents Imputation Balanced GAN (IB-GAN), a novel imputation-balancing approach to data augmentation that can pair any deep learning classifier with a generator-discriminator model to address class imbalance in MTSC.

1.1 Related Works. Modern MTSC methods involve combinations of Convolutional Neural Networks (CNNs), Recurrent Neural Nets (RNNs), and Long Short-Term Memory (LSTM) networks, with variants such as stacked CNNs and Squeeze-and-Excitation block [4], deep CNNs [16], and attention networks [17]. These works are not designed for imbalanced data. A 2019 survey paper [1] stated that very little has been done for imbalanced classes in MTSC other than class weights. Fawaz et al [13] and Tran et al [18] showed data augmentation with deep learning models improves image classification. We propose that any of these classifiers can be adapted for MTSC and paired with GANs to augment the training process, i.e., include synthetic data for under-sampled classes.

The GAN has two competing models: the *generator* creates synthetic data by learning the true data distribution, and the *discriminator* distinguishes real and synthetic samples. The vanilla GAN [15] has uses in many generative tasks [19, 20]. Deep-Convolutional GANs (DC-GAN) use convolution layers [21] to improve training stability via learning feature maps and object representations. Conditional GANs [22] extends DC-GAN and conditions on auxiliary information such as class labels. Additional techniques such as batch normalization [23], leaky-relu activation, max-pooling, and dropout all can improve GAN training. GANs specifically designed for imputation such as GAIN [24], Colla-GAN [25], Mis-GAN [26, 27], and for sequence generation [28, 29, 30] have also been introduced.

More recent works focus on generating better samples. The Balancing GAN (BAGAN) adopts an autoencoder-decoder structure to learn class conditioning in the latent space [14]. The Auxiliary Classifier GAN (AC-GAN) [31, 32] proposed a secondary classifier embedded in the discriminator in order to stabilize GAN training, but does not apply it directly for classification; a second-stage classifier is required. The Rumi formulation for GANs [33] trains the generator to specifically produce samples from a desired class. The Triple-GAN [34] uses a generator to produce pseudo-samples and classifier to produce pseudo-labels, and the discriminant tries to predict whether a pseudo-pair of (sample, label) is real or synthetic. All of these methods (BAGAN, AC-GAN, etc.) were designed for image data in mind, but could potentially be adapted for time-indexed data.

These GAN methods suffer from two drawbacks. First, they are specifically designed for data generation but requires a separate classification step, resulting in a two-stage training process where the generative model does not learn from classification feedback (see Figure 1). Second, these methods generate synthetic data from white or latent noise [28, 14], which risks

unrealistic or repetitive samples due to mode collapse. The Imputation Balanced GAN (IB-GAN) will address these two challenges through a joint training process and a two-pronged imputation-balancing approach.

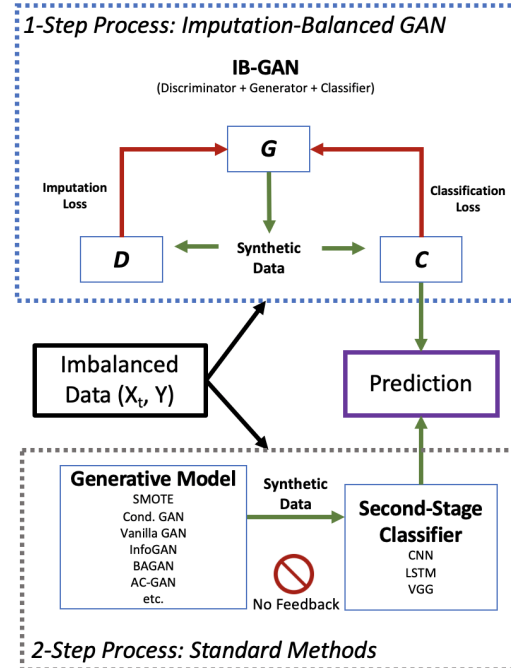


Figure 1: IB-GAN versus the two-step process for existing GAN data augmentation methods. IB-GAN directly utilizes C model in (D, G, C) triplet for classification; see Figure 2 for detailed set-up.

1.2 Contributions. We propose the Imputation Balanced GAN (IB-GAN) for imbalanced classification with three key contributions:

1. We present a unified one-step process (Figure 1) for joint data augmentation and classification under class imbalance, consisting of a triplet of generator G , discriminator D , and classifier C models that is model agnostic to different neural net architectures.
2. We adapt a novel imputation approach for generating synthetic samples from randomly masked vectors; data quality is improved via direct feedback from classification and imputation losses. A tuning parameter p_{miss} regulates innovations in generator imputations while preventing mode collapse common in standard GANs initialized by white noise.
3. We balance classifier training via resampling techniques and synthetic samples for under-observed classes, with significant performance gains against state-of-the-art baselines.

