

An Uncertainty Estimation Framework for Risk Assessment in Deep Learning-based Atrial Fibrillation Classification

James Belen, Sajad Mousavi, Alireza Shamsoshoara, Fatemeh Afghah
School of Informatics, Computing and Cyber Systems,
Northern Arizona University, Flagstaff, AZ 86011

Abstract—Atrial Fibrillation (AF) is among one of the most common types of heart arrhythmia afflicting more than 3 million people in the U.S. alone. AF is estimated to be the cause of death of 1 in 4 individuals. Recent advancements in Artificial Intelligence (AI) algorithms have led to the capability of reliably detecting AF from ECG signals. While these algorithms can accurately detect AF with high precision, the discrete and deterministic classifications mean that these networks are likely to erroneously classify the given ECG signal. This paper proposes a variational autoencoder classifier network that provides an uncertainty estimation of the network’s output in addition to reliable classification accuracy. This framework can increase physicians’ trust in using AI-based AF detection algorithms by providing them with a confidence score which reflects how uncertain the algorithm is about a case and recommending them to put more attention to the cases with a lower confidence score. The uncertainty is estimated by conducting multiple passes of the input through the network to build a distribution; the mean of the standard deviations is reported as the network’s uncertainty. Our proposed network obtains 97.64% accuracy in addition to reporting the uncertainty.

I. INTRODUCTION

Atrial Fibrillation (AF) is one of the most common types of heart arrhythmia which can put patients at an increased risk of ischemic strokes, pulmonary strokes, and cranial strokes [1]. In 2016, close to 25,000 people had AF as the underlying cause of death with the risk of AF estimated to be about 1 in 4 individuals in the US [2]. Monitoring the long ECG recordings to detect the AF cases is a time consuming task for the physicians and is subject to human errors; thus, precise and accurate automated AF detection methods can aid physicians to determine the best course of treatment for these patients.

Artificial Neural Networks (NNs) have seen increasing use in biomedical applications [3]–[8] and predicting cardiac arrhythmia such as AF. The “Computers in Cardiology Challenge 2001” obtained model accuracy around 80% [9]–[11] with some methods achieving up to 89%

specificity and sensitivity [12]. Several techniques have been developed to predict AF using deep learning approaches [13]–[16]. However, these deep learning methods report their predictions with complete confidence regardless of the data. Regarding a disease with high morbidity and mortality rate as AF, having deterministic predictions from these models can cause more harm than good. A major flaw in these NNs are their likelihood at mislabeling AF as a normal rhythm as certainly as labeling a normal rhythm as AF. Thus, providing an uncertainty metric alongside the model’s decision provides insight on the data and assists the physicians with reaching an accurate diagnosis [17].

Initially, preliminary methods of uncertainty estimation involved using a softmax output probabilities as model confidence. However, the distribution of the probability using the softmax function may not accurately reflect the model’s confidence as the model may report a high softmax classification, but it still retains low confidence [18]. Several methods such as Markov Chain Monte-Carlo and variational dropout have been developed to improve this estimation the uncertainty of deep learning methods [17]. These models report a metric to assess the amount of trust that should be placed on a network’s prediction. In general, there are two types of uncertainty, model or epistemic uncertainty and data or aleatoric uncertainty. Epistemic uncertainty relates to the prediction model and its parameters. Epistemic uncertainty can be trained away as more data is given to the network, however it is much more sensitive to insufficient data. Data or aleatoric uncertainty handles the uncertainty caused by the complexity or the amount of noise present in data. Noisier or new data is likely to provide higher aleatoric uncertainties. The objective of our proposed method is to measure the aleatoric uncertainty caused by the input data data. This method minimizes the risk of misidentifying AF signals as normal and vice versa caused by noisy data or misplaced or malfunctioning sensors.

