

# Data Augmentation for Electrocardiograms

**Aniruddh Raghu**

*Massachusetts Institute of Technology, USA*

ARAGHU@MIT.EDU

**Divya Shanmugam**

*Massachusetts Institute of Technology, USA*

DIVYAS@MIT.EDU

**Eugene Pomerantsev**

*Massachusetts General Hospital, USA*

EPOMERANTSEV@MGH.HARVARD.EDU

**John Guttag**

*Massachusetts Institute of Technology, USA*

GUTTAG@MIT.EDU

**Collin M. Stultz**

*Massachusetts Institute of Technology, USA*

CMSTULTZ@MIT.EDU

## Abstract

Neural network models have demonstrated impressive performance in predicting pathologies and outcomes from the 12-lead electrocardiogram (ECG). However, these models often need to be trained with large, labelled datasets, which are not available for many predictive tasks of interest. In this work, we perform an empirical study examining whether training time data augmentation methods can be used to improve performance on such data-scarce ECG prediction problems. We investigate how data augmentation strategies impact model performance when detecting cardiac abnormalities from the ECG. Motivated by our finding that the effectiveness of existing augmentation strategies is highly task-dependent, we introduce a new method, *TaskAug*, which defines a flexible augmentation policy that is optimized on a per-task basis. We outline an efficient learning algorithm to do so that leverages recent work in nested optimization and implicit differentiation. In experiments, considering three datasets and eight predictive tasks, we find that TaskAug is competitive with or improves on prior work, and the learned policies shed light on what transformations are most effective for different tasks. We distill key insights from our experimental evaluation, generating a set of best practices for applying data augmentation to ECG prediction problems.

**Data and Code Availability** We use three datasets: two are from Massachusetts General Hospital (MGH) and are not publicly available; the third is PTB-XL (Wagner et al., 2020), which is publicly

available on the PhysioNet repository (Goldberger et al., 2000). Code implementing our method is available here: [https://github.com/aniruddhraghu/ecg\\_aug](https://github.com/aniruddhraghu/ecg_aug).

## 1. Introduction

Electrocardiography is used widely in medicine as a non-invasive and relatively inexpensive method of measuring the electrical activity in an individual’s heart. The output of electrocardiography — the electrocardiogram (ECG) — is of great utility to clinicians in diagnosing and monitoring various cardiovascular conditions (Salerno et al., 2003; Fesmire et al., 1998; Blackburn et al., 1960).

In recent years, there has been significant interest in automatically predicting cardiac abnormalities, diseases, and outcomes directly from ECGs using neural network models (Hannun et al., 2019; Raghunath et al., 2020; Gopal et al., 2021; Diamant et al., 2021; Kiyasseh et al., 2021; Raghu et al., 2021a). Although these works demonstrate impressive results, they often require large labelled datasets with paired ECGs and labels to train models. In certain situations, it is challenging to construct such datasets. For example, consider inferring abnormal central hemodynamics (e.g., cardiac output) from the ECG, which is important when monitoring patients with heart failure or pulmonary hypertension (Schlesinger et al., 2021). Accurate hemodynamics labels are only obtainable through specialized invasive studies (Bajorath et al., 2006; Hiemstra et al., 2019) and hence it




Figure 1: **The effect of data augmentation on ECG prediction tasks is task-dependent.** We examine the mean/standard error of AUROC over 5 runs when applying SpecAugment (Park et al., 2019), a data augmentation method, to two different ECG prediction tasks. We observe performance improvement in one setting (left, Right Ventricular Hypertrophy), and performance reduction in another (right, Atrial Fibrillation).

is difficult to obtain large datasets with paired ECGs and hemodynamics variables.

Data augmentation (Hataya et al., 2020; Wen et al., 2020; Shorten and Khoshgoftaar, 2019; Iwana and Uchida, 2021a; Cubuk et al., 2019, 2020) during training is a useful strategy to improve the predictive performance of models in data-scarce regimes. However, there exists limited work studying data augmentation for ECGs. A key problem with applying standard data augmentations is that fine-grained information within ECGs, such as relative amplitudes of portions of beats, carry predictive signal: augmentations may worsen performance if such predictive signal is destroyed. Furthermore, the effectiveness of data augmentations with ECGs varies on a task-specific basis – applying the same augmentation for two different tasks could help performance in one case, and hurt performance in another (Figure 1).

In this work, we take steps towards addressing these issues. Our contributions are as follows:

- We propose *TaskAug*, a new task-dependent augmentation strategy. TaskAug defines a flexible augmentation policy that is optimized on a per-task basis. We outline an efficient learning algorithm to do so that leverages recent work in nested optimization and implicit differentiation. (Lorraine et al., 2020).
- We conduct an empirical study of TaskAug and other augmentation strategies on ECG predictive problems. We consider three datasets and eight dif-

ferent predictive tasks, which cover different classes of cardiac abnormalities.

- We analyze the results from our evaluation, finding that many augmentation strategies do not work well across all tasks. Given its task-specific nature, TaskAug is competitive with or improves on other methods for the problems we examined.
- We study the learned TaskAug policies, finding that they offer insights as to what augmentations are most appropriate for different tasks.
- We provide a summary of findings and best practices to assist future studies exploring data augmentation for ECG tasks.

## 2. Related Work

**Data augmentation for time-series.** Prior research on time-series data augmentation includes: (1) large-scale surveys exploring the impact of augmentation on various downstream modalities (Iwana and Uchida, 2021a,b; Wen et al., 2020); and (2) specific methods for particular modalities, including speech signals (Park et al., 2019, 2020), wearable device signals (Um et al., 2017), and time series forecasting (Bandara et al., 2021; Smyl and Kuber, 2016). There is relatively little work exploring how augmentation can impact performance for ECG-based prediction tasks, with prior studies mostly restricted to considering single tasks (Hatamian et al., 2020; Banerjee and Ghose, 2021). In contrast, in this paper, we evaluate a set of data augmentation methods on many different predictive tasks, studying when and why augmentations may help. In addition, the data augmentation strategy proposed in this work, *TaskAug*, can be readily adapted to new predictive tasks, unlike in existing works where the methods may be designed for a very specific downstream task.

There also exists related work on using data augmentation for contrastive pre-training with ECGs (Gopal et al., 2021; Kiyasseh et al., 2021; Raghu et al., 2021b; Mehari and Strodthoff, 2021). These works are complementary to ours; we focus specifically on supervised learning (rather than contrastive pre-training), and we hypothesize that our proposed augmentation pipeline could be used in these prior methods for improved contrastive learning.

**Designing and learning data augmentation policies.** The structure of *TaskAug*, our proposed augmentation strategy, was inspired by related work on flexible data augmentation policies in computer vi-

sion (Cubuk et al., 2019, 2020; Hataya et al., 2020). We extend these ideas to ECG predictive tasks by (1) selecting appropriate transformations for ECG data, and (2) allowing for class-specific transformation strengths. Since such policies introduce many hyperparameters, we use a bi-level optimization algorithm to enable scalable policy learning (Lorraine et al., 2020; Raghu et al., 2021c).

### 3. Problem Setup and Notation

We focus on supervised binary classification problems from ECG data. Let  $x \in \mathbb{R}^{12 \times T}$  refer to a 12-lead ECG of  $T$  samples and  $y \in \{0, 1\}$  refer to a binary target. We let  $\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N$  refer to a dataset of  $N$  ECG-label pairs.

Let  $f(x; \theta) \rightarrow \hat{y}$  be a neural network model with parameters  $\theta$  that outputs a predicted label  $\hat{y}$  given  $x$  as input. Network parameters are optimized to minimize the average binary cross entropy loss  $\mathcal{L}_{BCE}$  on the training dataset  $\mathcal{D}^{(\text{train})}$ .

We restrict our study to single label binary classification problems in this work in order to study the effect of data augmentation on a per-task basis. One can extend this to multilabel binary classification by letting  $y$  be a vector of several different binary labels and training the network to produce a vector of predictions.

**Training with Data Augmentation.** Let  $A(x, y; \phi) \rightarrow \tilde{x}$  refer to a data augmentation function with hyperparameters  $\phi$  that takes the input ECG  $x$  and its label  $y$  and outputs an augmented version  $\tilde{x}$ . Note that this formulation implicitly assumes that the augmentation is label preserving, since it does not also change the label  $y$ . Where relevant, the augmentation hyperparameters  $\phi$  may control the strength/probability of applying an augmentation.

The process of training with data augmentation<sup>1</sup> amounts to:

1. Sample a data point and label pair from the training set:  $(x, y) \sim \mathcal{D}^{(\text{train})}$ .
2. Apply the augmentation  $A : x \mapsto \tilde{x}$ , to transform the original input  $x$  to an augmented version  $\tilde{x}$ .
3. Use the pair  $(\tilde{x}, y)$  in training.

1. For the SMOTE baseline this process is slightly different; details are in Section 4.

## 4. Data Augmentation Methods

We now describe the data augmentation methods considered in our experiments. We also present our new, learnable data augmentation method that can be used to find task-specific augmentation policies, and an algorithm to optimize its parameters.

### 4.1. Existing Data Augmentation Methods

We evaluate the following set of existing data augmentation strategies, which includes operations in the signal (time-domain) space, frequency space, and interpolated signal space, providing good coverage of the possible space of augmentations.

**Time Masking.** This is a commonly used method in time-series and ECG data augmentation work (Iwana and Uchida, 2021a; Gopal et al., 2021). We mask out (set to zero) a contiguous fraction  $w \in [0, 1]$  of the original signal of length  $T$ . We choose a random starting sample  $t_s$  and set all samples  $[t_s, t_s + wT] = 0$ .

