Synthesis of Realistic ECG Waveforms Using a Composite Generative Adversarial Network for Classification of Atrial Fibrillation

Rohan Banerjee TCS Research Tata Consultancy Services rohan.banerjee@tcs.com Avik Ghose TCS Research Tata Consultancy Services avik.ghose@tcs.com

Abstract-In recent days, computer-aided diagnosis systems powered by artificial intelligence and machine learning have become an important part of medicine for assisting the doctors in critical decision making. They are popularly deployed in cardiology for early and automatic detection of various life-threatening diseases. However, a machine learning algorithm requires a large volume of training data to create the learning model which is an empirical problem in medical domain. Generating synthetic patient data has emerged as an important area of research to solve the issue. In this paper, we propose a novel Generative Adversarial Network (GAN) architecture using medical domain knowledge to create realistic Electrocardiogram (ECG) waveforms containing the signature of Atrial Fibrillation (AF), a common type of arrhythmia. Our composite architecture consists of a pair of GANs to simulate the disease-specific Heart Rate Variability (HRV) pattern and the unique signal morphology in the generated waveforms. The proposed architecture is applied on two public datasets for synthesis of AF-specific ECG to mitigate the class imbalance issue. Results show that the performance of existing AF classifiers significantly improves on both datasets by adding the synthetic data to the training set.

Index Terms—ECG, Deep learning, Data augmentation, Generative Adversarial Network

I. INTRODUCTION

Electrocardiogram (ECG) is a clinical test to record the electrophysiological activities of the heart. The recorded data is interpreted by the doctors for non-invasive diagnosis of cardiovascular diseases. However, it is practically impossible to manually analyse the large volume of ECG data generated each day in a hospital. An automatic diagnosis from the digitally recorded ECG is possible, thanks to the recent advancement in artificial intelligence, machine learning, and deep learning techniques. A supervised learning-based cardiac disease classification algorithm requires a large volume of annotated data to create the training model. Recording of large scale patient data is time consuming and often challenging due to privacy issues and associated risks in case of infectious diseases. The problem is addressed in data science by generating synthetic but realistic patient data.

Although normal ECG samples are relatively easy to record and are substantially available in various open-access databases, the quantity of abnormal recordings correspond-

ing to different heart diseases is often inadequate to train a machine learning or deep learning model. McSharry et al. [1] proposed a mathematical model to generate realistic ECG signals corresponding to normal people through a set of ordinary differential equations. Clifford et al. [2] proposed an approach to simulate abnormal ECG patterns using a sum of Gaussian kernels fitted to vectorcardiogram recordings and a hidden Markov model. The physiological model in [3] can simulate ECG based on a discretized reaction-diffusion system to mimic the main pacemakers in the heart. Apart from the physics-based models, pure statistical and deep learning approaches can also be found in literature. The deep learning algorithms aim to simulate newer artificial data by learning the distribution from a real-world training dataset. The Generative Adversarial Network (GAN) is a popular example of such generative modeling which is extensively used for creating realistic images and time-series data. Zhu et al. [4] proposed a GAN architecture using the MIT-BIH arrhythmia database. However, utility of the generated ECG data was not qualitatively evaluated on the existing AF classifiers. The GAN proposed by Hatamian et al. [5] can generate the spectrograms of ECG corresponding to Atrial Fibrillation (AF) patients, but not the time-series waveforms. Although the existing GAN models can successfully generate realistic ECG for normal people, their performance is relatively poor to generate disease-specific data. An ECG signal is complex in nature. Moreover, a disease-specific recording can have various anomalous components depending upon the underlying pathological condition. Hence, it is difficult to simulate using pure deep learning approaches.

In this paper, we propose a novel hybrid GAN architecture to generate relaistic disease-specific ECG waveforms, taking AF as an example. The architetcure is derived based on medical domain knowledge to incorpoarte the clinical markers of AF in generated data using multiple GANs. As shown in Fig.1, a complete cycle of a normal ECG waveform has three major components, the P wave, the QRS complex which contains the R peak, and the T wave. AF is a cardiac disorder where abnormal electrical impulses start firing in the atria, causing a faster heart rate [6]. It has two known clinical markers



Fig. 1. ECG pattern of a normal subject and an AF patient.

on ECG as shown. The P waves are either absent or are replaced by fibrillatory waves. Secondly, the irregular Heart Rate Variability (HRV) causes large variation in successive R-R interval distances. Our proposed architecture comprises two GANs to simulate these two AF markers:

- Long Short-Term Memory GAN (LSTM-GAN) to generate the HRV pattern of AF.
- Deep Convolutional GAN (DCGAN) to create the morphology of ECG cycle.

