

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 CLEF: CLINICALLY-GUIDED CONTRASTIVE LEARNING FOR ELECTROCARDIOGRAM FOUNDATION MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

The electrocardiogram (ECG) is a key diagnostic tool in cardiovascular health. Single-lead ECG recording is integrated into both clinical-grade and consumer wearables. While self-supervised pretraining of foundation models on unlabeled ECGs improves diagnostic performance, existing approaches do not incorporate domain knowledge from clinical metadata. We introduce a novel contrastive approach that utilizes an established clinical risk score to adaptively weight negative pairs: *clinically-guided contrastive learning*. It aligns the similarities of ECG embeddings to clinically meaningful differences between subjects, with an explicit mechanism to handle missing metadata. Using 12-lead ECGs from 161K patients in MIMIC-IV dataset, we pretrain single-lead *ECG foundation models* at three scales, collectively called *CLEF*, using only routinely-collected metadata without requiring per-sample ECG annotations. We evaluate *CLEF* on 18 clinical classification and regression tasks across 7 held-out datasets, and benchmark against 5 foundation model baselines and 3 self-supervised algorithms. When pretrained on 12-lead ECG data and tested on lead-I data, *CLEF* outperforms self-supervised foundation model baselines: the medium-sized *CLEF* achieves average AUROC improvements of at least 2.6% in classification and average reductions in MAEs of at least 3.2% in regression. Comparing with existing self-supervised learning algorithms, *CLEF* improves the average AUROC by at least 1.8%. Moreover, when pretrained only on lead-I data for classification tasks, *CLEF* performs comparably to the state-of-the-art ECGFounder, which has been trained in a supervised manner. Overall, *CLEF* allows more accurate and scalable single-lead ECG analysis, advancing remote health monitoring. We will publish our code and pretrained *CLEF* models.

1 INTRODUCTION

The electrocardiogram (ECG) (Geselowitz, 1989) captures the heart’s electrical activity as a sequence of voltage fluctuations (Lilly, 2012). While ECG interpretation is a clinical task, ECG-based AI started matching clinician performance, and enabled large-scale applications such as detecting heart conditions from wearables. Furthermore, Yao et al. (2021) found out that clinicians using AI were 30% more accurate in diagnosing left ventricular dysfunction from ECGs than those without AI support. The rise of wearables and mobile health devices with single-lead ECGs enables continuous monitoring beyond clinical settings, supporting early cardiac event detection, long-term tracking, and proactive interventions. Nevertheless, accurate analysis of single-lead ECGs remains a challenge due to the reduced spatial information and the increased noise and motion artifacts (Khamis et al., 2016; Halvaei et al., 2021; Khunte et al., 2023). Prior work has shown that single-lead ECG devices can be used to diagnose atrial fibrillation Svennberg et al. (2015), and that spatial correlations among ECG leads can help infer missing leads (Nelwan et al., 2004). Single-lead ECG is used for practical diagnoses, especially in wearable contexts (Qin et al., 2023; Li et al., 2024).

Collecting sufficiently large, human-annotated, single-lead ECGs is impractical, given the variability in downstream tasks across devices and health conditions. ECGs from healthcare organizations are usually provided in the standard 12-lead format but contain limited annotations. A more practical solution is to pretrain on 12-lead ECGs to learn generalizable representations, then fine-tune on smaller, task-specific data curated for the target health condition. Often, pretraining is performed using *contrastive learning* with data augmentation (Chen et al., 2020; Soltanieh et al., 2022), encouraging the representations of similar samples to be as close to one another while pushing apart those of

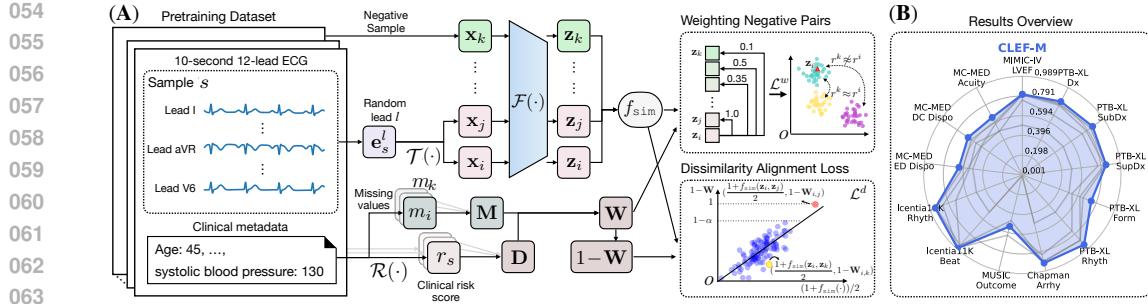


Figure 1: *CLEF*’s framework and performance overview (see §2 for notations). (A) Our clinically-guided contrastive pretraining. Key components include a negative weighting loss \mathcal{L}^w and a dissimilarity alignment loss \mathcal{L}^d that work in tandem to guide contrastive learning with clinical knowledge. (B) Spider plot on AUROC performance of *CLEF*-M (our medium-sized model) across 13 downstream classification tasks. Baseline performances are in gray lines (see §4).

This not only uses information from within samples (e.g. data augmentations), but also uses contextual factors, which can provide important cues for similarity that are not directly observable in the input. Contrastive learning for images is guided by the spatial proximity of image viewpoints or content (Thoma et al., 2020a;b), or by sample identities (Haslum et al., 2024). Prior work also incorporates subject and signal attributes in contrastive learning (Kim et al., 2025), but did not incorporate clinical knowledge, nor handle missing metadata.

Metadata recorded alongside ECGs, such as patient demographics, provides a potential source of contextual data. Some use metadata as model inputs (Erturk et al., 2025), while others have incorporated metadata prediction as part of the pretraining objective (Wang et al., 2024), or as downstream prediction targets (e.g., predicting gender or age from ECGs (Li et al., 2024)). Since gender or age is not a primary objective, this merely demonstrates the model’s ability to distinguish between patient groups. In this paper, we take a more integrative strategy to employ metadata as contextual information in contrastive pretraining, since different people might be associated with different risks of a particular health outcome. We leverage metadata in the form of *clinically validated risk scores* to pretrain a foundation model (FM) for single-lead ECGs. Risk scores estimate a person’s risk of experiencing an adverse health outcome in a specified time period, such as risk of developing cardiovascular disease in the next ten years (Conroy et al., 2003; working group and risk collaboration, 2021a;b). They take various metadata variables as inputs (e.g., patient demographics, past medical history, and comorbidities), and output a quantitative assessment of the risk of adverse outcomes. Risk scores can guide clinical decisions Hughes et al. (2023), such as whether to prescribe a drug (Lip et al., 2010), or perform invasive procedures (Fox et al., 2006). Our **contributions** are:

1. **Clinically-guided contrastive learning for ECG representation learning.** We leverage metadata to calculate clinically informative risk scores that provide flexible similarity relationships between unlabeled ECG samples. Unlike classic contrastive approaches that rely on binary notions of similarity/dissimilarity, our method incorporates clinical risk scores as soft guidance, enabling the model to learn richer and more nuanced representations.
2. **Extensive empirical analysis and benchmarking for ECG FMs.** We: (i) evaluate downstream performance on lead I/II ECG classification and regression, (ii) compare against 12-lead ECG approaches, (iii) assess robustness across different model architectures and pretraining methods, (iv) perform ablation studies on loss components and handling missing metadata, and (v) use linear probing analysis to assess representation quality.
3. **Effectiveness across diverse downstream tasks.** We evaluate on a wide range of applications. *CLEF* consistently outperforms strong baselines, achieving an average improvement of 3.1% in classification AUROC and a reduction of 2.9% in regression MAE.

2 CLINICALLY-GUIDED CONTRASTIVE PRETRAINING

Let $\mathcal{D} = \{(\mathbf{e}_s = (\mathbf{e}_s^1, \dots, \mathbf{e}_s^{12}), \mathbf{a}_s)\}_{s=1}^N$ denote the *pretraining dataset*, where $\mathbf{e}_s^l \in \mathbb{R}^t$, of length t , represents the ECG signal from lead l of the s -th sample, $\mathbf{a}_s = \{\mathbf{a}_s^1, \dots, \mathbf{a}_s^A\}$ denotes demographic or clinical *metadata* (i.e., attributes like age or blood pressure) associated with the subject and

not labels specific to the current ECG recording, and N is the number of 12-lead ECG samples. Importantly, certain attributes might be missing in some samples' metadata. Let \mathcal{R} denote a standard, clinically-validated function that combines a subject's metadata into a numerical estimate of the likelihood of a future health outcome (e.g. sudden cardiac arrest). We define $r_s = \mathcal{R}(\alpha_s)$ as a *risk score* associated to the sample \mathbf{e}_s . Our objective is to train a single-lead ECG FM $\mathcal{F}(\cdot)$ on \mathbb{D} such that the geometry of the *embedding space*, i.e., outputs $\mathbf{z} = \mathcal{F}(\mathbf{e}_s^l) \in \mathbb{R}^h$, with a dimension of h , reflects clinically meaningful similarities according to the risk score function $\mathcal{R}(\cdot)$.

To measure the success of our objective, we attach a linear layer $\mathcal{G}(\cdot)$ on top of the embeddings of condition-specific single-lead ECGs produced by $\mathcal{F}(\cdot)$. Let $\mathbb{C} = \{(\mathbf{e}_s^l, y_s)\}_{s=1}^M$ denote a *downstream evaluation dataset*, where $\mathbf{e}_s^l \in \mathbb{R}^t$, represents the ECG signal from lead l , and y_s indicates the ground-truth label associated to that ECG signal. Note that the lead is fixed across all samples in the test set (e.g. lead I). We evaluate $\mathcal{F}(\cdot)$ under two downstream scenarios: (1) *fine-tuning*, where all parameters of $\mathcal{F}(\cdot)$ and $\mathcal{G}(\cdot)$ are updated during training on the training split of \mathbb{C} , and (2) *linear probing*, where the pretrained $\mathcal{F}(\cdot)$ is kept frozen and only the linear head $\mathcal{G}(\cdot)$ is updated to map the embeddings $\mathbf{z} = \mathcal{F}(\mathbf{e}_s^l)$ to task-specific labels. To maintain a fair comparison, we adopt a consistent set of hyperparameters across all experiments; further details are provided in Appendix C.1.

For binary labels $y \in \{0, 1\}$, we report the AUROC metric, which measures the model's ability to discriminate between the positive and negative classes by computing the area under the Receiver Operating Characteristic (ROC) curve, representing the trade-off between true positive rate and false positive rate across all possible classification thresholds. For categorical labels $y \in \{1, \dots, K\}$ where K is the number of classes, we report the macro AUROC score using the one-vs-rest extension, where each class is compared against all others and the resulting scores are macro-averaged (i.e. unweighted mean). For numerical labels $y \in \mathbb{R}_0$, we report the mean absolute error (MAE) between the predicted and true values, providing an interpretable metric of prediction accuracy in the same units as the target variable. Performance is summarized as a percentage improvement over baseline.

2.1 CONTRASTIVE LEARNING

Contrastive Loss. We revisit contrastive learning (Chen et al., 2020). Given a sample $\mathbf{e}_s^l \stackrel{\text{i.i.d.}}{\sim} \mathbf{e}_s \stackrel{\text{i.i.d.}}{\sim} \mathbb{D}$, an independent stochastic transformation $\mathcal{T}(\cdot)$ is applied to \mathbf{e}_s^l to obtain an augmented *view*, denoted as $\mathbf{x}_i = \mathcal{T}(\mathbf{e}_s^l)$. Similarly, we generate another view $\mathbf{x}_j = \mathcal{T}(\mathbf{e}_s^l)$ forming a *positive pair* $(\mathbf{x}_i, \mathbf{x}_j)$ in contrastive learning. A batch of B samples yields $2B$ augmented views, such that any $k \in \{1, 2, \dots, 2B\} \setminus \{i, j\}$ forms *negative pairs* with both \mathbf{x}_i and \mathbf{x}_j , i.e., $(\mathbf{x}_i, \mathbf{x}_k)$ and $(\mathbf{x}_j, \mathbf{x}_k)$. Using a model $\mathcal{F}(\cdot)$, the augmented views are encoded into embeddings denoted as $\mathbf{z}_i, \mathbf{z}_j$, and $\mathbf{z}_k \in \mathbb{R}^h$. For each pair of $(\mathbf{x}_i, \mathbf{x}_j)$, the classic contrastive loss, termed NT-Xent (Chen et al., 2020), is given by

$$\mathcal{L}_{i,j} = -\log \frac{\exp(f_{\text{sim}}(\mathbf{z}_i, \mathbf{z}_j) / \tau)}{\sum_{k=1}^{2B} \mathbb{1}_{[k \neq i]} \exp(f_{\text{sim}}(\mathbf{z}_i, \mathbf{z}_k) / \tau)}, \quad (1)$$

where $f_{\text{sim}}(\mathbf{z}_i, \mathbf{z}_j) = \frac{\mathbf{z}_i^\top \mathbf{z}_j}{\|\mathbf{z}_i\| \cdot \|\mathbf{z}_j\|} \in [-1, 1]$ is cosine similarity, and τ is a temperature hyperparameter.

Weighted Contrastive Loss. While Eq. (1) treats all negative pairs equally, recent works highlight that negative pairs could contribute unequally to representation learning Robinson et al. (2021); Li et al. (2023); Zhuang et al. (2024); Yang et al. (2024). Conceptually, in the embedding space, negative samples are pushed away from a positive sample according to their similarity. Specifically, the less similar a negative sample is to the positive, the more it is pushed away. By introducing a weighting term \mathbf{W}_{ik} for each pair of negative embeddings $(\mathbf{z}_i, \mathbf{z}_k)$, we can adjust the relative importance of negatives according to their similarity to the positive, having

$$\mathcal{L}_{i,j}^w = -\log \frac{\exp(f_{\text{sim}}(\mathbf{z}_i, \mathbf{z}_j) / \tau)}{\sum_{k=1}^{2B} \mathbb{1}_{[k \neq i]} \mathbf{W}_{ik} \exp(f_{\text{sim}}(\mathbf{z}_i, \mathbf{z}_k) / \tau)}. \quad (2)$$

Eq. (2) generalizes Eq. (1) by replacing the uniform treatment of negatives with an adaptive weighting.

2.2 CLINICAL METADATA

Cardiovascular Risk Score. A risk score, quantified on clinical and demographic attributes, is the likelihood of experiencing an adverse outcome. They translate complex physiological factors into

interpretable values that guide clinical actions and interventions. The Systematic Coronary Risk Evaluation 2 score (SCORE2) commonly used by healthcare professionals, estimates the 10-year risk of cardiovascular disease events (death, myocardial infarction, or non-fatal stroke). SCORE2 is obtained based on 7 variables: age (a_s^1), gender (a_s^2), smoking status (a_s^3), systolic blood pressure (SBP, a_s^4), diabetes status (a_s^5), total cholesterol (a_s^6), and high-density lipoprotein cholesterol (a_s^7). Let $\mathbf{u}_s = [u_s^1, \dots, u_s^6] = g(\mathbf{a}_s)$ denote the vector after standardization of the ordered metadata feature $\mathbf{a}_s = [a_s^1, a_s^3, \dots, a_s^7]$. For each sample \mathbf{e}_s , we assign a risk score $r_s \in [0, 1]$, defined by the risk score function $\mathcal{R}(\cdot)$ as that SCORE2 value of \mathbf{e}_s :

$$r_s \stackrel{\text{def}}{=} \mathcal{R}(\mathbf{a}_s) = 1 - S_0(t)^{\exp(\mathbf{b}_1 \mathbf{u}_s^\top + u_s^1 \mathbf{b}_2 \mathbf{u}_s^\top)}. \quad (3)$$

where $S_0(t)$ is the baseline survival ¹ function at year t , and $\mathbf{b}_1, \mathbf{b}_2$ are the variable's coefficient estimated from the European cohort data covering over 600K individuals with more than 30K cardiovascular events (values varying by age and gender). Given that augmentations do not change metadata, for $\mathbf{x}_i, \mathbf{x}_j$ of sample s , we have $r^i = r^j = r_s$ (r^i, r^j are risk scores for $\mathbf{x}_i, \mathbf{x}_j$). We provide further details on calculating SCORE2 in Appendix C.6. **The CLEF framework supports any validated risk score. While we use SCORE2, we note that its development on European populations might limit its generalisability, and the optimal score may differ by application (see Appendix F).**

Handling Missing Metadata. MIMIC-IV-ECG (Gow et al., 2023) contains 161,352 unique subjects, each with multiple 12-lead ECG recordings. For each subject, the dataset reports only three variables: age, gender, and systolic blood pressure. This is a common challenge in real-world settings, where certain metadata variables are often unavailable either at a specific institution or for particular patients. Such variables are typically collected independently rather than at the time of the ECG recording. To address this issue, we assign different *multipliers* to negative pairs $(\mathbf{x}_i, \mathbf{x}_k)$ based on the number of metadata variables available when calculating their risk scores (r^i, r^k) . Specifically, we calculate

$$\mathbf{M}_{ik} = \exp\left(-\frac{A - m_i}{A} \times \frac{A - m_k}{A}\right), \quad (4)$$

where A is the number of variables (e.g., $A = 7$ in SCORE2), and m_i/m_k are the number of missing variables for $\mathbf{x}_i/\mathbf{x}_k$. Thus, $\mathbf{M}_{ik} \in (0, 1]$ indicates the level of missing metadata for a pair $(\mathbf{x}_i, \mathbf{x}_k)$, which functions as a relative reliability adjustment to ensure that only well-supported risk differences meaningfully shape the embedding geometry, we further elaborate this in section 2.3.

2.3 GUIDING REPRESENTATION LEARNING WITH CLINICAL RISK SCORES

Risk Score Dissimilarity. We aim to enhance weighted contrastive loss in Eq. (2) by using risk scores in Eq. (3) for pretraining our ECG FM $\mathcal{F}(\cdot)$. Our objective is to guide $\mathcal{F}(\cdot)$ toward learning embeddings that capture clinically relevant patterns, building a latent space for ECG signals where the distance between embeddings reflects dissimilarities in risk scores. For a negative pair $(\mathbf{x}_i, \mathbf{x}_k)$ with corresponding (r^i, r^k) , let $\delta_{ik} = (r^i - r^k)^2$. We define a negative pair's *dissimilarity*:

$$\mathbf{D}_{i,k} = (1 - \alpha) \frac{\delta_{ik} - \delta^{min}}{\delta^{max} - \delta^{min}} + \alpha, \quad \text{where } \delta^{max} = \max_{i,k}(\delta_{ik}) \quad \text{and } \delta^{min} = \min_{i,k}(\delta_{ik}). \quad (5)$$

Parameter $\alpha \in [0, 1]$ controls the minimum distance between the positive pair and all negative pairs. Dissimilarity to all negative pairs stays within $[\alpha, 1]$, having $\mathbf{D}_{i,j} = 0$ only for the positive pair. For $(\mathbf{x}_i, \mathbf{x}_k)$, even when their metadata indicates $r^i = r^k$, we have their dissimilarity $\mathbf{D}_{i,k}$ at least α .

