

3KG: Contrastive Learning of 12-Lead Electrocardiograms using Physiologically-Inspired Augmentations

Bryan Gopal*
Stanford University

BRYANG@CS.STANFORD.EDU

Ryan Han*
Stanford University

RYANHAN@CS.STANFORD.EDU

Gautham Raghupathi*
Stanford University

GAUTHAM@CS.STANFORD.EDU

Andrew Ng
Stanford University

ANG@CS.STANFORD.EDU

Geoff Tison†
University of California, San Francisco

GEOFF.TISON@UCSF.EDU

Pranav Rajpurkar†
Harvard Medical School

PRANAV_RAJPURKAR@HMS.HARVARD.EDU

Abstract

We propose 3KG, a physiologically-inspired contrastive learning approach that generates views using 3D augmentations of the 12-lead electrocardiogram. We evaluate representation quality by fine-tuning a linear layer for the downstream task of 23-class diagnosis on the PhysioNet 2020 challenge training data and find that 3KG achieves a 9.1% increase in mean AUC over the best self-supervised baseline when trained on 1% of labeled data. Our empirical analysis shows that combining spatial and temporal augmentations produces the strongest representations. In addition, we investigate the effect of this physiologically-inspired pre-training on downstream performance on different disease subgroups and find that 3KG makes the greatest gains for conduction and rhythm abnormalities. Our method allows for flexibility in incorporating other self-supervised strategies and highlights the potential for similar modality-specific augmentations for other biomedical signals.

Keywords: self-supervised learning, contrastive learning, electrocardiogram, spatiotemporal augmentations

1. Introduction

The 12-lead electrocardiogram (ECG) is a common, non-invasive test for the screening, diagnosis, and management of cardiovascular conditions (Salerno et al., 2003; Fesmire et al., 1998). Recent studies have demonstrated the ability of deep learning to predict conditions commonly diagnosed by physicians on the ECG as well as novel diagnoses beyond those traditionally identifiable by physicians with an ECG (Tison et al., 2019; Attia et al., 2019b,a). However, such supervised methods depend on tens of thousands of high quality labels to achieve strong generalization performance: we find that supervised learning is only marginally better than random for certain diagnoses at low label fractions. There is therefore a need for label-efficient learning strategies to reduce the resources required for algorithmic training.

Self-supervised contrastive learning presents a promising avenue for improving the label-efficiency of ECG interpretation strategies. Contrastive learning approaches leverage unlabeled data to pretrain models for fine-tuning on downstream tasks where labeled data is scarce (Hjelm et al., 2019; Wu et al., 2018). On computer vision tasks, contrastive learning approaches have reached performance comparable to that of supervised learning while using as little as 1% of the labeled data (Chen et al., 2020a; He et al., 2020; Vu et al., 2021). However, there still

* Equal Contribution

† Equal Contribution

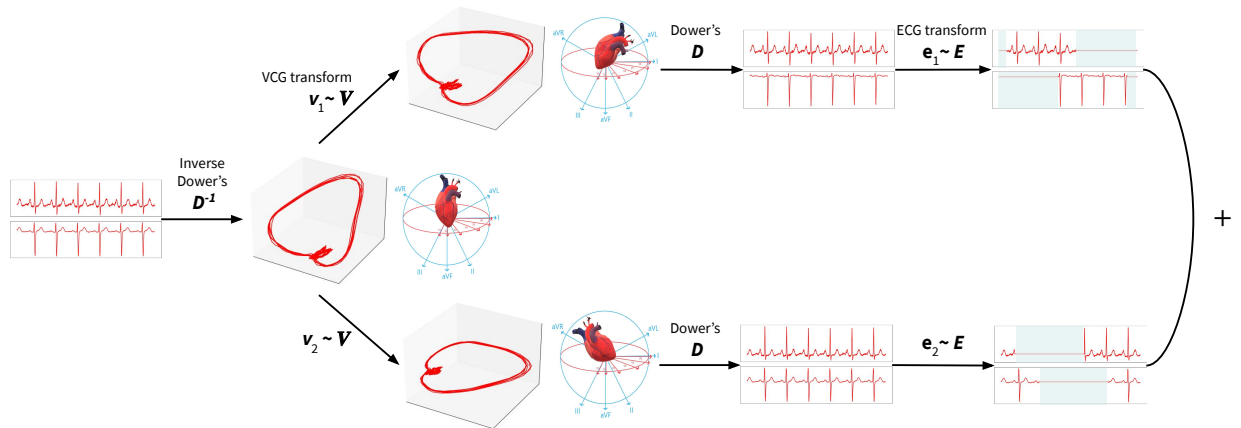


Figure 1: **3KG Overview.** Positive pairs for contrastive learning are generated by composing stochastic augmentations of the 12-lead ECG in the VCG and ECG spaces respectively. Leads are encoded independently and then treated as positive views if they derive from the same patient. Only leads II and V2 are shown in the figure for the sake of simplicity.

remains a large performance disparity between supervised learning and contrastive methods for ECG analysis that are fine tuned on small label fractions (Kiyasseh et al., 2020; Cheng et al., 2020; Mehari and Strodtzoff, 2021; Diamant et al., 2021).