The complete ECG waveforms are created by combining the output of the two GANs. The generated recordings are used to balance the AF and non-AF class ratio in the PhysioNet Challenge 2017 and the PhysioNet Challenge 2020 database and the utility is validated on two benchmark AF detection algorithms. In Section II, we provide a detailed description of our proposed architecture followed by experimental results in Section III and a conclusion in Section IV.

II. PROPOSED COMPOSITE GAN ARCHITECTURE

The GAN is a popular deep generative model containing two basic modules, generator and discriminator. The generator (G) takes an N-dimensional latent vector (z) as input that follows a Gaussian distribution and maps it to the generated data as its output, G(z). The discriminator (D) outputs D(G(z)), the probability to predict whether the generated data is real or fake based on a training set of real data, x. The generator and the discriminator reach a convergence state via a zero-sum game. The objective function of a GAN is expressed in terms of the min-max optimization process, as in eq. (1) [7].

$$\min_{G} \max_{D} V(D,G) = E_{x \sim P(x)}[logD(x)] + E_{z \sim P(z)}[log(1 - D(G(z)))]$$
(1)

The discriminator tries to maximize the probability to correctly classify real and fake data, and the generator tries to minimize the probability that the discriminator will predict its output as fake. Under an optimum state, distribution of the fake data becomes equivalent to the real data, and the discriminator classifies them at a probability of 0.5.

In this paper, we propose a novel GAN architecture using medical domain knowledge to generate realistic ECG waveforms similar to real AF patients. Block diagram of the proposed architecture is shown in Fig. 2. It also indicates



Fig. 2. Proposed architecture for generating AF like ECG.

the tensor dimension at the output of different layers. The architecture comprises an LSTM-GAN and a DCGAN which are effectively combined in a single composite structure. The LSTM-GAN is designed to generate the R-R interval distances time-series that follows the irregular HRV pattern of AF. Whereas, the DCGAN is responsible for generating the signal morphology between two adjacent R peaks in time domain. A set of 30-dimensional latent vectors, randomly sampled from a standard normal distribution is fed to both generators as input in batches to train the composite network. The proposed approach is separately evaluated on the PhysioNet Challenge 2017 and the PhysioNet challenge 2020 database in this study. A set of annotated AF recordings is selected from the two databases to form the real data for the discriminators.

A. LSTM-GAN for generating R-R interval distances

Irregular HRV is a known clinical marker for AF. An LSTM is a popular deep learning architecture that has its internal memory for sequential modeling of time-series data in terms of a hidden vector [8]. It can effectively learn the desired pattern from a very long sequence due to the unique cell structure, that enables to delete less important information from memory. The R-R interval distances extracted from an ECG data can be represented as a vector (rr_t) of k real numbers. Here, $rr_{t} = [rr_{1}, rr_{2}, ..., rr_{k}]$, where $rr_{i} = r_{i+1} - r_{i}$, and r_{i} is the location of i^{th} R peak in the ECG data on time axis. The R-R interval distances are computed from the two PhysioNet databases for creating the real data for the discriminator. Since the ECG recordings do not have a fixed length in the original databases, the number of points in the extracted R-R intervals varies accordingly. However, all instances of real data applied to the discriminator of a GAN are required to have the same dimension. Considering the median duration, length of an R-R intervals vector is selected as 50 in our architecture. The shorter recordings are repeated and merged to achieve the desired length, whereas the longer recordings are broken into multiple partially overlapping segments to increase the instances of real data in the training set.

The generator of the LSTM-GAN in our architecture contains two dense layers having 40 and 50 units respectively for mapping a 30-dimensional input latent vector to the desired dataspace of R-R intervals. Leaky Rectified Linear Unit (Leaky Relu) with negative slope coefficient = 0.2 and hyperbolic tangent (tanh) functions are used for non-linear activation of the units in the two dense layers. The discriminator module contains a pair of LSTM layers both having 64 units followed by a single-unit dense layer with a sigmoid activation function for classifying the generated sequences as real or fake.

