BENCHMARKING ECG FOUNDATIONAL MODELS: A REALITY CHECK ACROSS CLINICAL TASKS

Anonymous authors

Paper under double-blind review

ABSTRACT

The 12-lead electrocardiogram (ECG) is a long-standing diagnostic tool. Yet machine learning for ECG interpretation remains fragmented, often limited to narrow tasks or datasets. Foundation models promise broader adaptability, but their generalization across diverse ECG tasks is not well understood. We benchmarked eight ECG foundation models on 26 clinically relevant tasks using 12 public datasets comprising 1,650 regression and classification targets. Models were evaluated under fine-tuning and frozen settings, with scaling analyses across dataset sizes. Results show heterogeneous performance across domains: in the most widely studied domain, adult ECG interpretation, three foundation models consistently outperformed strong supervised baselines. In contrast, ECG-CPC, a compact structured state-space model pretrained on HEEDB, dominated other categories where most foundation models failed to surpass supervised learning. Foundation models also displayed distinct scaling behaviors with dataset size, which are critical for smallscale clinical applications. Overall, while foundation models show promise for adult ECG analysis, substantial gaps remain in cardiac structure, outcome prediction, and patient characterization. Notably, ECG-CPC's strong performance despite being orders of magnitude smaller and consuming minimal computational resources highlights untapped opportunities for advancing ECG foundation models.

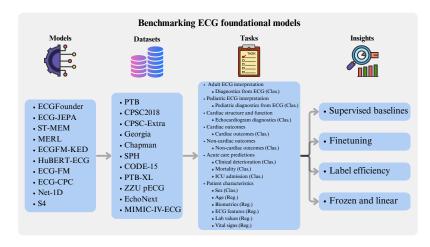


Figure 1: Overview of the benchmarking pipeline for ECG foundation models.

1 Introduction

Clinical relevance Electrocardiography (ECG) is a widely used non-invasive tool for assessing cardiac function and systemic physiology (Siontis et al., 2021). Accurate interpretation is essential for detecting myocardial infarctions (Strodthoff et al., 2020), evaluating cardiovascular risk (Bhatia & Dorian, 2018), and guiding clinical decisions (Rokos et al., 2010). ECG features also reflect systemic conditions such as electrolyte imbalances (Diercks et al., 2004), metabolic disorders (Wald, 2006), and physiological factors (Lopez Alcaraz et al., 2025), broadening its clinical utility. The growing availability of large-scale ECG datasets (Wagner et al., 2022; Gow et al., 2023; Ribeiro et al., 2021) has made machine learning increasingly pivotal for automated interpretation, enabling population-scale screening (Kalmady et al., 2024) and supporting clinical workflows (Graf et al., 2024).

Promise of foundational models for ECG Foundational models (FMs) have transformed fields such as natural language processing (Myers et al., 2024) and computer vision (Awais et al., 2025), where large-scale pretraining produces robust and transferable representations across tasks and domains (Subramanian et al., 2023). This approach is now emerging for biomedical time series, with ECG as a natural testbed, given its ubiquity and clinical relevance. Clinical FMs offer three key advantages: (i) higher predictive performance than training from scratch, (ii) improved label efficiency via pretrained representations, and (iii) the utility of FMs as frozen feature extractors for downstream tasks. These benefits could enable stronger predictive models, support rare-disease studies on limited data, and guide the choice of foundation models for specific applications.

Research gap Despite their promise, foundational models for ECG are under-evaluated. Prior studies often only compare FMs against weak baselines (Li et al., 2025; Kim, 2024; Na et al., 2024; Liu et al., 2024), and benchmarks are limited in scope, focusing on narrow datasets or tasks and hindering generalizable conclusions (Na et al., 2024; Liu et al., 2024). Key aspects such as label efficiency, cross-domain representation quality, and trade-offs between fine-tuning and frozen evaluation remain poorly understood (Li et al., 2025; Kim, 2024; Na et al., 2024; Liu et al., 2024). A systematic, like-for-like comparison of available FMs, leveraging a unified codebase, is still lacking.

Contributions of this work We present a comprehensive benchmark of ECG FMs, covering classification and regression across heterogeneous clinical open-source datasets. Our contributions are threefold: (1) FM comparison: We evaluate FMs in different evaluation modes against strong supervised baselines such as structured state-space (S4) models. While FMs often surpass weak baselines, only the strongest match S4 performance. FMs reveal heterogeneous, highly task-specific results. Certain FMs can effectively serve as frozen feature extractors for certain tasks with competitive performance. (2) Scaling: We investigate the prospects of finetuning FMs as a possible solution to the challenges of small sample sizes. Strong FMs can indeed increase label efficiency by factors of 2.5-9 compared to a strong supervised baseline. (3) FM training: We propose ECG-CPC, a lightweight foundation model using structured state-space models as the architecture and trained with minimal computational resources, which turns out as one of the strongest models in the benchmark.

2 BACKGROUND

AI-enhanced ECG Electrocardiography (ECG) is a widely used, non-invasive tool for assessing cardiac function and systemic health (Siontis et al., 2021). Machine learning and deep learning now underpin automated ECG interpretation, improving arrhythmia detection, risk stratification, and clinical decision support. Models of different architectural flavors, from CNNs to RNNs to transformers, trained through supervised learning, have been shown to achieve strong performance. However, they require large, curated datasets and often generalize poorly across populations, devices, and settings (Ribeiro et al., 2020; Hannun et al., 2019; Hong et al., 2020). Benchmarking robust supervised baselines across multiple datasets remains crucial for

assessing the benefits of representation learning and pretraining (Strodthoff et al., 2020; Nonaka & Seita, 2021; Hong et al., 2020).

ECG foundation models Inspired by the success of foundation models in language and vision, several approaches have been proposed for ECG. Architectures include CNNs (mainly ResNet variants), transformers (often with convolutional encoders), and structured state-space models. Pretraining methods vary from supervised and weakly supervised to contrastive and non-contrastive self-supervision. Pretraining datasets also differ widely, from Computing in Cardiology 2021 challenge subsets (Reyna et al., 2021b) to MIMIC-IV ECG (Gow et al., 2023) to the large-scale HEEDB dataset (Koscova et al., 2024).

Table 1: Summary of the eight ECG foundation models and two supervised baselines evaluated in this work, including their backbone architectures, training paradigms, pretraining datasets and sizes, and parameter sizes. Pretraining datasets: a: HEEDB; b: Chapman, Ningbo, CODE-15%; c: MIMIC-IV-ECG; d: CODE, PTB, CPSC2018, CPSC-Extra, PTB-XL, Georgia, Chapman, Ningbo, Hefei, SPH, MIMIC-IV-ECG; e: CPSC2018, CPSC-Extra, PTB-XL, Georgia, Ningbo, Chapman, MIMIC-IV-ECG

Name	Backbone	Pretraining	Pretraining Samples	Parameters
ECGFounder	CNN	Supervised	$10.7\mathrm{M}^a$	33.8M
ECG-JEPA	Transformer	JEPA	$174k^b$	87.2M
ST-MEM	Transformer	MAE	$174k^b$	90.3M
MERL	Transformer	Weak sup., contrastive	$800\mathrm{k}^c$	4.6M
ECGFM-KED	CNN	Weak sup., contrastive	$800k^c$	9.7M
HuBERT-ECG	Transformer	MLM	$9.1 \mathrm{M}^d$	97.2M
ECG-FM	Transformer	MLM+contrastive	$2.5\mathrm{M}^e$	93.9M
ECG-CPC	SSM	CPC	$10.7 \mathrm{M}^a$	3.8M
Net-1D	CNN	-	-	33.8M
S4	SSM	-	-	2.2M

Benchmarking foundation models The benchmarking of foundation models has a long tradition in other fields such as computer vision (Goldblum et al., 2023) and NLP (Chang et al., 2024). In the medical domain, most efforts have so far focused on medical imaging (Neidlinger et al., 2024; Lee et al., 2025). Recently, benchmarking results for EEG foundation models have been put forward (Xiong et al., 2025). However, no large-scale, comprehensive benchmark for ECG foundation models exists.

3 METHODS

3.1 Models

Table 1 summarizes the investigated models, including backbones, pretraining methods, and datasets, and parameter counts. As custom pretraining is infeasible, only foundation models with publicly available pretrained weights are included and integrated via wrapper modules into a common evaluation framework. Unfortunately, the lack of public weights excludes many proposed ECG FMs. We benchmark eight foundational models and two supervised baseline models as described below.

Foundation models We consider diverse architectures and pretraining strategies. ECGFounder (Li et al., 2025) uses a RegNet-inspired CNN pretrained on HEEDB with a supervised loss. ECG-JEPA (Kim, 2024) employs a transformer-based joint embedding predictive architecture (JEPA) (Assran et al., 2023). We use the multi-block variant. ST-MEM (Na et al., 2024) uses a ViT-1D transformer trained as a masked autoencoder (MAE) (He et al., 2022). MERL (Liu et al., 2024) applies contrastive text-signal alignment (Radford

et al., 2021), with the ResNet18 variant evaluated. ECGFM-KED (Tian et al., 2024) uses a ResNet backbone and minimizes a contrastive loss between signals and ECG report text. HuBERT-ECG and ECG-FM (McKeen et al., 2024) are transformer-based, pretrained via masked language modeling, with ECG-FM additionally incorporating a sequence-level contrastive loss. Finally, we pretrain a structured state-space (SSM) model on HEEDB using contrastive predictive coding (CPC) (van den Oord et al., 2018; Mehari & Strodthoff, 2023) (see Supplementary Material).

Supervised baselines We evaluate two backbones trained from scratch: Net-1D, from ECGFounder (Li et al., 2025), and S4 (Mehari & Strodthoff, 2023), a structured state-space model known for its ability to capture long-range dependencies (Gu et al., 2022). S4 has proven to be a strong baseline on PTB-XL and other datasets (Mehari & Strodthoff, 2023; Strodthoff et al., 2024).

3.2 Datasets and benchmarking tasks

We group the benchmark datasets into seven categories based on the qualitative nature of the underlying labels. See Table 3 in the Supplementary Material for details on sample sizes and task characteristics. Adult/pediatric ECG interpretation: These two categories, the most studied in the literature, involve predicting diagnostic ECG statements from cardiologists. Benchmarks include PTB (Bousseljot et al., 1995), Ningbo (Reyna et al., 2022), CPSC2018 and CPSC-Extra (Reyna et al., 2021a; 2022), Georgia (Reyna et al., 2021a; 2022), Chapman (Reyna et al., 2021a; 2022), SPH (Liu et al., 2022a;b), CODE-15% (Ribeiro et al., 2021), and PTB-XL (Wagner et al., 2022; 2020). A separate category includes ZZU pECG (Tan et al., 2025; Jian et al., 2025) for pediatric ECG interpretation. Cardiac structure and function: This category predicts outcomes from complementary modalities, here echocardiography. We use the EchoNext dataset (Elias & Finer, 2025), which is used to capture cardiac structure and function. Cardiac and non-cardiac outcomes: These categories cover the prediction of cardiac and non-cardiac discharge diagnoses from the first emergency department ECG (Strodthoff et al., 2024), distinguishing cardiac (ICD-10 chapter I) and non-cardiac diagnoses. Acute care predictions: ECG is a key modality for acute care decisions. We predict clinical deterioration, mortality at multiple time frames, and ICU admission (Alcaraz et al., 2025), training a joint model for cardiac, non-cardiac, and acute care outcomes to save computational resources. Patient characteristics: This category includes tasks where ECGs predict non-diagnostic patient characteristics, such as sex, age, biometrics, ECG features, laboratory values (Alcaraz & Strodthoff, 2024), and vital signs (Gow et al., 2023). The datasets span small specialized cohorts to large population studies, covering diverse classification and regression tasks. We train a single model for all tasks in this category using combined classification and regression losses.

3.3 METHODOLOGY

Finetuning/linear evaluation Pretrained models are supplemented with a linear classification head matching the downstream task outputs. We optimize binary cross-entropy or mean absolute error (MAE) using AdamW with a constant learning rate of 1e-3 and weight decay of 1e-3, performing model selection on the validation set via AUROC or MAE. In order to avoid negative interference for multiple regression targets at different scales, we use z-normalized targets (based on training set statistics). During finetuning, layer-dependent learning rates are applied: model architectures are divided into two parts, corresponding learning rates are scaled down by factors of 100 and 10 relative to that of the prediction head. Batch-norm statistics are frozen for linear evaluation and frozen evaluation. For frozen evaluation, we replace the linear prediction head with a learnable query attention head (Bardes et al., 2024) operating on the sequence of output tokens before pooling. If the model allows to adjust the input size, random 2.5s input crops are used for training (Mehari & Strodthoff, 2023), and predictions are averaged across non-overlapping sliding windows during test time.