Full overview of IB-GAN is outlined in Figure 2. Theoretically, a range of GAN variants (vanilla [15], conditional [22], Info-GAN [35] etc.) and deep learning classifiers (CNN, RNN, LSTM, VGG-16) is possible for IB-GAN given the model-agnostic framework; we explore these options in Section 3.1. Elaborating on contribution (2), IB-GAN generator initializes data generation with randomly masked data vectors using the Missing Completely At Random (MCAR) missing mechanism [36] where probability of missingness is independent of all data values, and thereby leverages existing information to produce better quality samples. Resulting imputations are analogous to image perturbations in computer vision applications [37, 38]. Empowered by these new weighted resampling and data masking steps, the final IB-GAN classifier generalizes well to hold-out sets with strong class imbalance.

2 Methods and Technical Solutions

We briefly discuss the challenge to train unbiased classifiers from highly imbalanced data. Methods such as class weights may cause numerical instability [39] and data augmentation, e.g., via GANs, poses a better solution. We then introduce the main IB-GAN framework and objective functions.

2.1 Data Augmentation for Imbalanced Classification. Without loss of generality, we assume all random variables involved are discrete with probability measure P . Suppose the observed data is coming from the realization of random pair: (\mathbf{X}, Y) . Y is the class label which is from a finite class label set \mathcal{Y} . Denote the prior label probability as $w_y = P(Y = y)$ for $y \in \mathcal{Y}$. \mathbf{X} is a k dimensional random vector (some characteristic features for class Y), $\mathbf{X} = (X_1, \dots, X_k)$ where $X_i \in \mathcal{X}$. Given $Y = y$, denote the conditional probability $P(\mathbf{X} = \mathbf{x} | Y = y)$ as $p_y(\mathbf{x})$. The task is to train a classifier C that maps \mathcal{X}^k to \mathcal{Y} or equivalently predicts the class label probability. We will use the latter definition of the task. That is, for $\mathbf{x} \in \mathcal{X}^k$, $C(\mathbf{x}) = \{c_y(\mathbf{x})\}_{y \in \mathcal{Y}}$, where $0 \leq c_y(\mathbf{x}) \leq 1$ and is the predicted probability that \mathbf{x} belongs to class y . We have $\sum_{y \in \mathcal{Y}} c_y(\mathbf{x}) = 1, \forall \mathbf{x} \in \mathcal{X}^k$. The classifier maximizing equal class weight negative cross-entropy loss:

$$(2.1) \quad \begin{aligned} L_{class}(C) &= \sum_{y \in \mathcal{Y}} \mathbb{E} [1\{Y = y\} \log c_y(\mathbf{X})] \\ &= \sum_{y \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}^k} w_y p_y(\mathbf{x}) \log c_y(\mathbf{x}) \end{aligned}$$

is known as the Bayes classifier and can be expressed as $C^*(x) = \{c_y^*(\mathbf{x})\}_{y \in \mathcal{Y}}$, where $c_y^*(\mathbf{x}) = \frac{w_y p_y(\mathbf{x})}{\sum_{y' \in \mathcal{Y}} w_{y'} p_{y'}(\mathbf{x})} = \frac{P(Y=y, \mathbf{X}=\mathbf{x})}{P(\mathbf{X}=\mathbf{x})} = P(Y = y | \mathbf{X} = \mathbf{x})$.

This is the optimal classifier for the data distribution (\mathbf{X}, Y) . However, our focus is on minimizing overall classification error when we treat each class equally. Then under this criterion, the Bayes classifier is suboptimal due to highly unbalanced prior label probability w_y and is biased towards labels with high prior probability. Let U be a uniform random variable over the class label set \mathcal{Y} . The joint distribution of (\mathbf{X}, U) is defined through letting $P(\mathbf{X} = \mathbf{x} | U = y) \equiv p_y(\mathbf{x})$. Then the optimal balanced class classifier is $\bar{C}^*(x) = \{\bar{c}_y^*(\mathbf{x})\}_{y \in \mathcal{Y}}$, where $\bar{c}_y^*(\mathbf{x}) = P(U = y | \mathbf{X} = \mathbf{x}) = \frac{p_y(\mathbf{x})}{\sum_{y' \in \mathcal{Y}} p_{y'}(\mathbf{x})}$.

In order to derive the optimal balanced class classifier for data distribution (X, Y) with class imbalance, a common method is inverse class weights in the loss. The classifier that maximizes the negative cross-entropy loss with inverse class weights:

$$(2.2) \quad \begin{aligned} L_{class}^w(C) &= \sum_{y \in \mathcal{Y}} w_y^{-1} \mathbb{E} [1\{Y = y\} \log c_y(\mathbf{X})] \\ &= \sum_{y \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}^k} p_y(\mathbf{x}) \log c_y(\mathbf{x}) \end{aligned}$$

is $\bar{C}^*(x)$. Solving this empirical optimization problem, e.g., with highly unbalanced inverse weights w_y^{-1} , can cause numerical instability issues.