B. DCGAN for generating signal morphology between two adjacent R peaks

Normal atrial activities of the heart become awry due to AF. This is reflected in the ECG morphology in terms of missing P waves or presence of abnormal fibrillatory waves before the QRS complex. We define a DCGAN structure that can generate an ECG cycle from an input latent vector incorporating such unique morphology. As shown in eq. (2), an ECG time-series, ecg_t can be represented as a vector of m real numbers.

$$ecg_t = [ecg_1, ecg_2, ...ecg_{r_1}, ...ecg_{r_2}, ...ecg_{r_i}, ...ecg_m]$$
 (2)

It contains a set of landmark points in terms of the R peaks whose locations are known. Here, r_i indicates the time location of i^{th} R peak, and ecg_{r_i} is the corresponding amplitude. An ECG cycle is defined in this paper as the segment between two adjacent R peaks. The p^{th} cycle is extracted as:

$$cycle_p = [ecg_{\{r_p\}}, ecg_{\{r_p+1\}}, ecg_{\{r_p+2\}}, ...ecg_{\{r_{p+1}-1\}}].$$
 (3)

In a DCGAN [9], the lower dimensional latent vector is converted to the desired space of realistic generated data using a series of convolution and transposed-convolution operations in the discriminator and the generator respectively through a set of filters (kernel). ECG signals in the two training databases are first set to a fixed sampling rate of 300 Hz. Subsequently, the ECG cycles are extracted with respect to the reference R peaks. The length of the extracted cycles is required to be fixed for applying them to the discriminator module. Mean heart rate of an AF patient is known to be higher than the normal range. The cycle length is selected as 200 (\approx 667 ms long, instantaneous heart rate = 90 bpm) to be generated in our architecture. Durations of the extracted real cycles are modified accordingly using cubic spline interpolation technique. The discriminator module of the proposed DCGAN architecture contains two 1D convolutional layers with associated batch-normalization and Leaky Relu activation layers. The convolutional layers contain 64 and 128 filters respectively with kernel dimension of 4. The output is downsampled at each convolutional layer by setting a stride length of 2. The resultant feature-map is flattened and applied to a single-unit dense layer to classify real and fake data through a sigmoid activation function. In order to mitigate the chance of over-fitting, 30% dropout is applied to the convolutional layers. The generator is comprised of a dense layer and a pair of strided transposed-convolutional layers (also known as deconvolutional layers) of having 128 and 64 filters (stride length = 2, kernel dimension = 4) with associated Leaky Relu layers to map the input latent vector to a higher dimensional space. There is a final convolutional layer, having a single filter of kernel dimension = 7 with a 'tanh' activation function to convert the feature-map to the desired shape of an ECG cycle.

C. Training of the proposed network

Separate mini-batches of real and fake data are used for training. The real and the fake samples are annotated as 1 and 0. The discriminator of a GAN aims to maximize the probability of correctly classifying an input as real or fake. The loss is expressed as $D_{loss} = log(D(x)) + log(1 - D(G(z)))$. These two terms are separately calculated on the mini-batches for real and generated fake data, providing a forward pass through the discriminator, and the gradients are calculated through a backward pass. For the generator, the loss term is $G_{loss} = log(D(G(z)))$. It tries to maximize log(D(G(z))), which is achieved by minimizing the term log(1 - D(G(z))). The loss is calculated based on the classification output of the generated data as predicted by the discriminator.

In order to ensure the ECG cycles generated by the DCGAN are close to real ECG morphology, the Mean Squared Error (MSE) between the real and the generated data is added to the generator loss function of the DCGAN as a penalty term to be minimized. Hence, the total generator loss of the DCGAN is:

$$G_{loss_{tot}}^{DCGAN} = G_{loss}^{DCGAN} + \lambda * \frac{1}{n} \sum (Y_i - \hat{Y}_i)^2$$
(4)

Y and \hat{Y} indicate real and generated data, n is the batch size, the constant λ controls the weight of the penalty term. A small value is set as $\lambda = 0.05$ so that unrestricted newer samples can be generated by the DCGAN keeping the overall morphology similar to real ECG cycles.

