

# Learning to Learn Personalized Neural Network for Ventricular Arrhythmias Detection on Intracardiac EGMs

Zhenge Jia<sup>1</sup>, Zhepeng Wang<sup>1</sup>, Feng Hong<sup>2</sup>, Lichuan Ping<sup>2</sup>, Yiyu Shi<sup>3</sup> and Jingtong Hu<sup>1</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, University of Pittsburgh

<sup>2</sup>Singular Medical Co., Ltd.

<sup>3</sup>Department of Computer Science and Engineering, University of Notre Dame

{zhenge.jia, zhepeng.wang}@pitt.edu, {hongfeng, pinglichuan}@singularmedical.net, yshi4@nd.edu, jthu@pitt.edu

## Abstract

Life-threatening ventricular arrhythmias (VAs) detection on intracardiac electrograms (IEGMs) is essential to Implantable Cardioverter Defibrillators (ICDs). However, current VAs detection methods count on a variety of heuristic detection criteria, and require frequent manual interventions to personalize criteria parameters for each patient to achieve accurate detection. In this work, we propose a one-dimensional convolutional neural network (1D-CNN) based life-threatening VAs detection on IEGMs. The network architecture is elaborately designed to satisfy the extreme resource constraints of the ICD while maintaining a high detection accuracy. We further propose a meta-learning algorithm with a novel patient-wise training tasks formatting strategy to personalize the 1D-CNN. The algorithm generates a well-generalized model initialization containing across-patient knowledge, and performs a quick adaptation of the model to the specific patient's IEGMs. In this way, a new patient could be immediately assigned with personalized 1D-CNN model parameters using limited input data. Compared with the conventional VAs detection method, the proposed method achieves 2.2% increased sensitivity for detecting VAs rhythm and 8.6% increased specificity for non-VAs rhythm.

## 1 Introduction

The Implantable Cardioverter Defibrillator (ICD) is a small device implanted to reduce Sudden Cardiac Death (SCD) risk and improve the survival rate for recipients by detecting ventricular arrhythmias (VAs) on IEGMs, which reflect electric changes measured by electrodes placed within the heart. Life-threatening VAs includes ventricular tachycardia (VT) and ventricular fibrillation (VF), which are the main cause of SCD [Zdarek and Israel, 2016]. ICDs are programmed to release defibrillation (i.e., electrical shock) therapy on VAs to restore rhythm back to normal. Therefore, VAs detection is a pivotal component to ICDs as detection accuracy severely affects the effectiveness and lifetime of ICDs.

Current VA detection methods count on a wide variety of criteria and there are hundreds of programmable criteria

parameters affecting the defibrillation delivery decision in these methods [Madhavan and Friedman, 2013]. To achieve the best detection performance for the patient, cardiologists have to periodically and manually tune the criteria parameters by learning from the cardiac rhythm of the patient [Senges-Becker *et al.*, 2005]. Current VA detection methods utilized in ICDs are often referred to as *criteria-based* VA detection.

However, existing criteria-based methods still suffer from a relatively high Inappropriate Shock Rate (ISR). *Inappropriate shock* is the shock delivered to a non-VA rhythm (i.e., the rhythm other than VT or VF) due to an incorrect detection and has been associated with proarrhythmia, intolerable pain, anxiety and depression on patients [Moss *et al.*, 2012]. Inappropriate shock has been reported in 12-23% ICD recipients [Greenberg *et al.*, 2004; Hofer *et al.*, 2017]. To alleviate the problem, designing new detection criteria and the personalized fine-tuning on criteria parameters are two vital approaches. However, both procedures demand a high level of expertise engaged and a long period of optimization.

Recently, deep learning based arrhythmias detection on electrocardiogram (ECG) has achieved outstanding performance in terms of accuracy [Hannun *et al.*, 2019; Elola *et al.*, 2019; Acharya *et al.*, 2018]. Compared with criteria-based detection methods, its most significant advantage is the reduction of the cardiological expertise required in the detection method design. The labeled ECG signal could be directly utilized as the input and the pre-trained CNN model outputs the prediction without cumbersome criteria design and selection.

However, there are two main challenges in applying the aforementioned CNN-based detection models to the scenarios of VA detection on IEGMs: 1) The existing CNNs for accurate arrhythmias detection cannot satisfy the resource constraints if deployed on the ICDs platform; 2) The detection accuracy of the pre-trained CNN model would degrade significantly on some patients due to their unique rhythm features, which is intolerable in the ICD scenarios. Fine-tuning the CNN model using the data from the specific patient could improve the detection accuracy. However, it is not practical since fine-tuning the CNN requires extensive amount of labeled IEGMs for each individual patient.

To address the challenges, we propose the one-dimensional convolutional neural network (1D-CNN) based VAs detection and a novel meta-learning algorithm to personalize the 1D-CNN for patient-specific arrhythmias detection. Specifically,

we develop a 1D-CNN to detect VAs on the one-channel IEGMs. The 1D-CNN only occupies tiny memory space and computational power such that it fits the resource-constrained ICDs. Since each patient has the unique characteristics in the rhythm, a pre-trained 1D-CNN cannot achieve the optimal detection performance for each individual. Thus we propose a novel meta-learning algorithm. Different from the *N-way-K-shot* classification targeted by conventional meta-learning approaches [Finn *et al.*, 2017a], our meta-learning algorithm accommodates the patient-specific VAs detection through a novel patient-wise training tasks formatting strategy. In this way, the 1D-CNN could learn across-patient knowledge and end up with a well-generalized model initialization. When a new patient is involved, a quick adaption would be performed on the well-generalized 1D-CNN to the specific patient's rhythm with only limited IEGMs of the patient. An *inner-loop update optimization* and a *dynamic task set mechanism* are devised to optimize the meta-learning process, and at the same time further improve the model generalization. The effectiveness of our meta-learning algorithm has also been examined by surface rhythm ECG, where the patient-specific arrhythmias detection are evaluated.