A better solution is in the form of data augmentation. Suppose we have additional samples for classifier training that are realizations of random pair (\mathbf{X}', Y') . For example, additional training samples can be derived from resampling of existing data, transformations and rotations, or model based data augmentation. Let the prior label probability of the additional data distribution to be $w'_y = P(Y' = y)$ for $y \in \mathcal{Y}$ and the conditional probability $p'_y(\mathbf{x}) := P(\mathbf{X}' = \mathbf{x} | Y' = y)$. We define the loss that combines the two source of data through a hyperparameter $0 < \alpha < 1$ as

$$(2.3) \quad \begin{aligned} L_{class}^\alpha(C) &= \alpha \sum_{y \in \mathcal{Y}} \mathbb{E} [1\{Y = y\} \log c_y(\mathbf{X})] \\ &\quad + (1 - \alpha) \sum_{y \in \mathcal{Y}} \mathbb{E} [1\{Y' = y\} \log c_y(\mathbf{X}')] \\ &= \sum_{y \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}^k} (\alpha w_y p_y(\mathbf{x}) + (1 - \alpha) w'_y p'_y(\mathbf{x})) \log c_y(\mathbf{x}) \end{aligned}$$

From Jensen's inequality, the optimal classifier that maximizes $L_{class}^\alpha(C)$ is $\bar{C}^*(\mathbf{x}) = \{\bar{c}_y^*(\mathbf{x})\}_{y \in \mathcal{Y}}$, where $\bar{c}_y^*(\mathbf{x}) = \frac{\alpha w_y p_y(\mathbf{x}) + (1 - \alpha) w'_y p'_y(\mathbf{x})}{\sum_{y' \in \mathcal{Y}} (\alpha w_{y'} p_{y'}(\mathbf{x}) + (1 - \alpha) w'_{y'} p'_{y'}(\mathbf{x}))}$. Data augmentation can control over the prior label probability w'_y . If $w'_y = \frac{|\mathcal{Y}|^{-1} - \alpha w_y}{1 - \alpha}$, $y \in \mathcal{Y}$ and we choose any α such that $\alpha < \frac{|\mathcal{Y}|^{-1}}{\max_y w_y}$, then under optimal condi-

tions where the augmented data has the same conditional distribution as the original data, $p'_y(\mathbf{x}) \equiv p_y(\mathbf{x})$, we have $\tilde{C}^*(x) \equiv \bar{C}^*(x)$. Thus with suitable choices of w'_y and hyperparameter α , we can effectively train the balanced class classifier through combination of true and GAN-generated samples. We now introduce the IB-GAN which achieves data augmentation and balanced classification in an unified process.

2.2 Imputation Balanced GAN. A standard GAN consists of two models [15], the generator G that learns the underlying data distribution p_{data} for data vector X from noise distribution Z , and the discriminator D that classifies whether the sample is from the true data distribution or a synthetic sample. For the IB-GAN [1] (Figure 2), we introduce a new component C , the classifier that predicts which class Y a sample belongs to, regardless of whether it is real or synthetic. Novel to prior work [31, 18, 34], the IB-GAN separates the discriminator and classifier into two different models with separate element-wise imputation loss and classifier loss, corresponding to the “imputation” and “balancing” aspect of IB-GAN respectively.

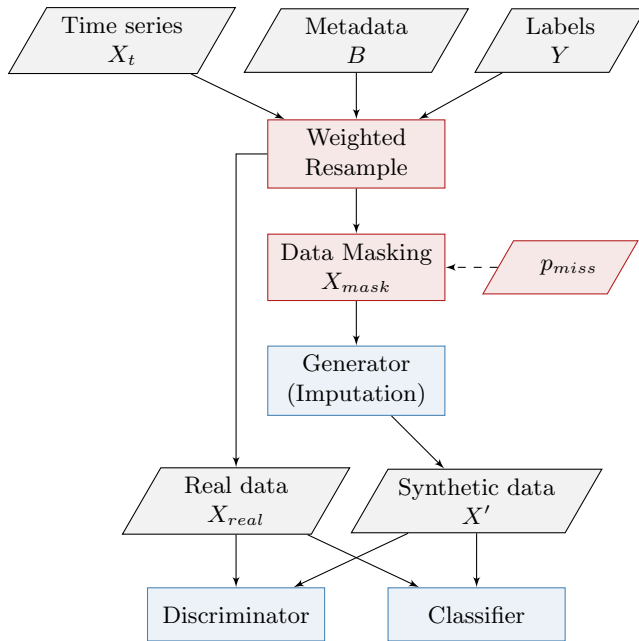


Figure 2: IB-GAN triple-model framework (blue) includes novel resampling and data masking steps (red) for joint data augmentation and classification.

2.2.1 Weighted Resample: Let X_t be a k -dimensional multivariate time series with sequence length m with optional metadata B and class labels Y . For a given mini-batch size n_{mb} , the weighted resampling step for $(X, Y) := (X_t, B, Y)$ will first sample real data X_{real} and corresponding labels from the true training data. For synthetic (X', Y') , we sample again X_{mask} (before masking) with weighted probability $\frac{1/w_y}{\sum_Y 1/w_y}$ from each class (denote the corresponding random variable as Y'). Together, X_{real} and X_{mask} have balanced samples from each class.