We use Adam optimizer with a learning rate of 0.0002 for the LSTM-GAN and the DCGAN. The mini-batch size is set as 64. The model weights are initialized from a normal distribution with zero mean and standard deviation of 0.02. Label smoothing is applied to modify the hard labels for real data slightly more or less than 1 and slightly more than 0 for fake data, where the variation for each label is done randomly. Additionally, we introduce some noise in the labels by randomly flipping the labels of a small fraction of real and fake data in each mini-batch. Both techniques have a regularization effect to avoid over-fit. The composite network is trained end to end up to 500 epochs applying the same set of 30-dimensional latent vectors to the generator modules of the LSTM-GAN and the DCGAN in every mini-batch.

D. Generating the complete ECG waveforms

Once the training is done, the generator of the LSTM-GAN and the DCGAN can generate a vector of R-R interval distances of length 50 and an ECG cycle of 200 points from a 30-dimensional input latent vector. The complete ECG waveform is created by modifying the length of the generated ECG cycle according to the R-R interval distances using cubic spline interpolation and merging them on time axis. The generated signal is finally applied to a 4^{th} order Butterworth bandpass filter, having cut-off frequencies of 0.5 Hz and 20 Hz to remove the undesired noise components.

III. EXPERIMENTAL RESULTS

The proposed architecture was implemented in python 3.6 using TensorFlow 1.15 library. The training was performed on a computer with having Intel i7-7820X processor, 16 GB primary memory, and a GeForce GTX 1080 Ti graphics processing unit. Few samples of generated ECG waveforms are shown in Fig. 3. Irregular HRV pattern, the primary marker for AF is clearly visible in all of them. Absence of P wave before the QRS complex can be found in the first two samples, whereas the remaining samples show traces of fibrillatory waves.

The efficacy of the proposed approach in data augmentation



Fig. 3. Sample ECG waveforms generated by the proposed composite GAN architecture.

to mitigate the class imbalance problem of a dataset in training a supervised learning-based AF classifier is also evaluated on the Physionet Challenge 2017 database [10] and the PhysioNet Challenge 2020 database [11]. Both databases contain ECG signals recorded from normal people and patients having AF and other cardiac diseases. Our objective is to study the performance improvement of the existing AF classifiers on a highly imbalnaced ECG dataset after applying various data augmentation approaches. Since the proposed architecture is designed to simulate AF-specific data, we focus only on classifying AF and non-AF recordings. Further classification of the non-AF recordings is beyond the scope of this study. For an effective analysis, we keep both normal as well as non-AF but other types of cardiac diseases in the non-AF class. The PhysioNet Challenge 2017 database [10] contains annotated single-lead ECG signals of 5154 normal, 771 AF, and 2557 other types of abnormal recordings. We convert the original database to a highly imbalanced database for binary classification by merging all types of non-AF recordings under a single class. The re-annotated database contains 7711 non-AF and 771 AF recordings. The recordings are sampled at

300 Hz. On the other hand, the PhysioNet Challenge 2020 database [11] is a very large multi-source corpus of 12-lead ECG data recorded from hospitals in different geographical locations. The signals in the original database are sampled at 500 Hz which are resampled at 300 Hz. A portion is randomly selected from the large original database for our study. The selected portion contains 400 AF and 4000 non-AF recordings. The non-AF class includes normal sinus rhythms and non-AF abnormal conditions like 1st degree AV block, left anterior fascicular block, complete right bundle branch block, bradycardia etc. The lead II data is considered in this study. In order to evaluate the impact of the porposed data augmentation technique on classification performance, we select two popular state-of-the-art open-source supervised learningbased AF classifiers to test on both the databases cosidered in this paper. The first algorithm is a classical machine learning approach by Datta et. al. [12], that trains a series of cascaded binary AdaBoost classifiers using more than 150 hand-crafted features, related to ECG morphology and shortterm HRV. The second algorithm by Zihlmann et. al. [13] proposes two separate deep learning AF-classifiers based on CNN and CRNN that take the 2-D spectrogram of ECG as input. The original algorithms are required to modify to binary classifier for the purpose of this study. From each database, 80% of data is randomly selected for training by maintaining the original class ratio, and the remaining portion is kept for testing. The AF portion in the training set of both databases is used for data augmentation to balance the class ratio of AF to non-AF data. Subsequently, the AF classifiers are trained on the balanced training set and evaluated on the test set. Classification performance is reported in this paper in terms of sensitivity (Se) and specificity (Sp) of detecting AF.