Weighting Negative Pairs. By combining Eq. (4) and (5) via the Hadamard product, we define the *weight matrix* $\mathbf{W} = \mathbf{D} \odot \mathbf{M}$, where each entry is given by $\mathbf{W}_{ik} = \mathbf{D}_{i,k} \cdot \mathbf{M}_{i,k} \in [0, 1]$. This value represents the weight assigned to pushing the embedding of \mathbf{x}_k away from that of \mathbf{x}_i . Intuitively, the weight increases when the risk scores of \mathbf{x}_i and \mathbf{x}_k differ more; conversely, it decreases when the risk scores are similar. **Importantly, when many metadata variables are missing, and SCORE2 would rely heavily on default imputations (see Appendix C.6), the resulting risk differences are thus less trustworthy. In such cases, the corresponding \mathbf{M}_{ik} between the negative pairs moves the weighting closer to a uniform SimCLR-like behavior, preventing uncertain or noisy risk differences from exerting undue influence on the contrastive objective.** For each batch of data, we calculate our *clinically-guided contrastive loss* by

¹The probability of not experiencing a cardiovascular event over 10 years for a reference individual.

$$216 \\ 217 \quad \mathcal{L}^w = -\frac{1}{2B} \sum_{i=1}^{2B} \mathcal{L}_{i,j}^w = -\frac{1}{2B} \sum_{i=1}^{2B} \log \frac{\exp(f_{\text{sim}}(\mathbf{z}_i, \mathbf{z}_j) / \tau)}{\sum_{k=1}^{2B} \mathbf{1}_{[k \neq i]} \mathbf{W}_{ik} \exp(f_{\text{sim}}(\mathbf{z}_i, \mathbf{z}_k) / \tau)} . \quad (6)$$

$$218 \\ 219 \\ 220 \\ 221 \\ 222$$

Dissimilarity Alignment Loss. We compute the mean squared error loss between the cosine similarity of each embedding and its corresponding weight, averaged across all pairs in the batch.

$$223 \quad \mathcal{L}^d = \frac{1}{B^2} \sum_{i,j} \left(\frac{1 + f_{\text{sim}}(\mathbf{z}_i, \mathbf{z}_j)}{2} - (1 - \mathbf{W}_{ij}) \right)^2 . \quad (7)$$

$$224 \\ 225$$

Note that cosine similarity $f_{\text{sim}} \in [-1, 1]$ is rescaled to $[0, 1]$. This encourages the model to map clinically similar ECG signals to nearby embeddings, while pushing dissimilar signals apart. This supports clinical use, where models must highlight differences between ECG signals to distinguish clinical categories (e.g., diagnoses or prognoses). Our final objective is: $\mathcal{L} = \mathcal{L}^w + \mathcal{L}^d$, where, in this paper, the contribution of the two losses are considered equal (details in Appendix D.6).

3 EXPERIMENTAL EVALUATION

We evaluate against (i) existing FMs and models pretrained on ECG data, (ii) widely-used self-supervised learning algorithms, and (iii) a state-of-the-art (SOTA) supervised foundation model.

Foundation Model \mathcal{F} . We use ResNeXt1D (Hong et al., 2020)², built on ResNeXt Xie et al. (2017), using one-dimensional convolutional filters, and widely used as a benchmark for ECG processing Li et al. (2024). We evaluate three model configurations: small (448K parameters), medium (30.7M parameters), and large (296M parameters), which differ in network depth, width, and complexity. A more detailed model structure description can be found in Appendix C.3 and Table S7.

Pretraining Dataset \mathcal{D} . We train $\mathcal{F}(\cdot)$ using the MIMIC-IV-ECG dataset (Gow et al., 2023), which contains 161,352 unique patients, each with multiple ECG recordings. To ensure that signals from the same patient are not considered as negative pairs during model pretraining, we use the first ECG recording of each patient. Each ECG record e is sampled at 500Hz over 10 seconds, resulting in a sequence length of $t = 5,000$. Following the common practice, we apply a Butterworth bandpass filter between 0.67 and 40 Hz, followed by z-score normalization for each sample.

Stochastic Data Augmentation \mathcal{T} . To simulate real-world signal perturbations, we apply some noise derived from free-living³ ECG recordings, including (i) muscle noise, (ii) movement artifacts, (iii) baseline wander, (iv) white noise, or (v) no perturbations. Function \mathcal{T} is a random selection from these five choices, each with equal probability of $p = 0.2$. For details, see Appendix C.2.

Random Lead Selection. Our *CLEF* can be fine-tuned to any lead, which is crucial since devices and health-monitoring scenarios vary. Also, many wearables only approximate a lead; e.g., smartwatches approximate lead I, and the target lead may not be known at pretraining time. With this motivation, we do not restrict training to a specific lead as pretraining exclusively on a specific lead (explored in section 4, Figure 4) requires one separate model per lead and is less flexible for cross-device use. Instead, for each sample in a batch, we randomly select one of the 12 leads and apply stochastic augmentation, thus every batch contains diverse samples from multiple leads. In this way, the model is collectively trained across all 12 leads, allowing \mathcal{F} to generalize and later be fine-tuned to any lead.

Downstream Datasets \mathcal{C} . We evaluate on various downstream tasks from cardiovascular conditions, including 3 well-established benchmarks used in (Li et al., 2024; Na et al., 2024), and 4 newly curated datasets from wearables or emergency visits. Tasks include multi-label diagnostic (e.g. heart block, myocardial infarction, hypertrophy), form (i.e. wave morphology), and rhythm statements on PTB-XL dataset (Wagner et al., 2020), left ventricular ejection fraction (LVEF) regression and classification at 50% threshold on MIMIC-IV-ECG dataset (Gow et al., 2023), and multi-label disease classification on the Chapman dataset (Zheng et al., 2020). We also include MC-MED (Kansal et al., 2025) and Aurora BP (Mieloszyk et al., 2022) for blood pressure estimation, the MUSIC dataset (Martin-Yebra et al., 2025) for long-term cardiovascular outcomes such as sudden cardiac death (SCD), and

²The PyTorch implementation of the model is adapted from github.com/hsd1503/resnet1d.

³The noise data is available at: physionet.org/content/ecg-ppg-simulator-arrhythmia.

270
 271
 272
 273
 274
 275
 276 Table 1: Results for finetuning on lead I (left) and lead II (right) ECGs. AUROCs are reported for 7
 277 clinical tasks: LVEF refers to left ventricular ejection fraction; Dx denotes diagnosis; SubDx denotes
 278 subdiagnosis; SupDx denotes superdiagnosis; Form refers to waveform morphology; Rhyt denotes
 279 rhythm; Arrhy denotes arrhythmia. Best result **bolded**. The second best underlined.

Dataset Task	Lead I						Lead II							
	MIMIC-IV LVEF	Dx	SubDx	SupDx	Form	Rhyt	MIMIC-IV LVEF	Dx	SubDx	SupDx	Form	Rhyt		
Moirai 2024 (91M)	.4968	.5003	.4966	.5011	.5158	.5031	.4992	.5000	.5014	.5001	.4993	.5003	.4988	.4987
Moment 2024 (125M)	.7763	.7780	.7627	.7700	.6322	.8243	.8150	.8024	.7896	.7788	.8260	.6164	.8754	.8905
ST-MEM 2024 (85M)	.7751	.7763	.7639	.7800	.5724	.7549	.8113	.7499	.6764	.7284	.7637	.5319	.7489	.8228
KED 2024 (8M)	<u>.8330</u>	.8390	.8332	.8302	.6696	.8887	.8897	.8073	.8072	.8304	.8407	.6492	.8736	.8941
<i>CLEF-S</i> (448K)	.8170	.8268	.8448	.8452	.6738	.9304	.9033	.8166	<u>.8205</u>	.8486	<u>.8445</u>	.6913	.9305	.9047
<i>CLEF-M</i> (30.7M)	.8083	.8292	.8566	.8430	.7409	.9361	.9061	.8079	.8307	<u>.8555</u>	.8446	.7378	.9326	.9089
<i>CLEF-L</i> (296M)	.7858	.8472	.8397	.8273	.7162	.9325	.9010	.8194	.8193	.8438	.8409	.7478	.9512	.9011
ECGFounder 2024 (76.3M)	.8512	<u>.8457</u>	.8500	.8376	.7626	.9501	.9090	.8312	.8024	.8244	.8437	.6823	.9229	.8755

283
 284 Icentia11K (Tan et al., 2019), a large-scale continuous wearable ECG dataset supporting beat and
 285 rhythm classification. More details are provided in Appendix B and C.1, respectively.

286
 287 **Foundation model baselines.** (1) ST-MEM (Na et al., 2024): a spatio-temporal masked autoencoder
 288 trained to reconstruct randomly masked ECGs. (2) KED (Tian et al., 2024): that aligns ECG
 289 signals with textual reports, enabling joint ECG–text representations. (3) Moirai (Woo et al., 2024):
 290 a forecasting model for time series, supporting both univariate and multivariate prediction. (4)
 291 Moment (Goswami et al., 2024): a transformer-based model designed for univariate time series tasks.

292
 293 **Self-supervised learning baselines.** (1) SimCLR (Chen et al., 2020): a canonical contrastive
 294 learning framework, as described in §2, Eq. (1). (2) BYOL (Grill et al., 2020): an augmentation-based
 295 approach that eliminates negative pairs by training another network to predict the target network’s
 296 representation of the same signal. (3) MoCo (He et al., 2020): maintains a dynamic memory bank to
 297 store representations from past batches, reducing the reliance on large in-batch negatives.

298
 299 **Supervised baseline.** Concurrent to this work, Harvard–Emory researchers released the HEEDB
 300 dataset (Ghanta et al., 2025) of 12-lead ECG recordings with 10M ECGs from 1.8M patients
 301 annotated with 150 diagnostic categories. While we could not get access to the dataset, they released
 302 a supervised ECG FM, ECGFounder (Li et al., 2024), trained for multi-label classification over the
 303 150 categories. We run ECGFounder on our evaluation datasets, and since it is trained on a large
 304 labeled dataset, we consider it as an upper bound relative to semi-supervised methods, including our
 305 proposed *CLEF*. Implementation details for all baselines are provided in Appendix C.4 and C.5.

306 4 RESULTS

307
 308 **Finetuning on Lead I & II (in-clinic datasets).** The AUROC results are detailed in Table 1 (with
 309 further details in Appendix D.1). *CLEF* outperforms the best baseline across both leads, with the best-
 310 performing *CLEF* variant for each task achieving an average AUROC improvement over the strongest
 311 baseline for that task of 3.1% on lead I, and 4.8% on lead II. On lead I, the *CLEF-S*, -M, and -L
 312 variants achieve improvements of 1.0%, 2.6%, and 1.3%, respectively, over the best performing
 313 baseline (1.6% on average), while on lead II, gains were even stronger, achieving 2.8%, 4.0%, and
 314 4.2% (3.7% average) performance gain. KED was the best-performing baseline in all but one task,
 315 and *CLEF-M* shows statistically significant superiority over KED ($p = 0.010$, by paired t-test). *CLEF*
 316 outperforms KED by $\geq 1.0\%$ on lead I and $\geq 2.8\%$ on lead II across all 3 variants, demonstrating
 317 its strong potential for pretraining. Additionally, the averaged confidence interval across all *CLEF*
 318 variants and tasks is ± 0.01 , showcasing *CLEF*’s consistent performance. *CLEF* performs particularly
 319 well on pattern identification tasks (form and rhythm tasks of PTB-XL, and the arrhythmia task of
 320 Chapman), with its best variant improving AUROC by 5.9% (lead I) and 8.5% (lead II) on average
 321 over the strongest baselines. These results suggest that *CLEF* excels at capturing both morphological
 322 features of individual beats and long-range rhythm dependencies across the signal.

323
 324 *CLEF* maintains robust performance across leads, with *CLEF-S*, -M, and -L differing by 0.3%, 0.0%,
 325 and 1.4% on average in AUROC across 7 tasks. In contrast, ST-MEM and KED which are specifically
 326 pretrained for ECG tasks, favored lead I, with AUROCs on lead II an average of 4.2% and 1.5%
 327 lower, respectively. When compared to the supervised model ECGFounder, *CLEF* performs less

well on lead I (AUROCs lower by 2.9%, 1.5%, and 2.7% for *CLEF*-S, -M, and -L, respectively) but better on lead II (AUROC increases by 1.3%, 2.5%, and 2.6%). This is primarily because ECGFounder’s performance drops 3.8% on lead II, compared to lead I (on which the single-lead version of ECGFounder was trained (Li et al., 2024)), while *CLEF* maintains performance across leads. Although the unsupervised *CLEF* did not perform as well as this SOTA supervised model on the lead for which that model was trained, it did achieve improved performance on another lead. Notably, our single-lead *CLEF* even outperforms some baselines when those baselines are pretrained and evaluated on 12-lead ECGs. Specifically, *CLEF*-S (evaluated on lead I & II) outperforms ST-MEM on 5 out of 7 tasks (Details are reported in Appendix D.2 and Table S10).

Finetuning for Single-lead ECG (out-of-clinic datasets). We use single-lead ECGs from MUSIC, Icentia11k, and AuroraBP collected by wearables, and MC-MED dataset of Emergency Department.

Classification. The AUROC results are detailed in Table 2 (for classification tasks) and Table 3 (for regression tasks). In classification, *CLEF* outperforms all semi-supervised baselines, with the best *CLEF* variant improving AUROC by 2.5% on average over the strongest baseline of the task. The best-performing *CLEF* variant for each task achieves an average AUROC improvement over the strongest baseline for that task of 2.5%. Individually, *CLEF*-S, -M, and -L achieve averaged AUROC improvements of 2.8%, 6.7%, and 2.5% over the best baseline (KED). On Icentia11K Beat and Rhythm tasks, the best-performing *CLEF* variant on each task achieves improved AUROC by 1.9% over the best baselines. When looking at each *CLEF* variant, *CLEF*-S, -M, and -L achieve average AUROC improvements of 1.5%, 2.5%, and 3.7% over the best baseline (KED), highlighting the effectiveness of *CLEF* in capturing quality ECG representations for wearable devices. However, Moment outperformed *CLEF* on MUSIC SCD prediction, the only case where a general time series model surpassed an ECG-specific model; likely because SCD risk depends on long-range and subtle temporal patterns in the ECG, making it more of a forecasting problem than a short-term classification.

Regression. We evaluate on the MIMIC-IV LVEF task (explored in prior work (Li et al., 2024)), and also systolic and diastolic blood pressure (SBP and DBP) prediction on the MC-MED and AuroraBP datasets. Target labels are first z-scored based on the statistics of the training set, and then at test time, predicted values are transformed back to the original scale of the labels for evaluation. See Appendix C.1 for details of our experimental setup. The mean absolute error (MAE) is reported in Table 3. The best-performing *CLEF* for each task achieves an average MAE reduction in comparison to the strongest baseline for that task of 2.9%. The *CLEF*-S and *CLEF*-M variants outperform all baselines, achieving average MAE reductions of 3.2% compared to the best baseline (ST-MEM). In contrast, *CLEF*-L outperforms 3 of the 4 baselines (yet MAE is 2.3% higher than ST-MEM on average). *CLEF* also outperforms the supervised model ECGFounder, with *CLEF*-S, -M, and -L achieving lower MAEs than ECGFounder by 5.5%, 5.6%, and 0.2% on average.

Linear probing. To evaluate the quality of the representations produced by *CLEF*, the parameters of FMs were kept frozen, and a linear classifier was trained on the output embeddings. The AUROCs are summarized in Figure 2 for *CLEF*-M and 4 baselines. Full results including those for the other 2 variants (*CLEF*-S and -L) are provided in Table S11 (Appendix D.3). The best-performing *CLEF* for each task outperforms the strongest baseline for that task of +7.3%. *CLEF*-M and *CLEF*-L both outperform all baselines, with average AUROC improvements of 8.5% and 9.9% respectively over

Table 2: Finetuning using single-lead ECGs. AUROCs are for: SCD outcome prediction, beat and rhythm recognition, disposition after emergency department visit (ED dispo) and after hospital stay (DC dispo), and ED triage acuity. Best result **bolded**. The second best underlined.

Model	MUSIC		Icentia11K		MC-MED		
	SCD	Beat	Rhythm	ED Dispo	DC Dispo	Acuity	
Moira	.0106	.4998	.4927	.5004	.4996	.4992	
Moment	.6223	.9764	.8540	.5488	.5042	.5782	
ST-MEM	.5389	.9452	.7115	.5323	.5645	.4943	
KED	.4701	.9424	.9228	.5992	.6200	.5978	
<i>CLEF</i> -S	.5493	<u>.9801</u>	.9135	.6065	.6261	.5650	
<i>CLEF</i> -M	.5304	.9792	<u>.9328</u>	<u>.6397</u>	.6607	.6510	
<i>CLEF</i> -L	<u>.5545</u>	.9800	<u>.9535</u>	.5805	.5800	.5940	
ECGFounder	.5420	.9822	<u>.9723</u>	<u>.6572</u>	<u>.6710</u>	<u>.6690</u>	

Table 3: ECG regression tasks. MAEs are reported for LVEF prediction (in %) on MIMIC-IV, and systolic and diastolic blood pressure prediction (in mmHg) on both the MC-MED and Aurora BP data sets. Best result **bolded**. The second best underlined.

Model	MIMIC-IV		Aurora BP		MC-MED	
	LVEF	SBP	DBP	SBP	DBP	
Moira	7.405	12.587	8.261	32.881	12.704	
Moment	7.277	12.707	8.150	18.623	12.576	
ST-MEM	7.149	12.707	8.157	18.623	12.552	
KED	8.313	<u>11.844</u>	8.032	18.598	12.647	
<i>CLEF</i> -S	6.569	11.667	8.653	17.901	12.314	
<i>CLEF</i> -M	<u>6.805</u>	12.110	8.099	17.880	12.318	
<i>CLEF</i> -L	<u>7.313</u>	12.819	8.865	18.763	12.441	
ECGFounder	7.900	13.577	8.332	17.880	12.378	

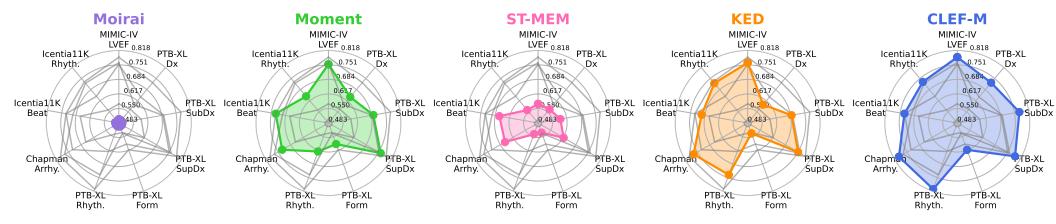
378
379
380
381
382 Table 4: Comparing self-supervised pretraining methods with AUROCs of single-lead ECG
383 classification. For datasets that originally have 12 leads, lead I was used for analysis. Instances
384 outperformed by *CLEF* are highlighted with cells having a grey background.
385
386
387
388

Model	MIMIC-IV			PTB-XL			Chapman Arrhy	MUSIC Outcome	Icential1K Beat	MC-MED		
	LVEF	Dx	SubDx	SupDx	Form	Rhyth				ED Dispo	DC Dispo	Acuity
MOCO (2020)	S	.8194	.5378	.5803	.7260	.5300	.5128	.7743	.5477	.9133	.4737	.6111
	M	.8112	.7664	.7885	.7757	.5021	.5163	.8119	.4546	.9764	.7115	.5926
	L	.6994	.5014	.4981	.5060	.5063	.4990	.5258	.4563	.8680	.5697	.5331
SimCLR (2020)	S	.8185	.7923	.7790	.8180	.5560	.8979	.8554	.5485	.9722	.9180	.6150
	M	.8243	.8247	.8352	.8210	.7338	.9248	.8888	.5531	.9756	.8936	.6380
	L	.8192	.8393	.8370	.8279	.7137	.9109	.8945	.4799	.9733	.9087	.5984
BYOL (2020)	S	.8147	.8134	.8282	.8259	.6883	.9296	.8789	.5448	.9729	.9065	.6166
	M	.8221	.8421	.8221	.8251	.7215	.9351	.8917	.4713	.9766	.8919	.6306
	L	.8221	.8409	.8378	.8239	.7144	.9257	.8978	.4719	.9819	.9032	.5515

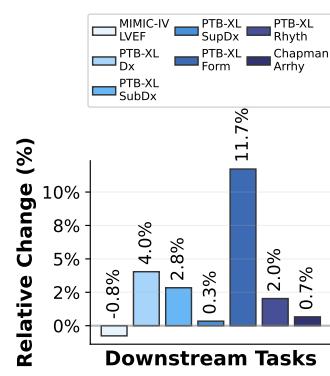
389
390 the best baseline (KED). However, *CLEF-S* does not outperform KED or Moment, which indicates
391 that the larger *CLEF* models produce better quality representations. Models pretrained specifically
392 for ECG tasks do not always yield better performance. Moment (which *CLEF-M* outperforms by
393 10.0%) performs better than ST-MEM (which *CLEF* outperforms by 28.1%).