In this work, we propose 3KG, a physiologically-inspired contrastive learning approach that leverages the unique spatiotemporal properties of the ECG. Our method generates views for contrastive learning using ECG-specific 3D augmentations in the vectorcardiogram (VCG) space. VCGs derive from ECGs and represent the electrical activity of the heart along three orthogonal spatial axes. We hypothesize that VCGs enable spatial manipulations of ECGs in a manner consistent with the underlying physiology they represent. After spatial manipulation, we project the VCG back in to the ECG space and apply temporal manipulations. We further extend this method by incorporating the spatial relationship between the leads by encoding them separately and treating them as positive views. An overview for this approach is provided in Figure 1.

We conduct pretraining and evaluate the quality of learned representations on the PhysioNet 2020 dataset (Alday et al., 2020) for the task of arrhythmia classification. Our contributions are as follows:

1. We demonstrate that 3KG outperforms previous self-supervised ECG strategies and closes the gap to achieving fully supervised performance when

only using a subset of the labeled data. Our best performing model achieves an AUROC of 0.826 when trained on 1% of labeled data compared to an AUROC of 0.898 for fully supervised learning on 100% of labeled data. For reference, we find that the previous best self-supervised learning approach achieves an AUROC of 0.757 when trained on the 1% label fraction.

2. We perform comparative empirical analyses with spatial augmentation strength and find that stronger spatial augmentations lead to better learned representations. Additionally, we observe that combining spatial augmentations increases performance over any individual spatial augmentation, and that combining spatial augmentations with temporal augmentations produces the strongest results of all.
3. We investigate how the introduction of an inductive bias through use of physiologically-inspired pretraining affects downstream performance on different disease subgroupings relative to a randomly initialized baseline. We find that 3KG makes the greatest gains for conduction and rhythm abnormalities and smaller gains for waveform abnormalities.

2. Related Work

Many contrastive approaches have been developed for computer vision tasks (Chen et al., 2020a; He et al., 2020; Grill et al., 2020). These methods define positive views, the representations of which are pushed together during pretraining, and negative views, the representations of which are pulled apart (Tian et al., 2020). The definition of views thus controls the semantic information encoded by the learned representations. Examples of definitions include Tian et al. (2019) who use different modal views of the same image as positive views, and Chen et al. (2020a) who define positive views as differently augmented views of the same image. Chen et al. (2020a) note that augmentation strength and composition are crucial for learning good representations; the development of augmentations applicable to ECGs is thus important.

Self-supervised methods have also been applied to the ECG domain. While Sarkar and Etemad (2020) pretrain their model on a pretext task of predicting which of six signal transformations (gaussian noise, scaling, temporal inversion, negation, permutation, and time-warping) was applied, Mehari and Strodthoff (2021) adapt SimCLR to the ECG domain through the use of random resized crops and time masking as augmentations. They additionally apply Contrastive Predictive Coding (van den Oord et al., 2018) to ECGs. Although these methods improve over randomly initialized baselines, they do not take advantage of the unique properties of the ECG as they are largely direct extensions from other domains.

Other contrastive ECG methods incorporate unique attributes of the ECG. Mehari and Strodthoff (2021) propose the use of physiological noise (baseline wander, powerline noise, electromyographic noise, and baseline shift) as augmentations for SimCLR. In Patient Contrastive Learning, Diamant et al. (2021) define positive pairs as different ECGs taken from the same patient at varying points in time. In contrast to our approach, these methods emphasize the clinical nature of the ECG recording without taking advantage of its spatiotemporal attributes. Another approach by Kiyasseh et al. (2020) proposes a trio of methods that leverage the spatiotemporal nature of the ECG by treating different segments, leads, and both segments and leads, as positive views. However, this approach limits itself to the ECG space, while our

approach takes advantage of natural augmentations to the ECG in the VCG space.

3. Methods

3.1. Problem Formulation

Given a batch of N 12-lead ECGs \mathbf{x} , a set of positive view pairs \mathcal{P} , a space of transformations \mathcal{T} , and an additional positive view condition \mathcal{C} , our goal during pretraining is to minimize a contrastive loss function

$$\begin{aligned} \mathcal{L}(\mathbf{x}) &= \frac{1}{|\mathcal{P}|} \sum_{i,j \in \mathcal{P}} \ell_{i,j}(\mathbf{x}) \\ \ell_{i,j}(\mathbf{x}) &= -\log \frac{s_{i,j}(\mathbf{x})}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \mathbb{1}_{[i,k \notin \mathcal{P}]} \mathbb{1}_{-\mathcal{C}(i,k)} s_{i,k}(\mathbf{x})} \\ s_{i,j}(\mathbf{x}) &= \exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j) / \tau) \end{aligned}$$

where the L_2 -normalized output embedding \mathbf{z}_i is produced by passing ECG sample \mathbf{x}_i through our pretraining model after it has been augmented with a transformation sampled from \mathcal{T} . The function sim denotes the cosine similarity function, and τ is a temperature hyperparameter.

Following Chen et al. (2020a), we do not sample negative views explicitly, instead defining a negative view for index i as all examples in the batch that are not part of a positive pair with i and fail to satisfy condition \mathcal{C} .