Evaluation For classification, we report macro-averaged AUROC, as a measure of overall discriminative performance. For regression, we report the average MAE across z-normalized predictions and targets. Individual label AUROC and MAE values are provided in the Supplementary Material. Statistical uncertainty is assessed via empirical bootstrapping on the test set (n=1,000 iterations). Pairwise comparisons are performed by bootstrapping performance differences, if the 95% confidence interval excludes zero, the difference is considered significant. Rankings are assigned based on these comparisons, with ties indicating no significant difference, capturing both relative performance, and statistical significance across task categories.

4 RESULTS

Results overview Table 2 shows an excerpt of finetuning results across tasks. Full finetuning, frozen evaluation, and linear evaluation results are given in Tables 4, 5, and 6 (Supplementary Material). Tables indicate statistically significant differences compared to the respective best method in this task. Ranked lists for all tasks and models are provided in Table 7. We further summarize this by reporting median ranks across all tasks of a given category in Figure 3. Appendix A.4 presents comparative predictions for the 10 best-performing conditions per task, illustrating both prediction quality and label diversity.

4.1 SUPERVISED BASELINES

Supervised baseline vs. literature Our supervised baseline performs on par with or surpasses the literature results. The S4-based model achieves AUROCs of 0.941 (vs. 0.9417 (Mehari & Strodthoff, 2023)) on PTB-XL, 0.908 (vs. 0.843 (Strodthoff et al., 2024)) for cardiac discharge, 0.849 (vs. 0.764 (Strodthoff et al., 2024)) for non-cardiac discharge, 0.863 (vs. 0.752 (Alcaraz et al., 2025)) for clinical deterioration, 0.747 (vs. 0.746 (Alcaraz et al., 2025)) for ICU admission, and 0.874 (vs. 0.816 (Alcaraz et al., 2025)) for mortality. Improvements on MIMIC-based tasks (except for ICU) are attributed to the multi-task training objective.

Impact of model architecture We also compare the supervised S4 model to the convolutional Net1D backbone used in ECGFounder, omitting supervised transformers, due to poor performance. Table 4 summarizes numerical results, and Table 7 shows statistically significant rankings: S4 consistently ranks first, while Net1D typically ranks at least four places lower. Although the backbone comparison was not the main goal, these results support the literature findings that CNNs are suboptimal for physiological time series (Strodthoff et al., 2024).

4.1.1 FINETUNING

Adult ECG interpretation: Across 11 tasks on 9 datasets, top performers are ECGFounder, ECG-JEPA, and ECG-CPC, often statistically surpassing the S4 baseline. ECG-FM ranks fourth overall, sometimes matching the top three but underperforming on Georgia, Chapman, and PTB-XL. MERL, ST-MEM, HuBERT-ECG, and ECGFM-KED generally fail to outperform the supervised baseline. Pediatric ECG interpretation: ECG-JEPA leads, followed by ECGFounder, ST-MEM, MERL, ECG-CPC, and S4, despite no pediatric pretraining data. Cardiac structure & function: ECG-CPC ranks first for echocardiography predictions, followed by ECGFounder, ECG-JEPA, ST-MEM, MERL, and S4. Cardiac and non-cardiac outcomes: ECG-CPC dominates, matching ECG-FM on non-cardiac and S4 on cardiac conditions. ECG-Founder performs relatively poorly, likely due to little overlap between diagnostic labels during pretraining and target labels for the task. Acute care predictions: ECG-CPC and ECG-FM perform best across three tasks, followed by ECGFounder, ECG-JEPA, and MERL; none of them outperform S4 significantly. Patient characteristics: ECG-CPC ranks first in 5 of 6 tasks, outperforming S4 in 3. MERL and ECG-FM generally match or slightly underperform S4, while ECGFounder and ECG-JEPA typically fall below the supervised baseline.

Table 2: Comparison of macro-AUROC (classification) and average z-normalized MAE (regression) under finetuning with linear prediction head. We use \uparrow/\downarrow to indicate whether higher or lower scores correspond to better model performance. Best results are bold and underlined; results not statistically worse are also bold. † indicates label subsets only used for evaluation. Full results involving all considered models are given in Table 4. See Table 5 and Table 6 for frozen and linear evaluation. Model abbreviations: ECGFounder (Founder), ECG-JEPA (JEPA), ECG-FM (FM), ECG-CPC (CPC).

	I	ounda	tion Model	s (Finetu	ned)		Supervise					
	Founder	JEPA	ST-MEM	MERL	FM	CPC	S4					
	Adu	lt ECG	interpretat	tion								
PTB ↑	0.656	0.679	0.694	0.717	0.725	0.702	0.654					
Ningbo ↑	0.974	0.973	0.954	0.955	0.971	0.973	0.972					
CPSC2018↑	0.966	0.974	0.946	0.936	0.972	0.969	0.962					
CPSC-Extra ↑	0.906	0.897	0.883	0.873	0.862	0.898	0.852					
Georgia ↑	0.920	0.918	0.888	0.912	0.912	0.913	0.903					
Chapman ↑	0.968	0.972	0.948	0.946	0.956	0.962	0.963					
-Chapman (rhythm) [†] ↑	0.991	0.989	0.985	0.975	0.993	0.987	0.986					
SPH ↑	0.983	0.980	0.964	0.944	0.966	0.981	0.981					
CODE-15% ↑	$\overline{0.987}$	0.991	0.974	0.982	0.986	0.989	0.991					
PTB-XL (all) ↑	0.934	0.940	0.908	0.925	0.927	0.949	0.941					
-PTB-XL (diag) [†] ↑	0.950	0.946	0.904	0.942	0.926	0.951	0.943					
-PTB-XL (form) [†] ↑	0.875	0.912	0.887	0.891		0.934	0.919					
-PTB-XL (rhythm) [†] ↑	0.965	0.956	0.951	0.912		0.959	0.956					
PTB-XL (sub) ↑	$\frac{0.943}{0.943}$	0.935	0.916	0.937		0.940	0.938					
PTB-XL (super) ↑	$\frac{0.945}{0.935}$	0.921	0.901	0.930			0.932					
PTB-XL (super) ↑ 0.935 0.921 0.901 0.930 0.916 0.934 0.932 Pediatric ECG interpretation												
ZZU pECG↑	0.898	0.911	0.893	0.886	0.887	0.892	0.897					
	Cardia	ac struc	ture & fun	ction								
EchoNext (Echo) ↑	0.817	0.817	0.816	0.822	0.772	<u>0.831</u>	0.819					
	(outcomes									
MIMIC (Cardiac) ↑	0.768	0.772	0.760	0.776	0.775	<u>0.781</u>	0.780					
	No	n-cardi	ac outcom	es								
MIMIC (Non-cardiac) ↑	0.701	0.711	0.688	0.712	<u>0.719</u>	0.719	0.714					
			e predictio									
MIMIC (Deterioration) ↑	0.717	0.747	0.714	0.743		0.764	0.756					
MIMIC (Mortality) ↑	0.810	0.792	0.784	0.800		0.803	0.793					
MIMIC (ICU) ↑	0.748	0.742	0.737	0.744	0.750	<u>0.753</u>	0.745					
	Pa	tient ch	aracteristi	cs								
$MIMIC(Sex)\!\uparrow$	0.913	0.904	0.883	0.916		<u>0.933</u>	0.919					
MIMIC (Age) ↓	0.461	0.463	0.504	0.449	0.412	0.437	0.455					
MIMIC (Biometrics) ↓	0.637	0.640	0.673	0.625	0.620	0.604	0.626					
MIMIC (ECG Features) \	0.458	0.460	0.500	0.463		0.451	0.452					
MIMIC (Lab Values) \(\)	0.679	0.677	0.688	0.676		$\overline{0.673}$	0.675					
MIMIC (Vital Signs) \(\)	0.704	0.703	0.715	0.702	0.704	$\overline{0.700}$	0.701					

4.2 Label efficiency

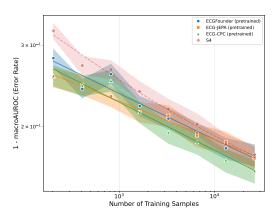


Figure 2: Scaling with dataset size on EchoNext (6 most-populated labels).

Setup To isolate the effect of dataset size from task difficulty, we perform a controlled scaling experiment for the cardiac structure & function prediction on EchoNext (incorporating only labels with >10,000 counts). Training and validation subsets are subsampled in powers of 2 down to 1/128, using multi-label stratification (Wagner et al., 2022) on diagnostic labels, age bins, and sex. We finetune ECGFounder, ECG-JEPA, and ECG-CPC on these subsets, comparing pretrained models to training from scratch, with the S4 supervised baseline as reference. Performance is plotted as $1-\max$ AUROC, and scaling curves are fitted as $CN^{-\alpha}+L_0$, where C is a constant, N is the size of the training set, α the scaling exponent, and L_0 the residual error.

Label efficiency Scaling curves for three FMs are shown in Figure 2 with scaling parameters listed in Table 31. All models stay below the supervised baseline (S4) throughout the entire range of considered training set sizes. We use the parametric form of the fits to work out a label efficiency ratio $r = N^*/N$, i.e., the fraction of samples N^* required for the pretrained model to reach the same performance as the supervised baseline forgivenN. For N in the range of 250 to 1000, this yields label efficiency ratios between 0.30-0.62 for ECGFounder, 0.11-0.42 for ECG-JEPA and 0.21-0.40 for ECG-CPC, see Table 32 for details. This positions ECG-JEPA as the most label-efficient model, in particular in the very low sample size regime, closely followed by ECG-CPC. This establishes the label efficiency as a relevant benchmark parameter for foundation models. Furthermore, the results show that ECG FMs fulfill the promise of an improved label efficiency not only in comparison to the respective model architecture trained from scratch but also against a strong supervised baseline, improving label efficiency by up to a factor of 9.

4.3 Frozen and linear evaluation

Frozen vs. supervised Model rankings under frozen evaluation largely mirror finetuning results, with some differences. ECGFounder and ECG-JEPA continue to perform strongly on adult ECG interpretation, matching the supervised baseline, while ECG-CPC ranks slightly lower, with a median rank of 3. ECG-JEPA still dominates pediatric ECG interpretation. In other categories, ECG-CPC continues to lead. Notably, ECGFounder and ECG-JEPA can serve as effective frozen feature extractors for adult ECG interpretation tasks, achieving supervised-level performance. In the remaining categories, this applies to selected tasks for ECG-CPC under frozen evaluation and for ST-MEM under linear evaluation.

Finetuning vs. frozen/linear Model rankings under finetuning largely mirror frozen/linear evaluation. Strong models like ECG-JEPA, ECGFounder, and especially ECG-CPC maintain strong performance across other evaluation modes. However, some models such as MERL, and ECG-FM rely more on finetuning to reach competitive rankings, Interestingly, ST-MEM shows a much better relative ranking under linear evaluation than finetuning. This highlights that finetuning and linear/frozen evaluation relate to largely similar, though not completely congruent, aspects of representational quality and should therefore both be considered.

Frozen vs. linear Most methods are only slightly affected when using a linear instead of a non-linear prediction head. Deviations from this pattern are ST-MEM, which ranks among the best FMs under linear evaluation unlike in frozen evaluation, and ECG-CPC, which consistently underperforms with a linear head. The latter effect likely relates to pretraining: supervised or global contrastive objectives encourage discriminative pooled representations, unlike purely token-level pretraining. These results can be seen as an incentive to combine token- and sequence-level objectives, as done in ECG-FM and also commonly observed in computer vision (Caron et al., 2021). However, we advocate frozen evaluation as a less biased measure of FM representational quality than finetuning.

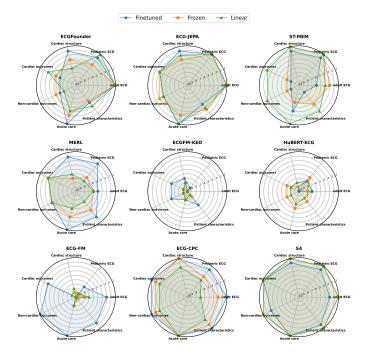


Figure 3: Radar plots summarizing model performance ranks (lower rank indicates statistically significantly better performance) for the eight FMs and one of the supervised baselines (S4). The plot is based on data from Table 8.