This paper proposes an AF prediction model using a variational autoencoder (VAE) approach to extract

¹This material is based upon work supported by the National Science Foundation under Grant Number 1657260.

high level features and sample an input from a random distribution that is passed to a classifier to obtain a prediction. The use of a VAE will yield more effective uncertainty estimation as the part of the network will approximate the posterior of the input which is then sampled several times to calculate the standard deviation of this latent distribution. VAEs take advantage of the encoder-decoder structure to obtain probabilistic features like uncertainty but acts as a generative network that retains the capability of comparison to the original input. Our network uses the VAE architecture to evaluate the decision and uncertainty as each input is run through this network multiple times to obtain a distribution of decisions. The most frequent decision is taken as the model’s final prediction and the standard deviation of the predictions are taken as the model’s uncertainty in the decision. Because aleatoric uncertainty will reflect the amount of signal noise, ECG signals will require less pre-processing compared to other methods that use the full signal, or heart rate variability (HRV) [19], morphological variability (MV), or wavelet features [20]. While deep neural networks trade off simplicity and time for better accuracy and clarity, the measure of confidence provided by more complex networks alleviate some of the trust issues present between AI and physicians. A physician may easily interpret a reasonable diagnosis with low uncertainty as correct, while a classification associated with high uncertainty may warrant a second opinion. These more complex networks provide measures to prevent misdiagnosis and clarity of decision to the physician at the cost of complexity and time.

The main contribution of this work is proposing an AF prediction model that combines state-of-the-art accuracy with uncertainty estimation that relays the confidence of our model in it’s decision of the input. The model’s uncertainty provides physicians a metric of trust on the predicted classification, alleviating the burden by indicating whether a case must be reviewed with expertise knowledge.

II. MODEL DESCRIPTION

The proposed model is composed of three main parts, including encoder, decoder and classifier networks. The encoder network employs a residual network (ResNet) architecture to automatically extract features from the raw signal data. The residual network consists of several CNNs with shortcut layers that passes a copy of the previous output. The ResNet architecture optimizes the convolutional layers such that their outputs achieve a residual close or at zero. Having this residual retains enough of the information about the input after each CNN layer and enables the capability to build deeper networks while minimizing over-fitting issues that are common with deep CNNs. The encoder (i.e., the ResNet

network) in parallel with several dense layers encode and extract the latent variables of the data that can be used by the decoder to reconstruct the input signal. Mathematically, the encoder approximates distribution, q , of input data over the latent space. The Kullback-Liebler (KL) divergence mechanism presented is the loss function used to shape the encoder to approximate the distribution (as shown in the next section). The KL divergence terms also acts as a regularizer, preventing our network from deviating too far from approximating the posterior distribution and prevents overfitting. Then, data points are sampled from the distribution to be fed into the decoder.

The decoder network consists of a network of transposed 2D convolution layers and is used to reconstruct the input data. To this end, the sampled data points are decoded by the decoder. Similar to the encoder network, the decoder network is in parallel with few dense layers and concatenated prior to a dense layer of the input size. Finally, the classifier network composing of two dense layers followed by a softmax layer employs the output of the previous step to perform the classification task. Figures 1, 2, and 3 show more details about the encoder, decoder, and the classifier that we used in our approach accordingly.

A. Uncertainty Estimation

We obtain an uncertainty estimation of the data by implementing a variational autoencoder (VAE) network consisting of separate encoder and decoder networks. Mathematically, VAE assumes that the data is part of a posterior distribution. Using the encoder output as an approximate distribution, by training the VAE, we are able to approximate the unknown prior distribution for Bayes’ theorem. The ResNet mentioned earlier acts as the encoder network for the VAE, thus we are able to infer the aleatoric uncertainty of the data by sampling features learned by the ResNet encoder as though they are random variables within a distribution. Because a layer with a random output cannot be backpropagated, discrete layers for the mean and standard deviation are used to allow backpropagation. Then the random sampling is obtained by combining the output of the random distribution output with the mean and standard deviation layers.