2.2.2 Data Masking with p_{miss} : The proportion of masked values for each sample is specified by hyperparameter p_{miss} . A low p_{miss} ensures imputed samples will be more realistic with greater similarity to original data; a higher p_{miss} encourages greater innovation and variety. It can be seen as a form of regularization on the generator imputation variability. At $p_{miss} = 1$, IB-GAN is analogous to combining a Conditional GAN generating data from white noise with a classifier (Naive GAN); at $p_{miss} = 0$, IB-GAN is equivalent to training on weighted bootstraps of the original data. Optimal choice for p_{miss} will likely depend on class ratio and sample size, and can be found via grid-search in practice.

Let \mathbf{I} be the component-wise indicator for whether each component of X_{mask} is masked (replaced with random white noise) or keeps the real value. \mathbf{I} has same dimension as the data and the randomly masked data vectors $X_{mask} = X \odot (\mathbf{1} - \mathbf{I}) + Z \odot \mathbf{I}$. By injecting randomness via noise Z into the generator input, we ensure diversity in the resulting synthetic samples and that the generator does not simply converge into a trivial time series predictor.

2.2.3 IB-GAN Triplet Training: The generator takes in (X_{mask}, Y') as input and impute synthetic samples (X', Y') , where Y' are not synthetic labels but rather the true class labels for the imputations, a key difference from Triple-GAN [34]. Let the prior label probability of the synthetic data be $w'_y = P(Y' = y)$ for $y \in \mathcal{Y}$ and the conditional probability $p'_y(x) := P(X' = x | Y' = y)$ which depends on the generator G . Both the true and synthetic samples (X_{real}, Y) and (X', Y') are now inputs to the discriminator and classifier. The classifier learns on a class-balanced sample, and the discriminator attempts to recover \mathbf{I} via the element-wise imputation loss. If the generator is performing well, the discriminator should find it difficult to distinguish between true and imputed values.

We then jointly optimize the triplet (D, C, G) through combined losses and hyper-parameter $0 < \alpha <$

¹<https://github.com/gdeng96/ibgan>

1 (Eq. 2.4), where we weigh the augmented samples during training of the classifier through a weight that depends on the current discriminator $D(\mathbf{x}) = \{d_y(\mathbf{x})\}_{y \in \mathcal{Y}}$, that is, $w_{D,y}(\mathbf{x}) = \frac{d_y(\mathbf{x})}{1-d_y(\mathbf{x})}$. Through these weights, the effect of augmented sample kicks in smoothly during the training of the classifier.

$$(2.4) \quad \min_G \max_C \max_D \{L_{GAN}(D, G) + L_{class}^\alpha(C, D, G)\}$$

$$(2.5) \quad \begin{aligned} L_{GAN}(D, G) &= \sum_{y \in \mathcal{Y}} w_y^{-1} \mathbb{E} [1\{Y = y\} \log D(\mathbf{X})] \\ &+ \sum_{y \in \mathcal{Y}} (w'_y)^{-1} \mathbb{E} [1\{Y' = y\} \log (1 - D(G(X_{mask})))] \\ &= \sum_{y \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}^k} (p_y(\mathbf{x}) \log d_y(\mathbf{x}) + p'_y(\mathbf{x}) \log(1 - d_y(\mathbf{x}))) \end{aligned}$$

$$(2.6) \quad \begin{aligned} L_{class}^\alpha(C, D, G) &= \alpha \sum_{y \in \mathcal{Y}} \mathbb{E} [1\{Y = y\} \log c_y(\mathbf{X})] + (1 - \alpha) \\ &\times \sum_{y \in \mathcal{Y}} \mathbb{E} [w_{D,y} 1\{Y' = y\} \log c_y(G(X_{mask}))] \\ &= \sum_{y \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}^k} [\alpha w_y p_y(\mathbf{x}) + (1 - \alpha) w'_y p'_y(\mathbf{x}) w_{D,y}(\mathbf{x}) \log c_y(\mathbf{x})] \end{aligned}$$

During initial epochs, the generator will impute lower-quality synthetic samples and the discriminator easily identifies fake vs. real components in X' . Hence, the weights $w_{D,y}(\mathbf{x})$ are approximately zero and the classifier is mainly updated through real samples. When the discriminator is close to the optimal D_* in the original GAN setting where $D_*(\mathbf{x}) = \{d_y^*(\mathbf{x})\}_{y \in \mathcal{Y}}$, $d_y^*(\mathbf{x}) = \frac{p_y(\mathbf{x})}{p_y(\mathbf{x}) + p'_y(\mathbf{x})}$. Then, $p'_y(\mathbf{x}) w_{D,y}(\mathbf{x}) \approx p'_y(\mathbf{x}) \frac{d_y^*(\mathbf{x})}{1-d_y^*(\mathbf{x})} = p_y(\mathbf{x})$, and the augmented samples contribute to IB-GAN classifier equivalently to real samples.