DISCUSSION

Details matter Reliable downstream performance depends on careful tuning. Layer-dependent learning rates consistently improve results, and some models (e.g., HuBERT-ECG) fail to train at all without them. Adjustable input size (where possible) also matters: using 2.5s crops with test-time averaging outperforms the full 10s inputs.

Comparable assessment ECGFounder, ECG-JEPA, and ECG-CPC are strong on adult ECG interpretation, sometimes surpassing the supervised baseline. Pediatric ECG is dominated by ECG-JEPA. Other categories are led by ECG-CPC, followed by ECG-FM, while ECGFounder and ECG-JEPA are weaker here. Several models fall clearly below the supervised baseline, showing that self-supervised pretraining does not necessarily yield effective downstream performance. No single model consistently excels across all tasks and evaluation modes, though ECG-CPC comes closest to this goal with finetuning. This underscores that foun-

dation model selection for downstream tasks should be informed by benchmarking results and similarity of the downstream task to the task categories considered here.

Pretraining strategies This benchmark cannot definitively identify the optimal pretraining method, as downstream performance depends on the interplay of the pretraining methodology, dataset, and model architecture (Table 1). On the one hand, superior supervised performance may indicate more suitable models, which would favor SSM models over CNNs. On the other hand, transformers are naturally compatible with masking strategies that are commonly used in self-supervised methods. Larger datasets help in some cases (ECGFounder, ECG-CPC) but are not sufficient (e.g., HuBERT-ECG). Large-scale models do not automatically generalize well. ECG-CPC's dominance outside ECG interpretation suggests self-supervised pretraining can be advantageous, though, its unidirectional backbone limits supervised performance compared to bidirectional models. These findings highlight the need for like-for-like comparisons of pretraining methods on a common dataset, such as HEEDB, using standardized architectures.

Model complexity Foundation models vary widely in complexity. This aspect is particularly relevant for clinical deployment or multimodal integration. Using the parameter count as an imperfect proxy (Dehghani et al., 2021), sizes range from 3.8M (ECG-CPC) to 97.2M (HuBERT-ECG). Notably, ECG-CPC matches ECGFounder and ECG-JEPA on adult ECG tasks with only 11% and 4% of their parameters respectively, and outperforms other FMs in most remaining categories.

Model insights Frozen and linear evaluation reveal the nature of learned representations and the implicit knowledge captured by different foundation models. Insights from probing, as used here, remain coarse, but methods from explainable AI such as analyzing representation structures and their alignment across layers and models (Vielhaben et al., 2025) could provide more detailed insights into the knowledge acquired by these models.

Limitations This work is subject to a number of limitations. First, the proposed tasks only includes indistribution tests, out-of-distribution tests would require further datasets with comparable label sets. While this is often challenging for diagnostic labels due to mismatching ontologies, it is feasible for demographic data. Second, multi-tasking models seem to improve performance for classification tasks. In the regression case, we tried to alleviate this issue through z-normalization of the targets. Still the multi-tasking model used for computational efficiency might underperform for certain task compared to single-task models. Third, the consistent positive impact of layer-dependent learning rates during finetuning suggests the potential for developing improved finetuning methodologies.

6 CONCLUSION

In this work, we present a comprehensive benchmark for ECG foundation models across seven different task categories. Results showed heterogeneous performance across domains: three FMs (ECGFounder, ECG-JEPA and ECG-CPC) were found to exhibit strong performance in the most widely considered domain of adult ECG interpretation, whereas ECG-CPC excels in most other categories where many of the other FMs fail to reach the performance of the supervised baseline. Overall, selected ECG FMs seem promising, outperforming strong supervised baselines during finetuning and/or serving as a frozen feature extractor, achieving performance on par with supervised models during frozen evaluation. ECG-CPC, proposed in this work, is a small-scale model based on a structured state-space model backbone. It was trained on limited computational resources (a single NVIDIA L40 GPU trained for three weeks). The fact that it turned out among the strongest FMs in the benchmark suggests significant opportunities for further improving ECG foundation models. Code and model weights are provided in the supplementary material and will be made publicly available. LLMs were exclusively used for language improvement. Our framework source code and model weights for the ECG-CPC model are publicly available at https://anonymous.4open.science/r/ecq-fm-benchmarking-D5E5.

REFERENCES

- Juan Miguel Lopez Alcaraz and Nils Strodthoff. Cardiolab: Laboratory values estimation and monitoring from electrocardiogram signals a multimodal deep learning approach. *arXiv* preprint 2411.14886, 2024. URL https://arxiv.org/abs/2411.14886.
- Juan Miguel Lopez Alcaraz, Hjalmar Bouma, and Nils Strodthoff. Enhancing clinical decision support with physiological waveforms a multimodal benchmark in emergency care. *Computers in Biology and Medicine*, 192:110196, June 2025. ISSN 0010-4825. doi: 10.1016/j.compbiomed.2025.110196. URL http://dx.doi.org/10.1016/j.compbiomed.2025.110196.
- Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- Muhammad Awais, Muzammal Naseer, Salman Khan, Rao Muhammad Anwer, Hisham Cholakkal, Mubarak Shah, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Foundation models defining a new era in vision: a survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- Antoine Bardes, Mehdi Mirza, Boris Oreshkin, Michael Auli, Ishan Misra, and Yann LeCun. V-jepa: Latent video prediction for visual representation learning. In *International Conference on Learning Representations (ICLR)*, 2024. URL https://openreview.net/forum?id=WFYbBOEOtv.
- R Sacha Bhatia and Paul Dorian. Screening for cardiovascular disease risk with electrocardiography. *JAMA Internal Medicine*, 178(9):1163–1164, 2018.
- R. Bousseljot, D. Kreiseler, and A. Schnabel. Nutzung der EKG-Signaldatenbank CARDIODAT der PTB über das Internet. *Biomedizinische Technik*, 40(Ergänzungsband 1):317, 1995.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9650–9660, 2021.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology*, 15(3):1–45, 2024.
- Mostafa Dehghani, Yi Tay, Anurag Arnab, Lucas Beyer, and Ashish Vaswani. The efficiency misnomer. In *International Conference on Learning Representations*, 2021.
- Deborah B Diercks, George M Shumaik, Richard A Harrigan, William J Brady, and Theodore C Chan. Electrocardiographic manifestations: electrolyte abnormalities. *The Journal of emergency medicine*, 27 (2):153–160, 2004.
- Pierre Elias and Joshua Finer. Echonext: A dataset for detecting echocardiogram-confirmed structural heart disease from ecgs. PhysioNet, 2025. URL https://doi.org/10.13026/r9pp-3y42.
- Micah Goldblum, Hossein Souri, Renkun Ni, Manli Shu, Viraj Prabhu, Gowthami Somepalli, Prithvijit Chattopadhyay, Mark Ibrahim, Adrien Bardes, Judy Hoffman, et al. Battle of the backbones: A large-scale comparison of pretrained models across computer vision tasks. *Advances in Neural Information Processing Systems*, 36:29343–29371, 2023.

Brian Gow, Tom Pollard, Larry A. Nathanson, Alistair Johnson, Benjamin Moody, Chrystinne Fernandes, Nathaniel Greenbaum, Jonathan W. Waks, Parastou Eslami, Tanner Carbonati, Ashish Chaudhari,

Press, 2024.

474 475

473

480 481

> 482 483 484

485 486 487

492 493 494

495 496 497

498 499 500

501

502 503 504

507 508 509

510

505 506

511 512 513

514

515

Elizabeth Herbst, Dana Moukheiber, Seth Berkowitz, Roger Mark, and Steven Horng. MIMIC-IV-ECG: Diagnostic Electrocardiogram Matched Subset (version 1.0). https://doi.org/10.13026/ 4ngg-sb35, 2023. PhysioNet. RRID:SCR_007345. Lennart Graf, Dagmar Krefting, Tibor Kesztyüs, Maximilian Oremek, Sven Zenker, and Nicolai Spicher. Interoperable integration of automatic ecg processing using dicomweb and the acuwave software suite.

Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. In International Conference on Learning Representations, 2022. URL https:// openreview.net/forum?id=uYLFoz1vlAC.

In Digital Health and Informatics Innovations for Sustainable Health Care Systems, pp. 1401–1405. IOS

- Awni Y Hannun, Pranav Rajpurkar, Masoumeh Haghpanahi, Geoffrey H Tison, Codie Bourn, Mintu P Turakhia, and Andrew Y Ng. Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. Nature medicine, 25(1):65–69, 2019.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 16000–16009, 2022.
- Shenda Hong, Yuxi Zhou, Junyuan Shang, Cao Xiao, and Jimeng Sun. Opportunities and challenges of deep learning methods for electrocardiogram data: A systematic review. Computers in biology and medicine, 122:103801, 2020.
- Tan Jian, Haoyi Fan, Jiawei Luo, Yanjie Zhou, Ning Wang, Xizheng Wang, Guizhi Liu, Chengyu Liu, and Zongmin Wang. A pediatric ECG database with disease diagnosis covering 11643 children, 5 2025. URL https://doi.org/10.6084/m9.figshare.27078763.v1.
- Sunil Vasu Kalmady, Amir Salimi, Weijie Sun, Nariman Sepehrvand, Yousef Nademi, Kevin Bainey, Justin Ezekowitz, Abram Hindle, Finlay McAlister, Russel Greiner, et al. Development and validation of machine learning algorithms based on electrocardiograms for cardiovascular diagnoses at the population level. npj Digital Medicine, 7(1):133, 2024.
- Sehun Kim. Learning general representation of 12-lead electrocardiogram with a joint-embedding predictive architecture. arXiv preprint arXiv:2410.08559, 2024. URL https://arxiv.org/abs/2410. 08559.
- Zuzana Koscova, Qiao Li, Chad Robichaux, Valdery Moura Junior, Manohar Ghanta, Aditya Gupta, Jonathan Rosand, Aaron Aguirre, Shenda Hong, David E Albert, et al. The harvard-emory ecg database. medRxiv, pp. 2024-09, 2024. doi: 10.1101/2024.09.27.24314503. URL https://doi.org/10. 1101/2024.09.27.24314503.
- Jaeung Lee, Jeewoo Lim, Keunho Byeon, and Jin Tae Kwak. Benchmarking pathology foundation models: Adaptation strategies and scenarios. Computers in Biology and Medicine, 190:110031, 2025.
- Jun Li, Aaron D. Aguirre, Valdery Moura Junior, Jiarui Jin, Che Liu, Lanhai Zhong, Chenxi Sun, Gari Clifford, M. Brandon Westover, and Shenda Hong. An electrocardiogram foundation model built on over 10 million recordings. NEJM AI, 2(7):AIoa2401033, 2025. doi: 10.1056/AIoa2401033. URL https://ai.nejm.org/doi/full/10.1056/AIoa2401033.