$$Se = \frac{TP}{TP + FN}, Sp = \frac{TN}{TN + FP}$$
 (5)

Here, TP, TN, FP and FN indicate the true positive, true negative, false positive, and false negative. Table I shows the performance improvement achieved by the machine learningbased AF classifier [12] on both databases when the training is done incorporating the proposed GAN model for data augmentation. Here, we show a comparison of the proposed approach with a number of existing data augmentaion approaches as well as the effect of no data augmentation. A classifier trained on a highly imbalanced database is expected to be biased towards the majority class. A significant number of AF recordings are misidentified on both databases when the classifier is trained on the original data where AF is the minority class, resulting in a high specificity and a low sensitivity. Synthetic Minority Oversampling Technique (SMOTE) [14] and Adaptive Synthetic (ADASYN) [15] are popularly used in machine learning for data augmentation. Instead of generating ECG waveforms, these methods can only simulate new datapoints in the feature-space based on local information from the hand-crafted features computed from the real ECG. The synthetically generated features are merged with the features extracted from the real ECG data to balance the class ratio in the feature matrix and applied to the classifier. Although the sensitivity of detecting AF improves due to them, there is a negative impact on specificity. On the other hand, the proposed GAN model generates newer ECG waveforms via learning the distribution of the original recordings. Hence, the hand-crafted features computed from the generated waveforms are found more effective than SMOTE and ADASYN, which significantly improves the classifier sensitivity without affecting the specificity. Table II shows that the proposed

 TABLE I

 QUANTITATIVE ANALYSIS OF THE MACHINE LEARNING-BASED AF

 CLASSIFIER IN [12], APPLYING VARIOUS DATA AUGMENTATION

 TECHNIQUES TO IMPROVE THE CLASS IMBALANCE OF THE TRAINING SET

 IN THE PHYSIONET CHALLENGE DATABASES

augmentation (not used)		augmentation (SMOTE)		augmentation (ADASYN)		augmentation (proposed GAN)					
The PhysioNet Challenge 2017 database											
Se	Sp	Se	Sp	Se	Sp	Se	Sp				
0.81	0.96	0.84	0.93	0.86	0.93	0.93	0.96				
The PhysioNet Challenge 2020 database											
Se	Sp	Se	Sp	Se	Sp	Se	Sp				
0.78	0.95	0.82	0.91	0.84	0.89	0.91	0.95				

data augmentation approach has a similar impact on the CNN and CRNN-based AF classifiers in [13]. Here, the proposed GAN architecture is compared with two other approaches. The first approach is a popular trick used in deep learning on unbalanced dataset, where the AF classifier is trained on the original imbalanced dataset by assigning 10 times higher class weight to the minority class, which pays more attention to AF. In spite of a significant improvement in sensitivity, it shows a negative impact on specificity. Secondly, we compare with the GAN architecture proposed in [5] that generates spectrogram of ECG without reconstructing the time-series. Unlike [5], our proposed method is designed based on clinical biomarkers of AF. It shows to be more capable of generating realistic ECG time-series to improve the diversity of the training set. Thus, it has the optimum impact on the AF classifier performance.

TABLE II

Quantitative analysis of the deep learning-based AF classifiers in [13], applying various data augmentation techniques to improve the class imbalance of the training set

	augmentation		augmentation		augmentation		augmentation					
	(not		(class		(GAN [5])		(proposed					
	used)		weight)				GAN)					
The PhysioNet Challenge 2017 database												
	Se	Sp	Se	Sp	Se	Sp	Se	Sp				
CNN	0.79	0.99	0.89	0.91	0.83	0.98	0.92	0.99				
CRNN	0.81	0.98	0.92	0.93	0.89	0.96	0.95	0.98				
The PhysioNet Challenge 2020 database												
	Se	Sp	Se	Sp	Se	Sp	Se	Sp				
CNN	0.69	0.90	0.72	0.86	0.81	0.94	0.90	0.97				
CRNN	0.75	0.92	0.89	0.90	0.90	0.93	0.92	0.97				

IV. CONCLUSION

Recording of large scale patient data can often be challenging due to various reasons. Synthesis of realistic data has become an important area of research in biomedicine to optimize the performance of a supervised learning classifier on a highly imbalanced dataset. In this paper, we propose a novel GAN architecture to generate disease-specific realistic ECG waveforms based on clinical biomarkers, considering AF as a use case. Results show that the proposed data augmentation approach can significantly improve the performance of existing AF detection algorithms on two highly imbalanced public databases. Although the proposed architecture is specific to AF, the concept can be extended to a generic framework to simulate realistic ECG waveforms for other types of cardiovascular diseases as well based on their clinical markers.