The main contributions summarized are as follows:

- We propose a novel 1D-CNN, which meets the requirements of ICDs in terms of memory, response time and power while maintaining a high detection accuracy.
- To the best of our knowledge, this paper is the first to propose a meta-learning method to accommodate the patient-specific arrhythmias detection through a novel patient-wise training tasks formatting strategy.
- We further improve the arrhythmias detection performance by optimizing the generalization of the meta-model initialization through a dynamic task set mechanism and an inner-loop update optimization.
- Our meta-learning algorithm is shown to be effective and generalized on various arrhythmias detection on surface rhythm from ECG standard databases.

## 2 Related Works and Motivations

### 2.1 Ventricular Arrhythmias and ICDs

VF and VT are called life-threatening ventricular arrhythmias (VAs) which are the main cause of Sudden Cardiac Death (SCD). More than 60% of deaths from cardiovascular disease are from the out-of-hospital SCD [Adabag *et al.*, 2010]. Unless the heart is shocked back into normal rhythm, the patient rarely survives.

There are various types of ICDs. In this paper, we focus on improving the VAs detection accuracy of single-chamber ICDs [Zanker *et al.*, 2016]. The single-chamber ICD has leads attached in right ventricle to sense ventricular heart rhythm and the rhythm data that the ICD normally utilizes to detect VAs is the one-channel EGMs sensed by right ventricular apex-bipolar (RVA-Bi) lead. Due to the constraints in physical size and battery lifetime, current ICDs can only provide 128-1024 KB for ROM and 8-32 KB for RAM to store data and execute program with 1-5 mA during VAs detection [Strydis *et al.*, 2013; Camara *et al.*, 2015].

### 2.2 VAs Detection Methods

Current VAs detection methods on ICDs utilize various criteria such as number of intervals to detect (NID), fast/slow interval determination, rhythm stability, and onset analysis on IEGMs to determine VAs [Zdarek and Israel, 2016; Thøgersen *et al.*, 2017]. To achieve the best detection performance for each patient, cardiologists have to periodically fine-tune the criteria parameters by learning the sensed rhythm of the patient, which requires massive manual intervention to achieve the optimal VAs detection.

Machine Learning (ML) based detection methods for arrhythmias classification become popular in recent years. The benefits of applying ML are the high detection accuracy and the degradation of expertise required in method design. Authors in [Li *et al.*, 2013] utilize SVM to discriminate VT/VF rhythm based on various combinations of features with the highest 97.52% accuracy. The convolutional neural network (CNN) proposed in [Hannun *et al.*, 2019] achieves an even better performance in 12 arrhythmias classification than cardiologists using single-lead ECG. However, both methods could not be directly deployed on ICDs since the computational resources are severely constrained. Authors in [Acharya *et al.*, 2018] propose a CNN with simple structure for VT/VF detection using single-lead surface ECG in intra-patient paradigms (i.e., the training and testing set may include data from the same patient). However, the detection accuracy on non-VAs rhythm is only 91.04%, which cannot be adopted in the ICDs scenarios.

### 2.3 Motivations

In this work, we are looking for a decent solution to personalize the VAs detection model with only limited rhythm data of each individual patient. We find the conventional meta-learning approaches (e.g. MAML [Finn *et al.*, 2017b] and Reptile [Nichol *et al.*, 2018]) provide a learning strategy that a meta-model with well-generalized initialization is obtained by training on the tasks containing support and query set. The meta-model can be further fine-tuned using the specific patient's rhythm data to obtain the personalized detection model. However, those methods mix training samples with different classes from various patients in a single training task during meta-learning process. The training tasks with such biased data distribution may cause gradient diminishing and training instability, which degrade the generalization of the model and leads to low VAs detection accuracy. Thus, it is necessary to devise a tasks formatting strategy for meta-learning approaches to optimize the training process and obtain a well-generalized model initialization.

## 3 Method

In this section, we first introduce the 1D-CNN based VAs detection. We then present the meta-learning algorithm which performs a quick and effective model personalization.

### 3.1 1D-CNN Based VAs Detection

We first demonstrate the 1D-CNN detection model applied on IEGMs in ICDs. As shown in Fig. 1, the input of the 1D-CNN is the set of 2s-length sampling points (i.e., total 500