**SpecAugment.** A highly popular method for augmenting speech signals (Park et al., 2019, 2020). We follow the approach from Kiyasseh et al. (2021), and apply masking (setting components to zero) in the time and frequency domains as follows. We take the Short-Time Fourier Transform (STFT) of the input signal, and independently mask a fraction  $w$  of the temporal bins and frequency bins (this involves setting the complex valued entries in these bins to  $0 + 0j$ ). The inverse STFT is then used to map the signal back to the time domain.

**Discriminative Guided Warping (DGW).** Introduced in Iwana and Uchida (2021b), this method uses Dynamic Time Warping (DTW) (Müller, 2007; Berndt and Clifford, 1994) to warp a source ECG to match a representative reference signal that is dissimilar to examples from other classes.

**SMOTE (Chawla et al., 2002).** A commonly used oversampling strategy, the SMOTE algorithm generates new synthetic examples of the minority class by interpolating minority class samples. Given that many ECG prediction problems are characterized by significant class imbalance, oversampling algorithms are important methods to consider. In contrast to the other methods, the SMOTE algorithm generates an augmented dataset prior to any training, based on a predefined training set size, rather than augmenting examples at each training iteration (as

presented in Section 4). We set this value to achieve a balanced number of the two classes.

## 4.2. TaskAug: A New Augmentation Policy

**Motivation.** The approaches mentioned so far are simple to implement and can be effective for various problems; however, they are fairly inflexible, given each individually uses only one or two fixed transformations. With ECGs, recall that it is unclear on a per-task basis which augmentations may help or worsen performance (Figure 1). Designing a more flexible augmentation strategy that is optimized on a per-task basis could help with this problem, and we now describe such an approach – *TaskAug*.

### 4.2.1. FORMALIZING TASKAUG

**High-level structure.** We define a set of operations  $\mathcal{S} = \{A_1, \dots, A_M\}$ , each of which is an augmentation function of the form  $A_i(x, y; \mu_0, \mu_1)$ , where  $x$  is the input data point to the augmentation function,  $y$  is the label, and  $\{\mu_0, \mu_1\}$  represent the augmentation strengths for datapoints of class label 0 and class label 1 respectively. We separately parameterize the augmentation strengths for each class because transformations may corrupt predictive information in the signal for one class but not the other.

The overall augmentation policy consists of a set of  $K$  stages, where at each stage we: (1) sample an augmentation function  $A_i$  to apply; and (2) apply it to the input signal to that stage. This allows composing combinations of operations in a stochastic manner. A high-level schematic is shown in Figure 2.

**Mathematical definition.** The policy is defined following Hataya et al. (2020). At each augmentation stage  $k \in \{1, \dots, K\}$  we have a set of operation selection parameters  $\pi^{(k)} \in [0, 1]^M$ , where  $\sum_i \pi_i^{(k)} = 1 \quad \forall k$ . Each vector  $\pi^{(k)}$  parameterizes a categorical distribution such that each entry  $\pi_i^{(k)}$  represents the probability of selecting operation  $i$  at augmentation stage  $k$ . We obtain a reparameterizable sample from this categorical distribution (using the Gumbel-Softmax trick, (Jang et al., 2016; Maddison et al., 2016)) at each stage to select the operation to use, as follows:

$$u \sim \text{Categorical}(\pi^{(k)}) \quad \# \text{ Note that } u \in \mathbb{R}^M \quad (1)$$

$$i = \arg \max u \quad (2)$$

$$\tilde{x} = \frac{u_i}{\text{stop\_grad}(u_i)} A_i(x, y; \mu_0, \mu_1). \quad (3)$$

The multiplicative factor  $\frac{u_i}{\text{stop\_grad}(u_i)}$  allows differentiation w.r.t the operation selection parameters  $\pi$ . This enables gradient-based optimization of  $\pi$  (see Section 4.2.2). The denominator is necessary because the reparameterized sample from the categorical distribution is not one-hot. Further details are in Appendix A.

Suppose a particular augmentation function  $A_i$  with strength parameters  $\mu_0$  and  $\mu_1$  is obtained following Eqns 1 and 2. Then, denoting the input to this augmentation stage as  $x$  with label  $y$ , the function  $A_i$  that computes the augmented output is defined as:

$$A_i(x, y; \mu_0, \mu_1) = t_i(x; s), \quad (4)$$

where  $t_i$  is the actual transformation applied to the signal (e.g., time masking), and  $s$  is the transformation strength, computed as follows:  $s = y\mu_1 + (1 - y)\mu_0$ . See Appendix A for a detailed example of the different steps in applying TaskAug.

**Extension to multiclass and multilabel settings.** Our instantiation of TaskAug is for the binary classification setting, since this is the scenario we consider in our experiments. The formulation can be extended to multiclass/multilabel problems by defining an operation selection probability matrix and strength matrix at each augmentation stage. The operation selection probabilities and operation strengths for a given example are then obtained by taking the matrix product of the relevant parameter matrix and the label vector  $y$ .

### 4.2.2. OPTIMIZING POLICY PARAMETERS

Although the defined policy is flexible, it introduces many new parameters – for a binary problem, there are  $M$  operation selection parameters for the categorical distributions at each stage, and 2 strength parameters at each stage, resulting in  $K \times (2 + M)$  total parameters. Finding effective values for these parameters with random/grid search or Bayesian optimization is computationally expensive since they require training models many times with different parameter settings. We therefore use a gradient-based learning scheme to learn these parameters online.

We optimize policy parameters to minimize a model’s validation loss, which is computed using non-augmented data. Following prior work (Lorraine et al., 2020; Hataya et al., 2020; Raghu et al., 2021c), we alternate gradient updates on the network parameters  $\theta$  and the augmentation parameters  $\phi$  by iter-

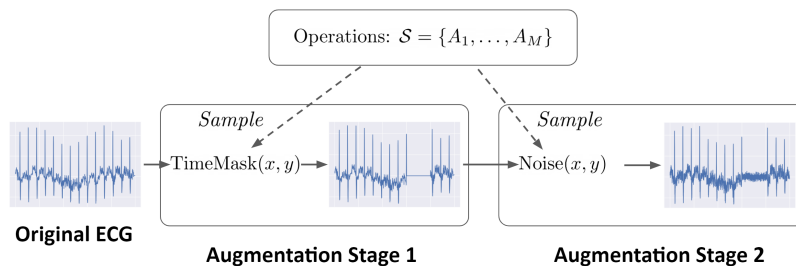


Figure 2: **Structure of TaskAug.** Augmentations to apply are sampled from a set of available operations, and applied in sequence. Here we show an example with  $K = 2$  stages of augmentation. We omit details relating to the per-class magnitudes and probabilities of sampling for clarity.

ating the following steps (details and full algorithm in Appendix A):

- Optimize the model parameters  $\theta$  for  $P$  steps: at each step, sample a batch  $(x, y)$  of data from  $\mathcal{D}^{(\text{train})}$ , augment the batch with the augmentation policy to obtain  $(\tilde{x}, y)$ , compute the predicted label  $\hat{y}$ , and update the model parameters using gradient descent:  $\theta \leftarrow \theta - \eta \nabla \mathcal{L}(y, \hat{y})$ .
- Compute the validation loss  $\mathcal{L}_V$  using an un-augmented batch from the validation dataset.
- Perform a gradient update on the augmentation parameters  $\phi$ . We use the chain rule to re-express the gradient wrt the augmentation parameters:

$$\frac{\partial \mathcal{L}_V}{\partial \phi} = \frac{\partial \mathcal{L}_V}{\partial \theta} \times \frac{\partial \theta}{\partial \phi},$$

and compute this as follows. The first term on the RHS is found exactly using straightforward back-propagation; the second term is approximated using the algorithm from Lorraine et al. (2020), leveraging implicit differentiation for efficient computation (since differentiating through training exactly is too memory-intensive). The augmentation parameters are then updated:  $\phi \leftarrow \phi - \eta \frac{\partial \mathcal{L}_V}{\partial \phi}$ .

By using this algorithm, augmentation parameters are learned on a per-task basis and analyzing the learned parameters may allow us to understand what augmentations are useful for different problems. We return to this in Section 5.2.2.

**Computational cost.** Optimizing policy parameters in this manner is significantly more computationally efficient than running a grid search over parameter values. With  $P = 1$ , running this algorithm has about  $2 - 3 \times$  the computational cost of training without any augmentations.

## 5. Experiments

We evaluate the data augmentation strategies on ECG prediction tasks. We have two main experimental questions: (1) in what settings can data augmentation be beneficial, and (2) when data augmentation does help, which augmentation strategies are most effective? To investigate these questions, we consider a range of settings that cover three different 12-lead ECG datasets and eight prediction tasks of varying difficulty, class imbalance, and training set sizes.

### 5.1. Experimental Setup

#### 5.1.1. DATASETS AND TASKS

We highlight key information about our datasets and tasks here, with a summary in Table 1.

**Dataset A** is from Massachusetts General Hospital (MGH) and contains paired 12-lead ECGs and labels for different cardiac abnormalities. Of the available labels in the dataset, we select Right Ventricular Hypertrophy (RVH) and Atrial Fibrillation (AFib) as two of the predictive tasks in our evaluation. These were chosen because (1) they have been previously studied as prediction targets from the ECG (Couceiro et al., 2008; Lin and Lu, 2020), and (2) they have low positive prevalence: 1% for RVH, and 5% for AFib, and therefore help to understand the impact of data augmentation in imbalanced prediction problems.

**Dataset B** is PTB-XL (Wagner et al., 2020; Goldberger et al., 2000), an open-source dataset of 12-lead ECGs. Each ECG has labels for four different categories of cardiac abnormality. This dataset has been used in prior work to evaluate ECG predictive models (Gopal et al., 2021; Kiyasseh et al., 2021).