3 Empirical Evaluation

3.1 UCR MTS Classification. We apply IB-GAN to 10 open-source multivariate time-series datasets from the popular UCR archive [40]. As two illustrative examples, CharacterTrajectories is a 20-class MTSC dataset with 3000 samples and Epilepsy is a 4-class dataset with 300 samples. Data imbalance is introduced by randomly dropping 75% of half of the classes.

Two classifier architectures are compared to showcase the flexibility of IB-GAN framework: a simple CNN with 2 Conv-1D layers and a state-of-the-art VGG-like network (3 VGG blocks with Conv2D and max pooling layers). Each classifier choice is paired with 3 GAN architectures: conditional GAN, vanilla GAN, and InfoGAN. Standard baselines are classifier only, class weights, upsampling, and downsampling. State-of-the-art GAN baselines are AC-GAN, BAGAN, SMOTE,

and Naive GAN (equivalent to generating data from white noise via Conditional GAN). Filter size for convolution layers is k , the time series dimension; $p_{miss} = 0.1$.

Findings. Each experiment is run for 5 replicates, and average classifier performance is listed in Table 1. The two metrics are Balanced Accuracy [41] and F1-score; both are macro-averaged across classes and penalizes for low performance in minority classes. The default IB-GAN classifier, with Conditional GAN for imputation as the generative model, outperforms against all baselines. IB-VanillaGAN and IB-InfoGAN demonstrate comparable high performance when utilizing the more powerful VGG classifier.

With both types of classifiers, SOTA GAN baselines such as AC-GAN and BAGAN did not show significant improvements compared to standard baselines, e.g., upsampling. The superior performance of the 3 IB-GAN variants compared to two-stage processes such as AC-GAN and BAGAN testifies that joint training for data augmentation and classification is a far more effective approach. In comparison, the folded classifier in AC-GAN and additional autoencoder-decoder in BAGAN did not generate better quality samples that translated to classification gains. Comparison with Naive GAN also indicates that samples generated via imputation based on randomly masked data vectors (IB-GAN) contributed to a greater performance boost than samples generated via white noise (Naive GAN). Even with a small injection of novelty with p_{miss} at 10%, the IB-GAN yields significantly higher Balanced Accuracy and F1-score. Finally, the imputation-balancing approach translates to IB-GAN classification with lower variance (see error bars), and stable prediction results are important for downstream applications.

The parametric SMOTE method also has relatively low performance, especially for smaller data size; SMOTE performs poorly for Epilepsy but has comparable performance to BAGAN and AC-GAN for CharacterTrajectories. SMOTE is constrained by forming new samples via linear combinations of existing samples, while IB-GAN functions consistently across sample sizes (Section 3.2.2). Details for additional UCR experiments and Kaggle EEG time series ($m > 1000$) with LSTM classifier are reported in Table 4 and 5 of Appendix. This testifies to the wide applicability of IB-GAN to time series of different lengths and flexibility of classifier choice.

3.2 Predicting Trending Products. We now apply IB-GAN to an empirical problem setting. The motivating application is predicting trending products on a large e-commerce website based on past time series metrics, which make up a minority class of all newly-launched products. This is known as the *cold-start*

Table 1: UCR Datasets - Average prediction metrics over 5 replicates for Imputation Balanced GAN and baselines with respect to standard multi-class CNN and VGG classifiers.

Experiment	CNN Classifier: CharacterTrajectories		VGG Classifier: CharacterTrajectories	
	Balanced Accuracy	F1-score	Balanced Accuracy	F1-score
IB-GAN	0.812 ± 0.032	0.810 ± 0.035	0.926 ± 0.009	0.905 ± 0.015
IB-InfoGAN	0.611 ± 0.037	0.535 ± 0.043	0.828 ± 0.039	0.776 ± 0.051
IB-VanillaGAN	0.651 ± 0.039	0.583 ± 0.043	0.862 ± 0.039	0.818 ± 0.049
Conditional (Naive) GAN	0.455 ± 0.069	0.416 ± 0.075	0.795 ± 0.068	0.786 ± 0.067
ACGAN + Second Stage Classifier	0.313 ± 0.075	0.276 ± 0.067	0.778 ± 0.040	0.772 ± 0.033
BAGAN + Second Stage Classifier	0.357 ± 0.066	0.326 ± 0.073	0.570 ± 0.087	0.545 ± 0.096
SMOTE + Second Stage Classifier	0.314 ± 0.095	0.257 ± 0.084	0.648 ± 0.083	0.595 ± 0.090
Baseline Class Weights	0.332 ± 0.024	0.284 ± 0.023	0.419 ± 0.110	0.344 ± 0.107
Baseline Upsample	0.457 ± 0.070	0.421 ± 0.080	0.683 ± 0.095	0.643 ± 0.103
Baseline Downsample	0.293 ± 0.062	0.246 ± 0.069	0.148 ± 0.049	0.092 ± 0.037
Baseline Classifier	0.357 ± 0.037	0.313 ± 0.038	0.393 ± 0.061	0.340 ± 0.067