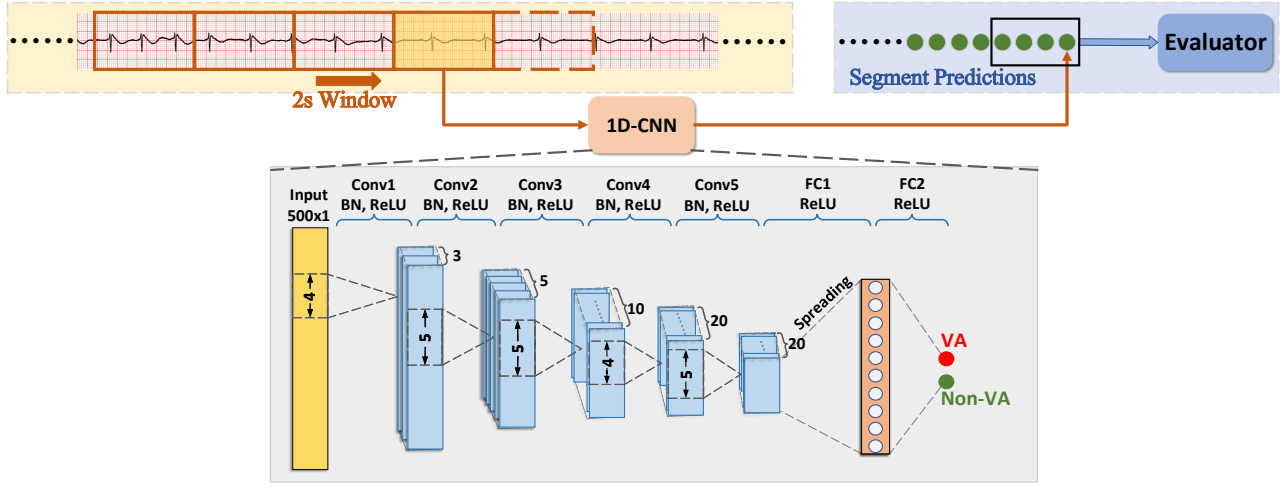


Figure 1: Illustrations of 1D-CNN based VAs detection.

points under 250Hz sampling rate) of the EGMs sensed by the ICD, and the output would be the prediction (i.e. VAs or non-VAs) on the given segment. The segment length is set to be 2s since it contains adequate rhythm information for accurate inference while we can maintain a relatively low computation overhead for the inference on the devices. The sampling rate of EGMs is 250Hz since it is the sampling rate applied in most ICDs [Simon *et al.*, 2007]. The 1D-CNN consists of five 1D-convolutions (stride 2) followed by two fully connected layers. The batch-normalization (BN) at each layer is applied to ensure the same magnitude of the numerical values.

The 1D-CNN based VAs detection mechanism leverages an evaluator to determine the VAs on the sensed IEGMs. As shown in Fig. 1, the evaluator receives the series of segment predictions and determines the VAs rhythm with an empirical criteria. The criteria is that the VAs rhythm would be reported if there are four consecutive VAs predictions on the 2s segments. In other words, the evaluator would consistently monitor the latest four inferences, and therefore the decision period is 8 seconds. The criteria is set since the detection period of criteria-based VAs detection in ICDs is usually 5 to 10 seconds [Madhavan and Friedman, 2013].

### 3.2 Learning to Learn Personalized 1D-CNN

There are two key steps to learn a personalized 1D-CNN for each individual patient’s VAs detection. 1) Meta-learning: Format the patient-wise training tasks and optimize learning process to obtain a well-generalized 1D-CNN initialization containing across-patient knowledge as *meta-model*. 2) Personalization: Perform a quick personalization on the meta-model to adapt to the specific patient’s intracardiac rhythm.

#### Meta-Learning

Formally, we denote the initial meta-model parameters of the 1D-CNN as  $\phi$ . The purpose of the learning process is to iteratively train the network over each *TaskSet*  $\mathcal{T}$  in the *TaskSet collection*  $\mathbf{T}$  to obtain a well-generalized meta-model parameters  $\phi^*$ . However, existing meta-learning algorithms focus on solving *N-way-K-shot* classification and are not applicable

in the VAs detection scenarios. To address the challenges, we propose the patient-wise training tasks and two optimizations to ensure that the across-patient knowledge is learned by the meta-model.

We first introduce **Patient-wise Training Tasks Formatting** along with the necessary formulation. For each task of the  $\mathcal{T}$ , we randomly select  $N$  patients from training patients dataset and randomly collect  $n$  VAs and  $n$  non-VAs labeled 2s segments from each selected patient to form a task  $\tau_i$ .

Once a task  $\tau_i$  is formed, the first  $p$  of every patient’s VAs and non-VAs segments would be extracted to form the *support set*, denoted as  $\tau_i^s$ . The rest  $n - p$  VAs and non-VAs segments of each patient in  $\tau_i$  would be formed as *query set*, denoted as  $\tau_i^q$ . The formal definition of both set are shown as follows:

$$\begin{aligned} \tau_i^s &= \{(x, y)_{i, M_j^{spt}}\} \text{ for } j = 1, \dots, 2p \cdot N, \\ \tau_i^q &= \{(x, y)_{i, M_j^{qry}}\} \text{ for } j = 1, \dots, 2(n - p) \cdot N, \end{aligned} \quad (1)$$

where  $(x, y)_i$  is the set of segment and label pairs in  $\tau_i$ , and  $M^{spt}$  and  $M^{qry}$  are the pair indices (in  $\tau_i$ ) set for  $\tau_i^s$  and  $\tau_i^q$  separately. There are total  $2p \cdot N$  indices for pairs in  $\tau_i^s$  and  $2(n - p) \cdot N$  indices for  $\tau_i^q$ . A TaskSet  $\mathcal{T}$  is constructed by repeating the aforementioned process to include a number of tasks  $\tau_i$  and a TaskSet collection  $\mathbf{T}$  contains multiple  $\mathcal{T}$ .