Dataset	Task name	Prevalence	Abnormality type	#ecgs/#patients
Dataset A	Right Ventricular Hypertrophy (RVH)	1%	Structural	705057/705057
	Atrial Fibrillation (AFib)	5%	Electrical	705057/705057
Dataset B (PTB-XL)	Hypertrophy (HYP)	12%	Structural	21837/18885
	ST/T Change (STTC)	22%	Ischemia	21837/18885
	Conduction Disturbance (CD)	24%	Electrical	21837/18885
	Myocardial Infarction (MI)	25%	Ischemia	21837/18885
Dataset C	Low Cardiac Output (CO)	4%	Hemodynamics	6290/4051
	High Pulmonary Capillary Wedge Pressure (PCWP)	26%	Hemodynamics	6290/4051

Table 1: Summary information about the datasets and tasks considered in our empirical evaluation.

**Dataset C** is from the same hospital (MGH) as Dataset A and contains paired ECGs and labels for two hemodynamics parameters, Cardiac Output (CO) and Pulmonary Capillary Wedge Pressure (PCWP). These measures of cardiac health are important in deciding treatment strategies for patients with cardiac disease (Yancy et al., 2013; Hurst et al., 1990; Solin et al., 1999). Typically, these parameters can only be measured accurately through an invasive cardiac catheterization procedure (Bajorat et al., 2006; Hiemstra et al., 2019). As a result, datasets with paired ECGs and hemodynamics measurements are relatively small. Considering the use of data augmentations to improve model performance in this limited data regime is therefore clinically relevant. We specifically consider inferring abnormally low Cardiac Output, and abnormally high Pulmonary Capillary Wedge Pressure.

Note that the tasks considered cover different classes of cardiac abnormalities: ischemia (MI, STTC), structural (HYP, RVH), electrical (CD, AFib), and abnormal hemodynamics (low CO, high PCWP).

**Dataset splitting.** Since the value of data augmentation can depend on the amount of training data, we train on different dataset sizes. For the non-hemodynamic tasks (Datasets A and B), we generate development datasets with 1000, 2500, and 5000 ECGs. On the more challenging hemodynamics inference tasks (Dataset C), for elevated PCWP, we consider two settings: using a development set of size 1000, and using the full dataset. For low CO, we only use the full dataset, since reducing the dataset size led to poor quality models.

In each setting, we split datasets into development and testing sets on a patient-level (no patient is in

both sets). We split the development set into an 80-20 training-validation split.

### 5.1.2. TASKAUG TRANSFORMATIONS

Based on prior work in time series and ECG data augmentation (Iwana and Uchida, 2021a; Mehari and Strodthoff, 2021) we use the following transformations in the TaskAug policy. Mathematical descriptions are in Appendix A.

- **Random temporal warp:** The signal is warped with a random, diffeomorphic temporal transformation. This is formed by sampling from a zero mean, fixed variance Gaussian at each temporal location in the signal to obtain a velocity field, and then integrating and smoothing (following Balakrishnan et al. (2018, 2019)) to generate a temporal displacement field, which is applied to the signal. The variance is the strength parameter, with higher variance indicating more warping.
- **Baseline wander:** A low-frequency sinusoidal component is added to the signal, with the amplitude of the sinusoid representing the strength.
- **Gaussian noise:** IID Gaussian noise is added to the signal, with the strength parameter representing the variance of the Gaussian.
- **Magnitude scale:** The signal amplitude is scaled by a number drawn from a scaled uniform distribution, with the scale being the strength parameter.
- **Time mask:** A random contiguous section of the signal is masked out (set to zero).
- **Random temporal displacement:** The entire signal is translated forwards or backwards in time by a random temporal offset, drawn from a uniform distribution scaled by a strength parameter.

Note that our instantiation of the augmentation policy could utilize many more operations, but we keep

it to this number for simplicity and to assist in interpreting the learned policies.

### 5.1.3. IMPLEMENTATION DETAILS

**Network architecture.** We standardize the network architecture to be a 1D convolutional network, based on the ResNet-18 architecture, since prior work has shown architectures of this form to be effective with ECG data (Diamant et al., 2021). Full architectural details are in the appendix.

**Training Details.** On Datasets A and B, all models are trained for 100 epochs, using early stopping based on validation loss. For the hemodynamics inference problems on Dataset C, we train models for 50 epochs with early stopping (since we observed significant overfitting after this point). We consider 15 random development/testing set splits for Datasets A and C (lower prevalences for some tasks meant that performance was more variable with fewer runs), and 5 splits for Dataset B. We train models using the Adam optimizer and a learning rate of  $1e-3$ . This value resulted in stable and effective training across all models (as compared to  $1e-4$ ,  $5e-4$ , and  $5e-3$ ). As evaluation, we compute the AUROC of the best performing model on the held-out testing set, and report mean/standard error across runs. We also report results for a baseline (NoAugs) that does not use any data augmentation.

**Augmentation Hyperparameters.** In TaskAug, we set the number of augmentation stages to  $K = 2$  (defined in Section 4.2.1), following prior work (Hataya et al., 2020). For the number of model optimization steps  $P$  (defined in Section 4.2.2), we evaluate both  $P = 1$  and  $P = 5$ , and select the best performing setting based on validation set loss. Further discussion on the choice of  $P$  is in Appendix A.

For Time Masking and SpecAugment, we search over the masking window, considering  $w \in \{0.1, 0.2\}$  for SpecAugment (range based on Kiyasseh et al. (2021)) and  $w \in \{0.1, 0.2, 0.5\}$  for Time Masking (range based on Gopal et al. (2021)).

## 5.2. Results

### 5.2.1. QUANTITATIVE RESULTS

**Non-hemodynamics tasks.** We first analyze performance of augmentation strategies on the non-hemodynamics tasks. Given that performance improvements are most evident in the lowest sample

regimes for both datasets ( $N = 1000$ ), we focus on this setting with results shown in Table 2. Results for the higher sample regimes are in the appendix. We summarize key findings here.

*The value of augmentation varies by task.* For some tasks such as RVH and MI, almost all augmentation strategies lead to performance improvements. On other tasks such as STTC and HYP, performance is the same or worse when applying augmentations. The improvement seen with RVH could be due to the fact that it is particularly low prevalence (1%), so all augmentation strategies have an oversampling effect and thus boost performance.

*TaskAug performs well on average.* TaskAug almost always improves on the NoAugs baseline, and even boosts performance on some tasks where other augmentations worsen performance (AFib). Although TaskAug does not always result in a statistically significant ( $p < 0.05$ ) improvement in AUROC, it is the only method to significantly improve AUPRC over NoAugs on the low-prevalence tasks, RVH and AFib (see Appendix C, Table 4).

When TaskAug results in lower performance than other augmentation strategies (e.g., TimeMasking for CD), it is still competitive with these methods and never causes a statistically significant reduction in performance compared to other methods. This suggests that for a new task, it may always be worth using TaskAug to see if performance is boosted. We hypothesise that TaskAug’s efficacy is due to its flexible and learned nature, examined in ablation studies (Section 5.2.3).

*Performance improvements are smaller on Dataset B.* The maximum improvement over the NoAugs baseline in Dataset A (5.8%) is greater than the maximum improvement in Dataset B (2.3%). We hypothesise two reasons for this. Firstly, the prevalence in Dataset B is higher, meaning that augmentations may not have as much of an effect at  $N = 1000$ . We study this in Appendix C, Table 11, where we examine performance at the  $N = 500$  data regime for Dataset B, and find that the maximum improvement (obtained with TaskAug for MI) goes up to 4%.

Secondly, Dataset A has narrower label definitions than Dataset B, and this affects performance, especially with TaskAug. The HYP, STTC, and CD classes of abnormalities in Dataset B aggregate many sub-categories together (see Appendix B), and these sub-categories may each benefit from different augmentations. In contrast, the labels in Dataset A are fine-grained, and so TaskAug, which optimizes aug-

	Dataset A		Dataset B			
	RVH	AFib	MI	HYP	STTC	CD
NoAugs	72.6 ± 2.7	<u>79.8 ± 1.4</u>	80.0 ± 0.8	<b>84.3 ± 1.4</b>	87.6 ± 0.8	82.2 ± 0.6
TaskAug	<b>78.4 ± 1.9</b>	<b>82.8 ± 1.0</b>	<b>82.3 ± 0.5*</b>	83.7 ± 0.5	<b>87.8 ± 0.4</b>	<u>83.1 ± 0.4</u>
SMOTE	75.9 ± 1.8	79.0 ± 1.4	<u>81.2 ± 0.6</u>	80.4 ± 0.6	87.0 ± 0.5	82.6 ± 0.8
DGW	73.6 ± 1.7	77.4 ± 1.5	81.1 ± 0.6	<u>83.9 ± 0.7</u>	87.5 ± 0.5	81.8 ± 1.0
SpecAug	<u>77.9 ± 1.7</u>	77.2 ± 2.1	81.1 ± 0.7	<u>83.5 ± 0.8</u>	<u>87.7 ± 0.4</u>	82.2 ± 0.7
TimeMask	<u>72.8 ± 2.1</u>	77.9 ± 1.9	81.1 ± 1.3	82.9 ± 0.7	<u>87.7 ± 0.7</u>	<b>83.8 ± 1.1</b>

Table 2: **Augmentation strategies improve AUROC on detecting most cardiac abnormalities in the low-sample regime ( $N = 1000$ ), and TaskAug is among the best-performing methods.** Table shows mean and standard error of AUROC (best-performing method bolded, second best underlined, statistically significant ( $p < 0.05$ ) improvement over NoAugs marked \*). The impact of augmentations is task-dependent, with some tasks (such as RVH, MI) showing improved performance on average with almost all strategies, and others (HYP) showing no improvement with any strategy. TaskAug is among the best methods across tasks, and improves performance on tasks such as AFib where no other augmentations help.