Because the VAE structure assumes that the input data has some posterior distribution $p(z|x)$ with z being a latent variable and x as the observed variable, we can obtain the posterior distribution through Bayes’ Theorem.

$$p(z|x) = \frac{p(x|z)p(z)}{p(x)} \quad (1)$$

$$p(x) = \int p(x|z)p(z)dz, \quad (2)$$

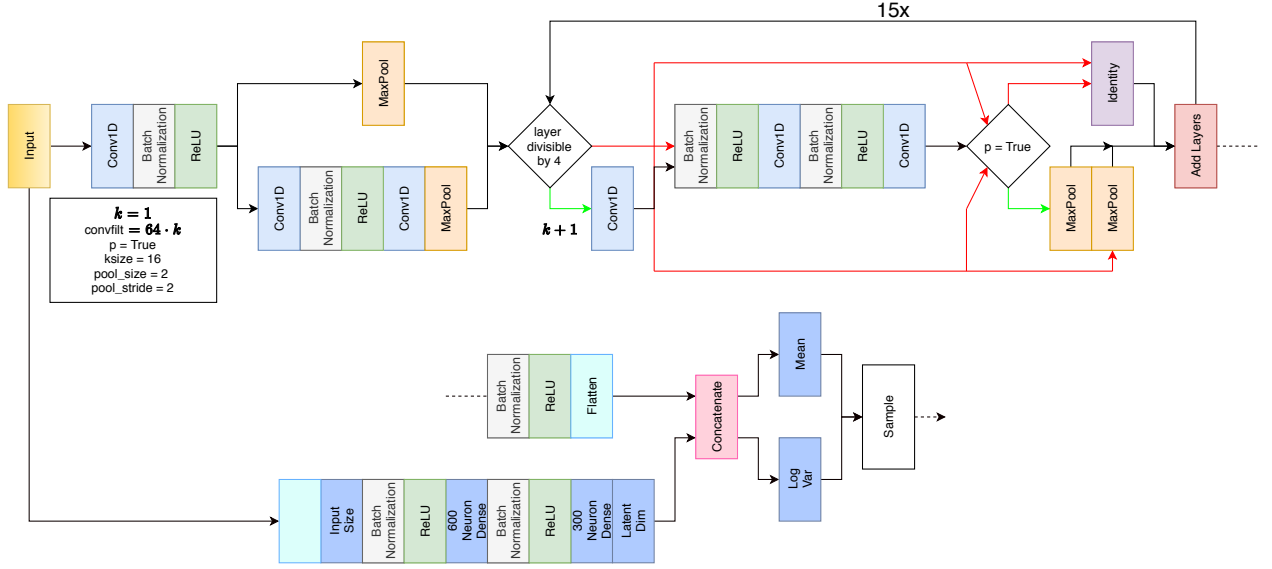


Fig. 1. Encoder Network

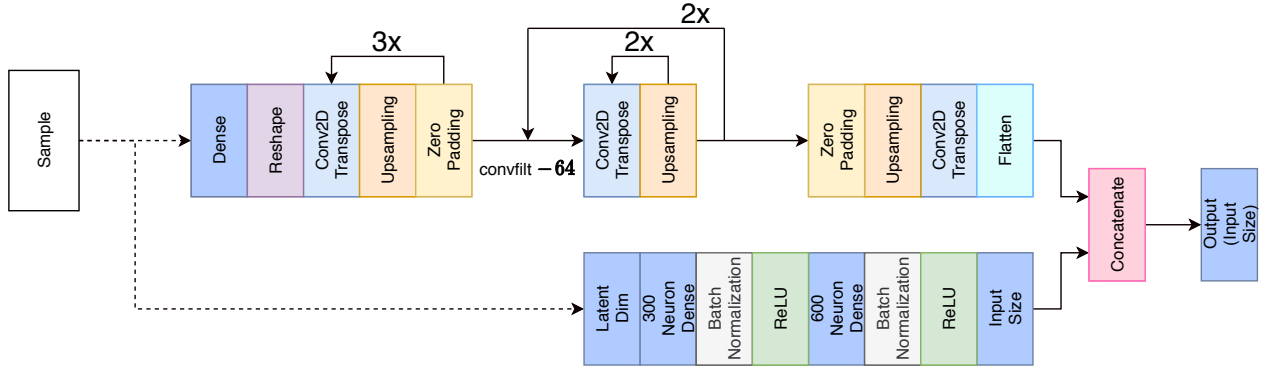


Fig. 2. Decoder Network

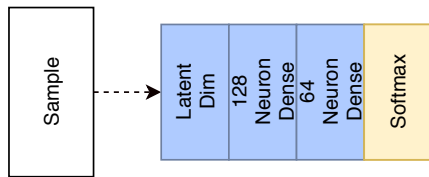


Fig. 3. Classifier Network

however, this requires knowing or computing $p(x)$ as in [2], the integral for which is intractable. On the other hand, a known distribution $q(z|x)$ can be used to approximate the posterior distribution. We use the KL Divergence to maximize the similarity. By minimizing the KL Divergence, we infer the characteristics of the posterior distribution of x , using our known Gaussian distribution

q . This can be rewritten as [21].

$$E_{q(z|x)} \log p(x|z) - KL(q(z|x)||p(z)) \quad (3)$$

where the first term shows the reconstruction likelihood, or Evidence Lower Bound (ELBO), and the second term is the KL Divergence of $q(z|x)$ in relation to the latent distribution $p(z)$. Because the KL Divergence is always non-negative, maximizing the full equation also minimize the KL Divergence, and approximates the posterior $p(x)$. This becomes the loss function for the VAE that drives the encoder to approximate the posterior and the decoder to successfully reconstruct the input, thereby minimizing the difference between the encoder's approximate distribution and true posterior.

While learning the posterior distribution using the VAE and the latent variable z poses an unsupervised learning task, adding a supervised classification output

merits that features extracted in z must be able to classify the inputs into one of several labels in addition to effectively reconstructing the input. To reflect this, we add a new term to the ELBO and KL Divergence such that the loss function becomes [21]:

$$loss = 10 * p(y|x) + E_{q(z|x)} \log p(x|z) - KL(q(z|x)||p(z)), \quad (4)$$