394
395 **Comparison with Self-supervised Algorithms.** Our proposed clinically-guided contrastive learning
396 method is compared against 3 widely used self-supervised pretraining methods: SimCLR, BYOL, and
397 MoCo. Each pretraining method is trained on all 3 size variants of *CLEF*. The AUROC results for the
398 comparator methods are presented in Table 4. A total of 100 out of 117 instances were outperformed
399 by *CLEF* (highlighted in gray). Overall, *CLEF* outperforms all 3 self-supervised pretraining methods,
400 with average AUROC improvements across all sizes of backbone models of 29.8%, 1.8%, and 2.3%
401 in comparison to MoCo, SimCLR, and BYOL, respectively.

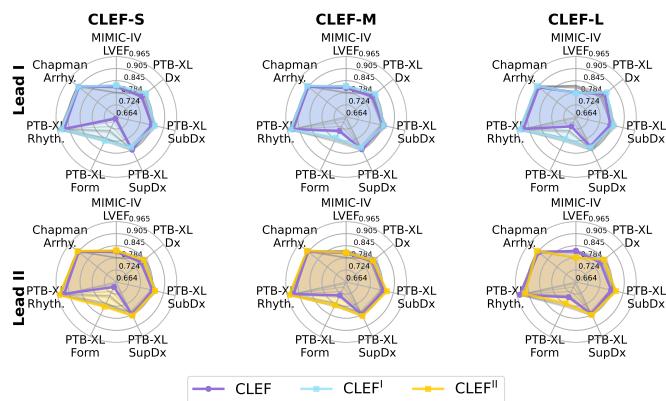
402
403 **Our Clinically-guided Contrastive Loss with Best Semi-supervised Baseline.** We initialize the
404 best-performing baseline for ECG, i.e. KED, with its original pretrained weights and apply our
405 proposed clinically-guided pretraining approach. AUROC results are presented in Figure 3. It can be
406



407
408
409
410
411 Figure 2: AUROC scores from linear probing on 9 classification tasks, comparing Moirai, Moment,
412 ST-MEM, KED, and our *CLEF*. Each subplot focuses on one model, with others shown in gray for
413 reference. Higher values indicate better performance (see further details in Table S11).
414



416
417
418
419
420
421
422
423
424
425
426
427 Figure 3: Changes in AUROC of KED after further training with *CLEF* objectives across
428 downstream ECG tasks.
429
430
431



428
429
430
431 Figure 4: Spider plot comparing *CLEF* with *CLEF* model pretrained on a specific lead (*CLEF*^I and *CLEF*^{II}). AUROC is
432 reported across 3 model variants and 7 downstream tasks.

432
 433 Table 5: Ablation results on handling missing metadata across *CLEF* models of different sizes, where
 434 $\neg M$ denotes the model trained without handling the missing metadata. We report both the results
 435 from the ablated models, and the performance changes relative to the corresponding *CLEF* models
 436 (shown in brackets). The \uparrow and \downarrow depict performance increase and drop, respectively. For better
 437 clarity, we highlighted cases where performance is inferior in blue.

Task	MIMIC-IV LVEF	Dx	SubDx	PTB-XL SupDx	Form	Rhyth	Chapman Arrhy	MUSIC Outcome	Icentia11K Beat	Icentia11K Rhyth
CLEF-S $\neg M$.8245 ($\uparrow .01$)	.8152 ($\downarrow .01$)	.8266 ($\downarrow .02$)	.8210 ($\downarrow .03$)	.6810 ($\uparrow .01$)	.9362 ($\uparrow .01$)	.8786 ($\downarrow .03$)	.4763 ($\downarrow .13$)	.9637 ($\downarrow .02$)	.9052 ($\downarrow .01$)
CLEF-M $\neg M$.8191 ($\uparrow .01$)	.8268 ($\downarrow .00$)	.8311 ($\downarrow .03$)	.8245 ($\downarrow .02$)	.7252 ($\downarrow .02$)	.9230 ($\downarrow .01$)	.8839 ($\downarrow .02$)	.5513 ($\uparrow .04$)	.9749 ($\downarrow .00$)	.8856 ($\downarrow .05$)
CLEF-L $\neg M$.7786 ($\downarrow .01$)	.8299 ($\downarrow .02$)	.8502 ($\uparrow .01$)	.8257 ($\downarrow .02$)	.7024 ($\downarrow .02$)	.9201 ($\downarrow .01$)	.8939 ($\downarrow .01$)	.4966 ($\downarrow .10$)	.9757 ($\downarrow .00$)	.8835 ($\downarrow .07$)

442
 443 observed that, apart from the LVEF classification task in MIMIC-IV, all tasks gained improvement,
 444 on average 3.0% higher AUROC. The form task in PTB-XL gets the greatest improvement (11.7%).
 445 More results on pretraining other ECG baseline models are provided in Appendix D.4.

446 **Pretraining *CLEF* on a Specific Lead.** To assess the potential upper bound of *CLEF*, we pretrain
 447 models exclusively on the corresponding lead of the downstream tasks (lead I or II). Pretraining was
 448 performed for 10 epochs, with the other downstream finetuning hyperparameters the same as the
 449 previous experiments. Results are summarized in Figure 4, and provided in full in Appendix D.5.
 450 On lead I, pretraining achieves average AUROC improvements over pretraining using all 12 leads of
 451 3.4%, 1.4%, and 2.4% for *CLEF-S*, *-M*, and *-L*, respectively. Moreover, this made the performance
 452 of *CLEF* comparable to that of the supervised ECGFounder model, with differences in average
 453 AUROC of only 0.2%, -0.1% , and -0.3% for *CLEF-S*, *-M*, and *-L*, respectively, on lead I (the same
 454 lead ECGFounder is trained on too). Improvements are also observed when pretraining on lead II, of
 455 3.4%, 2.0%, and 0.8% for *CLEF-S*, *-M*, and *-L*, respectively.

456 **Ablation on Handling Missing Metadata.** We assess the impact of our solution for handling missing
 457 metadata. As an ablation study, we report AUROC for all *CLEF* variants across 10 tasks including
 458 wearable datasets (MUSIC and Icentia11K) and single-lead ECG tasks (MIMIC-IV, PTB-XL, and
 459 Chapman) using lead I. Table 5 shows that without M in Eq. (4), the performance of all models
 460 degrades on the majority of tasks. On average, AUROC decreases by 1.9%, with a larger decrease on
 461 wearable datasets (4.0% drop) compared to lead I datasets (1.1% drop).

462 **Contribution of different loss components.** A study is conducted using CIFAR-100 to understand
 463 the contribution of different loss components. All models are initialized with ImageNet-1K pretrained
 464 ResNet-18 weights and further pretrained with: the standard contrastive loss \mathcal{L}_{nce} , the weighted
 465 contrastive loss \mathcal{L}^w , the dissimilarity alignment loss \mathcal{L}^d , $\mathcal{L}_{nce} + \mathcal{L}^d$, and our proposed $\mathcal{L}^w + \mathcal{L}^d$. A
 466 detailed experiment setup is provided in Appendix D.6. Our proposed $\mathcal{L}^w + \mathcal{L}^d$ achieves clearer
 467 separation with lower variance. Further details in Appendix D.6, Figure S2, and Table S15).

468 469 5 DISCUSSION AND CONCLUSION

470 We propose *CLEF*: clinically-guided contrastive learning to train single-lead ECG foundation models.
 471 On average, *CLEF* outperforms all baseline semi-supervised FMs across 18 clinical classification
 472 and regression tasks on 7 datasets. Furthermore, *CLEF* outperforms three self-supervised pretraining
 473 algorithms. When pretrained on the same lead as in downstream tasks, *CLEF* also performs
 474 comparably to ECGFounder, a state-of-the-art FM trained in a supervised manner on a labeled
 475 dataset. *CLEF* facilitates the creation of high-performance ECG FMs using only routinely-recorded
 476 metadata, without needing ECG-level annotations.

477 **Limitations.** We trained models on ECGs from a single hospital, which suit hospital-related tasks but
 478 are not representative of the general population. Performance on population-level tasks would likely
 479 improve with a more diverse cohort. The risk scores in this study are limited by only three out of
 480 seven SCORE2 input variables being available. Therefore, the scores cannot be considered a holistic
 481 assessment of cardiovascular risk, but instead are representative of real-world applications with
 482 incomplete metadata. Potentially richer embeddings could be obtained when using more complete
 483 metadata and, therefore, more precise risk scores. Finally, the single-lead ECGs used to assess
 484 performance in this study were measured using wet gel electrodes at the chest, and so are likely of
 485 higher quality than those typically measured by devices such as smartwatches and handheld ECGs.

486 **Future work.** One can investigate the utility of the proposed contrastive learning strategy for
487 other modalities and health conditions. For instance, it may also be useful for developing
488 photoplethysmography (PPG) FMs, since the PPG is also a cardiovascular signal and has been
489 found to be associated with cardiovascular risk Weng et al. (2024). It could also be applied to other
490 health conditions, such as incorporating the risk of deterioration in chronic respiratory conditions,
491 in contrastive learning for respiratory signals. The practical application of this strategy would be
492 aided by understanding which metadata variables contribute most to the quality of the representations
493 produced by *CLEF*. A further ablation study, preferably conducted on a dataset with complete
494 metadata, could identify the most important metadata variables and provide insight into how to apply
495 this approach across diverse datasets with different metadata. Finally, we can go beyond traditional
496 classification and regression tasks. Because the *CLEF*’s embedding space enables patient ECGs to
497 be analyzed without prior metadata. It is possible to assess the proximity of current embeddings
498 to prototypes of high-risk patients from the training set. This opens the door to clustering-based
499 decision-making and other novel applications.

500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540 ETHICS STATEMENT
541

542 This study makes exclusive use of publicly available electrocardiogram datasets, each of which
543 has independently undergone ethics approval by the respective data-providing institutions. All
544 datasets comply with relevant privacy regulations, and no additional data collection was conducted.
545 We acknowledge that biases may exist in the training data and recognize the importance of fair
546 performance across diverse populations. Future evaluations on broader and more representative
547 datasets will be essential for further improving the generalizability and fairness of our models. While
548 our work demonstrates potential for advancing remote health monitoring and clinical decision support,
549 we are mindful of the risks of misuse. Potential harmful applications could include unauthorized
550 health surveillance, discriminatory practices in insurance or employment, unfair credit assessments,
551 or the exploitation of sensitive health data for targeted marketing. We emphasize that this research
552 is intended solely for beneficial healthcare and scientific purposes and strongly advocate for its
553 responsible use. Our study adhered to established research ethics guidelines, and we declare no
554 conflicts of interest. We further encourage interdisciplinary dialogue to anticipate and mitigate
555 risks, develop appropriate governance frameworks, and safeguard individual rights. Recognizing
556 the broader societal implications of AI in healthcare, we encourage ethical AI research and foster
557 responsible deployment of these technologies.

558 REPRODUCIBILITY STATEMENT
559

560 We are committed to the reproducibility of our work. Models, code, and hyperparameters will be
561 publicly released, accompanied by user-friendly tutorials to facilitate adoption and extension. Our
562 work relies exclusively on publicly accessible datasets, which can be readily downloaded or requested
563 from the respective data-hosting institutions. To further support reproducibility, we provide detailed
564 descriptions of the datasets, model architectures, training and evaluation procedures, hyperparameters,
565 and preprocessing methods used throughout the paper. Together, these resources will enable other
566 researchers to replicate our results and build upon our contributions in future work.

567 REFERENCES
568

569 D.B. Geselowitz. On the theory of the electrocardiogram. *Proceedings of the IEEE*, 77(6):857–876,
570 1989.

572 Leonard S Lilly. *Pathophysiology of heart disease: a collaborative project of medical students and*
573 *faculty*. 2012.

574 Xiaoxi Yao, David R Rushlow, Jonathan W Inselman, Rozalina G McCoy, Thomas D Thacher,
575 Emma M Behnken, Matthew E Bernard, Steven L Rosas, Abdulla Akfaly, Artika Misra, et al.
576 Artificial intelligence–enabled electrocardiograms for identification of patients with low ejection
577 fraction: a pragmatic, randomized clinical trial. *Nature medicine*, 27(5):815–819, 2021.

579 Heba Khamis, Robert Weiss, Yang Xie, Chan-Wei Chang, Nigel H. Lovell, and Stephen J. Redmond.
580 QRS Detection Algorithm for Telehealth Electrocardiogram Recordings. *IEEE Transactions on*
581 *Biomedical Engineering*, 63(7):1377–1388, 2016.

582 Hesam Halvaei, Emma Svennberg, Leif Sörnmo, and Martin Stridh. Identification of Transient Noise
583 to Reduce False Detections in Screening for Atrial Fibrillation. *Frontiers in Physiology*, 12, 2021.

584 Akshay Khunte, Veer Sangha, Evangelos K. Oikonomou, Lovedeep S. Dhingra, Arya Aminorroaya,
585 Bobak J. Mortazavi, Andreas Coppi, Cynthia A. Brandt, Harlan M. Krumholz, and Rohan Khera.
586 Detection of left ventricular systolic dysfunction from single-lead electrocardiography adapted for
587 portable and wearable devices. *npj Digital Medicine*, 6(1):124, 2023.

589 Emma Svennberg, Johan Engdahl, Faris Al-Khalili, Leif Friberg, Viveka Frykman, and Mårten
590 Rosenqvist. Mass Screening for Untreated Atrial Fibrillation. *Circulation*, 131(25):2176–2184,
591 2015.

592 Stefan P Nelwan, Jan A Kors, Simon H Meij, Jan H van Bemmel, and Maarten L Simoons.
593 Reconstruction of the 12-lead electrocardiogram from reduced lead sets. *Journal of*
594 *Electrocardiology*, 37(1):11–18, 2004.

594 Yuzhen Qin, Li Sun, Hui Chen, Wenming Yang, Wei-Qiang Zhang, Jintao Fei, and Guijin Wang. Mvkt-
 595 ecg: Efficient single-lead ecg classification for multi-label arrhythmia by multi-view knowledge
 596 transferring. *Computers in Biology and Medicine*, 166:107503, 2023.

597

598 Jun Li, Aaron Aguirre, Junior Moura, Che Liu, Lanhai Zhong, Chenxi Sun, Gari Clifford, Brandon
 599 Westover, and Shenda Hong. An Electrocardiogram Foundation Model Built on over 10 Million
 600 Recordings with External Evaluation across Multiple Domains. *arXiv preprint arXiv:2410.04133*,
 601 2024.

602 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for
 603 contrastive learning of visual representations. In *International Conference on Machine Learning*,
 604 volume 119, pages 1597–1607, 2020.

605

606 Sahar Soltanieh, Ali Etemad, and Javad Hashemi. Analysis of Augmentations for Contrastive ECG
 607 Representation Learning. In *2022 International Joint Conference on Neural Networks (IJCNN)*,
 608 pages 1–10, 2022.

609

610 Janine Thoma, Danda Pani Paudel, Ajad Chhatkuli, and Luc Van Gool. Geometrically Mappable
 611 Image Features. *IEEE Robotics and Automation Letters*, 5(2):2062–2069, 2020a.

612

613 Janine Thoma, Danda Pani Paudel, and Luc V Gool. Soft Contrastive Learning for Visual Localization.
 614 In *Advances in Neural Information Processing Systems*, volume 33, pages 11119–11130. Curran
 615 Associates, Inc., 2020b.

616

617 Johan Fredin Haslum, Christos Matsoukas, Karl-Johan Leuchowius, Erik Müllers, and Kevin Smith.
 618 Metadata-guided consistency learning for high content images. In *Medical Imaging with Deep
 619 Learning*, volume 227, pages 918–936, 2024.

620

621 June-Woo Kim, Miika Toikkanen, Amin Jalali, Minseok Kim, Hye-Ji Han, Hyunwoo Kim, Wonwoo
 622 Shin, Ho-Young Jung, and Kyunghoon Kim. Adaptive metadata-guided supervised contrastive
 623 learning for domain adaptation on respiratory sound classification. *IEEE Journal of Biomedical
 624 and Health Informatics*, pages 1–13, 2025.

625

626 Eray Erturk, Fahad Kamran, Salar Abbaspourazad, Sean Jewell, Harsh Sharma, Yujie Li, Sinead
 627 Williamson, Nicholas J Foti, and Joseph Futoma. Beyond sensor data: Foundation models of
 628 behavioral data from wearables improve health predictions. *arXiv preprint arXiv:2507.00191*,
 629 2025.

630

631 Yue Wang, Xu Cao, Yaojun Hu, Haochao Ying, James Matthew Rehg, Jimeng Sun, Jian Wu, and
 632 Jintai Chen. AnyECG: Foundational Models for Electrocardiogram Analysis. *arXiv preprint
 633 arXiv:2411.17711*, 2024.

634

635 R M Conroy, K Pyörälä, A P Fitzgerald, S Sans, A Menotti, G De Backer, D De Bacquer,
 636 P Ducimetière, P Jousilahti, U Keil, I Njølstad, R G Oganov, T Thomsen, H Tunstall-Pedoe,
 637 A Tverdal, H Wedel, P Whincup, L Wilhelmsen, I M Graham, and SCORE project group.
 638 Estimation of ten-year risk of fatal cardiovascular disease in europe: the SCORE project. *Eur
 639 Heart J*, 24(11):987–1003, 2003.

640

641 SCORE2 working group and ESC Cardiovascular risk collaboration. Score2 risk prediction
 642 algorithms: new models to estimate 10-year risk of cardiovascular disease in europe. *European
 643 Heart Journal*, 42(25):2439–2454, 2021a.

644

645 SCORE2-OP working group and ESC Cardiovascular risk collaboration. Score2-op risk prediction
 646 algorithms: estimating incident cardiovascular event risk in older persons in four geographical risk
 647 regions. *European Heart Journal*, 42(25):2455–2467, 2021b.

648

649 J. Weston Hughes, James Tooley, Jessica Torres Soto, Anna Ostropolets, Tim Poterucha, Matthew Kai
 650 Christensen, Neal Yuan, Ben Ehler, Dhamanpreet Kaur, Guson Kang, Albert Rogers, Sanjiv
 651 Narayan, Pierre Elias, David Ouyang, Euan Ashley, James Zou, and Marco V. Perez. A deep
 652 learning-based electrocardiogram risk score for long term cardiovascular death and disease. *npj
 653 Digital Medicine*, 6(1):169, 2023.

648 Gregory Y. H. Lip, Robby Nieuwlaat, Ron Pisters, Deirdre A. Lane, and Harry J. G. M. Crijns.
 649 Refining Clinical Risk Stratification for Predicting Stroke and Thromboembolism in Atrial
 650 Fibrillation Using a Novel Risk Factor-Based Approach: The Euro Heart Survey on Atrial
 651 Fibrillation. *CHEST*, 137(2):263–272, 2010.

652
 653 Keith A. A. Fox, Omar H. Dabbous, Robert J. Goldberg, Karen S. Pieper, Kim A. Eagle, Frans Van
 654 de Werf, Álvaro Avezum, Shaun G. Goodman, Marcus D. Flather, Frederick A. Anderson, and
 655 Christopher B. Granger. Prediction of Risk of Death and Myocardial Infarction in the Six Months
 656 after Presentation with Acute Coronary Syndrome: Prospective Multinational Observational Study
 657 (GRACE). *BMJ*, 333(7578):1091, 2006.

658 Joshua David Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive Learning
 659 with Hard Negative Samples. In *International Conference on Learning Representations*, 2021.

660
 661 Haochen Li, Xin Zhou, Anh Tuan Luu, and Chunyan Miao. Rethinking Negative Pairs in Code
 662 Search. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.

663 Haojie Zhuang, Wei Emma Zhang, Jian Yang, Weitong Chen, and Quan Z Sheng. Not all negatives
 664 are equally negative: Soft contrastive learning for unsupervised sentence representations. In
 665 *Proceedings of the ACM International Conference on Information and Knowledge Management*,
 666 pages 3591–3601, 2024.