	Dataset C		
	Low CO	High PCWP: $N = 1000$	High PCWP: All Data
NoAugs	65.9 ± 1.2	66.7 ± 0.7	74.4 ± 0.5
TaskAug	<u>68.2 ± 1.0</u>	<b>67.9 ± 0.7</b>	<b>75.1 ± 0.4</b>
SMOTE	66.0 ± 1.4	67.2 ± 0.5	73.6 ± 0.5
DGW	<b>68.3 ± 0.9</b>	66.4 ± 0.6	74.9 ± 0.4
SpecAug	66.1 ± 0.9	66.4 ± 1.3	<u>75.0 ± 0.4</u>
TimeMask	66.8 ± 1.1	<u>67.3 ± 0.4</u>	74.6 ± 0.4

Table 3: **Training with data augmentation improves AUROC on two hemodynamics inference tasks, and TaskAug again is among the best-performing methods.** Table shows mean and standard error of AUROC (best-performing method bolded, second best underlined). All methods are comparable with or improve on the no augmentation baseline for Low CO prediction, possibly because of the low prevalence of the label (4%). The performance of methods on the High PCWP task is more variable across the two sample sizes. TaskAug obtains improvements in all three settings considered.

mentations on a per-task basis, learns more appropriate augmentation strategies. This hypothesis is supported by the fact that with MI (a more fine-grained label than HYP, CD, and STTC) we observe improvements over the NoAugs baseline (clearly seen in the  $N = 500$  regime, Appendix C, Table 11).

*Performance improvements at higher samples are lower*, as seen in the results in Appendix C. Augmentations do not worsen performance however, and some tasks (STTC, CD) benefit a small amount,  $\sim +1\%$  AUROC.

**Hemodynamics tasks.** Table 3 presents results for performance on the more challenging hemodynamics prediction tasks. All methods are comparable with or improve on the no augmentation baseline for low CO prediction, likely because of the low

prevalence of the positive label (4%). For inferring high PCWP, at both low sample and higher samples, TaskAug obtains improvements in performance (though not significant at the  $p < 0.05$  level); however, other methods do not consistently improve on the no augmentation baseline. Although improvements in AUROC are not statistically significant, we observe significant improvements with TaskAug in AUPRC for low CO detection (see Appendix C, Table 6). Again, we see that the benefit of augmentation varies with the task, prevalence, and dataset size, and that TaskAug is better than or competitive with other strategies.






Figure 3: **The TaskAug policy for Atrial Fibrillation detection.** We focus on the probability of selecting each transformation in both augmentation stages (left) and the optimized temporal warp strengths in the first stage (right). We show the mean/standard error of these optimized policy parameters over 15 runs. Given the characteristic features of AFib (e.g., irregular R-R interval), Time Masking is likely to be label preserving and therefore it is sensible that it has a high probability of selection. The temporal warp strength for positive samples is higher than that for negative samples, which makes sense since time warping a negative sample too strongly could change its label.




Figure 4: **The TaskAug policy for detecting elevated Pulmonary Capillary Wedge Pressure.** We focus on the probability of selecting each transformation in both augmentation stages (left) and the optimized magnitude scaling strengths in the second stage (right). We show the mean/standard error of these optimized policy parameters over 15 runs. There exists little domain knowledge about what features in the ECG may encode elevated PCWP, so examining the learned augmentations here could provide hypotheses of invariances in the data. Of interest is that the positive class is augmented with stronger magnitude scaling than the negative class, suggesting that scaling negative examples could affect their labels.

### 5.2.2. ANALYZING LEARNED POLICIES

We analyze the learned policies for three of the predictive tasks: AFib, PCWP, and RVH (appendix).

**AFib, Figure 3.** We see that time mask has a high probability of selection (Figure 3(a)). Since AFib is characterized in the ECG by an irregular R peak-R peak interval (Couceiro et al., 2008), which is often present regardless of which section of ECG is selected, time masking is likely label preserving, and is

a sensible choice. Considering the learned time warp strength in Figure 3(b), we observe that signals labelled negative for AFib are warped less strongly than those with AFib, again sensible since time warping may affect the label of a signal and introduce AFib in a signal where it was not originally present.

**PCWP, Figure 4.** We have limited domain understanding of what augmentations may be label preserving and help model performance, since detecting high PCWP from ECGs is not something clinicians




Figure 5: **Optimizing the TaskAug policy parameters results in performance improvements.** We show the mean/standard error of AUROC over 15 runs for AFib and over 5 runs for MI. Without optimizing policy parameters (InitAug), performance is comparable to not using augmentations at all, indicating the importance of learning the policy parameters.

are typically able to do (Schlesinger et al., 2021). Analyzing the augmentations could provide hypotheses about what features in the data encode the class label. Noise, displacement, and baseline wander all obtain higher weight in the first stage, and scaling obtains higher weight in the second stage. The high weight assigned to noise could be to help the model build invariance to it, and not use it as a predictive aspect of the signal. Studying the magnitude scaling in Figure 4(b), we see positive examples are scaled significantly more than negative examples. It is possible that negative examples are more sensitive to scale, and scaling them pushes them into positive example space. The positive examples may have more variance in scaling, and thus scaling them further has less of an effect.

### 5.2.3. ABLATION STUDIES

**How much does optimizing augmentations help?** Our results show that TaskAug offers improvements in performance. In Figure 5, we examine how the actual optimization of the augmentation policy parameters (operation selection probabilities and magnitudes, Section 4.2.1) affects performance, considering the AFib and MI detection tasks and  $N = 1000$ . We compare the performance of optimizing the policy parameters vs. keeping them fixed at their initialized values and training. We observe improvements in performance through the optimization process, suggesting that it is not only the range of augmentations that leads to improved performance, but also the optimization of the policy parameters. In




Figure 6: **Class-specific magnitude parameters in TaskAug lead to improvements in performance.** We show the mean/standard error of AUROC over 15 runs for AFib and over 5 runs for MI. This is particularly true for tasks such as AFib where some operations may not be label preserving.

Appendix C, we study this at different dataset sizes and find that performance is improved by optimization at each size.

### How much do class-specific magnitudes help?

TaskAug instantiates magnitude parameters for the augmentation operations on a per-class basis, as described in Section 4.2.1, allowing positive and negative examples to be augmented differently. We examine this further, considering the AFib and MI detection tasks and  $N = 1000$ . We compare performance using class-specific magnitude parameters (the positive and negative examples have independent augmentation magnitudes  $\mu_1$  and  $\mu_0$ ) vs. using global magnitude parameters (the positive and negative examples are forced to have the same augmentation magnitude:  $\mu = \mu_0 = \mu_1$ ). Results are shown in Figure 6. We observe noticeable improvements in performance with class-specific magnitude parameters, demonstrating the importance of independently specifying magnitudes for the two classes. In Appendix C, we study this at different dataset sizes and find that performance is improved at each size.

### 5.2.4. SUMMARY AND BEST PRACTICES

- Training with data augmentations does not always improve model performance, and may even hurt it. The impact of augmentation depends on nature of the task, positive class prevalence, and dataset size.
- Augmentations are most often useful in the low-sample regime. Where the prevalence is particularly low (see results for RVH detection) various

- augmentation strategies improve performance, perhaps by functioning as a form of oversampling.
- Data augmentations do not always improve performance at high sample sizes, but do not hurt it.
  - TaskAug, our proposed augmentation strategy, is the most effective method on average, and could therefore be the first augmentation strategy one tries on a new ECG prediction problem. TaskAug defines a flexible augmentation policy that is optimized on a task-dependent basis, which directly contributes to its effectiveness.
  - TaskAug also offers insights as to what augmentations are most effective for a given problem, which could be useful in novel prediction tasks (e.g., hemodynamics inference) to suggest what aspects of the ECG determine the class label.

## 6. Conclusion

In this work, we studied the use of data augmentation for prediction problems from 12-lead electrocardiograms (ECGs). We outlined *TaskAug*, a new, learnable data-augmentation strategy for ECGs, and conducted an empirical study of this method and several existing augmentation strategies.

In our experimental evaluation on three ECG datasets and eight distinct predictive tasks, we find that data augmentation is not always helpful for ECG prediction problems, and for some tasks may worsen performance. Augmentations can be most helpful in the low-sample regime, and specifically when the prevalence of the positive class is low. Our proposed learnable augmentation strategy, TaskAug, was among the strongest performing methods in all tasks. TaskAug augmentation policies are additionally interpretable, providing insight as to what transformations are most important for different problems. Future work could consider applying TaskAug to other settings (e.g., multiview contrastive learning) and modalities (e.g., EEGs) where flexible augmentation policies may be useful and could be interpreted to provide scientific insight.

## Institutional Review Board (IRB)

This study was approved by the Institutional Review Board (IRB) at Massachusetts General Hospital (protocol 2020P000132).