- Che Liu, Zhongwei Wan, Cheng Ouyang, Anand Shah, Wenjia Bai, and Rossella Arcucci. Zero-shot ecg classification with multimodal learning and test-time clinical knowledge enhancement. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.
 - Hui Liu, Dan Chen, Da Chen, Xiyu Zhang, Huijie Li, Lipan Bian, Minglei Shu, and Yinglong Wang. A large-scale multi-label 12-lead electrocardiogram database with standardized diagnostic statements. *Scientific data*, 9(1):272, 2022a. doi: doi.org/10.1038/s41597-022-01403-5. URL https://doi.org/10.1038/s41597-022-01403-5.
 - Hui Liu, Yinglong Wang, Da Chen, Xiyu Zhang, Huijie Li, Lipan Bian, Minglei Shu, and Dan Chen. A large-scale multi-label 12-lead electrocardiogram database with standardized diagnostic statements, 2022b. URL https://doi.org/10.6084/m9.figshare.c.5779802.v1.
 - Juan Miguel Lopez Alcaraz, Wilhelm Haverkamp, and Nils Strodthoff. Explainable machine learning for neoplasms diagnosis via electrocardiograms: an externally validated study. *Cardio-Oncology*, 11(1):70, 2025.
 - Kaden McKeen, Sameer Masood, Augustin Toma, Barry Rubin, and Bo Wang. Ecg-fm: An open electrocardiogram foundation model. *arXiv preprint arXiv:2408.05178*, 2024.
 - Temesgen Mehari and Nils Strodthoff. Towards quantitative precision for ecg analysis: Leveraging state space models, self-supervision and patient metadata. *IEEE Journal of Biomedical and Health Informatics*, 27(11):5326–5334, 2023.
 - Devon Myers, Rami Mohawesh, Venkata Ishwarya Chellaboina, Anantha Lakshmi Sathvik, Praveen Venkatesh, Yi-Hui Ho, Hanna Henshaw, Muna Alhawawreh, David Berdik, and Yaser Jararweh. Foundation and large language models: fundamentals, challenges, opportunities, and social impacts. *Cluster Computing*, 27(1):1–26, 2024.
 - Yeongyeon Na, Minje Park, Yunwon Tae, and Sunghoon Joo. Guiding masked representation learning to capture spatio-temporal relationship of electrocardiogram. *arXiv preprint arXiv:2402.09450*, 2024. URL https://arxiv.org/abs/2402.09450.
 - Peter Neidlinger, Omar SM El Nahhas, Hannah Sophie Muti, Tim Lenz, Michael Hoffmeister, Hermann Brenner, Marko van Treeck, Rupert Langer, Bastian Dislich, Hans Michael Behrens, et al. Benchmarking foundation models as feature extractors for weakly-supervised computational pathology. *arXiv preprint arXiv:2408.15823*, 2024.
 - Naoki Nonaka and Jun Seita. In-depth benchmarking of deep neural network architectures for ecg diagnosis. In *Machine Learning for Healthcare Conference*, pp. 414–439. PMLR, 2021.
 - Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
 - M. A. Reyna, N. Sadr, E. A. Perez Alday, A. Gu, A. J. Shah, C. Robichaux, A. B. Rad, A. Elola, S. Seyedi, S. Ansari, H. Ghanbari, Q. Li, A. Sharma, and G. D. Clifford. Will Two Do? Varying Dimensions in Electrocardiography: The PhysioNet/Computing in Cardiology Challenge 2021. In 2021 Computing in Cardiology (CinC), pp. 1–4, Brno, Czech Republic, 2021a. doi: 10.23919/CinC53138.2021.9662687.
 - M. A. Reyna, N. Sadr, A. Gu, E. A. Perez Alday, C. Liu, S. Seyedi, A. Shah, and G. D. Clifford. Will Two Do? Varying Dimensions in Electrocardiography: The PhysioNet/Computing in Cardiology Challenge 2021 (version 1.0.3). https://doi.org/10.13026/34va-7q14, 2022. PhysioNet. RRID:SCR_007345.

Matthew A Reyna, Nadi Sadr, Erick A Perez Alday, Annie Gu, Amit J Shah, Chad Robichaux, Ali Bahrami Rad, Andoni Elola, Salman Seyedi, Sardar Ansari, et al. Will two do? varying dimensions in electrocardiography: the physionet/computing in cardiology challenge 2021. In 2021 computing in cardiology (CinC), volume 48, pp. 1–4. IEEE, 2021b.

- Antônio H Ribeiro, Manoel Horta Ribeiro, Gabriela MM Paixão, Derick M Oliveira, Paulo R Gomes, Jéssica A Canazart, Milton PS Ferreira, Carl R Andersson, Peter W Macfarlane, Wagner Meira Jr, et al. Automatic diagnosis of the 12-lead ecg using a deep neural network. *Nature communications*, 11(1):1760, 2020.
- Antônio H. Ribeiro, Gabriela M. M. Paixao, Emilly M. Lima, Manoel Horta Ribeiro, Marcelo M. Pinto Filho, Paulo R. Gomes, Derick M. Oliveira, Wagner Meira Jr, Thömas B. Schon, and Antonio Luiz P. Ribeiro. CODE-15%: A Large Scale Annotated Dataset of 12-lead ECGs. https://doi.org/10.5281/zenodo.4916206, 2021. Zenodo, June 9, 2021.
- Ivan C Rokos, William J French, Amal Mattu, Graham Nichol, Michael E Farkouh, James Reiffel, and Gregg W Stone. Appropriate cardiac cath lab activation: optimizing electrocardiogram interpretation and clinical decision-making for acute st-elevation myocardial infarction. *American heart journal*, 160(6): 995–1003, 2010.
- Konstantinos C Siontis, Peter A Noseworthy, Zachi I Attia, and Paul A Friedman. Artificial intelligence-enhanced electrocardiography in cardiovascular disease management. *Nature Reviews Cardiology*, 18(7): 465–478, 2021.
- Nils Strodthoff, Patrick Wagner, Tobias Schaeffter, and Wojciech Samek. Deep learning for ecg analysis: Benchmarks and insights from ptb-xl. *IEEE journal of biomedical and health informatics*, 25(5):1519–1528, 2020.
- Nils Strodthoff, Juan Miguel Lopez Alcaraz, and Wilhelm Haverkamp. Prospects for artificial intelligence-enhanced electrocardiogram as a unified screening tool for cardiac and non-cardiac conditions: an explorative study in emergency care. *European Heart Journal Digital Health*, 5(4):454–460, May 2024. ISSN 2634-3916. doi: 10.1093/ehjdh/ztae039. URL http://dx.doi.org/10.1093/ehjdh/ztae039.
- Shashank Subramanian, Peter Harrington, Kurt Keutzer, Wahid Bhimji, Dmitriy Morozov, Michael W Mahoney, and Amir Gholami. Towards foundation models for scientific machine learning: Characterizing scaling and transfer behavior. *Advances in Neural Information Processing Systems*, 36:71242–71262, 2023.
- Jian Tan, Haoyi Fan, Jiawei Luo, Yanjie Zhou, Ning Wang, Xizheng Wang, Guizhi Liu, Chengyu Liu, and Zongmin Wang. A pediatric ecg database with disease diagnosis covering 11643 children. *Scientific Data*, 12(1):867, 2025. doi: 10.1038/s41597-025-05225-z. URL https://doi.org/10.1038/s41597-025-05225-z.
- Yuanyuan Tian, Zhiyuan Li, Yanrui Jin, Mengxiao Wang, Xiaoyang Wei, Liqun Zhao, Yunqing Liu, Jinlei Liu, and Chengliang Liu. Foundation model of ecg diagnosis: Diagnostics and explanations of any form and rhythm on ecg. *Cell Reports Medicine*, 5(12):101875, December 2024. ISSN 2666-3791. doi: 10.1016/j.xcrm.2024.101875. URL http://dx.doi.org/10.1016/j.xcrm.2024.101875.
- Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748, 2018. URL https://arxiv.org/abs/1807.03748.
- Johanna Vielhaben, Dilyara Bareeva, Jim Berend, Wojciech Samek, and Nils Strodthoff. Beyond scalars: Concept-based alignment analysis in vision transformers, 2025. URL https://arxiv.org/abs/2412.06639.

Patrick Wagner, Nils Strodthoff, Ralf-Dieter Bousseljot, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. Ptb-xl, a large publicly available electrocardiography dataset. Scien-tific data, 7(1):1–15, 2020. doi: 10.1038/s41597-020-0495-6. URL https://doi.org/10.1038/ s41597-020-0495-6. Philipp Wagner, Nils Strodthoff, Ralf Bousseljot, Wojciech Samek, and Tobias Schaeffter. PTB-XL, a large publicly available electrocardiography dataset (version 1.0.3). https://doi.org/10.13026/ kfzx-aw45, 2022. PhysioNet. RRID:SCR_007345. David A Wald. Ecg manifestations of selected metabolic and endocrine disorders. Emergency Medicine Clinics, 24(1):145–157, 2006. Wei Xiong, Jiangtong Li, Jie Li, and Kun Zhu. Eeg-fm-bench: A comprehensive benchmark for the system-atic evaluation of eeg foundation models. arXiv preprint arXiv:2508.17742, 2025.

A APPENDIX

A.1 BENCHMARKING TASKS

Table 3 provides a detailed overview of benchmarking tasks and the related datasets.

Table 3: This table provides an overview of the datasets and prediction tasks included in this study, covering routine cardiac diagnostics, clinical outcomes, and patient metadata, along with their sample sizes, patient counts, and label structures. For all tasks, we report effective sample sizes rather than the total number of available ECGs. In classification tasks, the effective sample size corresponds to the number of ECGs and patients with at least one positive label, since samples containing only negative labels do not contribute information about the presence of a condition. In regression tasks, it corresponds to the number of ECGs and patients with at least one valid numeric target value, as missing targets cannot be used for training or evaluation. Consequently, effective sample sizes may be substantially smaller than the total dataset size. We adopt this definition to provide a realistic estimate of usable data per task, ensure comparability across tasks, and avoid overstating the available training signal. †: evaluation only

Task	Dataset	Type	Samples	Patients	Outputs
	Adult ECG interpretar	tion			
ECG interpretation	PTB	Multi-label	549	290	22
ECG interpretation	Ningbo	Multi-label	34,808	unknown	68
ECG interpretation	CPSC2018	Multi-label	6,867	unknown	9
ECG interpretation	CPSC-Extra	Multi-label	3,441	unknown	33
ECG interpretation	Georgia	Multi-label	10,286	unknown	50
ECG interpretation	Chapman	Multi-label	10,646	10,646	42
ECG interpretation	Chapman (rhythm) [†]	Multi-label	10,646	10,646	9
ECG interpretation	SPH	Multi-label	25,770	24,666	35
ECG interpretation	CODE-15%	Multi-label	345,109	233,480	6
ECG interpretation	PTB-XL(all/sub/super)	Multi-label	21,799	18,869	71/23/5
ECG interpretation [†]	PTB-XL(diag/form/rhythm)	Multi-label	21,799	18,869	44/19/12
	Pediatric ECG interpret	ation			
ECG interpretation	12,328	10,350	58		
	Cardiac structure & fun	ction			
Echocardiogram findings	EchoNext	Multi-label	82,543	36,286	11
	Cardiac outcomes				
Cardiac discharge diagnoses	MIMIC-IV-ECG	Multi-label	114,355	49,400	158
	Non-cardiac outcom	es			
Non-cardiac discharge diagnoses	MIMIC-IV-ECG	Multi-label	178,163	81,930	918
	Acute care prediction	ns			
Clinical deterioration	MIMIC-IV-ECG	Multi-label	5,577	4,595	6
Mortality	MIMIC-IV-ECG	Multi-label	17,639	10,220	7
ICU admission	MIMIC-IV-ECG	Multi-label	18,690	13,868	2
	Patient characteristic				
Sex	MIMIC-IV-ECG	Binary	182,076	83,736	1
Age	MIMIC-IV-ECG	Regression		83,736	1
Biometrics	MIMIC-IV-ECG	Regression		53,702	3
ECG features	MIMIC-IV-ECG	Regression		83,721	7
Lab values	MIMIC-IV-ECG	Regression		53,293	18
Vital signs	MIMIC-IV-ECG	Regression	131,602	66,605	6

A.2 MODEL ARCHITECTURES

A.2.1 ECG-CPC MODEL

The model architecture largely follows (Mehari & Strodthoff, 2023) and is composed of four convolutional layers as the encoder, followed by four S4 layers (Gu et al., 2022) (state dimension 8 and model dimension 512) as the predictor. In contrast to (Mehari & Strodthoff, 2023), the model operates at a sampling frequency of 240 Hz, which is the minimal sampling frequency in the HEEDB dataset (Koscova et al., 2024) used for pretraining. To account for the deviation from the sampling frequency in the original publication, the model uses a kernel size of 3 and a stride of 2 in the first convolutional layer and predicts ahead 14 steps (as compared to 12 originally) in the CPC objective.

A.2.2 S4 MODEL

The S4-based supervised baseline follows the specification in (Strodthoff et al., 2024). It also uses four S4 layers (Gu et al., 2022) (state dimension 8 and model dimension 512) without a convolutional encoder. It operates at a sampling frequency of 100 Hz with an input size of 2.5 s.

A.3 PREDICTIVE PERFORMANCE

A.3.1 FINETUNING

Quantitative finetuning results are compiled in Table 4.

A.3.2 FROZEN EVALUATION

Frozen evaluation results are compiled in Table 5.

A.3.3 LINEAR EVALUATION

Linear evaluation results are compiled in Table 6.