667
 668 Zhen Yang, Ming Ding, Tinglin Huang, Yukuo Cen, Junshuai Song, Bin Xu, Yuxiao Dong, and Jie
 669 Tang. Does negative sampling matter? a review with insights into its theory and applications. *IEEE*
 670 *Transactions on Pattern Analysis and Machine Intelligence*, 46(8):5692–5711, 2024.

671 Brian Gow, Tom Pollard, Larry A Nathanson, Alistair Johnson, Benjamin Moody, Chrystinne
 672 Fernandes, Nathaniel Greenbaum, Jonathan W Waks, Parastou Eslami, Tanner Carbonati, et al.
 673 Mimic-iv-ecg: Diagnostic electrocardiogram matched subset. *Type: dataset*, 6:13–14, 2023.

674
 675 Shenda Hong, Yanbo Xu, Alind Khare, Satria Priambada, Kevin Maher, Alaa Aljiffry, Jimeng Sun,
 676 and Alexey Tumanov. HOLMES: Health OnLine Model Ensemble Serving for Deep Learning
 677 Models in Intensive Care Units. In *Proceedings of the ACM SIGKDD International Conference on*
 678 *Knowledge Discovery & Data Mining*, pages 1614–1624, 2020.

679
 680 Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual
 681 transformations for deep neural networks. In *Proceedings of the IEEE conference on computer*
 682 *vision and pattern recognition*, pages 1492–1500, 2017.

683
 684 Yeongyeon Na, Minje Park, Yunwon Tae, and Sunghoon Joo. Guiding Masked Representation
 685 Learning to Capture Spatio-Temporal Relationship of Electrocardiogram. In *International*
 686 *Conference on Learning Representations*, 2024.

687
 688 Patrick Wagner, Nils Strodtthoff, Ralf-Dieter Bousseljot, Dieter Kreiseler, Fatima I. Lunze, Wojciech
 689 Samek, and Tobias Schaeffter. PTB-XL, a large publicly available electrocardiography dataset.
 690 *Scientific Data*, 7(1):154, 2020.

691
 692 Jianwei Zheng, Jianming Zhang, Sidy Danioko, Hai Yao, Hangyuan Guo, and Cyril Rakowski. A
 693 12-lead electrocardiogram database for arrhythmia research covering more than 10,000 patients.
 694 *Scientific Data*, 7(1):48, 2020.

695
 696 Aman Kansal, Emma Chen, Tom Jin, Pranav Rajpurkar, and David Kim. Multimodal Clinical
 697 Monitoring in the Emergency Department (MC-MED). *PhysioNet*, v1, 10, 2025.

698
 699 Rebecca Mieloszyk, Hope Twede, Jonathan Lester, Jeremiah Wander, Sumit Basu, Gabe Cohn, Greg
 700 Smith, Dan Morris, Sidhant Gupta, Desney Tan, Nicolas Villar, Moni Wolf, Sailaja Malladi, Matt
 701 Mickelson, Lauren Ryan, Lindsey Kim, Jeffrey Kepple, Susanne Kirchner, Emma Wampler, Riena
 Terada, Joel Robinson, Ron Paulsen, and T. Scott Saponas. A Comparison of Wearable Tonometry,
 Photoplethysmography, and Electrocardiography for Cuffless Measurement of Blood Pressure in
 an Ambulatory Setting. *IEEE Journal of Biomedical and Health Informatics*, 26(7):2864–2875,
 2022.

702 Alba Martin-Yebra, Juan Pablo Martínez, and Pablo Laguna. Music (sudden cardiac death in chronic
 703 heart failure). 2025.

704

705 Shawn Tan, Guillaume Androz, Ahmad Chamseddine, Pierre Fecteau, Aaron Courville, Yoshua
 706 Bengio, and Joseph Paul Cohen. Iceni11k: An unsupervised representation learning dataset for
 707 arrhythmia subtype discovery. *arXiv preprint arXiv:1910.09570*, 2019.

708

709 Yuanyuan Tian, Zhiyuan Li, Yanrui Jin, Mengxiao Wang, Xiaoyang Wei, Liqun Zhao, Yunqing Liu,
 710 Jinlei Liu, and Chengliang Liu. Foundation model of ECG diagnosis: Diagnostics and explanations
 711 of any form and rhythm on ECG. *Cell Reports Medicine*, 5(12):101875, 2024.

712

713 Gerald Woo, Chenghao Liu, Akshat Kumar, Caiming Xiong, Silvio Savarese, and Doyen Sahoo.
 714 Unified training of universal time series forecasting transformers. In *International Conference on
 Machine Learning*, 2024.

715

716 Mononito Goswami, Konrad Szafer, Arjun Choudhry, Yifu Cai, Shuo Li, and Artur Dubrawski.
 717 MOMENT: A family of open time-series foundation models. In *International Conference on
 Machine Learning*, 2024.

718

719 Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena
 720 Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar,
 721 Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a
 722 new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*,
 723 volume 33, pages 21271–21284, 2020.

724

725 Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for
 726 unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on
 Computer Vision and Pattern Recognition (CVPR)*, 2020.

727

728 Manohar Ghanta, Reza Sameni, Aaron Aguirre, Qiao Li, Sahar Zafar, Gari Clifford, and M Brandon
 729 Westover. Harvard-Emory ECG Database. 2025.

730

731 Wei-Hung Weng, Sébastien Baur, Mayank Daswani, Christina Chen, Lauren Harrell, Sujay
 732 Kakarmath, Mariam Jabara, Babak Behsaz, Cory Y. McLean, Yossi Matias, Greg S. Corrado,
 733 Shravya Shetty, Shruthi Prabhakara, Yun Liu, Goodarz Danaei, and Diego Ardila. Predicting
 734 Cardiovascular Disease Risk Using Photoplethysmography and Deep Learning. *PLOS Global
 Public Health*, 4(6):e0003204, 2024.

735

736 Paul Kligfield, Leonard S. Gettes, James J. Bailey, Rory Childers, Barbara J. Deal, E. William
 737 Hancock, Gerard van Herpen, Jan A. Kors, Peter Macfarlane, David M. Mirvis, Olle Pahm, Pentti
 738 Rautaharju, and Galen S. Wagner. Recommendations for the Standardization and Interpretation of
 the Electrocardiogram. *JACC*, 49(10):1109–1127, 2007.

739

740 Tomas B Garcia and Neil E Holtz. *12 Lead ECG: The Art of Interpretation*. Jones & Bartlett Learning,
 741 2001.

742

743 Antônio H. Ribeiro, Manoel Horta Ribeiro, Gabriela M. M. Paixão, Derick M. Oliveira, Paulo R.
 744 Gomes, Jéssica A. Canazart, Milton P. S. Ferreira, Carl R. Andersson, Peter W. Macfarlane, Wagner
 745 Meira Jr., Thomas B. Schön, and Antonio Luiz P. Ribeiro. Automatic diagnosis of the 12-lead
 ECG using a deep neural network. *Nature Communications*, 11(1):1760, 2020. ISSN 2041-1723.

746

747 Awni Y. Hannun, Pranav Rajpurkar, Masoumeh Haghpanahi, Geoffrey H. Tison, Codie Bourn,
 748 Mintu P. Turakhia, and Andrew Y. Ng. Cardiologist-Level Arrhythmia Detection and Classification
 749 in Ambulatory Electrocardiograms Using a Deep Neural Network. *Nature medicine*, 25(1):65–69,
 2019.

750

751 Paul A Friedman. The electrocardiogram at 100 years: history and future. *Circulation*, 149(6):
 752 411–413, 2024.

753

754 Zachi I Attia, David M Harmon, Jennifer Dugan, Lukas Manka, Francisco Lopez-Jimenez, Amir
 755 Lerman, Konstantinos C Sontis, Peter A Noseworthy, Xiaoxi Yao, Eric W Klavetter, et al. Prospective
 evaluation of smartwatch-enabled detection of left ventricular dysfunction. *Nature
 medicine*, 28(12):2497–2503, 2022.

756 S P Preejith, R Dhinesh, Jayaraj Joseph, and Mohanasankar Sivaprakasam. Wearable ecg platform
 757 for continuous cardiac monitoring. In *Annual International Conference of the IEEE Engineering*
 758 *in Medicine and Biology Society (EMBC)*, pages 623–626, 2016.

759

760 Sumit Majumder, Leon Chen, Ognian Marinov, Chih-Hung Chen, Tapas Mondal, and M. Jamal Deen.
 761 Noncontact wearable wireless ecg systems for long-term monitoring. *IEEE Reviews in Biomedical*
 762 *Engineering*, 11:306–321, 2018.

763 Nicholar Clark, Edward Sandor, Calvin Walden, In Soo Ahn, and Yufeng Lu. A wearable ecg
 764 monitoring system for real-time arrhythmia detection. In *IEEE International Midwest Symposium*
 765 *on Circuits and Systems (MWSCAS)*, pages 787–790, 2018.

766 Yu-Jin Lin, Chen-Wei Chuang, Chun-Yueh Yen, Sheng-Hsin Huang, Peng-Wei Huang, Ju-Yi Chen,
 767 and Shuenn-Yuh Lee. Artificial intelligence of things wearable system for cardiac disease detection.
 768 In *IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS)*, pages
 769 67–70, 2019.

770 Nurul Huda, Sadia Khan, Ragib Abid, Samiul Based Shuvo, Mir Maheen Labib, and Taufiq Hasan. A
 771 low-cost, low-energy wearable ecg system with cloud-based arrhythmia detection. In *2020 IEEE*
 772 *Region 10 Symposium (TENSYMP)*, pages 1840–1843, 2020.

773

774 Peng Wang, Zihuai Lin, Xucun Yan, Zijiao Chen, Ming Ding, Yang Song, and Lu Meng. A wearable
 775 ecg monitor for deep learning based real-time cardiovascular disease detection. *arXiv preprint*
 776 *arXiv:2201.10083*, 2022.

777 Nassim Ravanshad, Hamidreza Rezaee-Dehsorkh, Reza Lotfi, and Yong Lian. A level-crossing
 778 based qrs-detection algorithm for wearable ecg sensors. *IEEE Journal of Biomedical and Health*
 779 *Informatics*, 18(1):183–192, 2014.

780 Paolo Melillo, Rossana Castaldo, Giovanna Sannino, Ada Orrico, Giuseppe de Pietro, and Leandro
 781 Peccchia. Wearable technology and ecg processing for fall risk assessment, prevention and detection.
 782 In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*
 783 *(EMBC)*, pages 7740–7743, 2015.

784

785 Dionisije Sopic, Amin Aminifar, Amir Aminifar, and David Atienza. Real-time event-driven
 786 classification technique for early detection and prevention of myocardial infarction on wearable
 787 systems. *IEEE Transactions on Biomedical Circuits and Systems*, 12(5):982–992, 2018.

788 Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx,
 789 Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson,
 790 Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel,
 791 Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano
 792 Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren
 793 Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter
 794 Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil
 795 Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar
 796 Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal
 797 Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu
 798 Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa,
 799 Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles,
 800 Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung
 801 Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu
 802 Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh,
 803 Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori,
 804 Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai
 805 Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi
 806 Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the
 807 Opportunities and Risks of Foundation Models. (arXiv:2108.07258), 2022.

808 Santiago Jiménez-Serrano, Miguel Rodrigo, Conrado J Calvo, José Millet, and Francisco Castells.
 809 From 12 to 1 ECG lead: multiple cardiac condition detection mixing a hybrid machine learning
 810 approach with a one-versus-rest classification strategy. *Physiological Measurement*, 43(6):064003,
 811 2022.

810 Eleni Angelaki, Georgios D Barmparis, Konstantinos Fragkiadakis, Spyros Maragkoudakis,
 811 Evangelos Zacharis, Anthi Plevritaki, Emmanouil Kampanieris, Petros Kalomoirakis, Spyros
 812 Kassotakis, George Kochiadakis, et al. Diagnostic performance of single-lead electrocardiograms
 813 for arterial hypertension diagnosis: a machine learning approach. *Journal of Human Hypertension*,
 814 39(1):58–65, 2025.

815 Elisa Ramirez, Samuel Ruiperez-Campillo, Ruben Casado-Arroyo, José Luis Merino, Julia E Vogt,
 816 Francisco Castells, and José Millet. The art of selecting the ECG input in neural networks to
 817 classify heart diseases: a dual focus on maximizing information and reducing redundancy. *Frontiers*
 818 *in physiology*, 15:1452829, 2024.

819 Weijie Sun, Sunil Vasu Kalmady, Amir Salimi, Nariman Sepehrvand, Eric Ly, Abram Hindle, Russell
 820 Greiner, and Padma Kaul. Ecg for high-throughput screening of multiple diseases: Proof-of-
 821 concept using multi-diagnosis deep learning from population-based datasets. *arXiv preprint*
 822 *arXiv:2210.06291*, 2022.

823 Mi-Na Kim, Yong Seok Lee, Youngmin Park, Ayoung Jung, Hanjee So, Joonwoong Park, Jin-Joo Park,
 824 Dong-Joo Choi, So-Ree Kim, and Seong-Mi Park. Deep learning for predicting rehospitalization
 825 in acute heart failure: Model foundation and external validation. *ESC Heart Failure*, 11(6):
 826 3702–3712, 2024.

827 Feifei Liu, Chengyu Liu, Lina Zhao, Xiangyu Zhang, Xiaoling Wu, Xiaoyan Xu, Yulin Liu, Caiyun
 828 Ma, Shoushui Wei, Zhiqiang He, et al. An open access database for evaluating the algorithms of
 829 electrocardiogram rhythm and morphology abnormality detection. *Journal of Medical Imaging*
 830 and *Health Informatics*, 8(7):1368–1373, 2018.

831 Antônio H Ribeiro, GM Paixao, Emilly M Lima, M Horta Ribeiro, Marcelo M Pinto Filho, Paulo R
 832 Gomes, Derick M Oliveira, Wagner Meira Jr, Thömas B Schon, and Antonio Luiz P Ribeiro.
 833 CODE-15%: A large scale annotated dataset of 12-lead ECGs. *Zenodo, Jun*, 9:10–5281, 2021.

834 Erick A Perez Alday, Annie Gu, Amit J Shah, Chad Robichaux, An-Kwok Ian Wong, Chengyu
 835 Liu, Feifei Liu, Ali Bahrami Rad, Andoni Elola, Salman Seyed, et al. Classification of 12-lead
 836 ecgs: the physionet/computing in cardiology challenge 2020. *Physiological measurement*, 41(12):
 837 124003, 2020.

838 Matthew A Reyna, Nadi Sadr, Erick A Perez Alday, Annie Gu, Amit J Shah, Chad Robichaux,
 839 Ali Bahrami Rad, Andoni Elola, Salman Seyed, Sardar Ansari, et al. Will two do? Varying
 840 dimensions in electrocardiography: the PhysioNet/Computing in Cardiology Challenge 2021. In
 841 *2021 computing in cardiology (CinC)*, volume 48, pages 1–4, 2021.

842 Che Liu, Zhongwei Wan, Cheng Ouyang, Anand Shah, Wenjia Bai, and Rossella Arcucci. Zero-Shot
 843 ECG Classification with Multimodal Learning and Test-time Clinical Knowledge Enhancement.
 844 In *International Conference on Machine Learning*, 2024.

845 Jiarui Jin, Haoyu Wang, Hongyan Li, Jun Li, Jiahui Pan, and Shenda Hong. Reading Your Heart:
 846 Learning ECG Words and Sentences via Pre-training ECG Language Model. In *International*
 847 *Conference on Learning Representations*, 2025.

848 Kaden McKeen, Laura Oliva, Sameer Masood, Augustin Toma, Barry Rubin, and Bo Wang. Ecg-fm:
 849 An open electrocardiogram foundation model. *arXiv preprint arXiv:2408.05178*, 2024.

850 Temesgen Mehari and Nils Strothoff. Self-supervised representation learning from 12-lead ECG
 851 data. *Computers in biology and medicine*, 141:105114, 2022.

852 Jungwoo Oh, Hyunseung Chung, Joon-myung Kwon, Dong-gyun Hong, and Edward Choi. Lead-
 853 agnostic self-supervised learning for local and global representations of electrocardiogram. In
 854 *Conference on Health, Inference, and Learning*, pages 338–353. PMLR, 2022.

855 Dani Kiyasseh, Tingting Zhu, and David A Clifton. CloCs: Contrastive learning of cardiac signals
 856 across space, time, and patients. In *International Conference on Machine Learning*, pages 5606–
 857 5615. PMLR, 2021.

864 Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised Feature Learning via
865 Non-Parametric Instance Discrimination. In *Proceedings of the IEEE Conference on Computer*
866 *Vision and Pattern Recognition (CVPR)*, June 2018.

867 Andrius Sološenko, Andrius Petrénas, Birutė Paliakaitė, Vaidotas Marozas, and Leif Sörnmo. Model
868 for simulating ecg and ppg signals with arrhythmia episodes. *PhysioNet*, 101:23, 2021a.

870 Ary L. Goldberger, Luis A. N. Amaral, Leon Glass, Jeffrey M. Hausdorff, Plamen Ch. Ivanov,
871 Roger G. Mark, Joseph E. Mietus, George B. Moody, Chung-Kang Peng, and H. Eugene Stanley.
872 Physiobank, physiotoolkit, and physionet. *Circulation*, 101(23):e215–e220, 2000.

873 Andrius Sološenko, Andrius Petrénas, Birutė Paliakaitė, Vaidotas Marozas, and Leif Sörnmo. Model
874 for simulating ecg and ppg signals with arrhythmia episodes. *PhysioNet*, 101:23, 2021b.

876 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.(2009).
877 2009.

878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

918

919

920

921

922

923

924

925

926

927

928

929

930

931

932

SUPPLEMENTARY MATERIAL

CLEF: CLINICALLY-GUIDED CONTRASTIVE LEARNING FOR ELECTROCARDIOGRAM FOUNDATION MODELS

A RELATED WORK

A.1 ECG AND ML

Utility of single-lead ECGs compared to 12-lead. Electrodes on the limbs and chest provide spatially diverse views of cardiac activity Kligfield et al. (2007), producing 12-lead ECGs to assess a range of heart conditions and abnormalities (Garcia and Holtz, 2001; Ribeiro et al., 2020). The growing usage of wearables and mobile health devices with single-lead ECG functionality (Hannun et al., 2019; Friedman, 2024; Attia et al., 2022) has opened new opportunities for continuous health monitoring beyond traditional care settings (Preejith et al., 2016; Majumder et al., 2018; Clark et al., 2018), with applications including early detection of cardiac events (Lin et al., 2019; Huda et al., 2020; Wang et al., 2022), long-term pattern tracking (Ravanshad et al., 2014; Majumder et al., 2018), and proactive health interventions (Melillo et al., 2015; Sopic et al., 2018; Bommasani et al., 2022).

Prior work Jiménez-Serrano et al. (2022) studied the Computing in Cardiology 2021 ECG dataset (covering 88,253 annotated 12-lead training recordings). They compared ML performance over 26 target cardiac conditions when processing only lead I versus all 12 leads. They found that 12 leads give the best overall performance, but using only lead I yields a modest average degradation rather than a collapse in accuracy: the average of G-metric (geometric mean of sensitivity and specificity) falls from 0.80 to 0.74. Specifically, among the 26 conditions, performance is similar (absolute G-metric change lower than 0.03) for 12 conditions; there is a moderate loss (G-metric drop from 0.04 to 0.07) for 7 conditions; and a significant drop (G-metric drop more than 0.07) for 7 conditions. On average, the drop in the G-metric is 0.06. Overall, this shows that single-lead (lead I) ECGs can reliably detect many rhythm-based abnormalities (making them attractive for wearables/screening), but they perform substantially worse for axis deviations, certain conduction blocks, and morphology-dependent findings that require precordial leads. Another observational study Angelaki et al. (2025) on 1,254 subjects showed that a single-lead ECG (lead I), when combined with other demographic features, can diagnose arterial hypertension with high accuracy (AUC 0.831, sensitivity 72%, specificity 82%), demonstrating that even single-lead ECGs contain sufficient diagnostic information when analyzed effectively. Additionally, a study by Ramirez et al. (2024) demonstrates that the standard 12-lead ECG contains redundant information when classifying cardiovascular diseases using CNNs. By selecting subsets of leads or applying transformations to the ECG signals, the authors evaluated how these adjustments influenced a CNN’s diagnostic performance. Their main finding is that carefully optimizing input configurations, hence reducing redundancy while preserving essential information, can improve deep learning model performance and enable efficient diagnostics even with fewer leads.