## References

- J Bajorat, R Hofmockel, DA Vagts, M Janda, B Pohl, C Beck, and G Noeldge-Schomburg. Comparison of invasive and less-invasive techniques of cardiac output measurement under different haemodynamic conditions in a pig model. *European journal of anaesthesiology*, 23(1):23–30, 2006.
- Guha Balakrishnan, Amy Zhao, Mert Sabuncu, John Guttag, and Adrian V. Dalca. An unsupervised learning model for deformable medical image registration. *CVPR: Computer Vision and Pattern Recognition*, pages 9252–9260, 2018.
- Guha Balakrishnan, Amy Zhao, Mert Sabuncu, John Guttag, and Adrian V. Dalca. Voxelmorph: A learning framework for deformable medical image registration. *IEEE TMI: Transactions on Medical Imaging*, 38:1788–1800, 2019.
- Kasun Bandara, Hansika Hewamalage, Yuan-Hao Liu, Yanfei Kang, and Christoph Bergmeir. Improving the accuracy of global forecasting models using time series data augmentation. *Pattern Recognition*, 120:108148, 2021.
- Rohan Banerjee and Avik Ghose. Synthesis of realistic eeg waveforms using a composite generative adversarial network for classification of atrial fibrillation. In *2021 29th European Signal Processing Conference (EUSIPCO)*, pages 1145–1149, 2021. doi: 10.23919/EUSIPCO54536.2021.9616079.
- Donald J Berndt and James Clifford. Using dynamic time warping to find patterns in time series. In *KDD workshop*, volume 10, pages 359–370. Seattle, WA, USA:, 1994.
- Henry Blackburn, Ancel Keys, Ernst Simonson, Pentti Rautaharju, and Sven Punsar. The electrocardiogram in population studies: a classification system. *Circulation*, 21(6):1160–1175, 1960.
- Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- Ricardo Couceiro, Paulo Carvalho, Jorge Henriques, Manuel Antunes, Matthew Harris, and Jörg Habetha. Detection of atrial fibrillation using model-based eeg analysis. In *2008 19th International Conference on Pattern Recognition*, pages 1–5. IEEE, 2008.
- Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 113–123, 2019.
- Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 702–703, 2020.
- Nathaniel Diamant, Erik Reinertsen, Steven Song, Aaron Aguirre, Collin Stultz, and Puneet Batra. Patient contrastive learning: a performant, expressive, and practical approach to eeg modeling. 2021.
- Francis M Fesmire, Robert F Percy, Jim B Bardoner, David R Wharton, and Frank B Calhoun. Usefulness of automated serial 12-lead eeg monitoring during the initial emergency department evaluation of patients with chest pain. *Annals of emergency medicine*, 31(1):3–11, 1998.
- A. Goldberger, L. A. Amaral, L. Glass, Jeffrey M. Hausdorff, P. Ivanov, R. Mark, J. Mietus, G. Moody, C. Peng, and H. Stanley. PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation*, 101 23:E215–20, 2000.
- Bryan Gopal, Ryan W. Han, Gautham Raghupathi, Andrew Y. Ng, Geoffrey H. Tison, and Pranav Rajpurkar. 3kg: Contrastive learning of 12-lead electrocardiograms using physiologically-inspired augmentations. 2021.
- 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):65–69, 2019.

- Faezeh Nejati Hatamian, Nishant Ravikumar, Sulaiman Vesal, Felix P Kemeth, Matthias Struck, and Andreas 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)*, pages 1264–1268. IEEE, 2020.
- Ryuichiro Hataya, Jan Zdenek, Kazuki Yoshizoe, and Hideki Nakayama. Meta approach to data augmentation optimization. *arXiv preprint arXiv:2006.07965*, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- Bart Hiemstra, Geert Koster, Renske Wiersema, Yoran M Hummel, Pim van der Harst, Harold Snieder, Ruben J Eck, Thomas Kaufmann, Thomas WL Scheeren, Anders Perner, et al. The diagnostic accuracy of clinical examination for estimating cardiac index in critically ill patients: the simple intensive care studies-i. *Intensive care medicine*, 45(2):190–200, 2019.
- J Hurst, C Rackley, E Sonnenblick, and N Wenger. *The heart, arteries and veins*, volume 1. McGraw-Hill, 1990.
- Brian Kenji Iwana and Seiichi Uchida. An empirical survey of data augmentation for time series classification with neural networks. *Plos one*, 16(7): e0254841, 2021a.
- Brian Kenji Iwana and Seiichi Uchida. Time series data augmentation for neural networks by time warping with a discriminative teacher. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 3558–3565. IEEE, 2021b.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Dani Kiyasseh, Tingting Zhu, and David A Clifton. Clocs: Contrastive learning of cardiac signals across space, time, and patients. In *International Conference on Machine Learning*, pages 5606–5615. PMLR, 2021.
- Gen-Min Lin and Henry Horng-Shing Lu. A 12-lead ecg-based system with physiological parameters and machine learning to identify right ventricular hypertrophy in young adults. *IEEE Journal of Translational Engineering in Health and Medicine*, 8:1–10, 2020. doi: 10.1109/JTEHM.2020.2996370.
- Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics*, pages 1540–1552. PMLR, 2020.
- Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Temesgen Mehari and Nils Strodthoff. Self-supervised representation learning from 12-lead ecg data. *arXiv preprint 2103.12676*, 2021.
- Meinard Müller. Dynamic time warping. *Information retrieval for music and motion*, pages 69–84, 2007.
- Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. Specaugment: A simple data augmentation method for automatic speech recognition. *arXiv preprint arXiv:1904.08779*, 2019.
- Daniel S Park, Yu Zhang, Chung-Cheng Chiu, Youzheng Chen, Bo Li, William Chan, Quoc V Le, and Yonghui Wu. Specaugment on large scale datasets. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6879–6883. IEEE, 2020.
- Aniruddh Raghu, John Guttag, Katherine Young, Eugene Pomerantsev, Adrian V Dalca, and Collin M Stultz. Learning to predict with supporting evidence: applications to clinical risk prediction. In *Proceedings of the Conference on Health, Inference, and Learning*, pages 95–104, 2021a.
- Aniruddh Raghu, Jonathan Lorraine, Simon Kornblith, Matthew McDermott, and David K Duvenaud. Meta-learning to improve pre-training. *Advances in Neural Information Processing Systems*, 34, 2021b.
- Aniruddh Raghu, Maithra Raghu, Simon Kornblith, David Duvenaud, and Geoffrey Hinton. Teaching

- with commentaries. In *International Conference on Learning Representations*, 2021c.
- Sushravya Raghunath, Alvaro E Ulloa Cerna, Linyuan Jing, Joshua Stough, Dustin N Hartzel, Joseph B Leader, H Lester Kirchner, Martin C Stumpe, Ashraf Hafez, Arun Nemani, et al. Prediction of mortality from 12-lead electrocardiogram voltage data using a deep neural network. *Nature medicine*, 26(6):886–891, 2020.
- Stephen M Salerno, Patrick C Alguire, and Herbert S Waxman. Competency in interpretation of 12-lead electrocardiograms: a summary and appraisal of published evidence. *Annals of Internal Medicine*, 138(9):751–760, 2003.
- Daphne Schlesinger, Nathaniel Diamant, Aniruddh Raghun, Erik Reinertsen, Katherine Young, Puneet Batra, Eugene Pomerantsev, and Collin M. Stultz. A deep learning model for inferring elevated pulmonary capillary wedge pressures from the 12-lead electrocardiogram. 2021.
- Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):1–48, 2019.
- Slawek Smyl and Karthik Kuber. Data preprocessing and augmentation for multiple short time series forecasting with recurrent neural networks. In *36th International Symposium on Forecasting*, 2016.
- Peter Solin, Peter Bergin, Meroula Richardson, David M Kaye, E Haydn Walters, and Matthew T Naughton. Influence of pulmonary capillary wedge pressure on central apnea in heart failure. *Circulation*, 99(12):1574–1579, 1999.
- Terry T Um, Franz MJ Pfister, Daniel Pichler, Satoshi Endo, Muriel Lang, Sandra Hirche, Urban Fietzek, and Dana Kulić. Data augmentation of wearable sensor data for parkinson’s disease monitoring using convolutional neural networks. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, pages 216–220, 2017.
- Patrick Wagner, Nils Strodthoff, Ralf-Dieter Bouseljt, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. PTB-XL, a large publicly available electrocardiography dataset. *Scientific data*, 7(1):1–15, 2020.
- Qingsong Wen, Liang Sun, Fan Yang, Xiaomin Song, Jingkun Gao, Xue Wang, and Huan Xu. Time series data augmentation for deep learning: A survey. *arXiv preprint arXiv:2002.12478*, 2020.
- Clyde W Yancy, Mariell Jessup, Biykem Bozkurt, Javed Butler, Donald E Casey, Mark H Drazner, Gregg C Fonarow, Stephen A Geraci, Tamara Horwich, James L Januzzi, et al. 2013 accf/aha guideline for the management of heart failure: executive summary: a report of the american college of cardiology foundation/american heart association task force on practice guidelines. *Journal of the American College of Cardiology*, 62(16):1495–1539, 2013.

## Appendix A. Augmentation Methods

In this section, we provide further details on the different augmentation strategies explored (existing and TaskAug), and visualize their operation.

### A.1. Existing methods

Figures 7-10 present examples following augmentation using the existing methods. We show only one lead for clarity; however, these operations will be applied to each lead.




Figure 7: Time Masking.




Figure 8: SpecAugment.




Figure 9: Discriminative Guided Warping (DGW).

### A.2. TaskAug

We provide more details about TaskAug: (1) further information about the mathematical formalism of the




Figure 10: SMOTE.

policy and an example of applying the different steps; (2) A more detailed description of the nested optimization algorithm used to learn TaskAug parameters, including a full algorithm; and (3) mathematical descriptions of the operations used in TaskAug in our experiments and a visualization of their effect on an ECG signal.