A.3.4 RANKING

In Table 7, we summarize model performance for each task in terms of a ranked list with ties, accounting for statistical significance. In Table 8, we summarize these results by reporting the median ranks for each of the seven categories considered.

A.4 LABEL-SPECIFIC MODEL PREDICTIONS

In Table 9-Table 30, we show model predictions for the 10 best predicted labels (sorted by the performance of the supervised S4 model). These tables reflect the high degree of specificity of the tasks that are covered by this benchmark.

A.5 SCALING ANALYSIS

Table 31 lists the fit parameters for the scaling curves of the form $CN^{-\alpha} + L_0$. Taking the supervised baseline (S4) as reference, one can use these parametric forms to work out for a given training set size N, what the required training set size would be N^* to reach the same level of performance with a given pretrained model. The ratio $r = N^*/N$ then characterizes the improved label efficiency of the pretrained

Table 4: Comparison of aggregated macro-AUROC for classification and MAE for regression under finetuning with a linear prediction head. We highlight with ↑ tasks where higher AUROC is better and ↓ tasks where lower standardized MAE values are better. The best-performing result is highlighted in boldface and underlined, while models that do not perform statistically significantly worse are also highlighted in boldface. † signifies evaluation of a model trained on another dataset (listed above).

	Foundation Models (Finetuned)										
	ECGFounder	ECG-JEPA	ST-MEM	MERL	ECGFM-KED	HuBERT-ECG	ECG-FM	ECG-CPC	S4	Net1D	
			Adult	ECG int	terpretation						
PTB ↑	0.656	0.679	0.694	0.717	0.612	0.699	0.725	0.702	0.654	0.564	
Ningbo ↑	0.974	0.973	0.954	0.955	0.940	0.958	0.971	0.973	0.972	0.968	
CPSC2018 ↑	0.966	0.974	0.946	0.936	0.930	0.956	0.972	0.969	0.962	0.949	
CPSC-Extra ↑	0.906	0.897	0.883	0.873	0.824	0.876	0.862	0.898	0.852	0.818	
Georgia ↑	0.920	0.918	0.888	0.912	0.877	0.883	0.912	0.913	0.903	0.884	
Chapman ↑	0.968	0.972	0.948	0.946	0.917	0.941	0.956	0.962	0.963	0.953	
-Chapman (rhythm) [†] ↑	0.991	0.989	0.985	0.975	0.963	0.982	0.993	0.987	0.986	0.983	
SPH ↑	0.983	0.980	0.964	0.944	0.932	0.953	0.966	0.981	0.981	0.967	
CODE-15% ↑	0.987	0.991	0.974	0.982	0.964	0.991	0.986	0.989	0.990	0.983	
PTB-XL (all) ↑	0.934	0.940	0.908	0.925	0.889	0.915	0.927	0.949	0.941	0.929	
-PTB-XL (diag) † ↑	0.950	0.946	0.904	0.942	0.913	0.925	0.926	0.951	0.943	0.940	
-PTB-XL (form) [†] ↑	0.875	0.912	0.887	0.891	0.829	0.860	0.904	0.934	0.919	0.893	
-PTB-XL (rhythm) [†] ↑	0.965	0.956	0.951	0.912	0.888	0.950	0.959	0.959	0.956	0.937	
PTB-XL (sub) ↑	0.943	0.935	0.931	0.912	0.908	0.930	0.939	0.939	0.938	0.937	
PTB-XL (super) ↑	0.935	0.933	0.910	0.930	0.905	0.908	0.932	0.934		0.924	
FTB-AL (super)	0.555	0.921	0.901	0.930	0.903	0.908	0.910	0.554	0.532	0.924	
			Pediatr	ic ECG i	nterpretation						
ZZU pECG ↑	0.898	<u>0.911</u>	0.893	0.886	0.861	0.883	0.887	0.892	0.897	0.868	
			Cardia	structu	re & function						
EchoNext (Echo) ↑	0.817	0.817	0.816	0.822	0.806	0.792	0.772	<u>0.831</u>	0.819	0.803	
			С	ardiac ou	itcomes						
MIMIC (Cardiac) ↑	0.768	0.772	0.760	0.776	0.767	0.719	0.775	0.781	0.780	0.747	
			Non	-cardiac	outcomes						
MIMIC (Non-cardiac) ↑	0.701	0.711	0.688	0.712	0.702	0.642	0.719	0.719	0.714	0.672	
			Acu	te care p	redictions						
MIMIC (Deterioration) ↑	0.717	0.747	0.714	0.743	0.728	0.664	0.767	0.764	0.756	0.731	
MIMIC (Mortality) ↑	0.810	0.792	0.784	0.800	0.768	0.722	0.811	0.803	0.793	0.744	
MIMIC (ICU) ↑	0.748	0.742	0.737	0.744	0.734	0.710	0.750	0.753	0.745	0.721	
			Pati	ent char	acteristics						
MIMIC (Sex) ↑	0.913	0.904	0.883	0.916	0.903	0.810	0.919	0.933	0.919	0.869	
MIMIC (Age) ↓	0.461	0.463	0.504	0.449	0.466	0.579	0.412	0.437	0.455	0.518	
MIMIC (Biometrics) ↓	0.637	0.640	0.673	0.625	0.647	0.723	0.620	0.604	0.626	0.684	
MIMIC (ECG Features) ↓	0.458	0.460	0.500	0.463	0.466	0.529	0.465	0.451	0.452	0.486	
MIMIC (Lab Values) ↓	0.679	0.677	0.688	0.676	0.685	0.709	0.688	0.673	0.675	0.691	
MIMIC (Vital Signs) \downarrow	0.704	0.703	0.715	0.702	0.705	0.717	0.704	0.700	0.701	0.712	

model. For each of the models, we evaluate r for different training dataset sizes. The results are compiled in Table 32.

Table 5: Comparison of aggregated macro-AUROC for classification and MAE for regression under the frozen evaluation mode. We highlight with \uparrow tasks where higher AUROC is better and \downarrow tasks where lower standardized MAE values are better. The best-performing result is highlighted in boldface and underlined, while models that do not perform statistically significantly worse are also highlighted in boldface. † signifies evaluation of a model trained on another dataset (listed above).

			Founda	ation Mo	dels (Frozen Eva	luation)			Supe	rvised		
	ECGFounder	ECG-JEPA	ST-MEM	MERL	ECGFM-KED	HuBERT-ECG	ECG-FM	ECG-CPC	S4	Net1D		
			Adult	ECG int	terpretation							
PTB ↑	0.681	0.681	0.686	0.661	0.592	0.592	0.639	0.716	0.654	0.564		
Ningbo ↑	0.961	0.971	0.956	0.942	0.833	0.911	0.927	0.953	0.972	0.968		
CPSC2018↑	0.966	0.975	0.956	0.956	0.874	0.919	0.930	0.959		0.949		
CPSC-Extra ↑	0.907	0.902	0.867	0.872	0.739	0.831	0.840	0.887	0.852	0.818		
Georgia ↑	0.924	0.910	0.893	0.890	0.787	0.836	0.858	0.894	0.903	0.884		
Chapman ↑	0.967	0.964	0.955	0.954	0.852	0.917	0.900	0.943	0.963	0.953		
-Chapman (rhythm) [†] ↑	0.983	0.988	0.984	0.972	0.849	0.951	0.974	0.983	0.986	0.983		
SPH ↑	0.966	0.980	0.946	0.958	0.885	0.939	0.938	0.961	0.981	0.967		
CODE-15% ↑	0.980	0.990	0.965	0.879	0.688	0.978	0.968	0.983	0.990	0.983		
PTB-XL (all) ↑	0.927	0.934	0.910	0.909	0.810	0.883	0.884	0.931	0.941	0.929		
-PTB-XL (diag) † ↑	0.940	0.943	0.910	0.918	0.827	0.898	0.890	0.934	0.943	0.940		
-PTB-XL (form) [†] ↑	0.876	0.889	0.894	0.879	0.783	0.815	0.838	0.905	0.919			
-PTB-XL (rhythm) † ↑	0.958	0.966	0.934	0.928	0.792	0.917	0.931	0.951	0.956			
PTB-XL (sub) ↑	0.939	0.934	0.931	0.917	0.816	0.903	0.912	0.934	0.938	0.926		
PTB-XL (super) ↑	0.928	0.917	0.923	0.915	0.865	0.892	0.877	0.919		0.924		
Pediatric ECG interpretation												
ZZU pECG ↑	0.891	0.905	0.899	0.870	0.789	0.857	0.845	0.879	0 807	0.868		
ZZC pEco	0.071	0.705				0.837	0.043	0.079	0.077	0.000		
					re & function							
EchoNext (Echo) ↑	0.803	0.811	0.817	0.801	0.791	0.778	0.772	0.822	0.819	0.803		
			С	ardiac ou	utcomes							
MIMIC (Cardiac) ↑	0.745	0.757	0.734	0.755	0.708	0.736	0.688	0.774	0.780	0.747		
			Non	-cardiac	outcomes							
MIMIC (Non-cardiac) ↑	0.671	0.688	0.657	0.681	0.630	0.659	0.616	0.703	0.714	0.672		
			Δen	te care n	redictions							
MIMIC (Deterioration) ↑	0.697	0.702	0.648	0.721	0.661	0.685	0.639	0.743	0.756	0.731		
MIMIC (Mortality) ↑	0.769	0.788	0.754	0.761	0.720	0.723	0.731	0.785	0.793			
MIMIC (ICU) ↑	0.731	0.734	0.715	0.727	0.698	0.719	0.691	0.750		0.721		
	0.751	0.75				0.717	0.071	<u> </u>	017 10	0.721		
					acteristics							
MIMIC (Sex) ↑	0.872	0.894	0.883	0.879	0.839	0.841	0.826	0.918		0.869		
MIMIC (Age) ↓	0.515	0.484	0.501	0.511	0.581	0.544	0.542	0.466	0.455	0.518		
MIMIC (Biometrics) ↓	0.702	0.700	0.681	0.683	0.715	0.702	0.751	0.640	0.626			
MIMIC (ECG Features) ↓		0.477	0.489	0.507	0.563	0.500	0.566	0.465	0.452	0.486		
MIMIC (Lab Values) ↓	0.703	0.694	0.740	0.694	0.712	0.703	0.763	0.676	0.675	0.691		
MIMIC (Vital Signs) ↓	0.719	0.716	0.739	0.716	0.729	0.722	0.747	0.703	0.701	0.712		

 Table 6: Comparison of aggregated macro-AUROC for classification and MAE for regression under the linear evaluation mode. We highlight with \uparrow tasks where higher AUROC is better and \downarrow tasks where lower standardized MAE values are better. The best-performing result is highlighted in boldface and underlined, while models that do not perform statistically significantly worse are also highlighted in boldface. † signifies evaluation of a model trained on another dataset (listed above).