Diagnostic potential and limitations of machine learning for ECG analysis. A proof-of-concept study by Sun et al. (2022) demonstrates that a ResNet model trained on large-scale, population-based ECG datasets can accurately predict a wide range of diseases, including many non-cardiovascular conditions. Using over 1.5 million ECGs linked to 11K unique WHO ICD codes from 240K patients across 26 hospitals in Canada, the authors identified 700 disease categories with sufficient data for modeling. Models achieved strong discriminative performance (AUROC > 80%) for 80 disease categories, with 18 categories exceeding AUROC > 90% (including non-cardiovascular conditions such as silicosis, type 1 diabetes mellitus, liver diseases, behavioral disorder due to drug use, and some maternal diseases during pregnancy). Despite excellent AUROC values, precision was limited for many conditions due to their low prevalence, suggesting greater utility for rule-out screening rather than definitive diagnosis. The findings highlight the untapped diagnostic potential of ECGs for diverse diseases, while also noting that predictions may reflect correlated comorbidities or patient

972
973
974
Table S1: Available metadata for each database. M-iv-ecg is short for MIMIC-IV-ECG, G12EC is
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
17210
17211
17212
17213
17214
17215
17216
17217
17218
17219
17220
17221
17222
17223
17224
17225
17226
17227
17228
17229
17230
17231
17232
17233
17234
17235
17236
17237
17238
17239
17240
17241
17242
17243
17244
17245
17246
17247
17248
17249
17250
17251
17252
17253
17254
17255
17256
17257
17258
17259
17260
17261
17262
17263
17264
17265
17266
17267
17268
17269
17270
17271
17272
17273
17274
17275
17276
17277
17278
17279
17280
17281
17282
17283
17284
17285
17286
17287
17288
17289
17290
17291
17292
17293
17294
17295
17296
17297
17298
17299
172100
172101
172102
172103
172104
172105
172106
172107
172108
172109
172110
172111
172112
172113
172114
172115
172116
172117
172118
172119
172120
172121
172122
172123
172124
172125
172126
172127
172128
172129
172130
172131
172132
172133
172134
172135
172136
172137
172138
172139
172140
172141
172142
172143
172144
172145
172146
172147
172148
172149
172150
172151
172152
172153
172154
172155
172156
172157
172158
172159
172160
172161
172162
172163
172164
172165
172166
172167
172168
172169
172170
172171
172172
172173
172174
172175
172176
172177
172178
172179
172180
172181
172182
172183
172184
172185
172186
172187
172188
172189
172190
172191
172192
172193
172194
172195
172196
172197
172198
172199
172200
172201
172202
172203
172204
172205
172206
172207
172208
172209
172210
172211
172212
172213
172214
172215
172216
172217
172218
172219
172220
172221
172222
172223
172224
172225
172226
172227
172228
172229
172230
172231
172232
172233
172234
172235
172236
172237
172238
172239
172240
172241
172242
172243
172244
172245
172246
172247
172248
172249
172250
172251
172252
172253
172254
172255
172256
172257
172258
172259
172260
172261
172262
172263
172264
172265
172266
172267
172268
172269
172270
172271
172272
172273
172274
172275
172276
172277
172278
172279
172280
172281
172282
172283
172284
172285
172286
172287
172288
172289
172290
172291
172292
172293
172294
172295
172296
172297
172298
172299
1722100
1722101
1722102
1722103
1722104
1722105
1722106
1722107
1722108
1722109
1722110
1722111
1722112
1722113
1722114
1722115
1722116
1722117
1722118
1722119
1722120
1722121
1722122
1722123
1722124
1722125
1722126
1722127
1722128
1722129
1722130
1722131
1722132
1722133
1722134
1722135
1722136
1722137
1722138
1722139
1722140
1722141
1722142
1722143
1722144
1722145
1722146
1722147
1722148
1722149
1722150
1722151
1722152
1722153
1722154
1722155
1722156
1722157
1722158
1722159
1722160
1722161
1722162
1722163
1722164
1722165
1722166
1722167
1722168
1722169
1722170
1722171
1722172
1722173
1722174
1722175
1722176
1722177
1722178
1722179
1722180
1722181
1722182
1722183
1722184
1722185
1722186
1722187
1722188
1722189
1722190
1722191
1722192
1722193
1722194
1722195
1722196
1722197
1722198
1722199
1722200
1722201
1722202
1722203
1722204
1722205
1722206
1722207
1722208
1722209
1722210
1722211
1722212
1722213
1722214
1722215
1722216
1722217
1722218
1722219
1722220
1722221
1722222
1722223
1722224
1722225
1722226
1722227
1722228
1722229
17222210
17222211
17222212
17222213
17222214
17222215
17222216
17222217
17222218
17222219
17222220
17222221
17222222
17222223
17222224
17222225
17222226
17222227
17222228
17222229
17222230
17222231
17222232
17222233
17222234
17222235
17222236
17222237
17222238
17222239
17222240
17222241
17222242
17222243
17222244
17222245
17222246
17222247
17222248
17222249
17222250
17222251
17222252
17222253
17222254
17222255
17222256
17222257
17222258
17222259
17222260
17222261
17222262
17222263
17222264
17222265
17222266
17222267
17222268
17222269
17222270
17222271
17222272
17222273
17222274
17222275
17222276
17222277
17222278
17222279
17222280
17222281
17222282
17222283
17222284
17222285
17222286
17222287
17222288
17222289
17222290
17222291
17222292
17222293
17222294
17222295
17222296
17222297
17222298
17222299
172222100
172222101
172222102
172222103
172222104
172222105
172222106
172222107
172222108
172222109
172222110
172222111
172222112
172222113
172222114
172222115
172222116
172222117
172222118
172222119
172222120
172222121
172222122
172222123
172222124
172222125
172222126
172222127
172222128
172222129
172222130
172222131
172222132
172222133
172222134
172222135
172222136
172222137
172222138
172222139
172222140
172222141
172222142
172222143
172222144
172222145
172222146
172222147
172222148
172222149
172222150
172222151
172222152
172222153
172222154
172222155
172222156
172222157
172222158
172222159
172222160
172222161
172222162
172222163
172222164
172222165
172222166
172222167
172222168
172222169
172222170
172222171
172222172
172222173
172222174
172222175
172222176
172222177
172222178
172222179
172222180
172222181
172222182
172222183
172222184
172222185
172222186
172222187
172222188
172222189
172222190
172222191
172222192
172222193
172222194
172222195
172222196
172222197
172222198
172222199
172222200
172222201
172222202
172222203
172222204
172222205
172222206
172222207
172222208
172222209
172222210
172222211
172222212
172222213
172222214
172222215
172222216
172222217
172222218
172222219
172222220
172222221
172222222
172222223
172222224
172222225
172222226
172222227
172222228
172222229
172222230
172222231
172222232
172222233
172222234
172222235
172222236
172222237
172222238
172222239
172222240
172222241
172222242
172222243
172222244
172222245
172222246
172222247
172222248
172222249
172222250
172222251
172222252
172222253
172222254
172222255
172222256
172222257
172222258
172222259
172222260
172222261
172222262
172222263
172222264
172222265
172222266
172222267
172222268
172222269
172222270
172222271
172222272
172222273
172222274
172222275
172222276
172222277
172222278
172222279
172222280
172222281
172222282
172222283
172222284
172222285
172222286
172222287
172222288
172222289
172222290
172222291
172222292
172222293
172222294
172222295
172222296
172222297
172222298
172222299
1722222100
1722222101
1722222102
1722222103
1722222104
1722222105
1722222106
1722222107
1722222108
1722222109
1722222110
1722222111
1722222112
1722222113
1722222114
1722222115
1722222116
1722222117
1722222118
1722222119
1722222120
1722222121
1722

1026 Table S2: ECG diagnostic labels in the PTB-XL dataset. The dataset contains 42 diagnostic labels
 1027 in a 3-level hierarchy: superclass, subclass, and specific diagnosis. The **Label** column denotes the
 1028 specific diagnostic label, while the **Subclass** and **Superclass** column indicates its corresponding
 1029 subclass and superclass categories.

1030

1031 Label	1032 Description	1033 Subclass	1034 Superclass
Conduction Disturbances (CD)			
1035 LAFB	1036 Left anterior fascicular block	1037 LAFB/LPFB	1038 CD
1036 LPFB	1037 Left posterior fascicular block	1038 LPFB	1039 CD
1037 IRBBB	1038 Incomplete right bundle branch block	1039 IRBBB	1040 CD
1038 ILBBB	1039 Incomplete left bundle branch block	1040 ILBBB	1041 CD
1039 CRBBB	1040 Complete right bundle branch block	1041 CRBBB	1042 CD
1040 CLBBB	1041 Complete left bundle branch block	1042 CLBBB	1043 CD
1041 AVB	1042 First degree AV block	1043 AVB	1044 CD
1042 3AVB	1043 Third degree AV block	1044 3AVB	1045 CD
1043 2AVB	1044 Second degree AV block	1045 2AVB	1046 CD
1044 IVCD	1045 Non-specific intraventricular disturbance	1046 conduction	1047 IVCD
1045 WPW	1046 Wolff-Parkinson-White syndrome	1047 WPW	1048 CD
Hypertrophy (HYP)			
1047 LVH	1048 Left ventricular hypertrophy	1049 LVH	1050 HYP
1048 LAO/LAE	1049 Left atrial overload/enlargement	1050 LAO/LAE	1051 HYP
1049 RVH	1050 Right ventricular hypertrophy	1051 RVH	1052 HYP
1050 RAO/RAE	1051 Right atrial overload/enlargement	1052 RAO/RAE	1053 HYP
1051 SEHYP	1052 Septal hypertrophy	1053 SEHYP	1054 HYP
Myocardial Infarction (MI)			
1052 IMI	1053 Inferior myocardial infarction	1054 IMI	1055 MI
1053 ILMCI	1054 Inferolateral myocardial infarction	1055 ILMCI	1056 MI
1054 IPLMI	1055 Inferoposterolateral myocardial infarction	1056 IPLMI	1057 MI
1055 IPMI	1056 Inferoposterior myocardial infarction	1057 IPMI	1058 MI
1056 INJIN	1057 Subendocardial injury in inferior leads	1058 INJIN	1059 MI
1057 INJIL	1058 Subendocardial injury in inferolateral leads	1059 INJIL	1060 MI
1058 ASMI	1059 Anteroseptal myocardial infarction	1060 ASMI	1061 MI
1059 AMI	1060 Anterior myocardial infarction	1061 AMI	1062 MI
1060 ALMI	1061 Anterolateral myocardial infarction	1062 ALMI	1063 MI
1061 INJAS	1062 Subendocardial injury in anteroseptal leads	1063 INJAS	1064 MI
1062 INJAL	1063 Subendocardial injury in anterolateral leads	1064 INJAL	1065 MI
1063 INJLA	1064 Subendocardial injury in lateral leads	1065 INJLA	1066 MI
1064 LMI	1065 Lateral myocardial infarction	1066 LMI	1067 MI
1065 PMI	1066 Posterior myocardial infarction	1067 PMI	1068 MI
ST/T Changes (STTC)			
1066 NDT	1067 Non-diagnostic T abnormalities	1068 STTC	1069 STTC
1067 DIG	1068 Digitalis effect	1069 DIG	1070 STTC
1068 LNGQT	1070 Long QT interval	1071 LNGQT	1072 STTC
1069 ANEUR	1072 ST-T changes compatible with ventricular aneurysm	1073 ANEUR	1074 STTC
1070 EL	1074 Electrolytic disturbance or drug effect	1075 EL	1076 STTC
1071 NST	1075 Non-specific ST changes	1076 NST	1077 STTC
1072 ISC	1076 Non-specific ischemic changes	1077 ISC	1078 STTC
1073 ISCIN	1077 Ischemic changes in inferior leads	1078 ISCIN	1079 STTC
1074 ISCIL	1078 Ischemic changes in inferolateral leads	1079 ISCIL	1080 STTC
1075 ISCAL	1079 Ischemic changes in anterolateral leads	1080 ISCAL	1081 STTC
1076 ISCAS	1080 Ischemic changes in anteroseptal leads	1081 ISCAS	1082 STTC
1077 ISCLA	1081 Ischemic changes in lateral leads	1082 ISCLA	1083 STTC
1078 ISCAN	1082 Ischemic changes in anterior leads	1083 ISCAN	1084 STTC
Normal (NORM)			
1079 NORM	1080 Normal ECG	1081 NORM	1082 NORM

1076

1077

1078 summarization of the current available ECG dataset is tabulated in Table S1. Among the datasets,
 1079 MIMIC-IV-ECG emerges as a particularly well-suited one for health risk assessment applications due
 to its combination of large-scale clinical data (161,352 patients) and a relatively inclusive metadata,

1080 Table S3: Form labels and label descriptions in the PTB-XL dataset. Labels represent morphological
 1081 and rhythm abnormalities detected in ECG recordings.

Label	Description	Label	Description
ABQRS	Abnormal QRS	LOWT	Low amplitude T-waves
DIG	Digitalis effect	LPR	Prolonged PR interval
HVOLT	High QRS voltage	LVOLT	Low QRS voltages
INVT	Inverted T-waves	NDT	Non-diagnostic T abnormalities
LNGQT	Long QT interval	NST	Non-specific ST changes
NT	Non-specific T-wave changes	PAC	Atrial premature complex
PRC(S)	Premature complex(es)	PVC	Ventricular premature complex
QWAVE	Q waves present	STD	Non-specific ST depression
STE	Non-specific ST elevation	TAB	T-wave abnormality
VCLVH	Voltage criteria for LVH		

1092 Table S4: Rhythm labels and label descriptions in the PTB-XL dataset. Labels represent various
 1093 cardiac rhythm patterns and arrhythmias detected in ECG recordings.

Label	Description	Label	Description
AFIB	Atrial fibrillation	PACE	Normal functioning artificial pacemaker
AFLT	Atrial flutter	PSVT	Paroxysmal supraventricular tachycardia
BIGU	Bigeminal pattern (unknown origin)	SARRH	Sinus arrhythmia
SBRAD	Sinus bradycardia	SR	Sinus rhythm
STACH	Sinus tachycardia	SVARR	Supraventricular arrhythmia
SVTAC	Supraventricular tachycardia	TRIGU	Trigeminal pattern (unknown origin)

1103 including gender, age, and SBP. In contrast, alternative datasets present significant limitations: smaller-
 1104 scale collections like Chapman and PTB-XL lack the statistical power for robust foundation model
 1105 training, while metadata-sparse datasets preclude comprehensive health risk modeling. Although
 1106 HEEDB offers a superior scale (1.8 million patients), its restricted accessibility limits reproducibility
 1107 and benchmarking capabilities.

B DETAILS OF THE DATASETS USED FOR PRETRAINING AND DOWNSTREAM EVALUATION

1113 **MIMIC-IV-ECG** dataset is a large-scale clinical ECG database containing 800,035 12-lead
 1114 diagnostic recordings collected from 161,352 unique patients.⁴ Each ECG is 10 seconds in length
 1115 and sampled at 500Hz. Corresponding metadata can be matched to each patient using the unique
 1116 patient code.⁵ A more detailed procedure of aligning metadata for each ECG sample is provided
 1117 in Appendix B.1. We use this dataset in 2 phases: (1) for pretraining, we use the first unique ECG
 1118 recording from each patient and train over all distinct patients, and (2) for downstream evaluation,
 1119 we follow prior work (Li et al., 2024) who obtained the labels for the LVEF estimation task from
 1120 the discharge section of MIMIC-IV-Notes. Specifically, LVEF values are provided as continuous
 1121 labels for the regression task, while an LVEF of 50% or higher is considered normal and below 50%
 1122 abnormal, thereby defining a binary classification task. We sequentially split the dataset into training
 1123 (60,329 samples, 80%), validation (7,541 samples, 10%), and test sets (7,542 samples, 10%).

1124 **MUSIC** dataset (MUerte Subita en Insuficiencia Cardiaca) focuses on assessing cardiac mortality
 1125 and sudden cardiac death (SCD) in ambulatory patients with chronic heart failure (CHF).⁶ It contains
 1126 992 patients with CHF consecutively enrolled from the specialized HF clinics of eight University
 1127 Spanish Hospitals between April 2003 and December 2004. All patients are measured with a 3-lead
 1128 resting electrocardiogram (ECG), or a 3-lead Holter ECG. In this study, we focus on using the first
 1129 lead of the collected Holter ECG signal. In the original dataset, outcome labels (non-cardiovascular
 1130 death, sudden cardiac death, or pump-failure death) are provided at the patient level. For each
 1131 patient, we extract the first 10 seconds of ECG recordings with non-zero signals. For the downstream

⁴MIMIC-IV-ECG dataset available at: physionet.org/content/mimic-iv-ecg.

⁵Metadata for MIMIC-IV available at: physionet.org/content/mimiciv.

⁶MUSIC dataset available at: physionet.org/content/music-sudden-cardiac-death.

1134

Table S5: Beat Symbol Definitions with Icentia11K data set

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

Symbol	Beat Description
N	Normal
S	ESSV (PAC): Premature or ectopic supraventricular beat, premature atrial contraction
V	ESV (PVC): Premature ventricular contraction, premature ventricular contraction

1157

Table S6: Rhythm symbol definitions with Icentia11K data set.

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

Symbol	Rhythm Type	Rhythm description
(N ...)	NSR (Normal sinus rhythm)	-
(AFIB ...)	AFib (Atrial fibrillation)	Irregular rhythm with absent P waves and irregular RR intervals
(AFL ...)	AFlutter (Atrial flutter)	Regular atrial arrhythmia with sawtooth flutter waves

1172

1173

1174

1175

1176

1177

1178

1179

PTB-XL dataset is a large-scale ECG dataset that has been widely used in prior research (cite) for evaluating model capacity in signal pattern and disease identification.⁷ We adopt the preprocessing and label alignment procedures described in the original dataset publication.⁸ Details of the tasks of PTB-XL are tabulated in Tables S2, S3, and S4. To ensure fair comparison, we follow prior work (Wagner et al., 2020) and adopt a data partitioning of [80/10/10]%, yielding 16,832, 2,100, and 2,098 samples for training, validation, and testing phases, respectively.

Icentia11K dataset is a wearable dataset that contains ECG signals collected from single-lead chest-mounted wearable devices.⁹ The sample rate for the dataset is 250 Hz. Beat and rhythm labels are extracted from the dataset’s annotation files, focusing on three beat types (Normal, Supraventricular, Ventricular) and three rhythm types (Normal Sinus Rhythm, Atrial Fibrillation, Atrial Flutter), see Tables S5 and S6 for a more detailed description. Note that the dataset also includes an undefined beat class, which we omit here as it does not correspond to a physiologically interpretable beat type. For both tasks, segments are generated by slicing 10-second windows starting from the annotated event and assigning numerical labels for classification. To ensure balanced representation across patients, we employ a patient-stratified sampling strategy: for each patient, we search their recordings for each target beat type and randomly select one representative instance, continuing until all types are found or available segments are exhausted. This results in relatively balanced distributions across beat classes (N: 10,866; S: 9,844; V: 9,287), whereas rhythm classes remain imbalanced due to the limited prevalence of atrial fibrillation and atrial flutter (N: 10,239; AFib: 743; AFL: 516). For experimental evaluation, the dataset is partitioned into training, validation, and test sets (90%/10%/10%), comprising 20,994/4,532/4,511 samples for beat classification and 8,042/1,723/1,733 samples for rhythm classification.

Chapman dataset is a collection of 12-lead ECG recordings from 45,152 patients labeled by clinical experts to support research in arrhythmia and cardiovascular disease detection.¹⁰ The dataset contains 10-second recordings sampled at 500 Hz, featuring 11 common cardiac rhythms and 67 additional cardiovascular conditions, all validated through a rigorous multi-physician review process with senior physician arbitration for diagnostic disagreements. We follow prior work (Jin et al., 2025) and use a refined version of the dataset, which contains 23,026 ECG recordings with 38 distinct labels. The samples are then split into training (16,546 samples, 70%), validation (1,860 samples, 10%), and test sets (4,620 samples, 20%).