where $p(y|x)$ is the error contributed by the classifier, this term is multiplied by 10 as an unscaled value had little impact on improving the classifier network resulting in a trained encoder-decoder structure with random, untrained weights on the classifier. This drives the VAE to extract features that can accurately reconstruct the data, and allow the added classifier network to discriminate the input into discrete classes.

After the network is trained, the output and uncertainty for the prediction is obtained by running the input through the network n times. The softmax probabilities of the output are then summed and appended to an array. The classification is calculated by averaging each class’s summed probability and the index of highest mean is taken as the model’s classification. The uncertainty is estimated by taking the softmax probability array and calculating the variance, the square root of this variance is then taken as the standard deviation that will be used as the model’s estimate of uncertainty.

III. EXPERIMENTS

In this section, we describe the dataset along with the results of experimental analysis to evaluate the performance of this proposed method.

A. Dataset Description

The performance of the proposed method is evaluated using the MIT-BIH AFIB database from Physionet [22]. This dataset contains 23 patients each with two 10-hour long ECG recordings sampled at 250 Hz over a range of $\pm 10mV$ with a 12-bit resolution. The recordings are then split into 167,422 5-second segments and are labeled based on a threshold parameter. A segment is labeled as AF if the percentage of annotated beats exceeds the threshold parameters, anything below the threshold is considered as normal. In our method, we used a 50% threshold, creating 66,939 AF signals and 100,483 normal signals. This caused an imbalance in the dataset, to avoid further complicating the network by creating class weights and other imbalanced data optimization, the full amount of AF signals were taken and an equal amount of normal signals were taken in order.

B. Experimental Setup

The performance of the proposed model is evaluated by comparing its classification accuracy against several state-of-the-art ANN-based AFib detection networks, and then evaluated while taking into account the uncertainty estimation.