#### A.2.1. STRUCTURE OF POLICY

**Mathematical definition.** As described in Section 4.2.1, the TaskAug policy is defined following Hataya et al. (2020). At each augmentation stage  $k \in \{1, \dots, K\}$  we have a set of operation selection parameters  $\pi^{(k)} \in [0, 1]^M$ , where  $\sum_i \pi_i^{(k)} = 1 \quad \forall k$ . Each vector  $\pi^{(k)}$  parameterizes a categorical distribution such that each entry  $\pi_i^{(k)}$  represents the probability of selecting operation  $i$  at augmentation stage  $k$ . We obtain a reparameterizable sample from this categorical distribution (using the Gumbel-Softmax trick, (Jang et al., 2016; Maddison et al., 2016)) at each stage to select the operation to use, as follows:

$$u \sim \text{Categorical}(\pi^{(k)}) \quad \# \text{ Note that } u \in \mathbb{R}^M \quad (5)$$

$$i = \arg \max u \quad (6)$$

$$\tilde{x} = \frac{u_i}{\text{stop\_grad}(u_i)} A_i(x, y; \mu_0, \mu_1). \quad (7)$$

**Why the multiplicative factor?** We use the multiplicative factor  $\frac{u_i}{\text{stop\_grad}(u_i)}$  to allow gradient flow to the operation selection parameters  $\pi$ . If we just selected  $i = \arg \max u$  and had no scaling in Eqn 7, then there would be no gradient flow to  $\pi$ , since the  $\arg \max$  operation is not differentiable.

The denominator of this scaling factor is necessary because  $u_i$ , obtained from the reparameterized sample from the categorical distribution, is not one-hot. The resulting fraction used as the scaling factor always has magnitude 1, since  $|\text{stop\_grad}(u_i)| = |u_i|$ .

When we take the gradient, we get:

$$\frac{\partial}{\partial \pi} \frac{u_i}{\text{stop\_grad}(u_i)} = \frac{1}{\text{stop\_grad}(u_i)} \frac{\partial u_i}{\partial \pi},$$

so the `stop_grad`( $u_i$ ) acts as a scaling term.

**Example application of TaskAug.** Suppose we have a one-stage TaskAug policy,  $K = 1$ , our augmentation set has two operations  $\mathcal{S} = \{A_1, A_2\}$  which are  $A_1 = \text{TimeMask}(x, y; \mu_0 = 0.2, \mu_1 = 0.1)$  and  $A_2 = \text{Noise}(x, y; \mu_0 = 2.1, \mu_1 = 5.3)$ , and the operation selection probability vector is  $\pi = [0.9, 0.1]$  (that is, we select TimeMask with probability 0.9, and noise with probability 0.1). Now consider applying TaskAug to a (data, label) pair  $(x, 1)$ , i.e., the label is 1. We follow these steps:

1. Obtain a reparameterizable sample  $u$  from `Categorical`([0.9, 0.1]): let this be  $u = [0.75, 0.25]$ .
2. Find  $i = \arg \max u$ ; in this case,  $i = 1$ .
3. Select the operation  $A_1$ , i.e. TimeMask.
4. Compute the masking strength based on the label. Recall this is defined as  $s = y\mu_1 + (1 - y)\mu_0$ , so  $s = 1 \times 0.1 + (1 - 1) \times 0.2 = 0.1$ .
5. Apply time-masking with strength 0.1 to  $x$ , generating  $\hat{x}$ .
6. Scale this by  $\frac{u_1}{\text{stop\_grad}(u_1)}$  to generate  $\tilde{x}$ .

### A.2.2. PARAMETER OPTIMIZATION

As detailed in the main text, there are many learnable parameters in TaskAug, and we use gradient-based optimization to learn these jointly with the base model parameters. Here, we provide some more details about the estimation of the gradient wrt the TaskAug parameters, and also include a full algorithm detailing the training procedure, Algorithm 1.

#### Estimating TaskAug parameter gradients.

Let the base model parameters after  $P$  update steps be denoted as  $\hat{\theta}(\phi)$ . We update the TaskAug policy parameters to minimize the base model’s validation loss  $\mathcal{L}_V$ , with the gradient of interest being:

$$\frac{\partial \mathcal{L}_V}{\partial \phi} = \frac{\partial \mathcal{L}_V}{\partial \hat{\theta}} \times \frac{\partial \hat{\theta}}{\partial \phi}.$$

The first term on the RHS can be found exactly using standard backpropagation. To compute the second term, we re-express it using the implicit function theorem (IFT) as in Lorraine et al. (2020). Using  $\mathcal{L}_T$  to denote the training loss, the IFT allows us to

re-express this second term as:

$$\frac{\partial \hat{\theta}}{\partial \phi} = - \left[ \frac{\partial^2 \mathcal{L}_T}{\partial \theta \partial \theta^T} \right]^{-1} \times \frac{\partial^2 \mathcal{L}_T}{\partial \theta \partial \phi^T} \Big|_{\hat{\theta}(\phi)}, \quad (8)$$

which is a product of an inverse Hessian and a matrix of mixed partial derivatives. Adopting the algorithm from Lorraine et al. (2020), we approximate this with a truncated Neumann series with 1 term, and implicit vector-Jacobian products.

**Training algorithm.** Incorporating this gradient estimator, the algorithm to jointly optimize base model parameters and TaskAug policy parameters is given in Algorithm 1, mirroring the approach used in Raghu et al. (2021c).

---

#### Algorithm 1 Optimizing TaskAug parameters.

---

- 1: Initialize base model parameters  $\theta$  and TaskAug parameters  $\phi$
  - 2: **for**  $t = 1, \dots, T$  **do**
  - 3:   Compute training loss,  $\mathcal{L}_T(\theta)$
  - 4:   Compute  $\frac{\partial \mathcal{L}_T}{\partial \theta}$
  - 5:   Update  $\theta \leftarrow \theta - \eta_\theta \frac{\partial \mathcal{L}_T}{\partial \theta}$
  - 6:   **if**  $t \% P == 0$  **then**
  - 7:     Set  $\hat{\theta} = \theta$
  - 8:     Compute the validation loss,  $\mathcal{L}_V(\hat{\theta})$
  - 9:     Compute  $\frac{\partial \mathcal{L}_V}{\partial \theta}$
  - 10:     Approximate  $\frac{\partial \hat{\theta}}{\partial \phi}$  using Equation 8.
  - 11:     Compute the derivative  $\frac{\partial \mathcal{L}_V}{\partial \phi} = \frac{\partial \mathcal{L}_V}{\partial \hat{\theta}} \times \frac{\partial \hat{\theta}}{\partial \phi}$  using the previous two steps.
  - 12:     Update  $\phi \leftarrow \phi - \eta_\phi \frac{\partial \mathcal{L}_V}{\partial \phi}$
  - 13:   **end if**
  - 14: **end for**
- 

**Choice of  $P$ .** The value of  $P$  influences how many ‘inner’ gradient steps (to the base model) we perform before an ‘outer’ gradient step (to the TaskAug parameters). There is a tradeoff here: if  $P$  is too small, then applying the IFT to approximate  $\frac{\partial \hat{\theta}}{\partial \phi}$  will result in a poor approximation (Lorraine et al., 2020); if  $P$  is too large, then updates to the policy parameters will have little effect on model parameters since the base model has already reached minimal training loss (and may start to overfit). In our experiments, we find that  $P > 5$  suffered from this second problem, and  $P = 1$  was sometimes unstable due to the first problem. In general,  $P = 1$  worked well at small sample sizes ( $N = 1000$ ), and  $P = 5$  worked better at  $N = 2500$  and  $N = 5000$ .



## A.2.3. AUGMENTATION OPERATIONS

Figure 11 shows the different operations used in TaskAug. We show only one lead for clarity; however, these operations will be applied to each lead. We now provide more details on the implementation of these operations in our experiments.

- **TimeMask.** As with the existing TimeMask strategies, we randomly select a contiguous portion of the signal to set to zero. We set 10% of the signal to zero in our implementation. This parameter is not optimized.
- **Gaussian Noise.** IID Gaussian noise is added to the signal. This is formed as follows. We first compute the standard deviation of each lead of the signal: let us denote this as  $\sigma$ . Then, the noise added to each sample of the signal is expressed as:  $\epsilon = 0.25 \times \sigma \times \text{sigmoid}(s) \times \mathcal{N}(0, 1)$ , where  $s$  is the learnable strength parameter, initialized to 0. The coefficient 0.25 was found by visual inspection of some augmented examples, and observing that this allowed flexible augmentations to be generated without overwhelming the signal with noise.
- **Temporal warping.** The signal is warped with a random, diffeomorphic temporal transformation. To form this, we sample from a Gaussian with zero mean, and a fixed variance  $100 \times s^2$ , where  $s$  is the learnable strength parameter (initialized to 1), at each temporal location, to generate a length  $T$  dimensional random velocity field. This velocity field is then integrated (following the scaling and squaring numerical integration routine used by Balakrishnan et al. (2018, 2019)). This resulting displacement field is then smoothed with a Gaussian filter to generate the smoothed temporal displacement field. This field represents the number of samples each point in the original signal is translated in time. The field is then used to transform the signal, translating each channel in the same way (i.e., the field is the same across channels).
- **Baseline wander.** We firstly form a wander amplitude by computing:  $A = 0.25 \times \text{sigmoid}(s) \times \text{Uniform}(0, 1)$ , where again  $s$  is a learnable strength parameter. Then, we compute the frequency and phase of the sinusoidal offset. The frequency is computed as:  $f = \frac{20 \times \text{Uniform}(0, 1) + 10}{60}$ , based on the approximate number of breaths per minute for an adult. The phase is:  $\phi = 2\pi \times \text{Uniform}(0, 1)$ . Then, the sinusoidal offset is computed as:  $A \sin(ft + \phi)$ .
- **Magnitude scaling.** We scale the entire signal by a random magnitude given by

$\text{sigmoid}(s) \times \text{Uniform}(0.75, 1.25)$ , where  $s$  is a learnable strength parameter, initialized to 0.