			Founda	ation Mo	dels (Linear Eva	luation)			Supe	ervised		
	ECGFounder	ECG-JEPA	ST-MEM	MERL	ECGFM-KED	HuBERT-ECG	ECG-FM	ECG-CPC	S4	Net1D		
			Adult	ECG int	terpretation							
PTB ↑	0.671	0.665	0.692	0.583	0.503	0.604	0.692	0.578	0.654	0.564		
Ningbo ↑	0.970	0.970	0.954	0.916	0.762	0.896	0.902	0.898	0.972	0.968		
CPSC2018 ↑	0.964	0.975	0.945	0.914	0.786	0.899	0.906	0.902	0.962	0.949		
CPSC-Extra ↑	0.910	0.902	0.885	0.858	0.553	0.855	0.842	0.794	0.852	0.818		
Georgia ↑	0.923	0.920	0.889	0.872	0.642	0.847	0.847	0.854	0.903	0.884		
Chapman ↑	0.968	0.962	0.949	0.916	0.745	0.904	0.891	0.868	0.963	0.953		
-Chapman (rhythm) [†] ↑	0.987	0.989	0.985	0.946	0.776	0.953	0.968	0.944	0.986	0.983		
SPH ↑	0.975	0.967	0.966	0.943	0.798	0.894	0.914	0.928	0.981			
CODE-15% ↑	0.976	0.984	0.975	0.716	0.568	0.965	0.976	0.968	0.990	0.983		
PTB-XL (all) ↑	0.931	0.928	0.908	0.883	0.706	0.867	0.841	0.904	0.941	0.929		
-PTB-XL (diag) [†] ↑	0.947	0.925	0.903	0.894	0.715	0.875	0.861	0.907	0.943	0.940		
-PTB-XL (form) [†] ↑	0.874	0.908	0.889	0.864	0.716	0.817	0.782	0.873	0.919			
-PTB-XL (rhythm) [†] ↑	0.961	0.969	0.951	0.877	0,666	0.902	0.865	0.938	0.956			
PTB-XL (sub) ↑	0.945	0.916	0.916	0.895	0.734	0.898	0.866	0.886	0.938			
PTB-XL (super) ↑	0.924	0.911	0.896	0.887	0.802	0.877	0.873	0.863		0.924		
Pediatric ECG interpretation												
ZZU pECG ↑	0.900	0.891	0.893	0.847	0.591	0.827	0.813	0.852	0.807	0.868		
ZZU PECU	<u>0.900</u>	0.691	0.093	0.847	0.391	0.827	0.813	0.832	0.097	0.808		
			Cardia		re & function							
EchoNext ↑	0.795	0.806	0.816	0.794	0.770	0.770	0.767	0.800	0.819	0.803		
			С	ardiac ou	utcomes							
MIMIC (Cardiac) ↑	0.751	0.751	0.761	0.751	0.683	0.719	0.675	0.751	0.780	0.747		
			Non	-cardiac	outcomes							
MIMIC (Non-cardiac) ↑	0.671	0.675	0.688	0.672	0.617	0.642	0.599	0.680	0.714	0.672		
			Acu	te care p	redictions							
MIMIC (Deterioration) ↑	0.713	0.720	0.717	0.704	0.627	0.664	0.616	0.733	0.756	0.731		
MIMIC (Mortality) ↑	0.774	0.782	0.788	0.744	0.681	0.722	0.676	0.769	0.793			
MIMIC (ICU) ↑	0.730	0.736	0.737	0.722	0.658	0.710	0.683	0.733	0.745			
			D. 4	4 .1	acteristics							
MIMIC (Sex) ↑	0.872	0.883	0.882	ent cnar: 0.853	0.801	0.810	0.853	0.873	0.919	0.869		
MIMIC (Sex) MIMIC (Age) ↓	0.511	0.883	0.882	0.833	0.639	0.579	0.833	0.873	0.455			
MIMIC (Age) ↓ MIMIC (Biometrics) ↓	0.700	0.485	0.503	0.711	0.740	0.723	0.837	0.686	0.626			
MIMIC (ECG Features) ↓		0.490	0.499	0.711	0.606	0.723	0.617	0.504	0.452	0.486		
MIMIC (ECG Features) ↓ MIMIC (Lab Values) ↓	0.693	0.490	0.499	0.543	0.713	0.709	0.017	0.504	0.432	0.480		
MIMIC (Vital Signs) ↓	0.716	0.093	0.090	0.719	0.738	0.709	0.834	0.708	0.701			
winvii€ (vitai Sigiis) ↓	0.710	0.714	0.711	0.719	0.736	0.717	0.654	0.708	0./01	0.712		

Table 7: Statistical ranking of foundation models across evaluation modes and datasets. Rankings (Fine-tuned/Frozen/Linear) are assigned based on statistical equivalence groups determined by bootstrap testing, where models not performing significantly worse than the best model share the same rank. Lower ranks indicate better performance. † signifies evaluation of a model trained on another dataset (listed above).

			Foundatio	n Models	s (Finetuned/Fro	zen/Linear)			Sup	ervised
	ECGFounder	ECG-JEPA	ST-MEM	MERL	ECGFM-KED	HuBERT-ECG	ECG-FM	ECG-CPC	S4	Net1D
			Adul	t ECG ir	terpretation					
PTB	1/1/1	1/1/1	1/1/1	1/1/6	8/9/10	1/6/6	1/6/1	1/1/6	8/6/1	10/9/6
Ningbo	1/4/1	1/1/1	7/4/5	7/7/6	10/10/10	7/8/8	1/8/6	1/6/8	1/1/1	6/1/1
CPSC2018	4/2/2	1/1/1	6/6/4	9/2/6	9/10/10	6/8/8	1/8/6	1/2/8	4/2/2	6/6/4
CPSC-Extra	1/1/1	1/1/1	1/3/3	5/3/5	9/10/10	5/6/5	5/6/5	1/3/9	5/6/3	9/6/5
Georgia	1/1/1	1/1/1	7/3/4	1/3/6	7/10/10	7/8/8	1/8/8	1/3/6	6/3/3	7/7/4
Chapman	1/1/1	1/1/1	6/4/4	6/4/6	10/10/10	9/8/6	3/8/8	3/7/8	3/1/1	6/4/4
-Chapman (rhythm) [†]	1/1/1	1/1/1	4/1/4	9/8/8	10/10/10	4/9/6	1/6/6	4/6/8	4/1/1	8/1/4
SPH	1/3/2	1/1/2	5/3/5	5/7/6	10/10/10	5/8/7	5/8/7	1/3/7	1/1/1	5/3/2
CODE-15%	5/5/4	1/1/1	5/5/1	5/9/9	10/10/10	1/5/6	5/5/6	1/1/6	1/1/1	5/1/4
PTB-XL (all)	2/2/2	2/2/2	8/6/5	5/6/7	10/10/10	8/8/8	5/8/9	1/2/5	2/1/1	5/2/2
-PTB-XL (diag) †	1/1/1	1/1/4	9/6/5	4/6/7	9/10/10	7/8/8	7/8/8	1/5/5	4/1/1	4/1/3
-PTB-XL (form) [†]	7/3/3	2/3/1	7/3/3	4/3/7	9/10/10	9/8/8	4/8/8	1/1/3	2/1/1	4/3/3
-PTB-XL (rhythm) [†]	1/1/1	1/1/1	5/4/3	9/7/7	10/10/10	5/7/7	1/4/9	1/4/5	5/1/3	8/7/5
PTB-XL (sub)	1/1/1	6/1/3	7/1/3	1/6/6	10/10/10	7/8/6	1/8/9	1/4/3	1/1/1	7/6/3
PTB-XL (super)	1/1/1	5/6/4	10/2/5	4/6/6	8/10/10	8/8/7	7/9/7	1/4/9	1/1/1	5/4/2
r 1 b-AL (super)	1/2/2	3/0/4	10/2/3	4/0/0	6/10/10	0/0//	11711	1/4/9	1/1/1	31412
					interpretation					
ZZU pECG	2/4/1	1/1/1	2/1/1	2/6/7	9/10/10	7/6/7	7/9/9	2/4/5	2/1/1	9/6/5
			Cardia	e structi	re & function					
EchoNext (Echo)	2/4/6	2/4/3	2/1/1	2/7/6	7/8/8	9/9/8	10/9/8	1/1/3	2/1/1	7/4/3
			(Cardiac o	outcomes					
MIMIC (Cardiac)	6/5/3	3/3/3	8/7/2	3/3/3	6/9/9	10/7/8	3/10/9	1/2/3	1/1/1	9/5/3
				n-cardia	coutcomes					
MIMIC (Non-cardiac)	6/5/7	4/3/4	8/7/2	4/4/4	6/9/9	10/7/8	1/10/10	1/2/3	3/1/1	9/5/4
			Ac	ute care j	predictions					
MIMIC (Deterioration)	6/4/5	1/4/1	6/8/5	1/4/5	6/8/8	10/4/8	1/8/8	1/1/1	1/1/1	6/1/1
MIMIC (Mortality)	1/1/1	1/1/1	1/1/1	1/6/7	8/6/7	8/10/7	1/6/7	1/1/1	1/1/1	8/6/1
MIMIC (ICU)	1/3/4	4/3/1	4/6/1	4/3/6	8/9/10	10/6/6	1/9/9	1/1/4	4/1/1	9/6/6
			Pat	tient cha	racteristics					
MIMIC (Sex)	5/6/4	6/3/2	8/4/2	2/4/7	6/8/9	10/8/9	2/10/7	1/1/4	2/1/1	9/6/4
MIMIC (Age)	5/5/4	5/3/2	8/4/3	3/5/7	5/10/10	10/8/8	1/8/8	2/2/6	4/1/1	9/7/5
MIMIC (Biometrics)	5/6/6	5/6/3	8/3/2	2/3/7	7/9/9	10/6/8	2/10/10	1/2/3	2/1/1	9/3/3
MIMIC (ECG Features)		4/3/4	9/5/5	5/8/8	7/9/9	10/7/7	5/9/10	1/2/6	2/1/1	8/4/2
MIMIC (Lab Values)	5/6/4	1/3/6	6/9/2	1/5/6	6/8/9	10/6/8	6/10/10	1/1/2	1/1/1	9/3/4
MIMIC (Vital Signs)	6/4/5	3/4/5	9/9/3	3/4/8	6/8/9	9/7/7	3/10/10	1/2/2	1/1/1	

Table 8: Median statistical rankings of foundation models across evaluation modes by categories. Rankings (Finetuned/Frozen/Linear) represent the median performance position across all datasets within each category. Lower values indicate better overall performance.

		Foundation Models (Finetuned/Frozen/Linear)											
	ECGFounder	ECG-JEPA	ST-MEM	MERL	ECGFM-KED	HuBERT-ECG	ECG-FM	ECG-CPC	S4	Net1D			
Adult ECG interpretation	1/1/1	1/1/1	6/3/4	5/6/6	10/10/10	7/8/7	3/8/7	1/3/7	3/1/1	6/4/4			
Pediatric ECG interpretation	2/4/1	1/1/1	2/1/1	2/6/7	9/10/10	7/6/7	7/9/9	2/4/5	2/1/1	9/6/5			
Cardiac structure & function	2/4/6	2/4/3	2/1/1	2/7/6	7/8/8	9/9/8	10/9/8	1/1/3	2/1/1	7/4/3			
Cardiac outcomes	6/5/3	3/3/3	8/7/2	3/3/3	6/9/9	10/7/8	3/10/9	1/2/3	1/1/1	9/5/3			
Non-cardiac outcomes	6/5/7	4/3/4	8/7/2	4/4/4	6/9/9	10/7/8	1/10/10	1/2/3	3/1/1	9/5/4			
Acute care predictions	1/3/4	1/3/1	4/6/1	1/4/6	8/8/8	10/6/7	1/8/8	1/1/1	1/1/1	8/6/1			
Patient characteristics	5/5.5/4	4.5/3/3.5	8/4.5/2.5	2.5/4.5/7	6/8.5/9	10/7/8	2.5/9.5/10	1/2/3.5	2/1/1	9/3.5/3.5			

Table 9: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the PTB dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
Healthy control	0.968	0.919	0.901	0.917	0.833	0.922	0.911	0.884	0.881	0.920
Arterial Hypertension	0.638	0.400	0.724	0.705	0.810	0.610	0.924	0.619	0.876	0.352
Atrial fibrillation	0.997	0.958	1.000	0.968	0.430	0.955	1.000	0.893	0.780	0.372
Cardiomyopathy	0.812	0.880	0.913	0.731	0.923	0.909	0.740	0.875	0.764	0.822
Ventricular fibrillation	0.465	0.684	0.528	0.786	0.547	0.709	0.584	0.564	0.736	0.664
Myocardial infarction acute	0.686	0.330	0.706	0.819	0.809	0.414	0.650	0.793	0.722	0.731
Myocardial infarction old	0.790	0.810	0.767	0.806	0.675	0.758	0.746	0.727	0.717	0.696
Myocardial infarction acute catheterized	0.778	0.760	0.828	0.736	0.652	0.767	0.824	0.710	0.715	0.696
Obesity	0.498	0.755	0.682	0.557	0.668	0.430	0.647	0.795	0.700	0.507
Arterial hypertension	0.552	0.713	0.621	0.743	0.654	0.704	0.674	0.680	0.693	0.556