We applied a 10-fold cross-validation training method to evaluate the proposed model. The data was split into 10 groups of data, each round involved separating the first fold for testing and training on the remaining nine.

The model was trained with a maximum of 10 epochs using a batch size of 128, with the Adam optimizer initialized at a learning rate of $a = 0.001$. The model utilized the full ResNet model using 15 intermediate layers and the full decoder structure. The classifier loss is scaled up as, during training, the respective loss calculated by Keras did not improve the classifier’s performance. Scaling this value by 10 solved the issue.

C. Results

Table I shows the performance of our proposed model compared to several previous state-of-the-art AFIB classification methods. The table shows that our proposed network provides comparable (and in several cases better) performance in terms of the *Sensitivity* and *Specificity* metrics. While the overall accuracy of our method is slightly lower than the algorithm proposed by Xia et al [23], its key advantage is providing a confidence metric on the network’s decision. We observe that our network achieves state-of-the-art classification capability, while also providing an estimation of the aleatoric uncertainty.

TABLE I
COMPARISON OF PERFORMANCE OF THE PROPOSED MODEL AGAINST OTHER ALGORITHMS ON THE MIT-BIH AFIB DATABASE WITH THE ECG SEGMENT OF SIZE 5-S.

Method	Best Performance (%)		
	Sensitivity	Specificity	Accuracy
Proposed VAE with Classifier	97.58	97.70	97.64
Xia et al. (2018) [23]	98.79	97.87	98.63
Asgari et al. (2015) [24]	97.00	97.10	—
Lee, et al. (2013) [25]	98.20	97.70	—
Jiang et al. (2012) [26]	98.20	97.50	—
Huang et al. (2011) [27]	96.10	98.10	—
Babaeizadeh, et al. (2009) [28]	92.00	95.50	—
Dash et al. (2009) [29]	94.40	95.10	—
Tateno et al. (2001) [30]	94.40	97.20	—

IV. UNCERTAINTY RESULTS

Our uncertainty algorithm was tested using the same MIT-BIH AFib dataset cross-validation train and test sets. The train set is the same dataset that was used to train that iteration of the model while the test set was the same used to evaluate the model’s validation metrics.

The uncertainty was estimated by running each sample 5 times from the dataset through the network to obtain

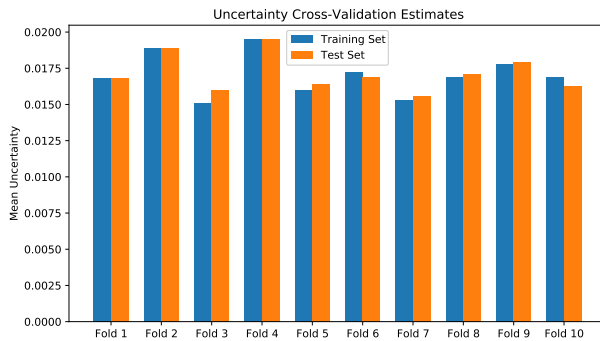


Fig. 4. The mean uncertainties for each Cross-Validation Set

an array of softmax probabilities. Each sample is run 5 times to compromise between effective uncertainty estimation and time to decision. More iterations may yield insignificant changes to the uncertainty at the cost of longer process times. The standard deviation is calculated from these probabilities and reported as the model’s measure of uncertainty. The test set included all samples from the cross-validation set, while the training set are randomly selected samples until the same amount of the test set is reached.

By utilizing the VAE architecture, we can construct a distribution of the output probabilities, and by calculating the standard deviation we can measure of the aleatoric uncertainty of the data. As shown in Fig. 4, over half the folds showed higher uncertainty on the newer, similar data present in the test set. The network is able to portray the difference in the learned training data and the new test set data with higher uncertainty. However, the aleatoric uncertainty is much less sensitive to the new data, as the changes from the outputs are much smaller. The similarity of the test data result in smaller differences in the standard deviation compared to the training set.