MC-MED dataset is a multimodal collection of emergency department visits from 118,385 adult patients at Stanford Health Care between 2020 and 2022.¹¹ The dataset combines continuous physiological monitoring (lead II ECG, photoplethysmography, respiration waveforms), clinical data

1180

1181

1182

1183

1184

1185

1186

1187

⁷PTB-XL dataset available at: physionet.org/content/ptb-xl.

⁸Data preprocessing code for PTB-XL available at: github.com/helme/ECG_PTB_XL_Benchmarking.

⁹Icentia11K dataset available at: physionet.org/content/centia11k-continuous-ecg.

¹⁰Chapman dataset available at: physionet.org/content/ecg-arrhythmia.

¹¹MC-MED dataset available at: physionet.org/content/mc-med.

(demographics, medical histories, laboratory results, medications, radiology reports), and temporal visit outcomes. It covers emergency department patients during and after the COVID-19 pandemic with granular physiological measurements.

The prediction targets encompass both classification and regression tasks across the clinical care continuum. For classification, there are 3 tasks: (1) emergency department (ED) disposition, which determines patient placement after ED assessment. This includes 4 classes: Discharge (outpatient), Inpatient (hospital admission), Observation (extended monitoring without full admission), and ICU (critical care); (2) Discharge (DC) disposition, classification of patient outcomes at hospital discharge. We categorized the outcome into 5 categories: home care, care facility, hospital transfer, psychiatric care, and death. and (3) Triage acuity assessment, which is predicting clinical urgency as determined by healthcare professionals during initial patient evaluation, including 5 tasks: Resuscitation, emergent, urgent, Semi-Urgent, and Non-Urgent. The regression tasks involve continuous prediction of systolic blood pressure (SBP) and diastolic blood pressure (DBP) values. We follow the chronological splitting provided in the original dataset based on patient admission dates, and allocate the first 78% (37,438 samples) for training, the following 11% (5,540 samples) for validation, and the most recent 11% (5,572 samples) for testing. Note that the exact number may differ from the original dataset, as we only include patients with lead II ECG recordings.

Aurora BP dataset is a collection of simultaneous multi-modal physiological recordings collected from 1,221 diverse participants,¹² serving the purpose for cuffless blood pressure research. The dataset contains synchronized tonometry, photoplethysmography (PPG), electrocardiography (ECG), accelerometry, and reference blood pressure measurements collected during both laboratory and 24-hour ambulatory monitoring phases, with participants spanning a wide range of ages and hypertensive status to ensure real-world applicability. The prediction targets are systolic blood pressure (SBP) and diastolic blood pressure (DBP), with the task formulated as the estimation of their actual values through regression. We use a 70%/15%/15% train/validation/test split on patient ID, resulting in 786/169/169 patients, which corresponds to 9,237/1,913/1,854 samples for training, validation, and testing, respectively.

B.1 METADATA ALIGNMENT WITH MIMIC-IV-ECG DATASET

The MIMIC-IV-ECG database provides raw electrocardiogram (ECG) waveforms together with patient identifiers and relative timestamps, but does not directly include demographic or clinical metadata. To match each ECG with patient information, we aligned it with the corresponding records in MIMIC-IV,¹³ which contains both static demographics (e.g., sex, age at admission) and timestamped clinical observations (e.g., blood pressure). Each ECG was first linked to the corresponding patient using the unique patient identifier, which is shared across MIMIC-IV and MIMIC-IV-ECG. Demographic variables that remain constant over time were directly assigned to all ECGs of the same patient. For time-varying clinical measurements, we selected the most recent value recorded at or before the ECG timestamp, ensuring that the metadata reflected the patient’s state at the time of acquisition. If no prior measurement was available, the variable was marked as missing, and no forward-filling across admissions or imputation beyond this step was performed. This procedure ensures that every ECG recording is annotated with the most temporally relevant metadata.

C SUPPLEMENTARY EXPERIMENT SETUP AND BASELINES

C.1 EXPERIMENT SETUP FOR DOWNSTREAM EVALUATION

This section supplements the experiment setup in section 3 of the main paper. All experiments were conducted using a Linux server (Ubuntu 22.04.5) with 4 NVIDIA L40S GPUs, 2 Intel Xeon Gold 6542Y CPUs. Code will be made available.

Pretraining: All experiments are conducted under fixed hyperparameter settings. All models are pretrained for 100 epochs with a batch size of 64 using the AdamW optimizer, except for the experiments for pretraining on a specific lead due to a reduced data diversity (section 4, pretraining

¹²A sample data of Aurora BP is available at: github.com/microsoft/aurorabp-sample-data. Access can be provided via application.

¹³MIMIC-IV database available at: physionet.org/content/mimiciv/3.1/

on a specific lead), where we pretrain for 10 epochs. A learning rate of 1×10^{-4} and a weight decay of 5×10^{-5} are employed, with a Cosine Annealing Learning Rate scheduler to adjust the learning rate. Training stops if the validation loss does not decrease for 20 consecutive epochs. Seed is fixed to 42 for all pretraining. For parameters specific to contrastive pretraining, we set the clinical-guidance coefficient to $\alpha = 0.2$ and temperature to $\tau = 0.07$, consistent with prior literature (Wu et al., 2018; He et al., 2020). The effect of α and τ is further examined in Appendix D.8, Table S17, and Table S18.

Downstream tasks: For downstream evaluation, training is also conducted for 100 epochs with a batch size of 64. We use binary cross-entropy loss for classification and L1 loss for regression tasks. Model parameters are optimized using Adam with a learning rate of 1×10^{-3} and a weight decay of 1×10^{-5} . The learning rate is adjusted using a Cosine Annealing with Warm Restarts scheduler. Training stops if the validation loss does not decrease for 5 consecutive epochs. Seed is fixed to 42 for all downstream evaluation, except for the seed robustness evaluation in Appendix D.1.

Downstream preprocessing: For all downstream tasks, the ECG signals are resampled to 500 Hz. For datasets containing recordings longer than 10 seconds, we keep the first 10 seconds as our input. We preprocessed all ECG signals using a 5th-order Butterworth bandpass filter with cutoff frequencies of 0.67 Hz and 40 Hz to remove low-frequency baseline wander (e.g. due to breathing) whilst maintaining heart rate frequencies, and to remove high-frequency noise (e.g. power-line interference) whilst maintaining fundamental frequencies of P, QRS, and T-waves. The input ECG samples are then standardized using z-scoring, an operation that transforms the data to zero mean and unit standard deviation.

Regression specific setup: In all regression tasks, the target labels are also z-scored using statistics computed from the training set. Predictions are de-normalized back to the original scale before comparison with the ground-truth labels. The errors are evaluated on the original scale.

C.2 DATA AUGMENTATION FOR CONTRASTIVE PRETRAINING

This section supplements the **stochastic data augmentation** in section 3.

Simulate physiological noise We inject physiological noise collected from the MIT-BIH noise stress test database (Sološenko et al., 2021a), as a form of data augmentation that encourages the model to learn representations invariant to electrode movement and motion-induced artifacts, etc.¹⁴ The noise injection is achieved by selecting the noise signal from the corresponding lead, given by $\hat{\mathbf{e}}_s^l = \mathbf{e}_s^l + \phi \times \mathbf{n}^l$, where ϕ controls the noise intensity, \mathbf{n}^l is the noise vector of the l -th lead. We provide visualization in Figure S1 of how different noises influence and augment the original signal. Note that we set $\phi = 0.02$ following prior work (Goldberger et al., 2000; Sološenko et al., 2021b), a parameter also used in the main experiment.

¹⁴The noise is available at: physionet.org/content/ecg-ppg-simulator-arrhythmia.

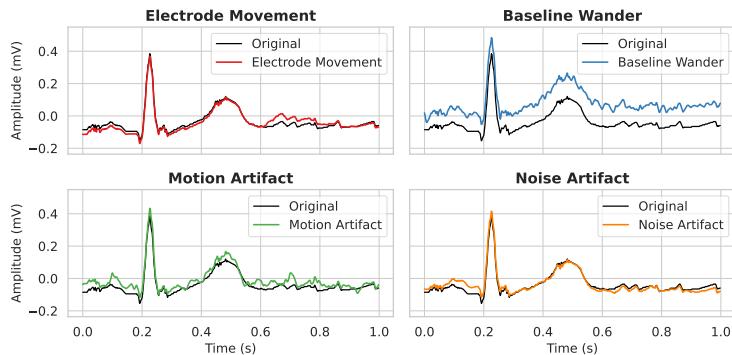


Figure S1: Physiological noise for single-lead ECG contrastive learning. We include four common sources of signal degradation in wearable ECG devices: electrode movement artifacts, baseline wander, motion-induced distortions, and additive noise. Original signals are shown in black, with augmented versions in color.

1296 Table S7: ResNeXt1D Model Architecture Configurations for *CLEF* variants. **Left:** Shared
 1297 hyperparameters of all three variants. **Right:** Detailed stage-wise model configurations.
 1298

Parameter	<i>CLEF-S</i>	<i>CLEF-M</i>	<i>CLEF-L</i>	Stage	<i>CLEF-S</i> h/γ	<i>CLEF-M</i> h/γ	<i>CLEF-L</i> h/γ
Model Size	448K	30.7M	296M	1	32 / 1	64 / 2	128 / 2
				2	64 / 1	160 / 2	256 / 3
Hidden Dimension d	32	64	128	3	64 / 2	160 / 2	256 / 3
				4	128 / 2	400 / 3	512 / 4
Ratio ι	0.5	1.0	1.5	5	128 / 2	400 / 3	512 / 4
				6	256 / 2	1024 / 4	1024 / 5
Groups Width g	8	16	32	7	—	1024 / 4	1024 / 5
				8	—	—	2048 / 6
Number of Stages ϖ	6	7	9	9	—	—	2048 / 6

1309 **Random mask** Randomly masking segments of the ECG signal encourages the model to learn
 1310 representations that are invariant to missing information. In our experiments, we set the masking
 1311 probability to $p = 0.2$, with each selected segment having 10% of its signal randomly masked.
 1312

1314 C.3 BACKBONE RESNET MODELS

1316 *CLEF* is built upon a multi-layer ResNeXt1D architecture, where each stage contains multiple
 1317 residual blocks. We implement three variants, i.e. small, medium, and large, for our proposed *CLEF*
 1318 model with different computational complexities. All variants follow the same architectural principles
 1319 but differ in depth, width, and parameter count.

1320 **Convolution layer.** The network begins with an initial convolutional layer with 1 channel, a kernel
 1321 size of 16, and a stride of 2. This layer takes the input ECG signal $\mathbf{e}_s^l \in \mathbb{R}^t$ from the s -th sample and
 1322 l -th lead, where t denotes the signal length. that takes input ECG signal $\mathbf{e}_s^l \in \mathbb{R}^t$ of s -th sample and
 1323 l -th lead, where t denotes the signal length. The input \mathbf{e}_s^l is first transformed by the layer and passed
 1324 through a swish activation, yielding the hidden state $\mathbf{H}_s^l \in \mathbb{D}^{d \times \ell}$, where $\ell = \frac{t}{2}$ is the sequence length,
 1325 and the hidden dimension d is set to 32, 64, or 128 for small, medium, and large variants, respectively.
 1326 For brevity, we omit the sample & lead indices s and l in the intermediate representations.

1327 **Residual stages.** The network then processes \mathbf{H}_s^l through a sequence of ϖ residual stages. Each
 1328 stage consists of γ residual blocks. Each residual block processes features through three convolution
 1329 layers sequentially:

1331 (1) A 1×1 convolution layer with stride 1 transforms \mathbf{H}_s^l to an intermediate dimension $\mathbf{H}^{(1)} \in \mathbb{R}^{d^{(1)} \times \ell}$,
 1332 where $d^{(1)} = d \times \iota$, with $\iota \in \{0.5, 1.0, 1.5\}$ is a predefined ratio parameter for small, medium, and
 1333 large models, respectively.

1334 (2) A 16×16 convolution layer with stride 1 is applied to $\mathbf{H}^{(1)}$, producing the layer output $\mathbf{H}^{(2)} \in$
 1335 $\mathbb{R}^{d^{(1)} \times \ell}$. Specifically, the input is first divided into $d^{(1)}/g$ groups along channel dimension, where
 1336 $g \in \{8, 16, 32\}$ specifies the group width for the small, medium, and large variants, respectively.
 1337 Each group is convolved independently with its own set of filters, and the resulting features are
 1338 concatenated along the channel dimension to form the output $\mathbf{H}^{(2)}$.

1339 (3) Finally, a 1×1 convolution layer maps $\mathbf{H}^{(2)}$ to the output tensor $\mathbf{H}^{(3)} \in \mathbb{R}^{h \times \ell}$, where h denotes
 1340 the final output channel dimension of the residual block.

1342 A swish activation is used between layers. Following the final convolutional output of the last block
 1343 of each stage, the output matrix $\mathbf{H}^{(3)}$ is first averaged along sequence dimension, forming vector
 1344 $\mathbf{h}_s \in \mathbb{R}^h$. This vector \mathbf{h}_s is passed through a swish-activated 2-layer perception (hidden dimension
 1345 is $\frac{h}{2}$, output dimension is h), followed by a sigmoid activation, forming channel-wise gating scores
 1346 $\mathbf{h}_r \in \mathbb{R}^h$. Each channel vector of $\mathbf{H}^{(3)}$ is multiplied by the corresponding element of \mathbf{h}_r , and the
 1347 reweighted channel vectors collectively form the residual tensor $\mathbf{H}_r \in \mathbb{R}^{h \times \ell}$. This residual term is
 1348 then added with $\mathbf{H}^{(3)}$ to form the output of the state, given by $\mathbf{H}_o = \mathbf{H}^{(3)} + \mathbf{H}_r$. The final output of
 1349 the last stage is averaged over the sequence dimension to obtain the feature embedding of the input
 signal, denoted as $\mathbf{z} \in \mathbb{R}^h$. Table S7 details the component parameters of each variant:

1350 Table S8: SCORE2 model coefficients (β_i) by gender and age group (<70 vs. ≥ 70 years).
1351

Coefficient	Male < 70	Female < 70	Male ≥ 70	Female ≥ 70
β_1	0.3742	0.4648	0.0634	0.0789
β_2	0.6012	0.7744	0.3524	0.4921
β_3	0.2777	0.3131	0.0094	0.0102
β_4	0.6457	0.8096	0.4245	0.6010
β_5	0.1458	0.1002	0.0850	0.0605
β_6	-0.2698	-0.2606	-0.3564	-0.3040
β_7	-0.0755	-0.1088	-0.0247	-0.0255
β_8	-0.0255	-0.0277	-0.0005	-0.0004
β_9	-0.0281	-0.0226	0.0073	-0.0009
β_{10}	0.0426	0.0613	0.0091	0.0154
β_{11}	-0.0983	-0.1272	-0.0174	-0.0107
$S_0(t)$	0.9605	0.9776	0.7576	0.8082
c	0	0	0.0929	0.2290

1364
1365
1366 C.4 BASELINE MODELS
13671368 **ST-MEM** (Na et al., 2024) is a self-supervised learning framework specifically designed to capture
1369 spatio-temporal dependencies in 12-lead ECG signals. It employs a spatio-temporal masked auto-
1370 encoder that reconstructs randomly masked signal segments, enabling the model to learn rich ECG
1371 representations.1372 **KED** (Tian et al., 2024) is an ECG foundation model pretrained with raw ECG signal and textual
1373 input (ECG reports). They used contrastive learning to align the representation of two modalities,
1374 enhancing diagnostic performance.1375 **ECGFounder** (Li et al., 2024) is a foundation model for ECG trained in a supervised manner on
1376 large-scale labeled datasets for multi-label classification. To address the challenge of missing labels
1377 that are commonly present in long-tailed clinical datasets, the authors propose a modified loss function
1378 that enables more robust and balanced representation learning despite label sparsity.1379 **Moirai** (Woo et al., 2024) is a foundation model for time series designed for general-purpose
1380 forecasting. It supports both univariate and multivariate forecasting. For multivariable input, to get an
1381 aligned sequence, they reshape the input sequence into a single, aligned sequence that integrates both
1382 temporal and variable dimensions, enabling unified processing through the transformer architecture.1383 **Moment** (Goswami et al., 2024) is a transformer-based model built on the T5 architecture, designed
1384 for univariate time series tasks. It supports both sequence-to-sequence forecasting and time series
1385 classification, offering a unified and scalable foundational framework for temporal modeling.1387
1388 C.5 BASELINE PRETRAINING METHODS
13891390 **SimCLR** (Chen et al., 2020) is a contrastive learning framework whose objective is to bring
1391 the representations of augmented views of the same sample closer together, while pushing apart
1392 representations of different samples. It relies on a large set of diverse negative samples, which are
1393 typically achieved through large batch sizes, to provide effective contrastive learning signals.1394 **BYOL** (Grill et al., 2020) is a self-supervised learning method that does not require negative samples.
1395 It proposed to have an online network and a target network. The online network is trained to predict
1396 the target network's representation of the same augmented input. The target network is updated via
1397 an exponential moving average of the online parameters.1398 **MoCo** (He et al., 2020) is also a modified contrastive learning framework. Instead of relying on
1399 negatives sampled directly from the dataset, they propose to have a dynamic memory bank that stores
1400 encoded representations from previous batches. This memory bank is updated using a momentum
1401 encoder, which ensures stable and consistent representations over time. The stored representations
1402 serve as negative samples for contrastive learning, providing a large and diverse set of negatives
1403 without requiring large batch sizes.

1404 C.6 DETAILS IN OBTAINING SCORE2 RISK SCORES
1405

1406 The SCORE2 risk score estimates the 10-year risk of fatal and non-fatal cardiovascular disease.¹⁵
 1407 Demographic information (e.g. sex and age) and clinical measurements (e.g. blood pressure) are
 1408 used to predict cardiovascular risk (working group and risk collaboration, 2021a;b). Specifically,
 1409 for sample s , the score uses 7 features, including: age (α_s^1 , in years), gender ($\alpha_s^2 \in \{\text{male, female}\}$),
 1410 smoking status ($\alpha_s^3 \in \{0, 1\}$), systolic blood pressure (SBP, α_s^4 in mmHg), diabetes status ($\alpha_s^5 \in \{0, 1\}$),
 1411 total cholesterol (α_s^6 in mmol/L), and high-density lipoprotein (HDL) cholesterol (α_s^7 in mmol/L).

1412 **Detailed risk score computation.** The 7 covariates are first scaled by subtracting their
 1413 corresponding mean values and then scaling by their pre-defined units (e.g., 5 years for age, 20
 1414 mmHg for SBP), except for the 3 binary features, gender, smoking status, and diabetes status. Given
 1415 sample s , the transformed variables are given as:

$$1417 u_s^1 := \frac{\alpha_s^1 - 60}{5}, \quad u_s^2 := \alpha_s^3, \quad u_s^3 := \frac{\alpha_s^4 - 120}{20}, \quad u_s^4 := \alpha_s^5, \quad u_s^5 := \alpha_s^6 - 6, \quad u_s^6 := \frac{\alpha_s^7 - 1.3}{0.5}. \quad (S1)$$

1419 These normalized covariates, collected in vector $\mathbf{u} = [u_s^1, u_s^2, \dots, u_s^6] \in \mathbb{R}^6$, are then transformed
 1420 with parameters $\mathbf{b}_1 = [\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6] \in \mathbb{R}^6$, and age-dependent transformation $\mathbf{b}_2 =$
 1421 $[0, \beta_7, \beta_8, \beta_9, \beta_{10}, \beta_{11}] \in \mathbb{R}^6$, given by:

$$1422 \chi_s = \mathbf{b}_1 \mathbf{u}_s^\top + u_s^1 \mathbf{b}_2 \mathbf{u}_s^\top, \quad (S2)$$

1424 where χ marks an intermediate score. Different parameter sets \mathbf{b}_1 and \mathbf{b}_2 are used for males and
 1425 females, and also for different age groups. Details are tabulated in Table S8.