We define the 1D-CNN inference as  $f_\theta(x)$ , where  $x$  is the input segment and  $\theta$  is the parameters of the network. When adapting the model to each task  $\tau_i$ ,  $k$ -step gradient update is applied to update the model parameter from  $\theta_i^0 = \phi$  to  $\theta_i^k$  using the segment and label pairs in the support set  $\tau_i^s$  of task  $\tau_i$ . The one example gradient update on step  $m$  (where  $0 < m \leq k$ ) is defined as follows:

$$\theta_i^m = \theta_i^{m-1} - \alpha \frac{1}{|\tau_i^s|} \sum_{(x, y) \in \tau_i^s} \nabla_{\theta_i^{m-1}} \mathcal{L}(f_{\theta_i^{m-1}}(x), y), \quad (2)$$

where  $\mathcal{L}$  is the loss function and  $\alpha$  is the learning rate. The gradient descent (GD) in Eqn. (2) would be processed from step 0 to  $k$  to obtain the task-specific model with  $\theta_i^k$  for  $\tau_i$ :

$$\theta_i^k = \text{GD}_k(\theta_i^0). \quad (3)$$

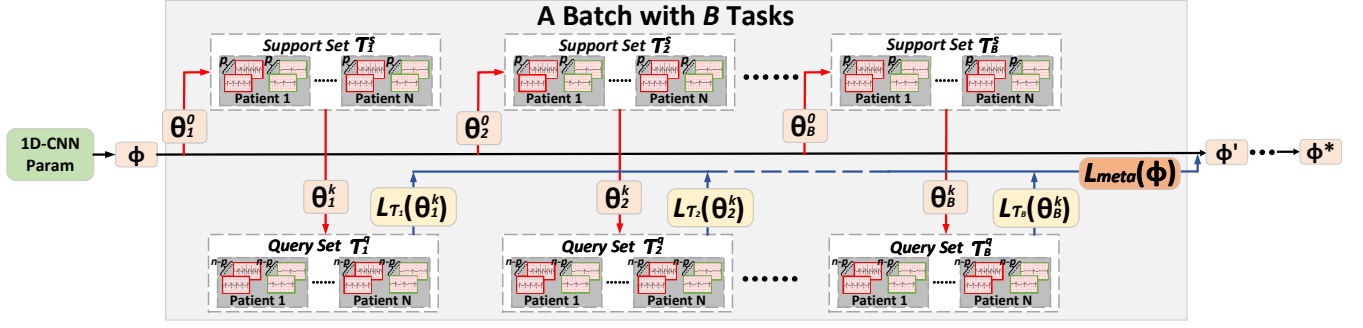


Figure 2: Illustration of the meta-learning over batches of  $N$ -patient-wise tasks. The red line indicates the inner-loop model update process, where each task-specific model is trained over its support set  $\tau_i^s$  to update its model parameters from  $\theta_i^0 = \phi$  to  $\theta_i^k$ . The blue line indicates the outer-loop meta-model update process, where the loss of each task-specific model over its query set  $\tau_i^q$  is calculated and accumulated as  $\mathcal{L}_{meta}(\phi)$ . The meta-model with parameters  $\phi$  learns the across-patient knowledge and iteratively updates its parameters batch by batch to obtain the well-generalized initialization  $\phi^*$ .

This process is defined as *inner-loop* model update.

Once obtaining  $|\mathcal{T}|$  number of task-specific models using Eqn. (3) for each task  $\tau_i \in \mathcal{T}$ , we then evaluate the generalization capability of the obtained models. We first calculate the loss of each task-specific model on the corresponding task's query set. The loss is calculated as follows:

$$L_{\tau_i}(\theta_i^k) = \frac{1}{|\tau_i^q|} \sum_{(x,y) \in \tau_i^q} \mathcal{L}(f_{\theta_i^k}(x), y). \quad (4)$$

Here, we fetch the tasks using a *tasks batch* methodology (i.e., the batch-size number of tasks are grouped as a batch) from the TaskSet  $\mathcal{T}$ . The meta objective of one batch of tasks is defined as follows:

$$\mathcal{L}_{meta}(\phi) = \sum_{b=1}^B L_{\tau_b}(\theta_b^k), \quad (5)$$

where  $b$  is the index of a particular task in the batch and  $B$  is the batch size. This objective in Eqn. (5) is utilized to evaluate the generalization of meta-model parameters  $\phi$  in terms of the total loss on all tasks in the batch. The objective should be minimized to optimize the meta-model parameters. Such optimization is defined as *outer-loop* meta-model update and could be solved using Adam or Stochastic Gradient Descent (SGD). The resulting meta-model parameters updating is shown as follows:

$$\phi' = \phi - \beta \nabla_{\phi} \sum_{b=1}^B L_{\tau_b}(\theta_b^k), \quad (6)$$

where  $\beta$  is the meta learning rate and  $\phi'$  is the updated parameters of meta-model which contains the across-patient knowledge and would be used as the initial parameters for next batch update. To obtain the well-generalized meta-model parameters  $\phi^*$ , we should iteratively update the meta-model over batches extracted from each TaskSet  $\mathcal{T} \in \mathbf{T}$ .

Fig. 2 demonstrates the meta-learning process from initial model parameters  $\phi$  to the well-generalized  $\phi^*$ . Each batch is formed by extracting different  $B$  tasks from the  $\mathcal{T}$ . As shown in Fig. 2, for each task  $\tau$  in a batch, the VAs segments

are denoted with red border and the non-VAs segments are denoted with green border.  $N$  selected patients' first  $p$  VAs and non-VAs segments form the support set  $\tau^s$  and the rest  $n - p$  segments form the query set  $\tau^q$  of  $\tau$ . The batch size  $B$  is fixed for all  $\mathcal{T} \in \mathbf{T}$ . Therefore, the meta-learning process is conducted batch by batch until every task from all  $\mathcal{T}$  in  $\mathbf{T}$  are utilized as a member of the batch.