V. CONCLUSION

In this paper, we proposed a Variational Autoencoder network algorithm that can distinguish and classify an AFib rhythm from a normal rhythm from an ECG signal input. The classification algorithm provides a measure of confidence in the aleatoric uncertainty of the input. Through experimentation and analysis, we showed how the state-of-the-art classification performance can be coupled with uncertainty estimation. Knowing how much confidence can be placed on the model’s decision leverages trusting the algorithm with a task, and requiring a physician’s re-evaluation, minimizing the misclassification risk posed by more deterministic networks.

REFERENCES

- [1] P. A. Wolf, R. D. Abbott, and W. B. Kannel, “Atrial fibrillation as an independent risk factor for stroke: the framingham study.” *Stroke*, vol. 22, no. 8, pp. 983–988, 1991.
- [2] E. J. Benjamin, P. Muntner, and M. S. Bittencourt, “Heart disease and stroke statistics-2019 update: a report from the american heart association,” *Circulation*, vol. 139, no. 10, pp. e56–e528, 2019.
- [3] O. Yildirim, M. Talo, E. J. Ciaccio, R. San Tan, and U. R. Acharya, “Accurate deep neural network model to detect cardiac arrhythmia on more than 10,000 individual subject ecg records,” *Computer methods and programs in biomedicine*, vol. 197, p. 105740, 2020.
- [4] S. Mousavi and F. Afghah, “Inter-and intra-patient ecg heartbeat classification for arrhythmia detection: a sequence to sequence deep learning approach,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 1308–1312.
- [5] S. Mousavi, F. Afghah, and U. R. Acharya, “Sleepnet: Automated sleep stage scoring with sequence to sequence deep learning approach,” *PLoS one*, vol. 14, no. 5, p. e0216456, 2019.
- [6] A. Shamsoshoara, M. Khaledi, F. Afghah, A. Razi, J. Ashdown, and K. Turck, “A solution for dynamic spectrum management in mission-critical uav networks,” in *2019 16th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. IEEE, 2019, pp. 1–6.
- [7] A. Shamsoshoara, M. Khaledi, F. Afghah, A. Razi, and J. Ashdown, “Distributed cooperative spectrum sharing in uav networks using multi-agent reinforcement learning,” in *2019 IEEE Annual Consumer Communications & Networking Conference (CCNC)*. IEEE, 2019, pp. 1–6.
- [8] S. Mousavi, A. Fotoohinasab, and F. Afghah, “Single-modal and multi-modal false arrhythmia alarm reduction using attention-based convolutional and recurrent neural networks,” *PLoS one*, vol. 15, no. 1, p. e0226990, 2020.
- [9] G. Moody, A. Goldberger, S. McClennen, and S. Swiryn, “Predicting the onset of paroxysmal atrial fibrillation: The computers in cardiology challenge 2001,” in *Computers in Cardiology 2001. Vol. 28 (Cat. No. 01CH37287)*. IEEE, 2001, pp. 113–116.
- [10] P. De Chazal and C. Heneghan, “Automated assessment of atrial fibrillation,” in *Computers in Cardiology 2001. Vol. 28 (Cat. No. 01CH37287)*. IEEE, 2001, pp. 117–120.
- [11] C. Maier, M. Bauch, and H. Dickhaus, “Screening and prediction of paroxysmal atrial fibrillation by analysis of heart rate variability parameters,” in *Computers in Cardiology 2001. Vol. 28 (Cat. No. 01CH37287)*. IEEE, 2001, pp. 129–132.
- [12] H. Costin, C. ROTARIU, and A. PĂSĂRICĂ, “A new method for paroxysmal atrial fibrillation automatic prediction,” *Buletinul Institutului Politehnic din Iași*, vol. 59, pp. 71–83, 2013.
- [13] S. Mousavi, F. Afghah, and U. R. Acharya, “Han-ecg: An interpretable atrial fibrillation detection model using hierarchical attention networks,” *arXiv preprint arXiv:2002.05262*, 2020.
- [14] S. Mousavi, F. Afghah, A. Razi, and U. R. Acharya, “Ecgnet: Learning where to attend for detection of atrial fibrillation with deep visual attention,” in *2019 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*. IEEE, 2019, pp. 1–4.
- [15] U. R. Acharya, H. Fujita, S. L. Oh, U. Raghavendra, J. H. Tan, M. Adam, A. Gertych, and Y. Hagiwara, “Automated identification of shockable and non-shockable life-threatening ventricular arrhythmias using convolutional neural network,” *Future Generation Computer Systems*, vol. 79, pp. 952–959, 2018.
- [16] S. Mousavi, F. Afghah, F. Khadem, and U. R. Acharya, “Ecg language processing (elp): a new technique to analyze ecg signals,” *arXiv preprint arXiv:2006.08841*, 2020.
- [17] H. Wang and D.-Y. Yeung, “Towards bayesian deep learning: A framework and some existing methods,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 28, no. 12, pp. 3395–3408, 2016.
- [18] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *international conference on machine learning*, 2016, pp. 1050–1059.