\mathcal{T} , \mathcal{C} , and \mathcal{P} all contribute to the definition of what a positive view of an ECG sample x can be. We use physiologically-based inductive biases to propose *3KG*, which describes novel formulations for all three.

3.2. Positive View Definition

3.2.1. AUGMENTATIONS

We seek to leverage the vectorcardiogram (VCG) in combination with time masking to build a space of transformations \mathcal{T} that preserves spatiotemporal features relevant for assessment of ECG signals.

The vectorcardiogram (VCG) represents the gold standard in visualizing the spatial propagation of a heart signal over time. In contrast to a 12-lead ECG (which provides 12 2D projections of a 3D signal), a VCG shows the electrical activity of the heart along three orthogonal spatial axes, tracing contours around a central point over time.

Since natural variability in cardiac structure and orientation exists without substantive clinical implications, we argue that patient VCGs can be similarly augmented through rotation and scaling without breaking the spatiotemporal invariances needed to make downstream diagnoses.

We also apply time masking on a per-lead basis as a temporal augmentation. Taking inspiration from self-supervised approaches in speech recognition (Baevski et al., 2020), we choose to implement time masking in order to mimic how ECG assessment is done in a clinical setting. Since cardiologists often only need a single beat on a single lead to ascertain a significant portion of ECG diagnoses, a self-supervised model should similarly be able to learn useful representations of ECGs that have some beats masked.

We formalize our transformation space as follows:

$$\mathcal{T} = \mathcal{E} \circ \mathcal{V},$$

where \mathcal{E} is the space of all temporal ECG augmentations and \mathcal{V} is the space of all spatial vectorcardiogram (VCG) perturbations. We define \mathcal{V} to be

$$\mathcal{V}(\mathbf{x}_i) = DSRD^{-1}\mathbf{x}_i,$$

where

- the Inverse Dowers transformation $D^{-1} \in \mathbb{R}^{3 \times 12}$ projects the 12-lead ECG into 3-dimensional vectorcardiogram space (Jaros et al., 2019).
- a rotation matrix $R \in \mathbb{R}^{3 \times 3}$ performs a 3-dimensional rotation of a VCG along all three orthogonal axes.
- a scaling matrix $S \in \mathbb{R}^{3 \times 3}$ scales the vectorcardiogram along all three orthogonal axes.
- the Dowers transformation $D \in \mathbb{R}^{12 \times 3}$ projects the VCG back into 12-lead ECG space (Jaros et al., 2019)

Note that because VCGs exist in a lower-dimensional space than ECGs, the conversion process between the two observes some loss since D approximates an inverse to the underdetermined system D^{-1} .

Similarly, we define \mathcal{E} to be the space of all possible time masks for each of the 12 ECG leads.

In practice, we apply our augmentation scheme in the following order:

1. VCG Augmentations \mathcal{V}

- Convert ECG to VCG.** We apply the Inverse Dowers transformation D^{-1} to compute a VCG from our 12-lead ECG.
- Random Rotation.** We rotate the VCG on all three axes in a random order, where each per-axis rotation is a random value uniformly sampled between $-r$ to $+r$ degrees for a user-specified value of r .
- Random Scale.** We scale the VCG on all three axes, where each per-axis scale is a random value uniformly sampled from a uniform distribution between $[1, s]$ with a 50% probability of inversion for a user-specified value of s . This means that on average, one half of the scale factors is between $[1/s, 1]$ and the other half is between $[1, s]$.
- Convert VCG to ECG.** We apply the Dowers transformation D to compute an ECG from our VCG.

2. ECG Augmentations \mathcal{E}

- Random Time Masking.** We uniformly sample 1 time step for each lead as a starting index, and replace the subsequent $[L \cdot p_t]$ time steps to be zero over all leads (wrapping around to the beginning of the recording as necessary), where L is the number of time steps in the sample and $0 \leq p_t \leq 1$ is the user-specified time mask percentage.

We sample two transformations $t_1, t_2 \sim \mathcal{T}$ at each training step and apply each to the batch, resulting in the positive pair $[t_1(\mathbf{x}_i), t_2(\mathbf{x}_i)]$ for each 12-lead ECG \mathbf{x}_i (see Figure 1 for more details).

3.2.2. PATIENT-AWARE VIEWS

We seek to ensure that ECG recordings that belong to the same patient are not considered as negative views of each other. We thus couple a recording number label to each ECG sample during pretraining and let our additional condition \mathcal{C} be "ECG samples that originate from the same recording are considered positive views of each other". This additional condition transforms our traditional contrastive loss function into a *Patient-Supervised Contrastive Loss* and is in line with previous works that have also taken a patient-aware approach in loss calculation (Kiyasseh et al., 2020; Diamant et al., 2021).

3.2.3. PER-LEAD VIEWS

We seek to leverage the inherent spatial relationship between individual ECG leads to learn better representations during pretraining. Some previous contrastive methods for 12-lead ECGs have treated different leads as independent channels akin to color channels in computer vision (Diamant et al., 2021; Mehari and Strodthoff, 2021). Because the 12 leads in an ECG provide different spatial views of the same underlying cardiac electrical activity, they can also be thought of as natural augmentations of each other (Kiyasseh et al., 2020).