Experiment	CNN Classifier: Epilepsy		VGG Classifier: Epilepsy	
	Balanced Accuracy	F1-score	Balanced Accuracy	F1-score
IB-GAN	0.568 ± 0.055	0.549 ± 0.045	0.767 ± 0.080	0.694 ± 0.111
IB-InfoGAN	0.428 ± 0.062	0.381 ± 0.072	0.746 ± 0.037	0.681 ± 0.043
IB-VanillaGAN	0.339 ± 0.098	0.308 ± 0.088	0.722 ± 0.046	0.652 ± 0.062
Conditional (Naive) GAN	0.343 ± 0.105	0.305 ± 0.100	0.487 ± 0.147	0.456 ± 0.154
ACGAN + Second Stage Classifier	0.318 ± 0.102	0.295 ± 0.095	0.441 ± 0.176	0.352 ± 0.190
BAGAN + Second Stage Classifier	0.224 ± 0.055	0.205 ± 0.052	0.280 ± 0.066	0.158 ± 0.103
SMOTE + Second Stage Classifier	0.248 ± 0.029	0.227 ± 0.030	0.275 ± 0.035	0.160 ± 0.059
Baseline Class Weights	0.307 ± 0.033	0.276 ± 0.038	0.284 ± 0.025	0.192 ± 0.043
Baseline Upsample	0.250 ± 0.043	0.235 ± 0.056	0.330 ± 0.083	0.222 ± 0.113
Baseline Downsample	0.266 ± 0.018	0.219 ± 0.053	0.260 ± 0.060	0.195 ± 0.065
Baseline Classifier	0.287 ± 0.039	0.237 ± 0.058	0.270 ± 0.034	0.156 ± 0.037




Figure 3: IB-GAN performance as p_{miss} increases, a form of regularization on the generator and level of novelty or innovations in synthetic samples.

problem in search, learning-to-rank, and recommendation systems. An additional challenge is the existence of time-invariant features based on product metadata, which will correspond to B in the IB-GAN set-up.

We can frame this problem as a large-scale imbalanced MTSC task with a 90K dataset for trending product metrics, split into 60K training and 30K testing. Each item has a binary label Y which indicates whether it is a top-ranked item, 15 time-invariant metadata features B and 8 daily time series X_t for 3 weeks. The full feature set is the joint vector $X = \{B, X_t\}$. B is different from class labels Y ; it is additional auxiliary information, e.g., product characteristics. Class imbalance is strong with only 11% samples in the 1-class. Standard accuracy is an inappropriate measure; an initial baseline CNN classifier with 89% testing accuracy has a precision of 0.29 and recall of 0.01 for the 1-class.

Given this large dataset, we conduct ablation experiments for IB-GAN training: (1) *How does choice of p_{miss} affect IB-GAN generator quality and classification accuracy?* (2) *How does IB-GAN perform across different training sample sizes?*

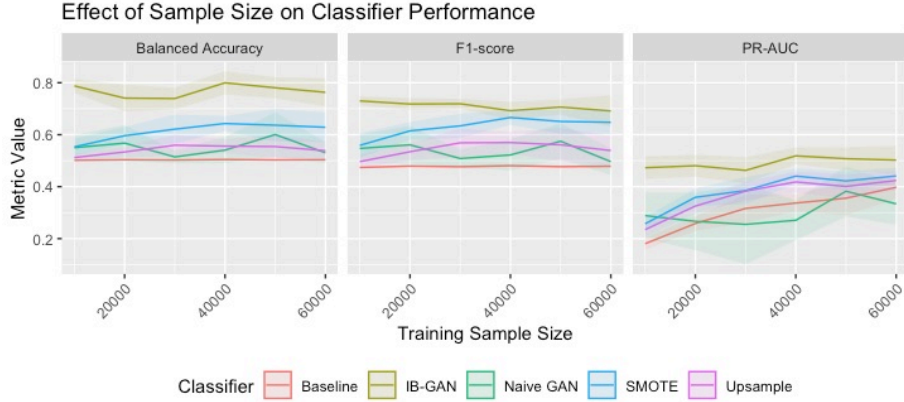


Figure 4: Average Balanced Accuracy, F1-score, and PR-AUC by training sample sizes for Imputation Balanced GAN and baselines evaluated on 30k test set for trending products data.

3.2.1 Novelty Parameter Tuning. Using the full product dataset, we compare IB-GAN classifier performance at different p_{miss} levels against benchmarks. We use a state-of-the-art VGG classifier with blocks of Conv1D, batch normalization, and max pooling layers; a similar architecture is adopted for the discriminator-generator duo, which conditions on class labels analogous to Conditional GAN. Each model is run for 100 epochs across 10 replicates. Figure 3 shows that 10% is a good default value, with an average Balanced Accuracy of 78.8%. The nearest benchmark is SMOTE with 64.1%. The three metrics are fairly consistent when p_{miss} is between 20-45%. As novelty in synthetic samples increase past 40-50%, classification performance declines on average. Higher p_{miss} dictates greater variability in generator output, which results in synthetic samples that are less similar to original data. This leads to greater variability in prediction accuracy; see error bands for F1-score and PR-AUC. See Table 2 in Appendix for results of standard and GAN baselines.