Table 10: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the Ningbo dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
Wolff-Parkinson-White	1.000	1.000	1.000	1.000	0.997	0.998	1.000	1.000	1.000	1.000
AAR	0.997	0.997	0.999	1.000	0.996	0.898	0.999	0.997	0.999	0.998
Ventricular fibrillation	1.000	1.000	1.000	0.994	0.990	1.000	1.000	0.999	0.999	0.999
Atrial flutter	0.999	0.999	0.998	0.998	0.996	0.997	0.999	0.998	0.999	0.998
Right atrial hypertrophy	0.996	0.999	0.984	1.000	1.000	0.905	0.994	0.997	0.998	0.998
Ventricular escape rhythm	0.999	0.999	0.999	0.998	0.998	0.999	0.998	0.999	0.998	0.999
Sinus tachycardia	0.998	0.997	0.995	0.996	0.995	0.996	0.997	0.998	0.998	0.997
Junctional tachycardia	0.985	1.000	0.996	0.972	0.997	1.000	0.999	0.992	0.998	0.997
Sinus bradycardia	0.998	0.998	0.996	0.994	0.998	0.998	0.999	0.998	0.998	0.998
Complete left bundle branch block	0.999	0.998	0.998	0.999	0.998	0.999	0.999	0.998	0.997	0.998



Table 11: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the CPSC2018 dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
Left bundle branch block	0.998	0.999	0.998	0.999	0.998	0.999	0.999	0.998	0.998	0.998
1st degree atrioventricular block	0.990	0.990	0.971	0.963	0.946	0.981	0.980	0.982	0.987	0.982
Atrial fibrillation	0.994	0.980	0.993	0.982	0.978	0.991	0.995	0.988	0.985	0.982
Right bundle branch block	0.989	0.981	0.971	0.983	0.974	0.981	0.983	0.983	0.985	0.980
NORMAL	0.966	0.973	0.954	0.948	0.930	0.947	0.963	0.973	0.961	0.961
ST depression	0.963	0.974	0.945	0.947	0.948	0.932	0.963	0.973	0.961	0.957
Premature ventricular contraction	0.941	0.957	0.945	0.891	0.874	0.914	0.969	0.946	0.948	0.931
ST elevation	0.925	0.962	0.849	0.885	0.887	0.922	0.938	0.951	0.943	0.873
Premature atrial contraction	0.927	0.950	0.884	0.828	0.835	0.941	0.957	0.928	0.890	0.881

Table 12: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the CPSC-Extra dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
Right atrial hypertrophy	1.000	0.979	0.993	0.956	0.979	0.996	0.975	0.929	1.000	0.912
Complete heart block	0.999	1.000	0.999	0.999	0.987	0.997	1.000	1.000	0.996	0.999
2nd degree atrioventricular block	0.997	1.000	0.996	0.996	0.972	1.000	1.000	0.996	0.994	0.994
Complete right bundle branch block	0.994	0.994	0.985	0.994	0.989	0.995	0.988	0.994	0.988	0.973
Sinus tachycardia	0.986	0.988	0.990	0.978	0.922	0.989	0.987	0.990	0.985	0.987
Atrial fibrillation and flutter	0.990	0.986	0.984	0.990	0.913	0.969	0.983	0.996	0.979	0.977
Bradycardia	0.978	0.975	0.970	0.970	0.932	0.976	0.942	0.981	0.978	0.980
Atrial flutter	0.985	0.990	0.990	0.990	0.919	0.948	0.902	0.975	0.977	0.962
Incomplete right bundle branch block	0.982	0.976	0.978	0.977	0.941	0.971	0.965	0.986	0.975	0.940
Atrial tachycardia	0.990	0.959	0.991	0.976	0.926	0.997	0.991	0.978	0.972	0.516

Table 13: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the Georgia dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
2nd degree atrioventricular block	0.999	0.997	0.999	0.998	0.998	0.997	0.998	1.000	1.000	1.000
Left bundle branch block	0.998	0.998	0.998	0.997	0.991	0.998	0.997	0.998	0.997	0.998
Sinus bradycardia	0.997	0.994	0.996	0.989	0.979	0.996	0.990	0.996	0.996	0.998
Sinus tachycardia	0.995	0.988	0.988	0.981	0.973	0.994	0.984	0.994	0.993	0.991
Supraventricular tachycardia	0.995	0.990	0.994	1.000	0.985	0.996	0.978	0.995	0.988	0.958
Ventricular pacing pattern	0.922	0.981	0.987	0.972	0.982	0.968	0.908	0.979	0.982	0.725
Left anterior fascicular block	0.986	0.974	0.880	0.978	0.956	0.965	0.963	0.971	0.981	0.982
Right bundle branch block	0.992	0.980	0.987	0.982	0.982	0.986	0.986	0.986	0.980	0.982
Bundle branch block	0.978	0.981	0.941	0.960	0.951	0.962	0.928	0.969	0.976	0.971
1st degree atrioventricular block	0.984	0.980	0.984	0.975	0.944	0.943	0.983	0.983	0.976	0.970

Table 14: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the Chapman dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
Myocardial infarction in the lower wall	0.999	0.997	1.000	0.997	0.995	1.000	0.999	1.000	1.000	0.998
Sinus bradycardia	1.000	1.000	0.999	0.996	0.994	0.999	0.999	0.999	0.999	0.999
Sinus tachycardia	0.999	0.998	0.994	0.995	0.991	0.996	0.999	0.999	0.998	0.997
Myocardial infarction in the front wall	0.997	0.995	0.994	0.998	0.999	0.993	0.997	0.997	0.997	0.987
Supraventricular tachycardia	0.998	0.997	0.996	0.994	0.983	0.994	0.998	0.997	0.996	0.997
Long RR interval	0.997	0.999	0.999	0.995	0.954	0.996	0.994	0.996	0.996	0.995
Atrial fibrillation	0.997	0.996	0.996	0.994	0.987	0.992	0.998	0.996	0.994	0.992
Atrial flutter	0.990	0.984	0.987	0.944	0.909	0.978	0.996	0.992	0.993	0.992
Left front bundle branch block	0.989	0.981	0.934	0.989	0.989	0.985	0.974	0.979	0.992	0.994
Right bundle-branch block	0.996	0.995	0.990	0.993	0.989	0.993	0.994	0.995	0.991	0.994

Table 15: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the SPH dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
AV block, complete (third-degree)	1.000	1.000	1.000	1.000	0.994	1.000	1.000	1.000	1.000	1.000
Left bundle-branch block	1.000	1.000	0.999	1.000	1.000	1.000	1.000	1.000	1.000	0.999
Junctional escape complex(es)	1.000	1.000	1.000	0.999	0.998	1.000	1.000	0.999	1.000	1.000
Right bundle-branch block	1.000	1.000	0.999	0.999	1.000	0.999	0.999	0.999	1.000	0.999
2:1 AV block	1.000	0.997	1.000	0.991	0.999	0.999	0.999	0.998	0.999	0.994
Atrial fibrillation	1.000	1.000	1.000	0.999	0.997	0.999	1.000	0.999	0.999	0.998
Prolonged QT interval	0.990	0.989	0.977	0.950	0.986	0.955	0.996	0.979	0.997	0.992
Anterior MI	0.995	0.995	0.996	0.991	0.986	0.998	0.999	0.994	0.995	0.981
Sinus bradycardia	0.995	0.995	0.994	0.989	0.991	0.994	0.995	0.995	0.995	0.994
Early repolarization	0.980	0.983	0.987	0.997	0.964	0.929	0.988	0.982	0.995	0.982

Table 16: Finetuning with a linear prediction head performance for the 6 labels, sorted by supervised S4 AUROC on the CODE-15% dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
Left bundle branch block	0.997	0.998	0.997	0.994	0.994	0.996	0.994	0.997	0.996	0.995
Sinus tachycardia	0.994	0.994	0.994	0.992	0.988	0.993	0.991	0.994	0.995	0.995
Right bundle branch block	0.969	0.986	0.955	0.991	0.951	0.991	0.979	0.984	0.990	0.977
Atrial fibrillation	0.992	0.992	0.988	0.992	0.958	0.991	0.981	0.992	0.989	0.982
1st degree atrioventricular block	0.986	0.985	0.984	0.967	0.959	0.987	0.980	0.984	0.984	0.976
Sinus bradycardia	0.985	0.994	0.929	0.958	0.935	0.990	0.990	0.980	0.984	0.969

Table 17: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the PTB-XL (all) dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
Subendocardial injury in inferior leads	0.997	0.998	1.000	0.999	0.994	0.993	0.992	0.998	1.000	0.996
Complete right bundle branch block	0.999	0.998	0.998	0.998	0.997	0.997	0.998	0.998	0.998	0.997
Complete left bundle branch block	0.998	0.998	0.997	0.998	0.990	0.993	0.999	0.997	0.998	0.998
Paroxysmal supraventricular tachycardia	1.000	0.998	0.997	0.998	0.993	0.998	1.000	0.998	0.998	0.995
Sinus tachycardia	0.996	0.994	0.991	0.991	0.987	0.986	0.990	0.996	0.995	0.994
Septal hypertrophy	0.976	0.999	0.972	0.994	0.982	0.996	0.976	0.999	0.994	1.000
Posterior myocardial infarction	0.996	0.981	0.981	0.942	0.948	0.940	0.938	0.987	0.991	0.994
Third degree AV block	0.999	0.996	0.989	0.992	0.989	0.962	0.998	0.992	0.990	0.973
Subendocardial injury in anteroseptal leads	0.992	0.994	0.984	0.991	0.963	0.989	0.983	0.993	0.990	0.987
Ventricular premature complex	0.988	0.991	0.991	0.954	0.875	0.992	0.984	0.986	0.988	0.975

Table 18: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the ZZU pECG dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
Ventricular escape complex(es)	1.000	0.999	1.000	0.995	0.999	0.999	1.000	1.000	0.999	0.999
Ventricular tachycardia	0.997	0.999	0.999	0.999	0.999	0.994	0.998	0.999	0.999	0.998
Sinus pause or arrest	0.963	0.968	0.956	0.989	0.892	0.988	0.872	0.975	0.998	0.992
AV block, advanced (high-grade)	0.989	0.998	0.987	0.989	0.998	0.996	0.985	0.985	0.998	0.983
Second-degree AV block, Mobitz type I(Wenckebach)	0.999	0.999	0.999	0.971	0.978	0.958	1.000	0.994	0.998	0.974
Right bundle-branch block	0.995	0.997	0.995	0.982	0.988	0.988	0.992	0.995	0.989	0.990
AV block, complete (third-degree)	0.996	0.999	0.999	0.997	0.987	0.996	0.998	0.997	0.987	0.998
Atrial flutter	0.989	0.993	0.968	0.894	0.967	0.972	0.959	0.968	0.987	0.979
TU fusion	0.981	0.957	0.960	0.997	0.958	0.912	0.947	0.989	0.984	0.921
Sinus bradycardia	0.984	0.974	0.981	0.982	0.979	0.984	0.976	0.986	0.984	0.985

Table 19: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the EchoNext dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
lvef lte 45 flag	0.901	0.897	0.896	0.902	0.883	0.878	0.873	0.910	0.897	0.879
rv systolic dysfunction moderate or greater flag	0.879	0.891	0.883	0.883	0.877	0.861	0.856	0.899	0.890	0.880
aortic stenosis moderate or greater flag	0.821	0.838	0.830	0.843	0.814	0.790	0.767	0.849	0.857	0.808
tricuspid regurgitation moderate or greater flag	0.832	0.845	0.840	0.845	0.825	0.806	0.779	0.851	0.844	0.827
mitral regurgitation moderate or greater flag	0.836	0.817	0.828	0.835	0.813	0.805	0.789	0.833	0.829	0.807
pulmonary regurgitation moderate or greater flag	0.875	0.857	0.880	0.872	0.853	0.825	0.861	0.870	0.828	0.851
pasp gte 45 flag	0.788	0.794	0.774	0.789	0.773	0.759	0.725	0.810	0.793	0.766
tr max gte 32 flag	0.779	0.792	0.757	0.781	0.764	0.748	0.722	0.804	0.784	0.748
aortic regurgitation moderate or greater flag	0.751	0.734	0.751	0.760	0.747	0.755	0.724	0.758	0.776	0.752
lvwt gte 13 flag	0.767	0.768	0.762	0.774	0.754	0.753	0.706	0.779	0.774	0.755