- **Temporal displacement.** We shift the entire signal in time, padding with zeros where required. Our implementation directly generates a displacement field (as with temporal warping) and uses the spatial transformation from Balakrishnan et al. (2018, 2019) to transform the signal. This allows the operation to be differentiable, and for us to learn the displacement strength  $s$ . The displacement magnitude is a Uniform distribution on  $[-100 \times s^2, 100 \times s^2]$ , with the strength being initialized to 0.5.

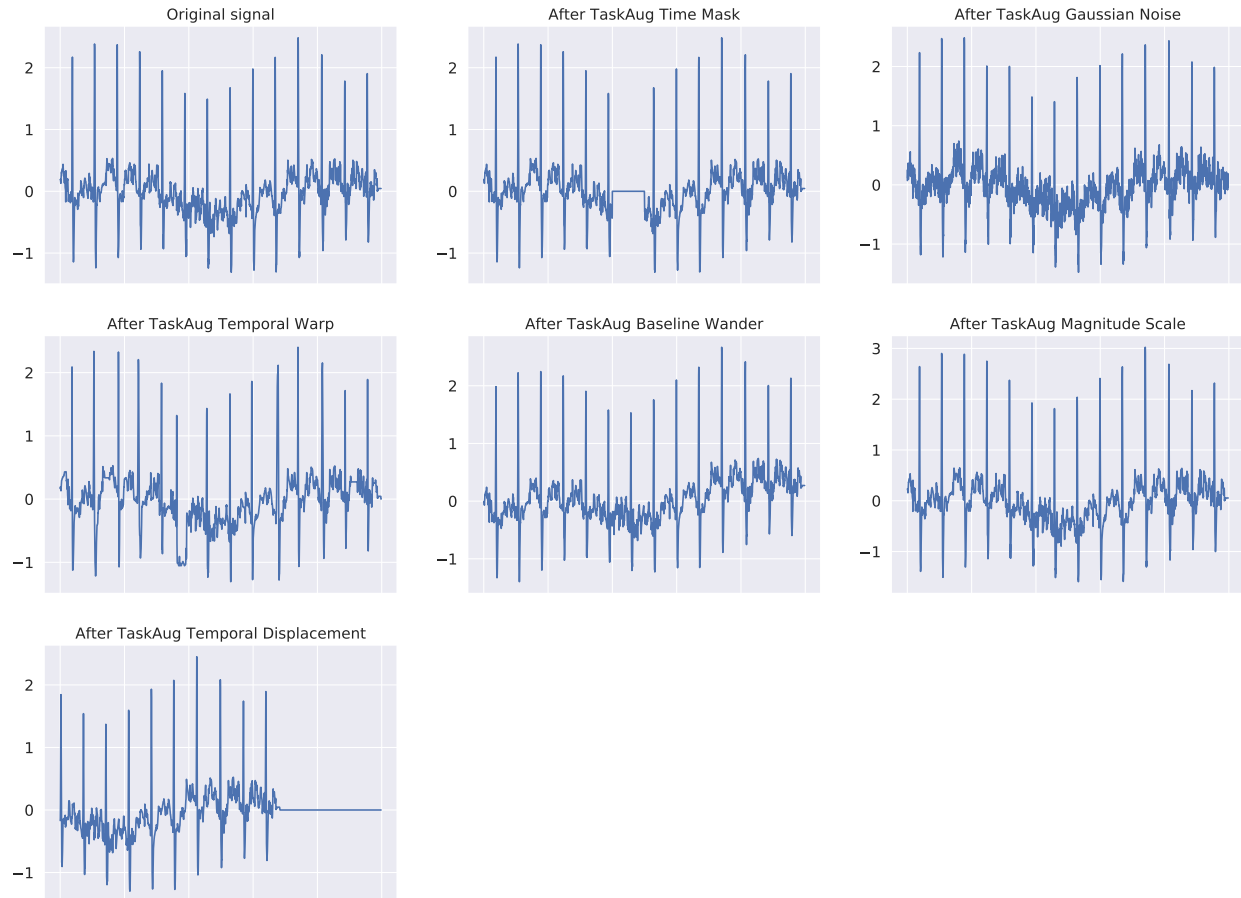


Figure 11: Examples of the different operations used in TaskAug.

## Appendix B. Dataset Details

We provide more details about the three datasets.

### B.1. Dataset A

The labels for RVH and AFib were assigned to each example based on whether relevant diagnostic statements were present in either a clinician’s read of the ECG, or a machine read of the ECG.

For RVH, there were six diagnostic statements that led to a positive label being assigned: “right ventricular hypertrophy”, “biventricular hypertrophy”, “combined ventricular hypertrophy”, “right ventricular enlargement”, “rightventricular hypertrophy”, “biventriclar hypertrophy”.

For AFib, there were nine such statements: “atrial fibrillation with rapid ventricular response”, “atrial

fibrillation with moderate ventricular response”, “fibrillation/flutter”, “atrial fibrillation with controlled ventricular response”, “afib”, “atrial fib”, “afibrillation”, “atrial fibrillation”, “atrialfibrillation”.

**Preprocessing.** ECGs were sampled at 250 Hz for 10 seconds, resulting in a  $2500 \times 12$  tensor for all 12 leads, per-ECG. We normalized the signals by dividing by 1000. Other forms of normalization for this dataset (e.g., z-scoring) resulted in some abnormally large/small values.

### B.2. Dataset B

The four labels are obtained by aggregating relevant sets of diagnostic statements – we refer the reader to the PTB-XL paper (Wagner et al., 2020) for further details. Of relevance here is that certain labels, such as MI, contain a small number of distinct diagnostic statements (3), potentially suggesting why many augmentation strategies can help – it is a fine-grained task. Others (such as CD) are much broader, covering many more diagnostic statements.

**Preprocessing.** ECGs in the dataset are sampled at 500 Hz for 10 seconds; we downsample these by a factor of 2 for consistency with Dataset A and C, resulting in a  $2500 \times 12$  tensor for all 12 leads, per-ECG. Normalization involved z-scoring, following the code provided with the dataset.

### B.3. Dataset C

The hemodynamics prediction cohort consists of patients who had an ECG and right heart catheterization procedure on the same day. The catheterization procedure measures hemodynamics variables including the pulmonary capillary wedge pressure (PCWP) and cardiac output (CO), and these are used to form the prediction targets. We consider inferring abnormally low Cardiac Output (less than 2.5 L/min), and abnormally high Pulmonary Capillary Wedge Pressure (greater than 20 mmHg).

**Preprocessing.** ECGs were sampled at 250 Hz for 10 seconds, resulting in a  $2500 \times 12$  tensor for all 12 leads, per-ECG. We normalized the signals by dividing by 1000. Other forms of normalization for this dataset (e.g., z-scoring) resulted in some abnormally large/small values, so we opted for the division-based normalization.

## Appendix C. Experiments

In this section, we provide further experimental details. We first provide implementation details, and then outline additional experimental results including: Results for AUPRC in the low-sample ( $N = 1000$ ) regime, performance on Datasets A and B in the high-sample regime, performance on Dataset B in an additional low sample regime ( $N = 500$  data points), interpretation of the TaskAug policy for RVH, a study of the impact of optimizing policy parameters across different sample size regimes, and a study of the impact of class-specific magnitudes across different sample size regimes.

### C.1. Implementation details

**Network architecture.** In all experiments, we use a 1D CNN based on a ResNet-18 (He et al., 2016) architecture. This model has convolutions with a kernel size of 15, and stride 2 (informed by the temporal window we want the convolutions to operate over). The blocks in the ResNet architecture have convolutional layers with 32, 64, 128, and 256 channels respectively. The output after the final block is average pooled in the temporal dimension, and then a linear layer is applied to predict the probability of the positive class.

**Optimization settings.** As discussed, we used Adam with a learning rate of 1e-3 for all methods, given that this resulted in stable training across all settings. When optimizing the TaskAug policy parameters, we used RMSprop with a learning rate of 1e-2, following Lorraine et al. (2020).

**Computational information.** All models and training were implemented in PyTorch and run on a single NVIDIA V100 GPU.

### C.2. Additional results

**AUPRC results at 1000 samples.** As discussed in Section 5.2, the improvements in AUROC are not always statistically significant. Given that some of the labels are very low prevalence (RVH: 1%, AFib: 5%, low CO: 4%), we evaluate the AUPRC in the low-sample regime, which provides additional information about model performance. Results are shown in Tables 4, 5, and 6. We observe that for the low prevalence RVH, AFib, and Low CO tasks, TaskAug obtains statistically significant improvements in performance. On Dataset A tasks (RVH and AFib), it is the only method to do so.

	RVH	AFib
NoAugs	7.4 ± 1.3	<u>21.2 ± 2.0</u>
TaskAug	<b>10.8 ± 0.8*</b>	<b>27.3 ± 1.8*</b>
SMOTE	9.7 ± 1.2	21.0 ± 2.2
DGW	7.1 ± 0.9	19.4 ± 2.3
SpecAug	<u>10.6 ± 1.2</u>	21.1 ± 2.0
TimeMask	10.1 ± 1.5	20.3 ± 2.3

Table 4: Mean and standard error of AUPRC for various data augmentation strategies when detecting cardiac abnormalities on Dataset A. We consider a low-sample regime with a development set of 1000 data points. The best-performing method is bolded, and the second best is underlined, and \* indicates statistically significant improvement at the  $p < 0.05$  level. TaskAug is the only method to obtain significant improvements in performance on both tasks.