3.2.2 Training Data Size. To understand IB-GAN performance across sample sizes, we fix $p_{miss} = 20\%$ and randomly sample subsets of the 60K training data. The same VGG classifier and IB-GAN set-up is used. Figure 4 shows average Balanced Accuracy, F1-score, and PR-AUC over 10 replicates as evaluated on the fixed 30K test set for IB-GAN and select baselines.

The IB-GAN classifier outperforms all benchmarks with consistent Balanced Accuracy and F1-score across sample sizes up to 60000, the maximum number of training samples. As sample size increases, upsampling and SMOTE become more effective due to increased examples from the minority class. When sample sizes are small, upsampling leads to overfitting and SMOTE is

unable to interpolate with variety to create meaningful new training samples. Meanwhile, GAN-based methods such as IB-GAN and even Naive GAN largely avoids this issue, and are recommended for applications with small data sizes. Error bars for IB-GAN metrics tend to be narrower than Naive GAN; the imputation mechanism and tuning parameter p_{miss} both regulate IB-GAN generator variability and ensure synthetic samples generated actually enables classifier to learn signals from minority classes. See Table 3 for details.

4 Significance and Impact

We have proposed a novel IB-GAN method for joint data augmentation and multivariate time series classification with highly imbalanced data in an unified one-step process. The framework consists of a triplet of generator G , discriminator D , and classifier C models that works seamlessly with choices of GAN architectures and deep learning classifiers. Compared to prior methods, IB-GAN does not require a cumbersome second stage classifier, and directly incorporates classification loss as feedback to improve synthetic data quality. IB-GAN uses a unique imputation and balancing approach that leverages existing training data to generate higher quality synthetic samples, and enables better classification performance on under-observed classes by learning on a class-balanced training set. The hyperparameter p_{miss} directly tunes the similarity vs. novelty in synthetic samples while side-stepping mode collapse from standard GAN training. Empirical results from UCR datasets show IB-GAN and its variants achieving significant performance against state-of-the-art parametric and GAN baselines. Ablation studies with a trending product dataset demonstrate IB-GAN performance is robust across sample sizes (narrower confidence inter-

vals) and p_{miss} levels up to 50% under MCAR.

The IB-GAN framework is quite versatile and can easily be extended to more complex sequence datasets such as images, sensor, audio, and video datasets; future work can identify the proper design of generator-discriminator and classifier architecture suitable for these tasks. IB-GAN can also be applied to mitigate data bias and fairness issues in many ML applications; see Section 5 in Appendix for a short note on fairness.

References

- [1] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery*, 33(4):917–963, 2019.
- [2] Fazle Karim, Somshubra Majumdar, Houshang Darabi, and Shun Chen. LSTM fully convolutional networks for time series classification. volume 6, pages 1662–1669. IEEE, 2017.
- [3] Bendong Zhao, Huanzhang Lu, Shangfeng Chen, Junliang Liu, and Dongya Wu. Convolutional neural networks for time series classification. volume 28, pages 162–169. BIAI, 2017.
- [4] Fazle Karim, Somshubra Majumdar, Houshang Darabi, and Samuel Harford. Multivariate LSTM-FCNs for time series classification. *Neural Networks*, 116:237–245, 2019.
- [5] Guozhong Li, Byron Choi, Jianliang Xu, Sourav S Bhowmick, Kwok-Pan Chun, and Grace LH Wong. Shapenet: A shapelet-neural network approach for multivariate time series classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 8375–8383, 2021.
- [6] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018.
- [7] G. Weiss, Kate McCarthy, and Bibi Zabbar. Cost-sensitive learning vs. sampling: Which is best for handling unbalanced classes with unequal error costs? In *DMIN*, 2007.
- [8] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- [9] Damien Dablain, Bartosz Krawczyk, and Nitesh V. Chawla. Deepsmote: Fusing deep learning and smote for imbalanced data, 2021.
- [10] Lara Lusa et al. Evaluation of SMOTE for high-dimensional class-imbalanced microarray data. In *2012 11th International Conference on Machine Learning and Applications*. IEEE, 2012.
- [11] Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):1–48, 2019.
- [12] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2020.
- [13] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. Data augmentation using synthetic data for time series classification with deep residual networks. *arXiv*, 2018.
- [14] Giovanni Mariani, Florian Scheidegger, Roxana Istrate, Costas Bekas, and Cristiano Malossi. BAGAN: Data augmentation with Balancing GAN. *arXiv*, 2018.
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [16] Yi Zheng, Qi Liu, Enhong Chen, Yong Ge, and J Leon Zhao. Exploiting multi-channels deep convolutional neural networks for multivariate time series classification. *Frontiers of Computer Science*, 10(1):96–112, 2016.
- [17] Xuchao Zhang, Yifeng Gao, Jessica Lin, and Chang-Tien Lu. TapNet: Multivariate time series classification with attentional prototypical network. In *AAAI Conference on Artificial Intelligence*, 2020.
- [18] Toan Tran, Trung Pham, Gustavo Carneiro, Lyle Palmer, and Ian Reid. A bayesian data augmentation approach for learning deep models. *arXiv*, 2017.
- [19] Manan Oza, Himanshu Vaghela, and Kriti Srivastava. Progressive generative adversarial binary networks for music generation. *arXiv*, 2019.
- [20] Lu Sheng, Junting Pan, Jiaming Guo, Jing Shao, Xiaogang Wang, and Chen Change Loy. Unsupervised bi-directional flow-based video generation from one snapshot. *arXiv*, 2019.
- [21] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv*, 2015.
- [22] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv*, 2014.
- [23] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.