- [19] S. Joo, K.-J. Choi, and S.-J. Huh, "Prediction of spontaneous ventricular tachyarrhythmia by an artificial neural network using parameters gleaned from short-term heart rate variability," *Expert Systems with Applications*, vol. 39, no. 3, pp. 3862 – 3866, 2012.
- [20] M. Alfaras, M. C. Soriano, and S. Ortín, "A fast machine learning model for ecg-based heartbeat classification and arrhythmia detection," *Frontiers in Physics*, vol. 7, p. 103, 2019.
- [21] D. P. Kingma and M. Welling, "An introduction to variational autoencoders," *arXiv preprint arXiv:1906.02691*, 2019.
- [22] "Mit-bih atrial fibrillation database v1.0.0," <https://physionet.org/content/afdb/1.0.0/>, Nov 2000, (Accessed on 03/15/2020).
- [23] Y. Xia, N. Wulan, K. Wang, and H. Zhang, "Detecting atrial fibrillation by deep convolutional neural networks," *Computers in Biology and Medicine*, vol. 93, 12 2017.
- [24] S. Asgari, A. Mehrnia, and M. Moussavi, "Automatic detection of atrial fibrillation using stationary wavelet transform and support vector machine," *Computers in Biology and Medicine*, vol. 60, 05 2015.
- [25] J. Lee, B. A. Reyes, D. D. McManus, O. Maitas, and K. H. Chon, "Atrial fibrillation detection using an iphone 4s," *IEEE Transactions on Biomedical Engineering*, vol. 60, no. 1, pp. 203–206, 2013.
- [26] K. Jiang, C. Huang, S.-M. Ye, and H. Chen, "High accuracy in automatic detection of atrial fibrillation for holter monitoring," *Journal of Zhejiang University. Science. B*, vol. 13, pp. 751–756, 10 2012.
- [27] C. Huang, S. Ye, H. Chen, D. Li, F. He, and Y. Tu, "A novel method for detection of the transition between atrial fibrillation and sinus rhythm," *IEEE Transactions on Biomedical Engineering*, vol. 58, no. 4, pp. 1113–1119, 2011.
- [28] S. Babaeizadeh, R. Gregg, E. Helffenbein, J. Lindauer, and S. Zhou, "Improvements in atrial fibrillation detection for real-time monitoring," *Journal of electrocardiology*, vol. 42, pp. 522–6, 08 2009.
- [29] S. Dash, K. Chon, S. Lu, and E. Raeder, "Automatic real time detection of atrial fibrillation," *Annals of biomedical engineering*, vol. 37, pp. 1701–9, 07 2009.
- [30] K. Tateno and L. Glass, "Automatic detection of atrial fibrillation using the coefficient of variation and density histograms of rr and δrrr intervals," *Medical & biological engineering & computing*, vol. 39, pp. 664–71, 12 2001.