	MI	HYP	STTC	CD
NoAugs	59.2±2.1	53.1±1.7	66.9±2.5	<b>67.3±1.1</b>
TaskAug	<b>63.1±1.7</b>	<b>55.2±0.9</b>	68.7±1.3	66.8±1.2
SMOTE	<u>62.0±1.6</u>	41.2±2.9	65.9±1.0	62.7±1.1
DGW	<u>61.1±1.2</u>	53.9±1.6	67.9±1.1	64.7±2.6
SpecAug	61.7±1.6	<u>54.5±1.5</u>	<u>68.8±1.5</u>	65.8±1.4
TimeMask	60.3±1.3	52.8±1.8	<b>68.8±1.2</b>	<b>70.1±1.3</b>

Table 5: Mean and standard error of AUPRC for various data augmentation strategies on detecting cardiac abnormalities on Dataset B. We consider a low-sample regime with a development set of 1000 data points. The best-performing method is bolded, and the second best is underlined, and \* indicates statistically significant improvement at the  $p < 0.05$  level.

	Low CO	High PCWP: $N = 1000$	High PCWP: All Data
NoAugs	7.2 ± 0.4	<u>42.5 ± 0.8</u>	49.7 ± 0.8
TaskAug	<b>8.8 ± 0.6*</b>	<b>43.5 ± 0.9</b>	<b>50.8 ± 0.8</b>
SMOTE	<b>8.8 ± 0.6*</b>	41.9 ± 0.7	46.9 ± 0.7
DGW	<u>8.1 ± 0.7</u>	41.2 ± 0.7	49.7 ± 1.0
SpecAug	7.8 ± 0.4	42.3 ± 1.1	<u>50.3 ± 0.8</u>
TimeMask	8.0 ± 0.5	42.4 ± 0.7	<u>50.1 ± 0.9</u>

Table 6: Mean and standard error of AUPRC for various data augmentation strategies for the hemodynamics inference task in Dataset C. We consider a low-sample regime with a development set of 1000 data points. The best-performing method is bolded, and the second best is underlined, and \* indicates statistically significant improvement at the  $p < 0.05$  level. TaskAug is the one of only two methods to obtain significant improvements in performance on the low CO detection task.

**Results at higher sample regimes.** Tables 7-10 show AUROC for the different augmentation methods on the tasks from Datasets A and B. We observe that augmentations are less effective at higher samples. Particularly when the development set sizes are 2500 and 5000 datapoints, we observe that the improvement with using augmentations (over the NoAugs baseline) with any of the methods is quite small, and nearly always less than 1% AUROC. This suggests that in general, augmentations are less useful at these higher data regimes.

	RVH	AFib
NoAugs	<u>86.1±0.9</u>	89.0 ± 0.4
TaskAug	<b>86.9±0.9</b>	89.1 ± 0.4
SMOTE	85.5±1.3	<u>89.1 ± 0.5</u>
DGW	84.8±1.3	88.4 ± 0.5
SpecAug	83.3±1.8	<b>89.1 ± 0.3</b>
TimeMask	85.8±1.1	88.2 ± 0.4

Table 7: Mean and standard error of AUROC for augmentation methods on Dataset A tasks with a development set of 2500 data points. The best performing method is bolded, and the second best is underlined.

	RVH	AFib
NoAugs	<u>90.6±0.6</u>	92.6±0.2
TaskAug	90.6±0.4	<b>92.8±0.1</b>
SMOTE	89.8±0.6	92.6±0.2
DGW	<b>90.8±0.5</b>	92.5±0.2
SpecAug	90.5±0.8	<u>92.7±0.1</u>
TimeMask	89.4±0.7	92.6±0.2

Table 8: Mean and standard error of AUROC for augmentation methods on Dataset A tasks with a development set of 5000 data points. The best performing method is bolded, and the second best is underlined.

	MI	HYP	STTC	CD
NoAugs	84.5±0.5	<u>86.4±0.4</u>	89.7±0.3	85.8±0.3
TaskAug	<b>86.1±0.5</b>	86.2±0.4	89.7±0.3	86.6±0.4
SMOTE	84.7±0.7	81.9±1.3	88.7±0.4	85.5±0.6
DGW	84.1±0.5	85.9±0.6	89.5±0.3	86.2±0.3
SpecAug	84.6±0.8	86.2±0.6	<b>90.2±0.3</b>	<u>86.8±0.6</u>
TimeMask	<u>85.7±0.4</u>	<b>86.6±0.3</b>	<u>90.1±0.1</u>	<b>87.0±0.7</b>

Table 9: Mean and standard error of AUROC for augmentation methods on Dataset B tasks with a development set of 2500 data points. The best performing method is bolded, and the second best is underlined.

	MI	HYP	STTC	CD
NoAugs	<u>89.4±0.3</u>	88.2±0.2	91.0±0.3	89.3±0.4
TaskAug	<u>89.4±0.3</u>	88.3±0.2	<b>91.6±0.2</b>	<b>90.0±0.2</b>
SMOTE	86.6±0.7	86.7±0.4	90.6±0.3	88.0±0.3
DGW	88.6±0.3	88.0±0.2	<u>91.3±0.1</u>	89.3±0.2
SpecAug	<b>89.5±0.2</b>	<u>88.4±0.4</u>	<b>91.6±0.2</b>	<u>89.9±0.2</u>
TimeMask	89.3±0.3	<b>88.6±0.2</b>	<b>91.6±0.2</b>	89.8±0.2

Table 10: Mean and standard error of AUROC for augmentation methods on Dataset B tasks with a development set of 5000 data points. The best performing method is bolded, and the second best is underlined.

**Results on Dataset B at  $N = 500$ .** Table 11 shows AUROC for the different augmentation methods in an additional low sample regime, with  $N = 500$ . We see that the maximum improvement over the NoAugs baseline by any augmentation strategy is greater in this regime than it was at  $N = 1000$  (see Table 2). Given that the prevalence of these tasks is relatively high, we see more significant performance improvements in the  $N = 500$  regime.



	MI	HYP	STTC	CD
NoAugs	$74.4 \pm 0.9$	<b><math>81.9 \pm 0.8</math></b>	$85.2 \pm 0.5$	$78.9 \pm 1.2$
TaskAug	<b><math>78.4 \pm 0.5</math></b>	<u><math>81.5 \pm 1.2</math></u>	$86.2 \pm 0.4$	<b><math>80.7 \pm 0.6</math></b>
SMOTE	$75.7 \pm 1.2$	$79.2 \pm 1.5$	$85.5 \pm 0.3$	$78.6 \pm 1.5$
DGW	<u><math>78.2 \pm 0.6</math></u>	$78.7 \pm 1.2$	$82.0 \pm 1.3$	$79.0 \pm 0.9$
SpecAug	$77.8 \pm 0.7$	$81.0 \pm 0.6$	<u><math>86.3 \pm 0.4</math></u>	$79.3 \pm 1.1$
TimeMask	$77.8 \pm 1.0$	$80.9 \pm 1.3$	<b><math>86.6 \pm 0.5</math></b>	<u><math>80.3 \pm 0.8</math></u>

Table 11: Mean and standard error of AUROC for augmentation methods on Dataset B tasks with a development set of 500 data points. The best performing method is bolded, and the second best is underlined.

**Interpreting the RVH policy.** We visualize the TaskAug policy for RVH in Figure 12. We observe high probability assigned to selecting two temporal operations in stage 1, namely masking and displacement. Relative magnitudes of different portions of the ECG affect the RVH label, so temporal operations having higher probability of selection is sensible since they are more likely to be label preserving than operations that change the relative magnitudes of different parts of the ECG. We examine the learned strengths for the displacement operation in Stage 1, Figure 12(b), and we see that there is little differentiation on a per-class basis. This is sensible, since we do not expect displacement of the signal in time to affect the RVH label for differently for the positive and negative classes.




Figure 12: **TaskAug policy for detecting Right Ventricular Hypertrophy.** The learned TaskAug policy: probability of selecting each transformation in both augmentation stages and the optimized displacement strengths in the first stage. We show the mean/standard error of the learned parameter values over 15 runs. Temporal operations (masking and displacement) have high probability of selection in Stage 1, which is sensible since these operations are likely to be label preserving (RVH is typically detected based on relative magnitudes of portions of beats in the ECG). We see that both positive and negative classes have similar optimized displacement augmentation strengths – we do not expect displacement to impact the class label differently for the two classes, so this is sensible.

**Further study on the impact of optimizing augmentations.** As shown in the main text, Figure 5, optimizing the policy parameters improves performance over keeping them fixed at their initial values. In Figure 13, we study this effect across different dataset sizes and find that the optimization has the most impact in the low sample regime, but still results in improvements even at higher samples. This could be due to the fact that at higher samples, augmentations boost performance less in general, so the specific parameter settings in TaskAug also have less impact.

**Further study on the impact of class-specific magnitudes.** As shown in the main text, Figure 6, optimizing class-specific magnitudes improves over learning one magnitude parameter for each class. Figure 14 studies this effect across different dataset sizes and we see that the class-specific parameters improve performance at all dataset sizes, but the improvement is most clearly seen at low samples. Similarly with the optimization of augmentation parameters, this could be due to the fact that at higher samples, augmentations boost performance less in general, so the class-specific parameterization in TaskAug has less impact.




Figure 13: **Studying performance when we do not optimize the policy parameters in TaskAug.** We show the mean/standard error of AUROC over 15 runs for AFib and over 5 runs for MI. We see that optimizing the policy parameters results in noticeable improvements in performance over keeping the policy parameters at their initial values (InitAugs). However, the impact of optimizing the parameters is reduced at larger dataset sizes, possibly due to the fact that augmentations are inherently less useful at higher sample regimes.




Figure 14: **Studying performance when we do not have class-specific magnitude parameters in TaskAug.** We show the mean/standard error of AUROC over 15 runs for AFib and over 5 runs for MI. Class-specific magnitude parameters improve performance most in the low sample regime. At higher samples, this impact is reduced, possibly due to the fact that augmentations are inherently less useful at higher sample regimes.