We thus propose modifying our set of positive pairs \mathcal{P} to consider individual leads as positive views. We evaluate the impact of this formulation by comparing models trained with and without per-lead views, keeping all other parts of experimental setup the same. We denote our two different methods for handling the 24 lead embeddings produced per positive pair $[t_1(\mathbf{x}_i), t_2(\mathbf{x}_i)]$ as follows:

- **SimCG-A**, which *does not* use a per-lead formulation of \mathcal{P} and instead averages the 12 embeddings produced per ECG. As such, there is only positive embedding pair per ECG sample \mathbf{x}_i , with each of the two component embeddings coming from $t_1(\mathbf{x}_i)$ and $t_2(\mathbf{x}_i)$.
- **SimCG-L**, which *does* use a per-lead formulation of \mathcal{P} , constructing positive pairs over all lead embeddings during loss computation. As such, we arrive at $\binom{24}{2} = 276$ positive embedding pairs per ECG \mathbf{x}_i .

We note that SimCG-L is similar to CMLC (Kiyasseh et al., 2020) in that different leads from the same patient are treated as positive pairs. However, notable differences include the fact that our approach incorporates spatiotemporal augmentations and uses a nonlinear projection head during pretraining to improve the representation quality of the final layer of the encoder. Additionally, we accumulate loss across all pairs of views of an ECG, which we did not find implemented in the released codebase.

Finally, we denote **3KG** to be the combination of our spatiotemporal augmentation strategy, Patient-Supervised Contrastive Loss, and SimCG-L.

3.3. Experimental Setup

3.3.1. DATASET

We run all experiments on the PhysioNet 2020 challenge training data (Alday et al., 2020). The dataset consists of 12-lead ECG recordings from four data sources with various distributions of patient characteristics, recording length, and diagnoses. There are a total of 43,134 recordings which we split into training, validation, and test sets in an 80/10/10 configuration; there was no patient overlap between these sets. The training set was further divided into 1%, 10%, and 100% label fraction splits to simulate label scarcity for the downstream task. Labels are at the recording level, and each recording could be associated with multiple labels. Although the training data had 111 classes present, we chose to classify only the 27 SNOMED CT codes that were included in the challenge evaluation metric. A board-certified cardiologist coalesced the following SNOMED CT labels after the determination that the diagnoses were substantially equivalent: complete right bundle branch block and right bundle branch block, premature atrial contraction and supraventricular premature beats, premature ventricular contractions and ventricular premature beats, 1st degree atrioventricular block and prolonged PR interval. Our final downstream task was a 23-class multilabel classification. F

Preprocessing. To ensure ECGs from different sources were standardized, all recordings were resampled to 500 Hz and then exhaustively cropped into disjoint 5-second segments. All crops were tagged with the file they originated from for later use in loss computation. ECGs were normalized to unit range, but were not translated so as to maintain the isoelectric line at 0.

3.3.2. ARCHITECTURES AND OPTIMIZATION

Our encoder is a 1-D convolutional neural network with a single channel input and 256-dimensional output (see Appendix A for more details). We chose to use 1 input channel and process each lead separately to increase the versatility of our model to settings where not all 12 leads are available, such as in the ambulatory (i.e. Holter) setting where the standard of care is a single lead ECG (Kennedy, 1992).

Following similar work in self-supervision, we leveraged a 2-layer multilayer perceptron (MLP) network in pretraining that projected this 256-dimensional output into 128-dimensional space (Chen et al.,

Augmentation	AUROC			F1		
	1%	10%	100%	1%	10%	100%
<i>Maximum Rotation (degrees)</i>						
5	0.771	0.835	0.851	0.286	0.366	0.395
10	0.783	0.839	0.858	0.295	0.374	0.404
25	0.800	0.854	0.873	0.312	0.386	0.418
45	0.801	0.856	0.877	0.315	0.386	0.422
<i>Maximum Scale Factor</i>						
1.05	0.771	0.831	0.849	0.273	0.360	0.392
1.1	0.764	0.834	0.851	0.260	0.366	0.398
1.5	0.795	0.852	0.868	0.303	0.379	0.408
2	0.789	0.846	0.866	0.299	0.381	0.412
<i>Combinations</i>						
Rotate 45 + Time Mask	0.812	0.870	0.881	0.306	0.397	0.428
Scale 1.5 + Time Mask	0.798	0.859	0.878	0.302	0.387	0.417
Rotate 45 + Scale 1.5	0.807	0.861	0.881	0.319	0.394	0.431
Rotate 45 + Scale 1.5 + Time Mask	0.812	0.869	0.883	0.320	0.400	0.432

Table 1: **Linear Evaluation for Different Augmentation Combinations on SimCG-A.** The time mask percentage was set to 50% for all experiments. Combinations were composed of best individual augmentations.

2020a). We used the Adam optimizer with a learning rate of 1×10^{-3} , a temperature $\tau = 0.1$, and a batch size of 512 for all pretraining experiments (Kingma and Ba, 2017).