- [24] Jinsung Yoon, James Jordon, and Mihaela Schaar. GAIN: Missing data imputation using generative adversarial nets. In *International Conference on Machine Learning*, pages 5689–5698. PMLR, 2018.
- [25] Dongwook Lee, Junyoung Kim, Won-Jin Moon, and Jong Chul Ye. CollaGAN: Collaborative gan for missing image data imputation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2487–2496, 2019.
- [26] Steven Cheng-Xian Li, Bo Jiang, and Benjamin Marlin. MisGAN: Learning from incomplete data with generative adversarial networks. *arXiv*, 2019.
- [27] Yonghong Luo, Xiangrui Cai, Ying Zhang, Jun Xu, et al. Multivariate time series imputation with generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 1596–1607, 2018.
- [28] Jinsung Yoon, Daniel Jarrett, and Mihaela van der Schaar. Time-series generative adversarial networks. 2019.
- [29] Ying Zhang, Baohang Zhou, Xiangrui Cai, Wenya Guo, Xiaoke Ding, and Xiaojie Yuan. Missing value imputation in multivariate time series with end-to-end generative adversarial networks. *Information Sciences*, 551:67–82, 2021.
- [30] Zijian Guo, Yiming Wan, and Hao Ye. A data imputation method for multivariate time series based on generative adversarial network. *Neurocomputing*, 360:185–197, 2019.
- [31] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with Auxiliary Classifier GANs. In *International conference on machine learning*, pages 2642–2651. PMLR, 2017.
- [32] Mingming Gong, Yanwu Xu, Chunyuan Li, Kun Zhang, and Kayhan Batmanghelich. Twin auxiliary classifiers GAN. *Advances in neural information processing systems*, 32:1328, 2019.
- [33] Siddarth Asokan and Chandra Seelamantula. Teaching a GAN what not to learn. In *Advances in Neural Information Processing Systems*, volume 33, pages 3964–3975, 2020.
- [34] Chongxuan Li, Kun Xu, Jun Zhu, and Bo Zhang. Triple generative adversarial nets. *arXiv*, 2017.
- [35] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. *arXiv*, 2016.
- [36] Fabrizia Mealli and Donald B Rubin. Clarifying Missing at Random and related definitions, and implications when coupled with exchangeability. *Biometrika*, 102(4):995–1000, 2015.
- [37] Stephan Zheng, Yang Song, Thomas Leung, and Ian Goodfellow. Improving the robustness of deep neural networks via stability training. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4480–4488, 2016.
- [38] Omid Poursaeed, Isay Katsman, Bicheng Gao, and Serge Belongie. Generative adversarial perturbations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4422–4431, 2018.
- [39] Chunlin Gong and Liangxian Gu. A novel smote-based classification approach to online data imbalance problem. *Mathematical Problems in Engineering*, 2016, 2016.
- [40] Hoang Anh Dau, Anthony Bagnall, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh Gharghabi, Chotirat Ann Ratanamahatana, and Eamonn Keogh. The UCR time series archive. *IEEE/CAA Journal of Automatica Sinica*, 6(6):1293–1305, 2019.
- [41] John D Kelleher, Brian Mac Namee, and Aoife D’arcy. *Fundamentals of machine learning for predictive data analytics: algorithms, worked examples, and case studies*. MIT press, 2020.
- [42] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *3rd Innovations in Theoretical Computer Science Conference*, 2012.
- [43] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017.
- [44] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *arXiv*, 2019.
- [45] Julia Dressel and Hany Farid. The accuracy, fairness, and limits of predicting recidivism. *Science advances*, 4(1):eaao5580, 2018.
- [46] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kiciman. Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in Big Data*, 2:13, 2019.
- [47] Vasileios Iosifidis and Eirini Ntoutsi. AdaFair: Cumulative fairness adaptive boosting. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2019.
- [48] Stephen Hutt, Margo Gardner, Angela L Duckworth, and Sidney K D’Mello. Evaluating fairness and generalizability in models predicting on-time graduation from college applications. *International Educational Data Mining Society*, 2019.