1426 The health risk score r of a patient sample is then obtained using an exponential transformation of χ :

$$1427 r_s = 1 - S_0(t)^{\exp(\chi_s - c)}, \quad (S3)$$

1429 where $S_0(t)$ is the baseline survival function at time t (10 years) and c is an offset. Different $S_0(t)$
 1430 and c are also used for different gender and age groups. Details are tabulated in Table S8. Note that
 1431 the original SCORE2 calculation includes a step for regional calibration. As regional information is
 1432 unavailable in our dataset, this step was omitted. We directly used the uncalibrated 10-year risk score,
 1433 with no further region-specific adjustment applied.

1434 The SCORE2 risk score initially uses 7 variables, but because MIMIC-IV-ECG only records the age,
 1435 gender, and systolic BP, we infer missing values with simple imputation strategies. Smoking and
 1436 diabetes status were assumed to be absent if not recorded (i.e., set to 0). The missing total cholesterol
 1437 and HDL cholesterol values were estimated using population-based reference values stratified by sex,
 1438 with added Gaussian noise to reflect natural biological variation. Specifically, the total cholesterol
 1439 was imputed as 5.2 ± 0.5 mmol/L and HDL cholesterol as 1.3 ± 0.2 mmol/L, based on the mean
 1440 values of age group 40-45 (working group and risk collaboration, 2021a). Additionally, there are 15
 1441 records with missing age, which we impute as 40. The number of missing values is accounted for in
 1442 the **handling missing metadata** step described in section 2.2.

1443 ¹⁵SCORE2 risk score is obtained following the implementation available at:
 1444 github.com/dvicencio/RiskScorescvd.

1446 Table S9: Results for finetuning on lead I (upper table) and lead II (lower table) ECGs with confidence
 1447 interval. AUROCs are reported for seven clinical tasks: LVEF refers to left ventricular ejection
 1448 fraction; Dx denotes diagnosis; SubDx denotes subdiagnosis; SupDx denotes superdiagnosis; Form
 1449 refers to waveform morphology; Rhyth denotes rhythm; Arrhy denotes arrhythmia.

Dataset	MIMIC-IV			PTB-XL			Chapman
Task	LVEF	Dx	SubDx	SupDx	Form	Rhyth	Arrhy
Lead I							
CLEF-S (448K)	0.8276 \pm 0.007	0.8473 \pm 0.011	0.8515 \pm 0.009	0.8311 \pm 0.007	0.7632 \pm 0.047	0.9515 \pm 0.011	0.9001 \pm 0.003
CLEF-M (30.7M)	0.8230 \pm 0.010	0.8499 \pm 0.011	0.8581 \pm 0.003	0.8335 \pm 0.005	0.7708 \pm 0.019	0.9415 \pm 0.006	0.9074 \pm 0.003
CLEF-L (296M)	0.7875 \pm 0.015	0.8533 \pm 0.005	0.8543 \pm 0.008	0.8320 \pm 0.004	0.7732 \pm 0.032	0.9442 \pm 0.007	0.9041 \pm 0.004
Lead II							
CLEF-S (448K)	0.8186 \pm 0.004	0.8284 \pm 0.008	0.8515 \pm 0.004	0.8475 \pm 0.003	0.7727 \pm 0.042	0.9510 \pm 0.011	0.9026 \pm 0.003
CLEF-M (30.7M)	0.8115 \pm 0.005	0.8361 \pm 0.004	0.8615 \pm 0.005	0.8482 \pm 0.004	0.7849 \pm 0.025	0.9531 \pm 0.011	0.9098 \pm 0.003
CLEF-L (296M)	0.7850 \pm 0.032	0.8333 \pm 0.009	0.8604 \pm 0.011	0.8489 \pm 0.007	0.7836 \pm 0.020	0.9538 \pm 0.005	0.9094 \pm 0.005

1458 Table S10: Results for finetuning on 12-lead ECGs. AUROCs are reported, with numbers in brackets
 1459 denoting the change relative to the corresponding single-lead baselines using lead I (in left hand
 1460 side of Table 1), where (\uparrow) and (\downarrow) denote a performance gain and degradation, respectively. A grey
 1461 background denotes that the method was outperformed by *CLEF*. LVEF refers to left ventricular
 1462 ejection fraction; Dx denotes diagnosis; SubDx denotes subdiagnosis; SupDx denotes superdiagnosis;
 1463 Form refers to waveform morphology; Rhyth denotes rhythm; Arrhy denotes arrhythmia.

Model	MIMIC-IV		PTB-XL			Chapman	
	LVEF	Dx	SubDx	SupDx	Form	Rhyth.	Arrhy.
ST-MEM	.8095 (\uparrow .08)	.8340 (\uparrow .23)	.8261 (\uparrow .13)	.8459 (\uparrow .11)	.6716 (\uparrow .26)	.8419 (\uparrow .12)	.8462 (\uparrow .03)
KED	.8704 (\uparrow .08)	.9293 (\uparrow .15)	.9079 (\uparrow .09)	.9250 (\uparrow .10)	.8549 (\uparrow .32)	.9683 (\uparrow .11)	.9419 (\uparrow .05)
Moment	.7884 (\downarrow .02)	.8995 (\uparrow .14)	.8727 (\uparrow .12)	.8926 (\uparrow .08)	.7874 (\uparrow .28)	.9528 (\uparrow .09)	.9298 (\uparrow .04)
ECGFounder	.8748 (\uparrow .05)	.9107 (\uparrow .14)	.9106 (\uparrow .10)	.9085 (\uparrow .08)	.8046 (\uparrow .18)	.9321 (\uparrow .01)	.9362 (\uparrow .07)

D SUPPLEMENTARY EVALUATIONS

D.1 DETAILED CONFIDENCE INTERVAL RESULTS

This section supplements the results reported in section 4. To quantify uncertainty for the proposed *CLEF* model, we conducted 6 rounds of experiments with the prediction head initialized with different seeds ‘10, 42, 111, 123, 1111, 1234’, and computed 95% confidence intervals for all reported metrics. Table S9 enumerates the results, where for the majority of the tasks, most confidence intervals remain tight (0.003 – 0.011), apart from Form classification in the PTB-XL dataset, which shows larger variability across seeds. It can also be observed that the selected seed ‘42’ yields relatively conservative results compared to other initializations.

D.2 COMPARISON WITH FINE-TUNING ON 12-LEAD ECGS

To better quantify the performance upper bounds for downstream tasks, we compare single-lead results for *CLEF* against baseline methods that use full 12-lead ECG data. Table S10 summarizes the finetuning results for baseline methods on 12-lead ECGs. Several 12-lead baselines fail to surpass the performance of the single-lead *CLEF* model (shown as gray-highlighted cells). Performance changes relative to the corresponding single-lead baselines are shown in parentheses. All models benefited from multi-lead inputs for all tasks, except for Moment, where there is a small reduction in performance on the LVEF classification task when using 12-lead data. This confirms that additional leads provide valuable complementary information.

D.3 LINEAR PROBING

This section supplements the linear probing experiments in section 4 in the main paper, where we analyzed the performance of *CLEF-M* linear probing against the baselines. Table S11 presents the detailed performance for *CLEF-M* and further results for *CLEF-S* and *CLEF-L*, comparing with the

1498 Table S11: Results for linear probing on lead I ECG representations expressed as AUROCs. LVEF
 1499 refers to left ventricular ejection fraction; Dx denotes diagnosis; SubDx denotes subdiagnosis; SupDx
 1500 denotes superdiagnosis; Form refers to waveform morphology; Rhyth denotes rhythm; Arrhy denotes
 1501 arrhythmia. Best result **bolded**. The second best underlined.

Model	MIMIC-IV		PTB-XL			Chapman		Icentia11K	
	LVEF	Dx	SubDx	SupDx	Form	Rhyth	Arrhy	Beat	Rhyth
Moirai	.5000	.4989	.5000	.5000	.5000	.5000	.4988	.5002	.4927
Moment	.7537	.6392	.6932	.7611	.5863	.6227	.7292	.7298	.6434
ST-MEM	.5709	.5652	.5872	.6189	.5300	.5374	.6587	.6640	.5605
KED	.7607	.5935	.6883	.7514	.5333	.7372	.7712	.6984	.7221
<i>CLEF-S</i>	.7728	.6005	.6509	.7503	.5266	.5872	.6913	.6581	.6761
<i>CLEF-M</i>	.7864	.7256	.7736	.7903	.6146	.8077	.7937	.7302	.7302
<i>CLEF-L</i>	.7939	.7388	.7713	.7971	.6110	.7966	.8010	.7843	.7479
ECGFounder	.8228	<u>.7945</u>	<u>.8423</u>	<u>.8443</u>	<u>.7412</u>	<u>.9422</u>	<u>.8428</u>	<u>.9692</u>	<u>.9721</u>

Table S12: Results when retraining the backbone models with clinically-guided pretraining and finetuning, using lead I ECG. Numbers in brackets denote the change relative to the corresponding single-lead baselines using lead I, with (\uparrow) denoting a performance gain and (\downarrow) a degradation. LVEF refers to left ventricular ejection fraction; Dx denotes diagnosis; SubDx denotes subdiagnosis; SupDx denotes superdiagnosis; Form refers to waveform morphology; Rhyth denotes rhythm; Arrhy denotes arrhythmia. For better clarity, we highlighted cases where performance is better in **teal**.

Model	MIMIC-IV LVEF	Dx	SubDx	PTB-XL SupDx	Form	Rhyth	Chapman Arrhy
ST-MEM	.7788 (\uparrow .00)	.7690 (\downarrow .01)	.7568 (\downarrow .01)	.7878 (\uparrow .01)	.5468 (\downarrow .04)	.8078 (\uparrow .07)	.8046 (\downarrow .01)
KED	.8267 (\downarrow .01)	.8729 (\uparrow .04)	.8569 (\uparrow .03)	.8330 (\uparrow .00)	.7481 (\uparrow .12)	.9067 (\uparrow .02)	.8956 (\uparrow .01)
ECGFounder	.8263 (\downarrow .03)	.8515 (\uparrow .01)	.8649 (\uparrow .02)	.8322 (\downarrow .01)	.7957 (\uparrow .04)	.9516 (\uparrow .00)	.9038 (\downarrow .01)

baseline models. Among the 3 variants, *CLEF*-L achieves the highest linear probing performance with a 7.4% improvement over best baselines on average across 9 downstream classification tasks. The second best being *CLEF*-M (6.0% improvement), while *CLEF*-S exhibits a 7.3% performance deficit. This descending trend indicates that larger models develop more expressive and generalizable representations during pretraining, which can be effectively leveraged through simple linear classification heads.

Notably, comparing results in Table S11 with Table 1, *CLEF*-M generally achieves the best balance of capacity and data efficiency during fine-tuning, where it outperforms *CLEF*-L on many tasks, particularly those with limited labeled data. This potentially reflects a capacity–data mismatch under shared fine-tuning hyperparameters. The 296M-parameter *CLEF*-L requires stronger regularization and/or more labeled data for stable end-to-end fine-tuning. When it is given with relatively small downstream datasets (especially on wearable and ED tasks), it is more prone to overfitting or reaching optimization plateaus when using the same fine-tuning recipe as *CLEF*-M. Nevertheless, the stronger linear-probe performance of *CLEF*-L indicates that its representations are of higher quality.

Additionally, the supervised model ECGFounder outperformed all three variants of *CLEF*, with *CLEF*-M AUROCs an average of 12.7% lower than those of ECGFounder. This is expected, as ECGFounder was pretrained on a much larger labeled dataset, whereas *CLEF* relies purely on self-supervised learning.

D.4 EFFECTIVENESS OF RISK-GUIDED PRETRAINING ON BASELINE MODELS

This section supplements the experiment in section 4 on applying our clinically-guided contrastive loss to the best semi-supervised baseline, KED. We further extend our analysis by applying the proposed pretraining approach to other baseline models, thereby testing its generalizability across architectures. Each baseline model has been initialized with its original pretrained weights and pretrained with our proposed clinically-guided pretraining approach. The AUROC results are presented in Table S12, with values in the brackets indicating the changes relative to the finetuning results of the original model weights, as reported in Table 1. On average, pretraining with our proposed clinically-guided method yields an overall performance gain of 0.7%. Among the three baselines, KED shows the largest improvement, with its AUROC score increasing by 3.0%. This substantial improvement likely stems from KED’s reliance on purely self-supervised learning without explicit clinical guidance,

Table S13: Results on finetuning *CLEF* variants pretrained on the corresponding lead. The parentheses indicate the lead on which the model was pretrained. AUROCs are shown for Lead I (left) and Lead II (right). A **teal** text denotes improved performance compared to Table 1. LVEF refers to left ventricular ejection fraction; Dx denotes diagnosis; SubDx denotes subdiagnosis; SupDx denotes superdiagnosis; Form refers to waveform morphology; Rhyth denotes rhythm; Arrhy denotes arrhythmia.

Task	MIMIC-IV LVEF	PTB-XL				Chapman Arrhy	Task	MIMIC-IV LVEF	PTB-XL				Chapman Arrhy		
		Dx	SubDx	SupDx	Form	Rhyth			Dx	SubDx	SupDx	Form	Rhyth		
<i>CLEF</i> ^I -S	.8215	.8536	.8604	.8324	.7958	.9449	.9057	<i>CLEF</i> ^{II} -S	.8207	.8386	.8620	.8497	.8001	.9552	.9107
<i>CLEF</i> ^I -M	.8170	.8505	.8575	.8336	.7805	.9476	.9128	<i>CLEF</i> ^{II} -M	.8102	.8349	.8719	.8523	.7921	.9552	.9135
<i>CLEF</i> ^I -L	.7821	.8606	.8564	.8327	.7870	.9513	.9133	<i>CLEF</i> ^{II} -L	.7862	.8421	.8680	.8464	.7829	.9257	.9134

Table S14: “Risk score” of CIFAR-100 superclasses and subclasses. Subclasses within the same superclass have similar risk scores, while more distant superclasses have larger dissimilar scores.

Superclass	Subclass	Risk Score
Natural Environment (0.05–0.20)	Large natural outdoor scenes	0.05
	Trees	0.10
	Flowers	0.15
	Fruit and vegetables	0.20
Aquatic Life (0.25–0.30)	Fish	0.25
	Aquatic mammals	0.30
Invertebrates (0.35–0.40)	Insects	0.35
	Non-insect invertebrates	0.40
Land Mammals (0.45–0.70)	Small mammals	0.45
	Medium mammals	0.50
	Reptiles	0.55
	Large omnivores and herbivores	0.60
	Large carnivores	0.65
	People	0.70
Human-Made Objects (0.75–1.00)	Food containers	0.75
	Household furniture	0.80
	Household electrical devices	0.85
	Large man-made outdoor things	0.90
	Vehicles 1	0.95
	Vehicles 2	1.00

making it particularly receptive to our clinically-guided pretraining strategy that provides essential domain-specific knowledge.

D.5 SUPPLEMENTS IN PRETRAINING WITH A SPECIFIC LEAD

This section supplements section 4 on **pretraining on a specific lead**, which explores the upper bound performance by pretraining exclusively on lead I or lead II data, corresponding to the leads used in our downstream evaluation tasks. Although less generalizable, this method reduces potential domain shift compared with our primary pretraining approach using all 12 leads.

Table S13 present the results. The ones colored in red with an uparrow (\uparrow) denote a performance improvement compared to the performance of the corresponding-sized *CLEF* model pretrained on all 12 leads, while those with a downarrow (\downarrow) indicate a performance degradation. Overall, lead-specific pretraining yields performance gains across all tasks, with an average improvement of 2.0%. Specifically, *CLEF^I-S*, *CLEF^I-M*, and *CLEF^I-L* achieve improvements of 3.4%, 1.6%, and 1.0%, respectively, compared to their corresponding models pretrained on all leads across all tasks and both leads. Interestingly, the **Small** variant gives the largest performance gain. This may be because the reduced capacity of the smaller model helps prevent overfitting on the relatively limited data, allowing it to capture the essential patterns more effectively than the larger counterparts.

D.6 EFFECTIVENESS OF DIFFERENT LOSS COMPONENTS

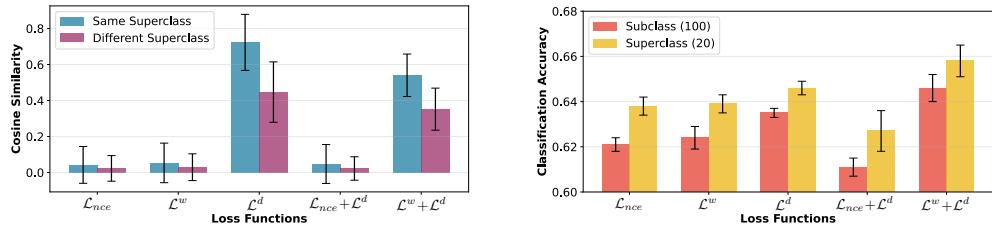
An ablation study is conducted using CIFAR-100 (Krizhevsky et al., 2009) (the representative computer vision benchmark) to understand the contribution of different loss components. This dataset consists of 100 subclasses, each grouped into a higher-level superclass. While instances within the same superclass may belong to different subclasses, they are generally more similar to each other than to instances from different superclasses. To simulate the setup in informing the contrastive learning process with risk scores, we assign each CIFAR-100 superclass a risk score within the range $[0, 1]$. The assigned scores are decided in a way such that subclasses belonging to the same superclass receive relatively close risk scores, while subclasses from more distinct superclasses are assigned increasingly dissimilar scores. The detailed scores given to each subclass are provided in Table S14.