3.3.3. EVALUATION ON DOWNSTREAM TASK

We perform a downstream linear evaluation of our pretrained encoder across 1%, 10%, and 100% label fractions. After pretraining, we freeze our encoder’s parameters and drop our MLP, attaching a single linear layer to predict diagnosis labels from the learned output representations. We average the normalized single lead embeddings generated by passing a 12-lead ECG crop into our single-lead encoder. This final averaged embedding is sent downstream to the linear classifier. The linear classifier is trained using a binary cross-entropy loss averaged over all classes. For the 1% and 10% label fractions, we trained 5 different linear classifiers on 5 different fractional splits and used all of their output predictions in metric calculation. We report the AUROC and F1 scores averaged across all 23 classes in the test set for our experimental results.

We used the Adam optimizer and a batch size of 512 for all downstream experiments (Kingma and Ba, 2017). For each experiment, we performed an automatic learning rate search over the space $[10^{-8}, 1]$

following the methodology outlined by Smith (2017). In practice, we found that the optimal learning rate for the linear layer landed within the range $[0.05, 0.2]$ for most experiments.

Statistical Testing. For each experiment, we evaluate the statistical significance of each of our experiments against a stated baseline model using a non-parametric bootstrapping setup. We randomly sample N_T test set predictions with replacement (where N_T is the size of the test set) from both the baseline and the model of interest. We then compute the metric difference between the two resampled prediction sets. We repeat this entire process 1000 times, and construct a 95% confidence interval using all the resulting differences. If 0 lies outside this interval, we consider the model of interest to be significantly different than the baseline model.

4. Experiments

We use AUROC as our primary metric to evaluate performance across the 23 classes, and also report F1 scores. Our results are shown in Table 1 and 2.

Method	AUROC			F1		
	1%	10%	100%	1%	10%	100%
<i>Supervised Baseline</i>						
Random Initialization	0.694	0.836	0.898	0.170	0.349	0.444
<i>Self-Supervised Baselines</i>						
Patient Contrastive Learning (Diamant et al., 2021)	0.757	0.843	0.857	0.259	0.335	0.362
Contrastive Multi-Segment Coding (Kiyasseh et al., 2020)	0.728	0.782	0.792	0.210	0.277	0.308
Contrastive Multi-Lead Coding (Kiyasseh et al., 2020)	0.730	0.769	0.779	0.205	0.274	0.321
Contrastive Multi-Segment/Lead Coding (Kiyasseh et al., 2020)	0.711	0.774	0.788	0.201	0.263	0.299
SimCLR + RRC/TO (Mehari and Strodthoff, 2021)	0.594	0.634	0.740	0.121	0.153	0.242
SimCLR + Physio. (Mehari and Strodthoff, 2021)	0.608	0.641	0.744	0.141	0.163	0.229
Contrastive Predictive Coding (Mehari and Strodthoff, 2021)	0.554	0.573	0.677	0.099	0.093	0.170
Signal Transformation (Sarkar and Etemad, 2020)	0.548	0.562	0.564	0.085	0.085	0.099
<i>Our Methods</i>						
SimCG-A + VCG + Time Mask	0.812	0.869	0.883	0.320	0.400	0.432
SimCG-L	0.789	0.851	0.869	0.304	0.389	0.420
3KG (SimCG-L + VCG + Time Mask)	0.826	0.876	0.890	0.309	0.411	0.438

Table 2: **PhysioNet2020 performance of linear classifiers trained on representations learned with self-supervised methods.** Augmentation parameters (Rotate 45, Scale 1.5) were chosen based on highest performance on SimCG-A. Differences between the top baseline (Patient Contrastive Learning) and all our methods are statistically significant for all metrics and fractions displayed. The top performer among baselines and top performer among our methods are bolded. Abbreviations: Random Resized Crop (RRC), Timeout (TO), Physiological Noise (Physio).

4.1. Spatial VCG Augmentations

Design We compare the performance of different augmentation parameters and examine the improvement in performance due to a combination of augmentations. Tian et al. (2020) argue that a good view in contrastive learning retains task-relevant information while reducing the mutual information between views. We are motivated to generate positive views of an ECG signal by applying physiologically-inspired augmentations under this principle. Chen et al. (2020b) show that the combination of augmentations in a contrastive learning setup could amount to gains in performance greater than the gains in performance of any individual augmentation. Therefore, we are also motivated to combine our physiologically-inspired augmentations with time masking, an augmentation commonly used for signal data (Baevski et al., 2020). We choose to evaluate our physiologically-inspired augmentations with SimCG-A, since we want to isolate the impact of our augmentations to better optimize the augmentation parameters. We default to 50% for the user-specified Time Mask percentage as that is the optimal percentage found by Baevski et al. (2020).

Single Augmentation Results. We find Rotate 45 and Scale 1.5 to be the strongest augmentation

parameters (as seen in Table 1). Among Rotate 5, 10, 25, and 45, we find that Rotate 45 is our best performer at the 1% (0.801), 10% (0.856), and 100% (0.877) label fractions. Among Scale 1.05, 1.1, 1.5, and 2, we find that Scale 1.5 is our best performer at the 1% (0.795), 10% (0.852), and 100% (0.868) label fractions.