To increase the generalization of the final meta-model parameters and speed up the learning process, we further propose the following optimizations:

**Dynamic Task Set Mechanism.** For each TaskSet  $\mathcal{T} \in \mathbf{T}$ , we apply a dynamic setting to construct the support and query set for the tasks. In conventional  $N$ -way- $K$ -shot classification problems, the settings of tasks are fixed. The task must consist of samples from different  $N$  classes and its support set must have  $K$  elements to fit the classification problem. On the other hand, in our scenarios of VAs detection, we could have a flexible setting for the tasks in one  $\mathcal{T}$  to increase the universality of the TaskSet collection  $\mathbf{T}$  without violating the classification formulation. Here, we extract  $N$  patients and a dynamic number of VAs and non-VAs segments of each patient to form the support and query set among different  $\mathcal{T}$ . The value of  $N$  is set to be a random integer value in a preset range  $[N_{lower}, N_{upper}]$ . The support and query set defined in Eqn. (1) are formed with a dynamic  $p \in [U_{low}, U_{up}]$ , where  $U_{low}, U_{up}$  are the lower and upper bound of the random integer  $p$ . Note that inside one  $\mathcal{T}$ , the value of  $N$  and  $p$  are kept the same. This optimization technique could increase the universality of the data collection so that we could obtain a more generalized meta-model parameters.

**Inner-loop Update Optimization.** The training of the conventional meta-learning method is time-consuming and may come with training instability problem [Antoniou *et al.*, 2018]. To speed up the training process while maintaining a high quality of the meta-model parameters, we propose an inner-loop update steps annealing strategy. This strategy is based on the observation that the meta-model would rapidly learn the across-patient knowledge through the preceding part of the TaskSet collection  $\mathbf{T}$ . As a consequence, for the  $\mathcal{T}_a$  from the latter of  $\mathbf{T}$ , the losses calculated on each update step

of the task-specific model are relatively stable or even the same. Therefore, we propose an annealing strategy based on a step function to evenly reduce the inner-loop update steps  $k$  in Eqn. (3) for each  $\mathcal{T}$  in  $\mathbf{T}$ . The strategy is defined as follows:

$$\theta_i^k = \text{GD}_k(\theta_i^0), \text{ for } k = \max(k - a \bmod \mathbf{w}, K), \quad (7)$$

where  $a$  is the index of a particular  $\mathcal{T}_a \in \mathbf{T}$ ,  $\mathbf{w}$  is the preset updating stride, and  $K$  is the minimal inner-loop update steps. To stabilize the gradients, we further invoke the multi-step loss optimization to take the loss of task-specific model based on the loss from *every* inner-loop update step. The Eqn. (4) can be re-written as:

$$L_{\tau_i}(\theta_i^k) = \sum_{m=0}^k v_k(m) \left( \frac{1}{|\tau_i^q|} \sum_{(x,y) \in \tau_i^q} \mathcal{L}(f_{\theta_i^m}(x), y) \right), \quad (8)$$

where  $k$  is the inner-loop steps calculated in Eqn. (7) and  $v_k(m)$  is the weight of the loss at step  $m$ .

### Personalization

The personalization of the 1D-CNN for the given new patient would start with the well-generalized model parameters  $\phi^*$  obtained from the meta-learning. Here, we take advantage of a routine process in the ICD implantation surgery where limited IEGMs containing VAs and non-VAs of the patient would be recorded to assist the cardiologist to tune the criteria parameters before implanting the ICD [Theuns and Gold, 2010]. Instead of having the cardiologists to optimize the criteria parameters manually, we utilize those limited VAs and non-VAs rhythm to automatically personalize the 1D-CNN by fine-tuning the model with  $\phi^*$ . The 1D-CNN could quickly adapt to the specific patient’s rhythm feature since  $\phi^*$  contains across-patient knowledge through meta-learning. Specifically, the parameters of the 1D-CNN with  $\phi^*$  would be updated by SGD and back-propagation with those limited but labeled VAs and non-VAs segments of the new patient. The 1D-CNN based VAs detection model would be personalized to adapt to the specific patient with a few training iterations.

## 4 Experiments

### 4.1 Dataset and Data Preprocessing

Our dataset is retrieved from volume I & II of Ann Arbor Electrogram Libraries (AAEL), which is one of the largest dataset for IEGMs and used by all manufacturers developing implantable defibrillators to test their methods [AnnArbor, 2003]. The sampling rate of all recordings is 1,000 Hz. Different periods of recordings have been annotated and reviewed by cardiac electrophysiologists to ensure an accurate interpretation of arrhythmia.

Here, we select 217 recordings over 78 patients to form the dataset by excluding the patients who do not experience VAs and the recordings with length less than 20s. Each selected recording contains the one-channel IEGMs sensed by RVA-Bi lead. All recordings are resampled to 250 Hz as the sampling rate is widely utilized in implantable devices [Simon *et al.*, 2007]. As the raw physiological signals contain noise, we perform very common filtering techniques on the IEGMs recordings by applying a band-pass FIR filter with a

Data	2s-Segments		Events
	Non-Overlapping	Overlapping	
<b>VAs</b>	2,197	10,542	145
<b>Non-VAs</b>	5,097	10,206	214

Table 1: Data profile for the dataset.

pass-band frequency of 0.5 Hz and stop-band frequency of 50 Hz. Next, the recording is participated into various VA or non-VA *events* based on the rhythm diagnostic annotations on the IEGMs periods. Then, each event is segmented into *segments* using a 2-second sliding window (250 Hz  $\times$  2 s = 500 samples) with and without a overlap. The overlap is set to perform data augmentation for training only. The detailed segments and events statistics are illustrated in Table ??.