Table 20: Finetuning with a linear prediction head performance for the 10 best-predicted labels, sorted by supervised S4 AUROC on the MIMIC (cardiac) dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
I447 Left bundle block	0.938	0.935	0.928	0.923	0.932	0.921	0.921	0.929	0.939	0.920
I4510 Right bundle block	0.936	0.925	0.931	0.937	0.930	0.912	0.915	0.919	0.918	0.926
I5023 Acute/chronic systolic HF	0.900	0.900	0.899	0.900	0.894	0.891	0.892	0.907	0.906	0.894
I428 Other cardiomyopathies	0.881	0.892	0.884	0.885	0.883	0.861	0.882	0.894	0.901	0.878
I255 Ischemic cardiomyopathy	0.902	0.882	0.896	0.895	0.888	0.879	0.870	0.892	0.900	0.873
I132 HTN heart+CKD w/HF, ESRD	0.897	0.899	0.873	0.893	0.913	0.847	0.886	0.912	0.896	0.876
I482 Chronic AF	0.895	0.894	0.892	0.888	0.862	0.865	0.882	0.899	0.893	0.872
I211 STEMI, inferior wall	0.896	0.881	0.867	0.876	0.838	0.776	0.794	0.885	0.887	0.754
I078 Rheumatic tricuspid disease	0.869	0.867	0.879	0.867	0.877	0.817	0.842	0.873	0.878	0.873
I44 AV + LBBB	0.876	0.869	0.876	0.866	0.868	0.855	0.869	0.879	0.875	0.850

Table 21: Finetuning with a linear prediction head performance for the 10 best predicted-labels, sorted by supervised S4 AUROC on the MIMIC (non-cardiac) dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
L9740 Chronic ulcer heel/midfoot	0.841	0.864	0.762	0.908	0.897	0.764	0.886	0.904	0.941	0.805
Z4502 ICD defibrillator management	0.945	0.918	0.938	0.929	0.910	0.950	0.919	0.935	0.936	0.919
K767 Hepatorenal syndrome	0.909	0.886	0.812	0.885	0.813	0.631	0.856	0.924	0.903	0.801
Z681 BMI 19.9 or less, adult	0.885	0.896	0.861	0.896	0.881	0.812	0.890	0.918	0.901	0.824
V850 Driver, construction vehicle accident	0.850	0.871	0.850	0.891	0.912	0.711	0.877	0.913	0.893	0.810
Z992 Renal dialysis dependence	0.878	0.886	0.876	0.888	0.867	0.823	0.886	0.891	0.886	0.850
V433 Car occupant, collision nontraffic	0.863	0.894	0.885	0.881	0.858	0.805	0.890	0.915	0.885	0.871
E660 Obesity, excess calories	0.860	0.855	0.810	0.868	0.874	0.757	0.886	0.897	0.882	0.804
N186 End-stage renal disease	0.875	0.882	0.866	0.884	0.867	0.819	0.883	0.890	0.882	0.849
V422 Outside car, motorcycle collision	0.868	0.897	0.873	0.886	0.855	0.809	0.894	0.892	0.882	0.869

Table 22: Finetuning with a linear prediction head performance for the clinical deterioration labels, sorted by supervised S4 AUROC on the MIMIC (deterioration) dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
cardiac arrest	0.868	0.861	0.871	0.882	0.852	0.825	0.895	0.887	0.858	0.801
vasopressors	0.771	0.768	0.764	0.760	0.739	0.719	0.771	0.800	0.800	0.774
ecmo	0.724	0.798	0.759	0.748	0.712	0.586	0.841	0.819	0.788	0.783
mechanical ventilation	0.786	0.777	0.784	0.781	0.765	0.741	0.785	0.802	0.785	0.759
inotropes	0.654	0.727	0.646	0.716	0.711	0.642	0.721	0.720	0.776	0.706
severe hypoxemia	0.498	0.553	0.458	0.573	0.591	0.470	0.591	0.555	0.530	0.559

Table 23: Finetuning with a linear prediction head performance for the mortality labels, sorted by supervised S4 AUROC on the MIMIC (mortality) dataset.

1	084
1	085
1	086

Label	ECGFounder	ECGJEPA	ECGFMKED	MERL	STMEM	Hubert	ECGFM	ECGCPC	S4	Net1D
mortality 1d	0.829	0.837	0.895	0.843	0.807	0.825	0.849	0.876	0.835	0.794
mortality 7d	0.816	0.809	0.815	0.813	0.801	0.760	0.826	0.831	0.815	0.790
mortality 28d	0.811	0.811	0.795	0.808	0.791	0.739	0.814	0.823	0.807	0.780
mortality 365d	0.798	0.798	0.776	0.800	0.791	0.731	0.805	0.809	0.802	0.770
mortality 90d	0.791	0.795	0.770	0.791	0.779	0.719	0.800	0.803	0.796	0.765
mortality 180d	0.792	0.791	0.766	0.791	0.781	0.723	0.796	0.802	0.795	0.763
mortality stay	0.836	0.699	0.671	0.755	0.626	0.552	0.787	0.679	0.705	0.548

Table 24: Einstuning w

Table 24: Finetuning with a linear prediction head performance for the ICU labels, sorted by supervised S4 AUROC on the MIMIC (icu) dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
icu stay	0.751	0.747	0.741	0.749	0.738	0.713	0.755	0.757	0.748	0.726
icu 24h	0.744	0.738	0.734	0.739	0.729	0.706	0.745	0.749	0.742	0.716

Table 25: Finetuning with a linear prediction head performance for the sex label, sorted by supervised S4 AUROC on the MIMIC (sex) dataset.

Label	ECGFounder	ECG-JEPA	ECGFM-KED	MERL	ST-MEM	Hubert-ECG	ECG-FM	ECG-CPC	S4	Net1D
sex	0.913	0.904	0.883	0.916	0.903	0.810	0.919	0.933	0.919	0.869

Table 26: Finetuning with a linear prediction head performance for the age label based on MAE on the MIMIC (age) dataset.

Label	Units	ECGFounder	ECGJEPA	ECGFMKED	MERL	STMEM	Hubert	ECGFM	ECGCPC	S4	Net1D	Baseline
Age	Years	8.78	8.82	8.88	8.54	9.59	11.03	7.835	8.32	8.66	9.86	15.07

Table 27: Finetuning with a linear prediction head performance for the biometrics labels, sorted by supervised S4 MAE on the MIMIC (biometrics) dataset.

Label	Units	ECGFounder	ECGJEPA	ECGFMKED	MERL	STMEM	Hubert	ECGFM	ECGCPC	S4	Net1D	Baseline
Height	Inches	2.69	2.70	2.71	2.66	2.78	3.02	2.804	2.58	2.63	2.86	3.30
Weight	Lbs	27.85	28.27	28.34	27.33	29.79	32.04	26.417	26.38	27.53	30.16	35.73
BMI	kg/m2	4.14	4.12	4.23	4.05	4.40	4.67	3.943	3.90	4.06	4.43	5.30

Table 28: Finetuning with a linear prediction head performance for the ECG features labels, sorted by supervised S4 MAE on the MIMIC (ECG features) dataset

Label	Units	ECGFounder	ECGJEPA	ECGFMKED	MERL	STMEM	Hubert	ECGFM	ECGCPC	S4	Net1D	Baseline
RR	ms	107.01	106.59	106.78	108.28	108.98	108.93	112.913	105.73	105.94	109.96	154.38
QRS	ms	7.32	7.59	7.65	7.27	7.80	8.79	7.190	7.18	7.13	8.22	15.48
QT	ms	26.33	26.51	26.85	26.63	26.86	28.58	27.153	26.14	26.16	27.36	37.84
QTc	ms	18.09	17.89	18.16	18.36	17.91	20.67	17.657	17.41	17.60	19.48	26.84
P wave axis	0	16.38	16.40	16.55	16.21	18.44	20.03	16.200	15.81	15.97	17.46	22.47
QRS axis	0	15.07	15.41	16.01	15.86	20.59	18.22	15.849	15.35	15.35	16.41	36.06
T wave axis	0	24.36	24.19	24.75	24.59	28.31	30.95	24.648	24.05	23.94	25.34	35.41
PT	sec	2.84	2.84	2.94	2.86	2.92	3.04	2.824	2.80	2.82	2.93	3.85

Table 29: Finetuning with a linear prediction head performance for the laboratory values predictions, sorted by supervised S4 MAE on the MIMIC (laboratory values) dataset

Label	Units	ECGFounder	ECGJEPA	ECGFMKED	MERL	STMEM	Hubert	ECGFM	ECGCPC	S4	Net1D	Baseline
Albumin	g/dL	0.44	0.43	0.44	0.43	0.45	0.49	0.442	0.44	0.44	0.44	0.49
Anion Gap	mEq/L	2.35	2.32	2.34	2.32	2.38	2.39	2.330	2.32	2.31	2.34	2.39
Bicarbonate	mEq/L	2.66	2.65	2.67	2.68	2.69	2.72	2.993	2.64	2.67	2.68	2.71
Bilirubin, Total	mg/dL	0.55	0.56	0.56	0.55	0.55	0.56	0.551	0.56	0.55	0.55	0.67
Calcium, Total	mg/dL	0.49	0.49	0.49	0.50	0.50	0.51	0.547	0.50	0.49	0.49	0.51
Creatinine	mg/dL	0.39	0.39	0.40	0.39	0.40	0.40	0.399	0.39	0.38	0.39	0.48
Ferritin	ng/mL	325.91	214.00	239.76	232.05	548.12	582.04	826.690	308.41	206.29	239.27	297.79
Urea Nitrogen	mg/dL	8.15	8.11	8.25	8.08	8.30	8.45	8.415	7.98	8.05	8.33	9.94
Hematocrit	%	3.88	3.90	3.96	3.87	3.93	4.12	3.901	3.87	3.88	4.01	4.34
Hemoglobin	g/dL	1.39	1.39	1.42	1.38	1.41	1.48	1.442	1.38	1.39	1.44	1.59
Lymphocytes	%	8.55	8.57	8.63	8.57	8.70	8.89	8.450	8.48	8.56	8.78	9.37
MCHC	%	1.10	1.08	1.09	1.08	1.10	1.10	1.082	1.07	1.07	1.08	1.13
RDW	%	1.22	1.23	1.23	1.22	1.23	1.28	1.218	1.21	1.21	1.24	1.42
Red Blood Cells	m/uL	0.47	0.46	0.47	0.46	0.47	0.49	0.456	0.46	0.46	0.48	0.52
RDW-SD	fL	4.48	4.47	4.52	4.50	4.58	4.71	4.465	4.55	4.49	4.61	5.31
Creatine Kinase	IU/L	264.66	232.67	229.90	237.43	272.06	249.10	238.532	243.92	242.04	225.86	275.63
NTproBNP	pg/mL	3769.34	3729.16	4037.84	3766.55	3699.87	3730.16	3683.700	3593.63	3655.23	3564.98	4538.09

Table 30: Finetuning with a linear prediction head performance for the vital signs predictions, sorted by supervised S4 MAE on the MIMIC (vital signs) dataset

Label	Units	ECGFounder	ECGJEPA	ECGFMKED	MERL	STMEM	Hubert	ECGFM	ECGCPC	S4	Net1D	Baseline
dbp	mmHg	11.41	11.38	11.43	11.41	11.48	11.65	11.353	11.36	11.41	11.58	11.83
heartrate	bpm	11.61	11.67	11.68	11.68	11.72	11.81	11.788	11.57	11.57	11.74	15.00
o2sat	%	1.54	1.54	1.55	1.53	1.55	1.57	1.535	1.53	1.54	1.56	1.64
resprate	bpm	2.20	2.18	2.19	2.18	2.25	2.20	2.232	2.18	2.18	2.19	2.29
sbp	mmHg	17.26	17.32	17.31	17.21	17.81	17.98	17.102	17.16	17.27	17.65	18.36
temperature	°F	0.62	0.61	0.61	0.62	0.63	0.62	0.615	0.61	0.61	0.62	0.63

Table 31: Fit parameters for the scaling analysis

Model	C	α	L_0	R^2
ECGFounder (pretrained)	0.462	0.109	0.018	0.933
ECGFounder (from scratch)	0.887	0.270	0.120	0.998
ECG-JEPA (pretrained)	0.402	0.083	1.32×10^{-13}	0.989
ECG-CPC (pretrained)	0.463	0.104	4.35×10^{-7}	0.946
ECG-CPC (scratch)	0.501	0.101	9.13×10^{-10}	0.957
S4	0.677	0.206	0.089	0.983

Table 32: Label efficiency for different training datasets

Model	r(N=250)	r(N=500)	r(N=1000)	r(N=2000)
ECGFounder	0.30	0.40	0.51	0.62
ECG-JEPA	0.11	0.17	0.27	0.40
ECG-CPC	0.21	0.27	0.34	0.40