Combination Augmentation Results. We find Rotate 45 + Scale 1.5 + Time Mask to be the best performing combination overall as seen in Table 1. The combination of Rotate 45 + Scale 1.5 performs better than Rotate 45 or Scale 1.5 alone at the 1% (0.807 vs 0.801 & 0.795), 10% (0.861 vs 0.856 & 0.852), and 100% (0.881 vs 0.877 & 0.868) label fractions. When adding Time Mask, we find that Rotate 45 + Time Mask outperforms Rotate 45 at the 1% (0.812 vs 0.801), 10% (0.870 vs 0.856), and 100% (0.881 vs 0.877) label fractions. Similarly, Scale 1.5 + Time Mask outperforms Scale 1.5 at the 1% (0.798 vs 0.795), 10% (0.859 vs 0.852), and 100% (0.878 vs 0.868) label fractions. Rotate 45 + Scale 1.5 + Time Mask underperforms Rotate 45 + Time Mask at the 10% (0.869 vs 0.870) label fraction. However, it is still the best performer at the 1% (0.812) and 100% (0.883) label fractions.




Figure 2: **Per-Diagnosis AUROC Comparisons.** We show the per-diagnosis improvement in AUROC that 3KG (linear evaluation) has over a fully supervised, randomly-initialized model, both trained on 1% of labels. 95% confidence intervals for the improvements on each diagnosis are also shown. A board-certified cardiologist classified the 23 abnormalities into four types: multilead waveform, single-lead waveform, rhythm, and conduction. All improvements on all diagnoses are statistically significant.

4.2. Comparison to Previous Methods

We compare our best performer, 3KG (SimCG-L + VCG + Time Mask), with 8 prior self-supervised learning methods applied to the ECG domain.

Reproduction. To represent previous works as faithfully as possible, we adapt their publicly released codebases to accommodate the PhysioNet 2020 challenge dataset and modify their encoders to use our standardized encoder architecture. The following methods use a 12-channel encoder: Patient Contrastive Learning (PCLR) (Diamant et al., 2021), SimCLR + RRC/TO (Random Resized Crop/Timeout) (Mehari and Strodthoff, 2021), SimCLR + Physio. (Physiological Noise) (Mehari and Strodthoff, 2021), and Contrastive Predictive Coding (CPC) (Mehari and Strodthoff, 2021). In contrast, the following methods use a 1-channel encoder: Signal Transformation (Sarkar and Etemad, 2020), Contrastive Multi-Segment Coding (CMSC) (Kiyasseh et al., 2020), Contrastive Multi-Lead Coding (CMLC) (Kiyasseh et al., 2020), and Contrastive Multi-Segment/Lead Coding (CMSMLC) (Kiyasseh

et al., 2020). Where possible, we pretrain models using their out of the box training procedures and configurations (more details specified in Appendix B). We loaded all pretrained models into our unified codebase for downstream evaluation to ensure that output metric results are identical.

Results. Each of our proposed methods (SimCG-A + VCG + Time Mask, SimCG-L, and 3KG) significantly outperforms all 8 self-supervised baselines, as shown in Table 2. Among our baselines, we find PCLR to be the top performer. SimCG-A + VCG + Time Mask and SimCG-L individually outperform PCLR at the 1% (0.812 & 0.789 vs 0.757), 10% (0.869 & 0.851 vs 0.843), and 100% (0.883 & 0.869 vs 0.857) label fractions. 3KG outperforms PCLR by an even greater difference at the 1% (0.826 vs 0.757), 10% (0.876 vs 0.843), and 100% (0.890 vs 0.857) label fractions. This experiment shows the independent value of SimCG-A + VCG + Time Mask and SimCG-L over our baselines, as well as the additional value of combining these methods in 3KG.

4.3. Per-Pathology Performance Analysis

We explore 3KG at the per-diagnosis level to better understand how our physiologically-inspired self-supervised learning scheme affects the quality of the learned representations. We compare the linear evaluation of our best performer 3KG to a fully supervised, randomly-initialized model where both are trained on the same 1% label fraction. A board-certified cardiologist classified the diagnosis into four different abnormality types: multilead waveform, single-lead waveform, rhythm, and conduction.

Results. As shown in Figure 2, 3KG generally makes greater AUROC gains for conduction (mean = 0.22) and rhythm (mean = 0.17) abnormalities and smaller AUROC gains for single-lead waveform (mean = 0.12) and multilead waveform (mean = 0.04) abnormalities. All improvements across all diagnoses are statistically significant.

5. Discussion

We introduce 3KG, a method to use the spatiotemporal properties of the ECG to generate and select positive views for contrastive learning, and demonstrate the utility of this method on an arrhythmia diagnosis dataset.

Can we improve performance by incorporating physiological constraints into our positive pairs? Yes. Our best pretrained strategy obtains a 19.0% increase in AUROC on the 1% fraction over a supervised baseline, as well as a 9.1% increase over the approach implemented in [Diamant et al. \(2021\)](#), which we found to be the strongest previous self-supervised approach on our dataset.

Our empirical analysis shows that stronger spatial augmentations lead to better learned representations with diminishing returns at higher augmentation strengths. Performance eventually decreases for scaling at the higher end of the range (from 1.5 to 2) but no decrease is observed for rotation, although this may be due to the limited range of sampled strengths. Additionally, we observe that combining spatial augmentations increases performance over any individual spatial augmentation, and that combining spatial augmentations with temporal augmentations produces the strongest results of all. A similar result has been observed by [Chen et al. \(2020a\)](#) in the computer vision domain demonstrating augmentation strength

and composition are both crucial for learning good representations.