### 4.2 Inter-Patient Evaluation Paradigms

The dataset is partitioned patient-wisely into 5 folds to ensure that the patients data is not mixed between the training and testing sets. That is, once the patient is selected for training or testing, all related events and segments would be utilized only for training or testing. The split is done randomly on the patients and we perform 5-fold cross-validation.

For each validation, four folds of patients’ overlapping 2s segments serve as the *training set* and are used to perform learning to learn a well-generalized 1D-CNN model initialization in our method or training in the baseline methods. For each patient in the last fold for testing, the first 10s rhythm (five consecutive 2s non-overlapping segments) of a VAs and a non-VAs event are extracted to serve as the *personalizing set* and used to personalize or fine-tune the model for the specific patient. The rest 2s non-overlapping segments of the patients serve as *testing set* and are used to report the detection performance. In the testing fold, each time we extract a different specific patient with the corresponding segments from the personalizing set to personalize the model (if applicable) and the testing set to evaluate the model.

### 4.3 Methods Evaluated

We present performance comparisons for the following VAs detection methods:

**Criteria-based VAs detection.** We implement a VAs detection method used in single-chamber ICDs [Zanker *et al.*, 2016], denoted as *Classic*. This method continuously monitors each heart beat and report VAs if criteria is satisfied. We set two detection zones for VT and VF respectively. The heart rate boundary of VT/VF zone and fast/slow interval threshold are fine-tuned for each testing patient to simulate the manual intervention such that the best discrimination performance could be achieved.

**SVM-based VAs detection.** We implement an existing machine learning based VAs detection method using support vector machine (SVM) [Li *et al.*, 2013], denoted as *SVM*. The features extracted in SVM are *Count2* and *Leakage*.

Methods	F1	Se/Sp	BAC/ACC	PPV/NPV
Classic	.923	.972/.910	.941/.935	.880/.979
SVM	.902	.951/.892	.921/.916	.861/.964
DL1	.926	.923/.954	.939/.940	.931/.947
DL1-FT	.953	.974/.952	.963/.961	.935/.981
DL2	.911	.903/.947	.925/.929	.922/.935
DL2-FT	.955	.950/.974	.962/.965	.960/.968
MAML	.959	.967/.973	.969/.943	.952/.953
Reptile	.943	.955/.941	.951/.948	.932/.976
Meta-Static	.973	.982/.985	.983/.987	.977/.976
Meta	<b>.993</b>	<b>.993/.996</b>	<b>.995/.995</b>	<b>.994/.996</b>

Table 2: Performance of evaluated methods on events.

Methods	F1	Se/Sp	BAC/ACC	PPV/NPV
SVM	.851	.937/.906	.922/.911	.788/.970
DL1	.879	.910/.944	.927/.931	.860/.960
DL1-FT	.925	.955/.958	.957/.957	.900/.981
DL2	.865	.890/.937	.914/.922	.844/.953
DL2-FT	.926	.936/.967	.951/.957	.916/.972
MAML	.929	.934/.923	.933/.951	.921/.969
Reptile	.935	.946/.922	.936/.945	.929/.966
Meta-Static	.954	.961/.959	.959/.948	.948/.961
Meta	<b>.965</b>	<b>.968/.960</b>	<b>.962/.953</b>	<b>.964/.977</b>

Table 3: Performance of evaluated methods on segments.

**CNN-based VAs detection.** We implement two deep learning based VAs detection methods. One is based on the CNN demonstrated in Fig. 1 and the other is an existing CNN-based VAs detection model [Acharya *et al.*, 2018]. Both models would be trained on the training set following the mini-batch training process. We denote these two detection approaches as *DL1* and *DL2*. Moreover, the pre-trained models would be firstly fine-tuned with the personalizing set and then evaluated on the testing set. We denote these two approaches as *DL1-FT* and *DL2-FT*.

**Meta learning-based VAs detection.** We implement our meta-learning based 1D-CNN VAs detection model and denote it as *Meta*. We also implement the same meta-learning based model but without inner-loop update optimization and dynamic task set mechanism, denoted as *Meta-Static*. We further implement a conventional meta-learning approach with and without first-order approximation [Finn *et al.*, 2017b; Nichol *et al.*, 2018], denoted as *MAML* and *Reptile*. Entropy loss is applied for all CNN models.

#### 4.4 Results

##### VAs Detection Performance Evaluation

We evaluate our meta-learning based 1D-CNN VAs detection against other methods in terms of various metrics including F1 score (F1), Sensitivity (Se), Specificity (Sp), balanced

accuracy (BAC), accuracy (ACC), positive predictive value (PPV), and negative predictive value (NPV). The condition positive is VAs and condition negative is non-VAs.

Table 2 illustrates the detection performance on events. The performance on events is more practical since the defibrillation therapy should be delivered based on the rhythm period but not a single 2s segment. Compared with the Classic, DL1 achieves 4.9% deduction from a baseline of 97.2% on VAs rhythm detection rate represented Se but 4.4% increase on non-VAs rhythm detection rate from a baseline of 91.0% represented by Sp. As for DL2, it achieves 6.9% lower on VAs rhythm detection rate and 3.7% higher on non-VAs rhythm detection rate compared with Classic. Both detection rates increase to 95.0% and 97.4% respectively after fine-tuning. As for SVM, it does not perform well since the features extracted in the method are not applicable to the IEGMs. The two conventional meta-learning approaches, MAML and Reptile, achieve the similar performances on VAs detection as the DL1-FT and DL2-FT. As for the Meta, it achieves the best performance on all testing metrics. It has the near-optimal detection accuracy on VAs rhythm (99.3%) and non-VAs rhythm (99.6%), and the highest F1 score (0.993).

Table 3 illustrates the detection performance on segments. Meta still achieves the best performance on all metrics compared with other evaluated methods. The fine-tuning is necessary for CNN-based methods to increase the detection accuracy. Therefore, the generalization of the CNN initialization is critical in the VAs detection scenarios and the proposed meta-learning provides an effective solution to the problem. The detailed experimental setup and the performance of the proposed meta learning method for arrhythmias detection on ECG would be demonstrated in the Appendix.

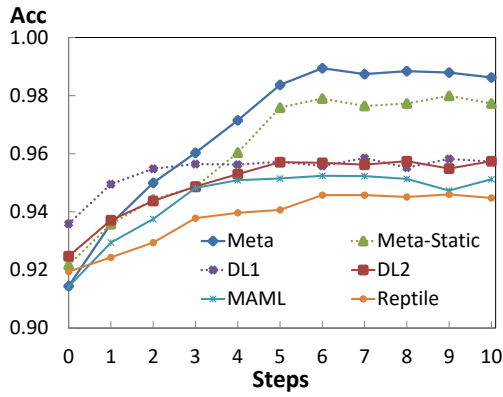
##### Model Generalization and Training Overhead

To demonstrate the importance of the proposed dynamic task set mechanism and inner-loop update optimization, we present the average accuracy and loss curve for each patient’s personalization process of Meta, Meta-Static, MAML, Reptile, DL1 and DL2 together with the training overhead.

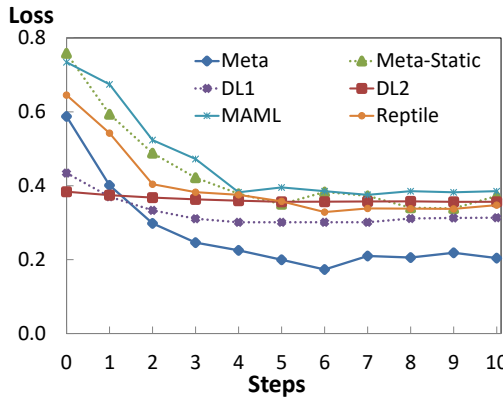
The 10-step gradient descent (GD) is applied to personalize each pre-trained model through fine-tuning on the patient-specific IEGMs. As shown in Fig. 3(a), at step 0, the initial model of our Meta method do not perform well on the patients compared with the models of DL1 and DL2. However, the model initialization of Meta could rapidly adapt to the specific patient’s rhythm and achieves higher accuracy through 10-step GD personalization. As shown in Fig. 3(b), the convergence speed of Meta is much faster than that of DL1, DL2, MAML and Reptile. The accuracy of Meta improves from 91.4% to 98.9% while that of Meta-Static improves from 92.2% to 97.8%. It indicates that the dynamic task set mechanism improves the quality of the model initialization.

As for training overhead, the inner-loop optimization of Meta significantly reduces the training overhead of meta-learning. The average training hours of Meta is 4.3, while the Meta-Static requires 8.4 training hours on average.





(a) Accuracy curve.



(b) Loss curve.

Figure 3: Comparison of the Meta methods and other baselines during personalization.

**Performance on Hardware**

We deploy the 1D-CNN on the Apollo3 Blue [Sparkfun, 2018] board to test its performance in terms of energy, latency and memory overhead on real hardware. The model requires only 36.2 KB to store its model parameters and occupies at most 5.32 KB to process the intermediate data. The average latency on one inference over a 2s segment is 368 ms and the power is 5.445 mW (supplied with 3.3 V). It turns out that the model could satisfy the performance requirements of the implantable devices on VAs detection [Strydis *et al.*, 2013].

**5 Conclusion**

In this paper, we propose a 1D-CNN based VAs detection mechanism and a novel meta-learning algorithm for patient-specific arrhythmias detection. The 1D-CNN is applied to detect VAs by inputting segments of intracardiac EGMs. The meta-learning is then utilized to generate personalized hyperparameters setting on the 1D CNN for each patient. The 1D-CNN meets the resources constraints and real-time requirements of ICDs while maintaining a high detection accuracy. The proposed patient-wise tasks formatting strategy improves the applicability of the meta-learning methods. The 1D-CNN personalized by our method achieves 2.2% higher on VAs rhythm detection rate and 8.6% higher on non-VAs

rhythm detection rate compared with the classic VAs detection method in ICDs.

**References**

[Acharya *et al.*, 2018] U Rajendra Acharya, Hamido Fujita, Shu Lih Oh, U Raghavendra, Jen Hong Tan, Muhammad Adam, Arkadiusz Gertych, and Yuki Hagiwara. Automated identification of shockable and non-shockable life-threatening ventricular arrhythmias using convolutional neural network. *Future Generation Computer Systems*, 79:952–959, 2018.