References

- P. E. McSharry, G. D. Clifford, L. Tarassenko, and L. A. Smith, "A dynamical model for generating synthetic electrocardiogram signals," *IEEE transactions on biomedical engineering*, vol. 50, no. 3, pp. 289– 294, 2003.
- [2] G. Clifford, S. Nemati, and R. Sameni, "An artificial multi-channel model for generating abnormal electrocardiographic rhythms," in 2008 *Computers in Cardiology*. IEEE, 2008, pp. 773–776.
- [3] M. Quiroz-Juárez, O. Jiménez-Ramírez, R. Vázquez-Medina, V. Breña-Medina, J. Aragón, and R. Barrio, "Generation of ecg signals from a reaction-diffusion model spatially discretized," *Scientific Reports*, vol. 9, no. 1, pp. 1–10, 2019.
- [4] F. Zhu, F. Ye, Y. Fu, Q. Liu, and B. Shen, "Electrocardiogram generation with a bidirectional lstm-cnn generative adversarial network," *Scientific reports*, vol. 9, no. 1, pp. 1–11, 2019.
- [5] F. N. Hatamian, N. Ravikumar, S. Vesal, F. P. Kemeth, M. Struck, and A. Maier, "The effect of data augmentation on classification of atrial fibrillation in short single-lead ecg signals using deep neural networks," in *ICASSP 2020-2020 IEEE International Conference on Acoustics*, Speech and Signal Processing (ICASSP). IEEE, 2020, pp. 1264–1268.
- [6] D. with the special contribution of the European Heart Rhythm Association (EHRA), E. by the European Association for Cardio-Thoracic Surgery (EACTS), A. F. Members, A. J. Camm, P. Kirchhof, G. Y. Lip, U. Schotten, I. Savelieva, S. Ernst, I. C. Van Gelder *et al.*, "Guidelines for the management of atrial fibrillation: the task force for the management of atrial fibrillation of the european society of cardiology (esc)," *European heart journal*, vol. 31, no. 19, pp. 2369–2429, 2010.
- [7] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672– 2680.
- [8] H. Sak, A. W. Senior, and F. Beaufays, "Long short-term memory recurrent neural network architectures for large scale acoustic modeling," 2014.
- [9] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," arXiv preprint arXiv:1511.06434, 2015.
- [10] G. D. Clifford, C. Liu, B. Moody, H. L. Li-wei, I. Silva, Q. Li, A. Johnson, and R. G. Mark, "Af classification from a short single lead ecg recording: The physionet computing in cardiology challenge 2017," *Proceedings of Computing in Cardiology*, vol. 44, p. 1, 2017.
- [11] E. A. P. Alday, A. Gu, A. J. Shah, C. Robichaux, A.-K. I. Wong, C. Liu, F. Liu, A. B. Rad, A. Elola, S. Seyedi *et al.*, "Classification of 12-lead ecgs: the physionet/computing in cardiology challenge 2020," *Physiological measurement*, vol. 41, no. 12, p. 124003, 2020.
- [12] S. Datta, C. Puri, A. Mukherjee, R. Banerjee, A. D. Choudhury, R. Singh, A. Ukil, S. Bandyopadhyay, A. Pal, and S. Khandelwal, "Identifying normal, af and other abnormal ecg rhythms using a cascaded binary classifier," *Computing*, vol. 44, p. 1, 2017.
- [13] M. Zihlmann, D. Perekrestenko, and M. Tschannen, "Convolutional recurrent neural networks for electrocardiogram classification," *Computing*, vol. 44, p. 1, 2017.
- [14] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "Smote: synthetic minority over-sampling technique," *Journal of artificial intelligence research*, vol. 16, pp. 321–357, 2002.
- [15] H. He, Y. Bai, E. A. Garcia, and S. Li, "Adasyn: Adaptive synthetic sampling approach for imbalanced learning," in 2008 IEEE international joint conference on neural networks (IEEE world congress on computational intelligence). IEEE, 2008, pp. 1322–1328.