To the best of our knowledge, we are the first to investigate how the introduction of an inductive bias through contrastive pretraining affects downstream performance on different disease subgroupings. Compared to a randomly initialized baseline, we find that 3KG makes the greatest gains for conduction and rhythm abnormalities and smaller gains for waveform abnormalities. However, it is hard to attribute these results to a specific component of 3KG due to the various components of our best performing model.

Limitations. Some limitations of this work should be noted. First, we do not perform an exhaustive hyperparameter search to determine the ideal time mask percentage, rotation, and scale amounts. Second, even though our dataset was comprised of multiple independent sets of ECG recordings collected with different institutions and included additional metadata, we did not explore the utility of this additional metadata in our method ([Vu et al., 2021](#)). Finally, we did not formally investigate the relationship between encoder size and performance, instead opting to use a relatively small network in both pretraining and fine-tuning due to computing resource constraints.

In closing, our work demonstrates the benefits of incorporating spatiotemporal augmentations into contrastive learning for ECGs and highlights the potential for similar modality-specific augmentations for other biosignals. The introduction of an inductive bias through physiologically-inspired pretraining is especially relevant for clinicians, who may be more inclined to trust algorithms trained with intuitive procedures that reflect well-understood physiological principles. Strengths of our work include the fact that our approach is easily extensible in that our framework allows flexibility in incorporating other self-supervised strategies; additionally, all data used for experiments are publicly available, as will the codebase be that was used to implement our methods.

References

- Erick A Perez Alday, Annie Gu, Amit J Shah, Chad Robichaux, An-Kwok Ian Wong, Chengyu Liu, Feifei Liu, Ali Bahrami Rad, Andoni Elola, Salman Seyedi, Qiao Li, Ashish Sharma, Gari D Clifford, and Matthew A Reyna. Classification of 12-lead ecgs: the physionet/computing in cardiology challenge 2020. *Physiological Measurement*, 2020. URL <http://iopscience.iop.org/article/10.1088/1361-6579/abc960>.
- Zachi I Attia, Suraj Kapa, Francisco Lopez-Jimenez, Paul M McKie, Dorothy J Ladewig, Gaurav Satam, Patricia A Pellikka, Maurice Enriquez-Sarano, Peter A Noseworthy, Thomas M Munger, et al. Screening for cardiac contractile dysfunction using an artificial intelligence-enabled electrocardiogram. *Nature medicine*, 25(1):70–74, 2019a.
- Zachi I Attia, Peter A Noseworthy, Francisco Lopez-Jimenez, Samuel J Asirvatham, Abhishek J Deshmukh, Bernard J Gersh, Rickey E Carter, Xiaoxi Yao, Alejandro A Rabinstein, Brad J Erickson, et al. An artificial intelligence-enabled ecg algorithm for the identification of patients with atrial fibrillation during sinus rhythm: a retrospective analysis of outcome prediction. *The Lancet*, 394(10201):861–867, 2019b.
- Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations, 2020.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations, 2020a.
- Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners, 2020b.
- Joseph Y. Cheng, Hanlin Goh, Kaan Dogrusoz, Oncel Tuzel, and Erdrin Azemi. Subject-aware contrastive learning for biosignals, 2020.
- Nathaniel Diamant, Erik Reinertsen, Steven Song, Aaron Aguirre, Collin Stultz, and Puneet Batra. Patient contrastive learning: a performant, expressive, and practical approach to ecg modeling, 2021.
- Francis M Fesmire, Robert F Percy, Jim B Bardoner, David R Wharton, and Frank B Calhoun. Usefulness of automated serial 12-lead ecg monitoring during the initial emergency department evaluation of patients with chest pain. *Annals of emergency medicine*, 31(1):3–11, 1998.
- Jean-Bastien Grill, Florian Strub, Florent Althé, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning, 2020.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning, 2020.
- R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization, 2019.
- Rene Jaros, Davide Radek Martinek, and Lukas Danys. Comparison of different electrocardiography with vectocardiography transformations. *Sensors*, 19(14), July 2019. doi: 10.3390/s19143072. URL <https://dx.doi.org/10.3390/S19143072>.
- Harold L. Kennedy. Ambulatory (holter) electrocardiography technology. *Cardiology Clinics*, 10(3):341–359, 1992. ISSN 0733-8651. doi: [https://doi.org/10.1016/S0733-8651\(18\)30218-2](https://doi.org/10.1016/S0733-8651(18)30218-2). URL <https://www.sciencedirect.com/science/article/pii/S0733865118302182>.
- Ambulatory Electrocardiography: Current Clinical Concepts.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- Dani Kiyasseh, Tingting Zhu, and David A. Clifton. Clocs: Contrastive learning of cardiac signals across space, time, and patients, 2020.
- Temesgen Mehari and Nils Strodthoff. Self-supervised representation learning from 12-lead ecg data, 2021.

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.

Pritam Sarkar and Ali Etemad. Self-supervised ecg representation learning for emotion recognition. *IEEE Transactions on Affective Computing*, pages 1–1, 2020. doi: 10.1109/TAFFC.2020.3014842.