[Adabag *et al.*, 2010] A Selcuk Adabag, Russell V Luepker, Véronique L Roger, and Bernard J Gersh. Sudden cardiac death: epidemiology and risk factors. *Nature Reviews Cardiology*, 7(4):216, 2010.

[AnnArbor, 2003] AnnArbor. Ann Arbor Electrogram Libraries, Chicago IL, USA, 2003.

[Antoniou *et al.*, 2018] Antreas Antoniou, Harrison Edwards, and Amos Storkey. How to train your maml. *arXiv preprint arXiv:1810.09502*, 2018.

[Camara *et al.*, 2015] Carmen Camara, Pedro Peris-Lopez, and Juan E Tapiador. Security and privacy issues in implantable medical devices: A comprehensive survey. *Journal of biomedical informatics*, 55:272–289, 2015.

[Elola *et al.*, 2019] Andoni Elola, Elisabete Aramendi, Unai Irusta, Artzai Picón, Erik Alonso, Pamela Owens, and Ahamed Idris. Deep neural networks for ECG-based pulse detection during out-of-hospital cardiac arrest. *Entropy*, 21(3):305, 2019.

[Finn *et al.*, 2017a] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.

[Finn *et al.*, 2017b] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.

[Greenberg *et al.*, 2004] Henry Greenberg, Robert B Case, Arthur J Moss, Mary W Brown, Elizabeth R Carroll, Mark L Andrews, MADIT-II Investigators, et al. Analysis of mortality events in the multicenter automatic defibrillator implantation trial (MADIT-II). *Journal of the American College of Cardiology*, 43(8):1459–1465, 2004.

[Hannun *et al.*, 2019] Awni Y Hannun, Pranav Rajpurkar, Masoumeh Haghpanahi, Geoffrey H Tison, Codie Bourn, Mintu P Turakhia, and Andrew Y Ng. Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. *Nature Medicine*, 25(1):65, 2019.

[Hofer *et al.*, 2017] Daniel Hofer, Jan Steffel, David Hürlimann, Laurent Haegeli, Thomas F Lüscher, Firat Duru, Urs Eriksson, and Nazmi Krasniqi. Long-term incidence of inappropriate shocks in patients with implantable cardioverter defibrillators in clinical practice—An underestimated complication. *Journal of Interventional Cardiac Electrophysiology*, 50(3):219–226, 2017.

- [Li *et al.*, 2013] Qiao Li, Cadathur Rajagopalan, and Gari D Clifford. Ventricular fibrillation and tachycardia classification using a machine learning approach. *IEEE Transactions on Biomedical Engineering*, 61(6):1607–1613, 2013.
- [Madhavan and Friedman, 2013] Malini Madhavan and Paul A Friedman. Optimal programming of implantable cardiac-defibrillators. *Circulation*, 128(6):659–672, 2013.
- [Moss *et al.*, 2012] Arthur J Moss, Claudio Schuger, Christopher A Beck, Mary W Brown, David S Cannom, James P Daubert, NA Mark Estes III, Henry Greenberg, W Jackson Hall, David T Huang, et al. Reduction in inappropriate therapy and mortality through ICD programming. *New England Journal of Medicine*, 367(24):2275–2283, 2012.
- [Nichol *et al.*, 2018] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- [Senges-Becker *et al.*, 2005] Julia C Senges-Becker, Martina Klostermann, Ruediger Becker, Alexander Bauer, Karl E Siegler, Hugo A Katus, and Wolfgang Schoels. What is the “Optimal” follow-up schedule for ICD patients? *EP Europace*, 7(4):319–326, 2005.
- [Simon *et al.*, 2007] Fernando Simon, Juan Pablo Martinez, Pablo Laguna, Bart van Grinsven, Cyril Rutten, and Richard Houben. Impact of sampling rate reduction on automatic ecg delineation. In *2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 2587–2590. IEEE, 2007.
- [Sparkfun, 2018] Sparkfun. SparkFun Edge Development Board - Apollo3 Blue, 2018.
- [Strydis *et al.*, 2013] Christos Strydis, Robert M Seepers, Pedro Peris-Lopez, Dimitrios Siskos, and Ioannis Sourdis. A system architecture, processor, and communication protocol for secure implants. *ACM Transactions on Architecture and Code Optimization (TACO)*, 10(4):1–23, 2013.
- [Theuns and Gold, 2010] Dominic AMJ Theuns and Michael R Gold. Defibrillation threshold testing at implantation: Can we predict the patient with a high defibrillation threshold? *EP Europace*, 12(3):309–310, 2010.
- [Thøgersen *et al.*, 2017] Anna Margrethe Thøgersen, Jacob Moesgaard Larsen, Jens Brock Johansen, Moeen Abedin, and Charles D Swerdlow. Failure to treat life-threatening ventricular tachyarrhythmias in contemporary implantable cardioverter-defibrillators: Implications for strategic programming. *Circulation: Arrhythmia and Electrophysiology*, 10(9):e005305, 2017.
- [Zanker *et al.*, 2016] Norbert Zanker, Diane Schuster, James Gilkerson, and Kenneth Stein. Tachycardia detection in icds by Boston Scientific. *Herzschrittmachertherapie+ Elektrophysiologie*, 27(3):186–192, 2016.
- [Zdarek and Israel, 2016] Jan Zdarek and Carsten W Israel. Detection and discrimination of tachycardia in ICDs manufactured by St. Jude Medical. *Herzschrittmachertherapie+ Elektrophysiologie*, 27(3):226–239, 2016.