1620
 1621 Table S15: Cosine similarity statistics between class pairs. A spade symbol (\spadesuit) indicates that the
 1622 two classes belong to the same superclass in the CIFAR-100 dataset, whereas a diamond symbol
 1623 (\diamond) represents pairs from different superclasses. Values are reported as mean \pm standard deviation
 1624 of cosine similarity computed over the sampled pairs. We also report the classification accuracy on
 1625 both sub and superclass classification, which has 100 and 20 classes in total, with mean and standard
 1626 deviation reported on 6 different random seed initializations. Best classification results are **bolded**.
 1627

Class Pair	Groups	\mathcal{L}_{nce}	\mathcal{L}^w	\mathcal{L}^d	$\mathcal{L}_{nce} + \mathcal{L}^d$	$\mathcal{L}^w + \mathcal{L}^d$
Palm – Pine Tree	\spadesuit	.058 \pm .125	.081 \pm .139	.699 \pm .214	.067 \pm .132	.556 \pm .153
Palm Tree – Baby	\diamond	.018 \pm .060	.025 \pm .064	.548 \pm .131	.016 \pm .052	.432 \pm .075
Pine Tree – Baby	\diamond	.012 \pm .052	.016 \pm .053	.564 \pm .136	.010 \pm .045	.438 \pm .075
Shark – Trout	\spadesuit	.022 \pm .081	.024 \pm .078	.738 \pm .104	.019 \pm .077	.542 \pm .074
Shark – Bicycle	\diamond	.014 \pm .060	.026 \pm .067	.527 \pm .146	.012 \pm .051	.385 \pm .094
Trout – Bicycle	\diamond	.023 \pm .073	.038 \pm .081	.501 \pm .137	.019 \pm .062	.380 \pm .100
Apple – Pear	\spadesuit	.053 \pm .125	.062 \pm .148	.761 \pm .110	.083 \pm .162	.559 \pm .116
Apple – Bottle	\diamond	.015 \pm .055	.017 \pm .059	.544 \pm .107	.030 \pm .060	.411 \pm .081
Pear – Bottle	\diamond	.036 \pm .099	.038 \pm .094	.589 \pm .121	.035 \pm .083	.452 \pm .089
Bridge – House	\spadesuit	.065 \pm .129	.087 \pm .136	.762 \pm .135	.057 \pm .119	.561 \pm .115
Bridge – Sea	\diamond	.044 \pm .102	.056 \pm .113	.311 \pm .222	.046 \pm .107	.284 \pm .165
House – Sea	\diamond	.018 \pm .064	.021 \pm .067	.307 \pm .217	.016 \pm .062	.262 \pm .154
Cloud – Forest	\spadesuit	.013 \pm .052	.013 \pm .048	.654 \pm .215	.011 \pm .050	.482 \pm .133
Cloud – Bus	\diamond	.015 \pm .060	.018 \pm .055	.220 \pm .196	.014 \pm .053	.201 \pm .153
Forest – Bus	\diamond	.033 \pm .086	.041 \pm .091	.357 \pm .264	.025 \pm .074	.275 \pm .181
<i>ResNet18</i> pretrained						
Subclass (100)		.621 \pm .003	.624 \pm .005	.635 \pm .002	.611 \pm .004	.646 \pm .006
Superclass (20)		.638 \pm .004	.639 \pm .004	.646 \pm .003	.627 \pm .009	.658 \pm .007

1629
 1630 For the experiments, we pretrained the ResNet18 model (initialized with ImageNet-1K pretrained
 1631 weights) using different loss functions, which are the standard contrastive loss \mathcal{L}_{nce} , the weighted
 1632 contrastive loss \mathcal{L}^w , the dissimilarity alignment loss \mathcal{L}^d , $\mathcal{L}_{nce} + \mathcal{L}^d$, and our proposed $\mathcal{L}^w + \mathcal{L}^d$. The
 1633 models were trained for 50 epochs with a batch size of 256, optimized using the AdamW optimizer
 1634 with an initial learning rate of 1×10^{-3} . A Cosine Annealing learning rate schedule was used, and
 1635 we stop the training when the training loss does not decrease over 10 epochs.
 1636

1637
 1638 **Inter-class similarity analysis.** We first analyzed the similarity relationships between groups by
 1639 randomly selecting 5 sets of triplet classes. Each triplet consisted of 2 subclasses from the same
 1640 superclass and one subclass from a different superclass. For each triplet, we obtain the similarity
 1641 scores between all pairs of samples (resulting in 10,000 scores per subclass pair, given 100 samples
 1642 per subclass), and report the mean and standard deviation across groups. Results are summarized on
 1643 the left of Figure S2 and detailed in the upper part of Table S15. It can be observed that training the
 1644 model with only \mathcal{L}_{nce} yields minimal distinction between similarity scores of samples within the same
 1645 superclass (intra-class) compared to those from different superclasses (inter-class). When training
 1646 with \mathcal{L}^w , the inter-class differentiation becomes marginally more pronounced, although it remains
 1647 insufficiently distinct. Meanwhile, training with \mathcal{L}^d produces substantially more distinguishable
 1648 inter-class similarity differences. However, this loss function enforces hard mapping of inter-class
 1649 distances to align with risk scores without constraining intra-class clustering, consequently increasing
 1650 the intra-class similarity variance. This indicates that samples of the same class are scattered rather
 1651 than clustered together. The proposed loss functions $\mathcal{L}^w + \mathcal{L}^d$ and $\mathcal{L}^w + \mathcal{L}^d$ achieve the best
 1652 performance in terms of both intra-class similarity and inter-class differentiation.
 1653



1654
 1655 Figure S2: Ablation study on CIFAR-100 dataset. **Left:** Cosine similarity between samples from the
 1656 same vs. different superclasses. Results averaged across 5 randomly selected triplet groups, with 2 of
 1657 them from the same superclass. **Right:** Classification accuracy on CIFAR-100 subclass (100 classes)
 1658 and superclass (20 classes) tasks.
 1659

Table S16: Effect of weighting over loss. All combinations are normalized by dividing by the sum of their coefficients to maintain comparable loss scales. Experiments were conducted over all classification tasks. For the 12-lead ECG datasets, we use lead I as input. Table cells are colored where darker blue indicates better performance and lighter blue indicates worse performance. Best results are **bolded**.

Loss	MIMIC-IV			PTB-XL			Chapman Arrhy	MUSIC Outcome	Iceni11K		MC-MED		
	Dx	SubDx	SupDx	Form	Rhyth	Beat			Rhyth	ED Dispo	DC Dispo	Acuity	
\mathcal{L}^d	0.8051	0.7990	0.8135	0.8363	0.6451	0.9206	0.8808	0.4767	0.9803	0.9251	0.6016	0.6475	0.6222
$\mathcal{L}^w + 5\mathcal{L}^d$	0.8096	0.8195	0.8536	0.8467	0.6917	0.9581	0.9028	0.4747	0.9780	0.9224	0.6070	0.6341	0.6179
$\mathcal{L}^w + 2\mathcal{L}^d$	0.7965	0.8253	0.8518	0.8396	0.7365	0.9325	0.9088	0.5513	0.9755	0.9289	0.6148	0.6416	0.5861
$\mathcal{L}^w + \mathcal{L}^d$	0.8083	0.8292	0.8566	0.8430	0.7409	0.9361	0.9061	0.5304	0.9792	0.9328	0.6397	0.6607	0.6510
$\mathcal{L}^w + .5\mathcal{L}^d$	0.8204	0.8271	0.8503	0.8473	0.7471	0.9360	0.9130	0.5496	0.9858	0.9350	0.6026	0.6492	0.6221
$\mathcal{L}^w + .2\mathcal{L}^d$	0.8118	0.8158	0.8520	0.8484	0.7325	0.9593	0.9032	0.4747	0.9794	0.9211	0.6082	0.6369	0.6216
\mathcal{L}^w	0.8018	0.7990	0.8188	0.8387	0.6494	0.9215	0.8838	0.4701	0.9786	0.9241	0.6017	0.6446	0.6153

than forming compact clusters, which is a less desirable property for downstream classification. Meanwhile, the combination of \mathcal{L}_{nce} with \mathcal{L}^d offered limited gains compared to \mathcal{L}_{nce} . In contrast, our proposed combination $\mathcal{L}^w + \mathcal{L}^d$ achieves notable differentiation in similarity scores whilst maintaining consistency, which can be inferred from the reduced standard deviation.

Classification accuracy on CIFAR-100. We also provide classification accuracy over the 100 sub-classes and also 20 super-classes, experiment initialized with 6 random seeds, which are ‘10, 42, 111, 123, 1111, 1234’, and we present the average and standard deviation of the classification accuracy. Results are summarized on the right of Figure S2 and enumerated in the lower part of Table S15. For subclass classification (100 classes), our proposed $\mathcal{L}^w + \mathcal{L}^d$ combination achieves the highest accuracy of $64.6\% \pm 0.6\%$, outperforming all baseline approaches by at least (1.7%). Notably, whilst \mathcal{L}^d alone yields the second best performance ($63.5\% \pm 0.2\%$), the combination of $\mathcal{L}_{nce} + \mathcal{L}^d$ exhibits degraded classification accuracy by (3.8%) ($61.1\% \pm 0.4\%$), indicating that simply combining the InfoNCE loss with the difference alignment loss would contrarily harm the model performance, which is consistent with our similarity score observations. For super-class classification (20 classes), the performance remains consistent, with our proposed method achieving $65.8\% \pm 0.7\%$, bringing a substantial improvement by (1.9%) over the best-performance baseline ($64.6\% \pm 0.3\%$) and by 3.0% compared to the naïve contrastive baseline ($63.8\% \pm 0.4\%$).

D.7 ANALYSIS OVER DIFFERENT LOSS WEIGHTING

In Appendix D.6, we used CIFAR-100 to investigate how different combinations of losses influence the cosine similarity between representation pairs and classification accuracy. In this section, we extend this analysis to ECG tasks to examine the effect of having different weighting between the two proposed losses, \mathcal{L}^w and \mathcal{L}^d . Specifically, we evaluated four configurations: (1) \mathcal{L}^w only, (2) \mathcal{L}^d only, (3) our default balanced weight $\mathcal{L}^w + \mathcal{L}^d$, and (4) unbalanced mixes with coefficients $\lambda = \{5, 2, 0.5, 0.2\}$ applied to $\mathcal{L}^w + \lambda\mathcal{L}^d$. All combinations have been normalized by dividing the sum of their corresponding coefficients to maintain comparable loss scales. It can be observed that *CLEF* is not particularly sensitive to the loss weight ratio (with an averaged standard deviation of 0.016 across all tasks), and the two components provide complementary guidance. Using balanced weights for \mathcal{L}^w and \mathcal{L}^d consistently provides strong performance, confirming our initial motivational analysis using CIFAR100.

D.8 ANALYSIS OVER PRETRAINING PARAMETERS

Effect of parameter α : We provide experiments on the effect of α , using the medium variant of the *CLEF* model backbone, and pretrained the model with $\alpha = \{.1, .2, .3, .4, .5\}$. Results are enumerated in Table S17. The model remains robust when α is within the range $0.1 \sim 0.3$, with 0.007 standard deviation averaged across all tasks. Increasing α beyond 0.3 results in a performance degradation, indicating that over-compressing the weighting factor can lead to suboptimal representation learning. Interestingly, if we compare the performance with Table 4, it can be observed that the performance when $\alpha = 0.5$ is similar to SimCLR, implying the weighting factor is over-compressed and becomes less effective. Overall, our results confirm that while α can affect the shape of the embedding space,

1728 Table S17: Sensitivity analysis regarding α . Experiments were conducted over all classification tasks.
 1729 For the 12-lead ECG datasets, we use lead I as input. Table cells are colored where darker blue
 1730 indicates better performance and lighter blue indicates worse performance. Best results are **bolded**.
 1731

α	MIMIC-IV LVEF	PTB-XL				Chapman Arrhy	MUSIC Outcome	Icentia11K		MC-MED				
		Dx	SubDx	SupDx	Form	Rhyth		Beat	Rhyth	ED	Dispo	DC	Dispo	Acuity
0.5	0.8050	0.8084	0.8452	0.8357	0.7111	0.9212	0.8904	0.4737	0.9662	0.9381	0.6053	0.6112	0.5694	
0.4	0.7983	0.8117	0.8550	0.8419	0.7291	0.9213	0.8950	0.4767	0.9550	0.9251	0.6019	0.6475	0.6407	
0.3	0.8030	0.8295	0.8699	0.8515	0.7524	0.9338	0.9023	0.5513	0.9848	0.9475	0.6260	0.6491	0.6548	
0.2	0.8083	0.8292	0.8566	0.8430	0.7409	0.9361	0.9061	0.5304	0.9792	0.9328	0.6397	0.6607	0.6510	
0.1	0.8176	0.8203	0.8596	0.8484	0.7625	0.9360	0.9119	0.5476	0.9733	0.9397	0.6186	0.6593	0.6427	

1737
 1738 it does not directly drive the performance gains. The main improvements in the performance are due
 1739 to using clinically-guided relative risk scores rather than the exact margin offset defined by α .
 1740

1741 **Effect of parameter τ** In our experiment, τ is set to 0.07 following prior work in contrastive
 1742 learning (Wu et al., 2018; He et al., 2020). To further validate the sensitivity of the model to τ we
 1743 conducted further experiments using the small and medium variants of the *CLEF* model backbone,
 1744 with τ set to $\{0.01, 0.1, 0.5\}$. Results are enumerated in Table S18. Performance is best at small
 1745 τ ($\tau = \{0.07, 0.1\}$) and degrades by little as τ increases. This indicates that lower τ enhances the
 1746 discrimination between positives and negatives and is beneficial for learning clinical-informed ECG
 1747 representations.

1748 D.9 COMPUTATIONAL COMPLEXITY ANALYSIS

1749 The proposed clinically-guided pretraining framework introduces minimal computational overhead
 1750 compared to the baseline SimCLR. Specifically, risk scores are computed once per patient from
 1751 clinical metadata during the preprocessing stage, prior to model training. Following the SCORE2
 1752 algorithm detailed in Appendix C.6, each patient’s risk score is derived from 7 clinical covariates
 1753 through a series of operations (Equations S1–S3). For a single patient, this requires $O(1)$ operations,
 1754 as the number of covariates and parameters is fixed.
 1755

1756 Additionally, during training, the clinical guidance adds two operations per mini-batch: (i) obtaining
 1757 weight matrix $\mathbf{W}_{i,k}$ from pairwise risk score dissimilarities $\mathbf{D}_{i,k}$ and missing metadata matrix \mathbf{M} ,
 1758 and (ii) computing the dissimilarity alignment loss L^d (Eq. 7). Both operations require $O(B^2)$
 1759 computations for batch size B . In contrast, the standard SimCLR contrastive loss requires $O(B^2d)$
 1760 operations for d -dimensional embeddings. Since for *CLEF*, $d \in \{256, 1024, 2048\}$, the additional
 1761 $O(B^2)$ cost results in no significant increase in complexity.
 1762

1763 E ALGORITHMS

1764 Algorithm 1 summarizes our clinically-guided contrastive pretraining and downstream evaluation
 1765 procedure. In the pretraining phase, we sample minibatches of ECG embeddings and generate
 1766 stochastic augmentations, which are encoded by the foundation model $\mathcal{F}(\cdot)$. Risk scores from clinical
 1767 metadata are used to compute pairwise dissimilarities, which are combined with multipliers weighted
 1768 by the number of missing metadata to form the final pairwise weights. The foundation model \mathcal{F} is
 1769 updated by minimizing a combination of weighted contrastive loss \mathcal{L}^w and dissimilarity alignment
 1770 loss \mathcal{L}^d . In the evaluation phase, embeddings of downstream labeled ECGs are extracted and used
 1771 to train a linear classifier or regressor $\mathcal{G}(\cdot)$. Model selection is performed by choosing the model
 1772

1773 Table S18: Sensitivity analysis regarding τ . Experiments were conducted over all classification tasks.
 1774 For the 12-lead ECG datasets, we use lead I as input. Table cells are colored where darker blue
 1775 indicates better performance and lighter blue indicates worse performance. Best results are **bolded**.
 1776

τ	MIMIC-IV LVEF	PTB-XL				Chapman Arrhy	MUSIC Outcome	Icentia11K		MC-MED				
		Dx	SubDx	SupDx	Form	Rhyth		Beat	Rhyth	ED	Dispo	DC	Dispo	Acuity
.07	.8083	.8292	.8566	.8430	.7409	.9361	.9061	.5304	.9792	.9328	.6397	.6607	.6510	
.1	.7809	.8405	.8568	.8519	.7988	.9604	.9024	.4737	.9440	.9265	.6033	.6432	.6463	
.5	.7760	.8416	.8440	.8429	.7742	.9220	.9022	.4708	.9345	.9255	.6038	.6421	.6308	

1782 that achieves the lowest validation loss. The final model is evaluated on the held-out test set using
 1783 AUROC (classification) or MAE (regression).
 1784

1785 F DISCUSSION OF CHOICE OF RISK SCORE

1786 In this work, we used the SCORE2 risk score to guide contrastive learning. SCORE2 was selected
 1787 because it is a standard cardiovascular risk score relevant to ECG data, which has been externally
 1788 validated on data from over one million individuals from 15 countries, and its required input variables
 1789 overlap with the metadata in the MIMIC-IV-ECG pretraining dataset. However, we do not claim
 1790 that SCORE2 is the optimal risk score for *CLEF* as we did not test other risk scores. Furthermore,
 1791 the optimal risk score might vary across applications. For instance, SCORE2 was developed using
 1792 data from European subjects, for adults aged 40-69, and for long-term cardiovascular risk prediction.
 1793 Alternative risk scores may be more appropriate for non-European populations, for different age
 1794 groups, and for non-cardiovascular risk assessment. In addition, the optimal risk score will likely
 1795 vary depending on the available metadata in a pretraining dataset.
 1796

1797 LLM USAGE STATEMENT

1798 LLMs are used to aid in the writing of this paper. We use LLMs to improve the phrasing clarity
 1799 to better convey our contributions and research findings to the readers. Specifically, this involves
 1800

1801 Algorithm 1: Clinically-Guided Contrastive Pretraining and Evaluation

1802 Pretraining Phase:

1803 **Input:** Unlabeled ECG dataset $\mathbb{D} = \{(\mathbf{e}_s = (\mathbf{e}_s^1, \dots, \mathbf{e}_s^{12}), \mathbf{a}_s)\}_{s=1}^N$, risk score function $\mathcal{R}(\cdot)$,

1804 **Output:** Foundation ECG model $\mathcal{F}(\cdot)$

1805 **for** each minibatch $\{\mathbf{e}_s^l\}_{s=1}^B$ **do**

1806 **for** each \mathbf{e}_s^l **do**

1807 Sample stochastic augmentations $\mathbf{x}_i = \mathcal{T}(\mathbf{e}_s^l)$, $\mathbf{x}_j = \mathcal{T}(\mathbf{e}_s^l)$;

1808 Compute embeddings $\mathbf{z}_i = \mathcal{F}(\mathbf{x}_i)$, $\mathbf{z}_j = \mathcal{F}(\mathbf{x}_j)$;

1809 Obtain risk score $r_s = \mathcal{R}(\mathbf{a}_s)$ from clinical metadata (Eq. (3)) ;

1810 **for** each pair $(\mathbf{x}_i, \mathbf{x}_k)$ in batch **do**

1811 Compute risk dissimilarity $\delta_{ik} = (r^i - r^k)^2$;

1812 Normalize to $\mathbf{D}_{ik} \in [\alpha, 1]$ (Eq. (5)) ;

1813 Compute metadata missingness weight \mathbf{M}_{ik} (Eq. (4)) ;

1814 Set final weight $\mathbf{W}_{ik} = \mathbf{D}_{ik} \cdot \mathbf{M}_{ik}$;

1815 Compute weighted contrastive loss \mathcal{L}^w (Eq. (6)) ;

1816 Compute dissimilarity alignment loss \mathcal{L}^d (Eq. (7)) ;

1817 Update encoder $\mathcal{F}(\cdot)$ by minimizing $\mathcal{L} = \mathcal{L}^w + \mathcal{L}^d$;

1818 Evaluation Phase:

1819 **Input:** downstream labeled dataset $\mathbb{C} = \{(\mathbf{z}_s = \mathcal{F}(\mathbf{e}_s^l), y_s)\}_{s=1}^M$

1820 **Output:** downstream linear model $\mathcal{G}(\cdot)$

1821 **for** each labeled ECG (\mathbf{e}_s, y_s) in \mathbb{C}_{train} **do**

1822 Extract embedding $\mathbf{z}_s = \mathcal{F}(\mathbf{e}_s)$;

1823 Train/update a linear classifier/regressor $\hat{y}_s = \mathcal{G}(\mathbf{z}_s)$;

1824 **for** each labeled ECG (\mathbf{e}_s, y_s) in \mathbb{C}_{val} **do**

1825 Extract embedding $\mathbf{z}_s = \mathcal{F}(\mathbf{e}_s)$;

1826 Predict discrete/continuous values $\hat{y}_s = \mathcal{G}(\mathbf{z}_s)$;

1827 Select \mathcal{G} with best validation loss ;

1828 **for** each labeled ECG (\mathbf{e}_s, y_s) in \mathbb{C}_{test} **do**

1829 Extract embedding $\mathbf{z}_s = \mathcal{F}(\mathbf{e}_s)$;

1830 Predict discrete/continuous values $\hat{y}_s = \mathcal{G}(\mathbf{z}_s)$;

1831 Evaluate $\{\hat{y}_s\}$ against $\{y_s\}$ using AUROC (classification) or MAE (regression) ;

1836 grammar correction, wording polishing, and sentence structure refinement. The extent of LLM
1837 assistance was minimal relative to the overall manuscript content. We note that LLMs are not
1838 involved in any aspect of data analysis, the development of methods, the design of the experiment, or
1839 results interpretation. The usage of LLMs is strictly limited to linguistic refinement. All scientific
1840 content originates entirely from the author's research work. The authors take full responsibility for
1841 all content presented in the paper and have verified that the usage of LLMs has not resulted in any
1842 form of plagiarism. All claims, findings, and contributions presented are the authors' own work and
1843 have been thoroughly validated for originality and accuracy.

1844

1845

1846

1847

1848

1849

1850

1851

1852

1853

1854

1855

1856

1857

1858

1859

1860

1861

1862

1863

1864

1865

1866

1867

1868

1869

1870

1871

1872

1873

1874

1875

1876

1877

1878

1879

1880

1881

1882

1883

1884

1885

1886

1887

1888

1889