Leslie N. Smith. Cyclical learning rates for training neural networks, 2017.

Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding, 2019.

Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning?, 2020.

Geoffrey H Tison, Jeffrey Zhang, Francesca N Delling, and Rahul C Deo. Automated and interpretable patient ecg profiles for disease detection, tracking, and discovery. *Circulation: Cardiovascular Quality and Outcomes*, 12(9):e005289, 2019.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding, 2018.

Yen Nhi Truong Vu, Richard Wang, Niranjana Balachandrar, Can Liu, Andrew Y. Ng, and Pranav Rajpurkar. Medaug: Contrastive learning leveraging patient metadata improves representations for chest x-ray interpretation, 2021.

Zhirong Wu, Yuanjun Xiong, Stella Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance-level discrimination, 2018.

Appendix A. Model Architectures

The convolutional neural network we used for our experiments is described below, with all convolutions being 1-dimensional and all activations being Leaky ReLUs with a negative slope of 0.2:

Number	Layer Info (K : Kernel Size, C : Output Channels, S : Stride)
1	Conv: $K = 7, C = 16, S = 4$ Activation BatchNorm
2	Conv: $K = 7, C = 32, S = 3$ Activation BatchNorm
3	Conv: $K = 5, C = 64, S = 2$ Activation BatchNorm
4	Conv: $K = 3, C = 64, S = 1$ Activation BatchNorm
5	Conv: $K = 3, C = 128, S = 1$ Activation BatchNorm
6	Conv: $K = 3, C = 256, S = 1$ Activation BatchNorm
7	Adaptive Average Pool Flatten

Thus, given an input shape of $[N, 1, 2500]$ (representing a batch of N 5-second samples of 1-lead ECGs sampled to 500Hz), our output shape will be $[N, 256]$.

Appendix B. Baseline Implementation Notes

We made the following changes to baseline methods to allow for the standardization of pretraining data and encoders:

Patient Contrastive Learning (Diamant et al., 2021) The released codebase for this publication did not have training code. As such, the pretraining method was reproduced to the best of our ability and trained on our own codebase.

Contrastive Learning of Cardiac Signals (Kiyasseh et al., 2020) We used the original codebase for this method, changing the encoder used in the original publication to our own.

Contrastive Predictive Coding (Mehari and Strodtzoff, 2021) We used the original codebase for this method. However, this model originally used a convolutional neural network encoder which passed

its outputs into an LSTM predictor. To normalize output embedding dimension across methods, we reduced the input size of the LSTM layer to 256. Additionally, we reduced the number of hidden units in the LSTM layer to 85 in order to match the input dimensionality required for the downstream task. We did not use the original models for SimCLR + RRC/TO, SimCLR + Physio. Noise, and CPC. We use

SimCLR + RRC/TO and Physio. (Mehari and Strodt Hoff, 2021) We used the original codebase for this method, changing the encoder used in the original publication to our own.

Signal Transformation (Sarkar and Etemad, 2020) We reimplemented the codebase in PyTorch, changing the encoder to our own.

- Complete right bundle branch block and right bundle branch block
- Premature atrial contraction and supraventricular premature beats
- Premature ventricular contractions and ventricular premature beats
- 1st degree AV block and prolonged PR interval

This turns the original Physionet 2020 challenge task into the 23-arrhythmia classification mentioned in section 3.3.3.

Appendix C. Detailed Dataset Information

C.1. Class Distribution

Class	Train	Val	Test
1st degree AV block	0.0453	0.0451	0.0454
Atrial fibrillation	0.0576	0.0573	0.0577
Atrial flutter	0.0052	0.0053	0.0052
Bradycardia	0.0048	0.0048	0.0048
Complete right bundle branch block	0.0509	0.0505	0.0511
Incomplete right bundle branch block	0.0267	0.0265	0.0266
Left anterior fascicular block	0.0300	0.0300	0.0294
Left axis deviation	0.1009	0.1024	0.0990
Left bundle branch block	0.0173	0.0171	0.0173
Low QRS voltages	0.0092	0.0092	0.0091
Nonspecific intraventricular conduction disorder	0.0165	0.0171	0.0163
Pacing Rhythm	0.0050	0.0049	0.0050
Premature atrial contraction	0.0322	0.0321	0.0322
Premature ventricular contractions	0.0091	0.0092	0.0091
Prolonged QT interval	0.0251	0.0249	0.0251
Q wave abnormal	0.0166	0.0181	0.0170
Right axis deviation	0.0069	0.0082	0.0070
Sinus arrhythmia	0.0206	0.0204	0.0206
Sinus bradycardia	0.0391	0.0389	0.0392
Sinus rhythm	0.3454	0.3432	0.3467
Sinus tachycardia	0.0398	0.0395	0.0399
T wave abnormal	0.0774	0.0768	0.0776
T wave inversion	0.0184	0.0183	0.0185

C.2. Equivalent Classes

We evaluated our pre-trained encoders on the PhysioNet 2020 challenge task. Although the training data had 111 classes, we chose to classify only SNOMED CT codes that were included in the challenge evaluation metric.

The original Physionet 2020 challenge task was 27-arrhythmia classification. However, we coalesced the following diagnoses after a cardiologist determined them to be substantially equivalent: