

Echo2ECG: Enhancing ECG Representations with Cardiac Morphology from Multi-View Echos

Michelle Espranita Liman^{1(✉)}, Özgün Turgut¹,
Alexander Müller², Eimo Martens²,
Daniel Rueckert^{1,3,4}, and Philip Müller¹

¹ Chair for AI in Healthcare and Medicine, Technical University of Munich (TUM)
and TUM University Hospital, Germany

{michelle.liman,oezguen.turgut,daniel.rueckert,philip.j.mueller}@tum.de

² Department of Cardiology, TUM University Hospital, Germany
{Alexander.Mueller,Eimo.Martens}@mri.tum.de

³ Department of Computing, Imperial College London, UK

⁴ Munich Center for Machine Learning (MCML), Germany

Abstract. Electrocardiography (ECG) is a low-cost, widely used modality for diagnosing electrical abnormalities like atrial fibrillation by capturing the heart’s electrical activity. However, it cannot directly measure cardiac morphological phenotypes, such as left ventricular ejection fraction (LVEF), which typically require echocardiography (Echo). Predicting these phenotypes from ECG would enable early, accessible health screening. Existing self-supervised methods suffer from a representational mismatch by aligning ECGs to *single-view* Echos, which only capture local, spatially restricted anatomical snapshots. To address this, we propose **Echo2ECG**, a multimodal self-supervised learning framework that enriches ECG representations with the heart’s morphological structure captured in *multi-view* Echos. We evaluate **Echo2ECG** as an ECG feature extractor on two clinically relevant tasks that fundamentally require morphological information: (1) classification of structural cardiac phenotypes across three datasets, and (2) retrieval of Echo studies with similar morphological characteristics using ECG queries. Our extracted ECG representations consistently outperform those of state-of-the-art unimodal and multimodal baselines across both tasks, despite being 18× smaller than the largest baseline. These results demonstrate that **Echo2ECG** is a robust, powerful ECG feature extractor. Our code is accessible at: <https://github.com/michelleespranita/Echo2ECG>.

Keywords: Multimodal Learning · ECG · Echocardiography

1 Introduction

Electrocardiography (ECG) is a cost-effective and ubiquitous tool for cardiac assessment. While it is highly effective at diagnosing electrical heart abnormalities, such as atrial fibrillation [1], it cannot directly quantify cardiac morphological phenotypes like left ventricular ejection fraction (LVEF). Traditionally,

these structural measurements require echocardiography (Echo), a modality often limited by cost, medical expertise, and accessibility. Recently, AI-based ECG models have demonstrated the ability to detect morphological phenotypes such as left-sided valvular heart disease [4] and hypertrophic cardiomyopathy [8], enabling early and accessible health screening.

Current methods for transferring morphological information from Echo to ECG primarily rely on supervised training using Echo-derived labels [2, 4, 8, 13]. However, they fail to leverage large-scale unlabeled data through self-supervised learning, which is critical for building robust ECG representations that could be used for other prediction tasks.

Recent multimodal self-supervised methods, most notably EchoingECG [5], attempt to bridge this gap by pre-training on ECG, Echo, and text-based discharge summaries. By aligning ECGs to *single-view* Echos (e.g., apical four-chamber/A4C), the approach attempts to map global electrical signals to local, spatially restricted anatomical snapshots. This creates a representational mismatch: a single-view Echo captures only a subset of the heart’s structure, whereas the ECG reflects the heart’s global electrical activity. Furthermore, by requiring paired text reports, the model’s scalability is restricted to datasets where clinical reports are available.

To address these limitations, we propose **Echo2ECG**, a multimodal self-supervised learning framework that enriches ECG representations with morphological information from *multi-view* Echo studies. By leveraging powerful unimodal encoders pre-trained on large datasets, our method enables distillation of a complete picture of the heart’s morphological structure into ECG representations without depending on paired text data. The resulting ECG encoder serves as a powerful ECG feature extractor for downstream tasks requiring morphological information. Our contributions are as follows:

- We propose **Echo2ECG**, a multimodal self-supervised learning framework that aligns global ECGs to multi-view Echo studies to learn rich, structure-aware ECG representations.
- We evaluate **Echo2ECG** as an ECG feature extractor for the classification of structural cardiac phenotypes across three datasets. Our extracted ECG features consistently outperform those of state-of-the-art unimodal and multimodal baselines, despite our model being 18× smaller than the largest baseline. This demonstrates that morphological information can be effectively distilled into a lightweight yet powerful ECG feature extractor.
- We demonstrate the alignment of our learned representation space by successfully retrieving Echo studies with similar morphological characteristics using only ECG queries.

2 Related Work

The increasing accessibility of large-scale multimodal datasets, such as MIMIC-IV [7] and the UK Biobank [17], has enabled research into cross-modal knowledge

transfer in the medical domain. Most multimodal learning frameworks are inspired by CLIP [14], which aligns modalities by pulling paired representations together and pushing unpaired ones apart in a joint embedding space. In the cardiac domain, MMCL [18] and PTACL [16] use this approach to distill cardiac morphological information from Cardiac MRI (CMR) into ECG. Most similar to our work, EchoingECG [5] aligns ECGs with single-view Echos and discharge reports. These studies consistently demonstrate that aligning ECGs with "gold standard" imaging modalities significantly improves the prediction of structural cardiac phenotypes solely based on ECG.

Concurrently, large-scale ECG foundation models such as xECG [11] and ECGFounder [9], and time-series foundation models such as OTIS [19], have emerged, utilizing self-supervised learning on massive ECG datasets like MIMIC-IV-ECG [6] and CODE [15]. We argue that while these unimodal models are powerful, their performance on structural cardiac phenotype prediction can be further enhanced through multimodal alignment. To this end, we initialize the ECG encoder in our framework with OTIS [19].

Our work diverges from existing literature in two key ways. First, we align ECGs with complete multi-view Echo studies, solving the representational mismatch that occurs when aligning ECGs with individual Echo views. Secondly, unlike EchoingECG, we focus strictly on ECG-Echo relationship by omitting the usage of text reports during pre-training. To our knowledge, this is the first study to investigate the direct alignment of ECGs with full multi-view Echo studies.

3 Method

The goal of our proposed method is to transfer rich cardiac morphological information from echocardiography into ECG representations. This enables the prediction of Echo-derived structural cardiac phenotypes, such as left ventricular ejection fraction (LVEF), using the widely accessible and low-cost ECG modality alone. As shown in Figure 1, our framework comprises four main components: (1) an ECG encoder that maps an ECG signal to an embedding; (2) an Echo encoder that maps each Echo view to an embedding; (3) an Echo view aggregation module that combines multiple view-level embeddings from the same study into a single study-level representation; and (4) modality-specific projection layers.

3.1 ECG and Echo Encoders

We initialize the ECG and Echo encoders using strong, pre-trained unimodal models to leverage large-scale prior knowledge. The ECG encoder is a 12-layer tiny transformer [20], initialized from OTIS [19], which was pre-trained on a multi-domain time-series corpus including MIMIC-IV-ECG [6]. Each ECG is represented by the mean token embedding. The Echo encoder is a MViTv2 [10], initialized from EchoPrime [21], which was trained via multimodal contrastive learning between Echo views and clinical reports from an in-house Cedars-Sinai dataset. Each Echo view is represented by the [CLS] token embedding.




Fig. 1. Overview of Echo2ECG. To address the representational mismatch in prior methods aligning ECGs with *single-view* Echos, our framework uses contrastive learning to align ECGs with *multi-view* Echo studies. This multi-view alignment distills comprehensive cardiac morphological information from Echos into ECG representations. We aggregate view-level Echo embeddings from a frozen, powerful Echo encoder into a study-level Echo embedding. During pre-training, we optimize only the ECG encoder, the Echo view aggregator, and modality-specific projection layers. The ECG encoder is then used to extract ECG representations for subsequent downstream tasks.

The ECG encoder, the Echo view aggregator, and the projection layers are fully trainable, while the Echo encoder is kept frozen. Overall, the framework has 12.5M trainable parameters and is trained for 50 epochs. Hyperparameter details are available in the code.

3.2 Echo View Aggregation

A single Echo view captures only part of the heart, e.g., the A4C view captures the four chambers of the heart but not the aortic valve, which is visible in the PSAX view. In clinical practice, cardiologists do not base diagnoses on a single view; instead, they integrate information across multiple views from a complete Echo study to obtain a comprehensive cardiac assessment. Similarly, by distilling information from the full Echo study rather than isolated views into an ECG, we ensure that the resulting ECG features are informed by the entire heart.

To achieve this, each Echo view is first encoded independently by the Echo encoder to produce view-level embeddings. These embeddings are projected from 512- to 1024-dim and aggregated via attention pooling into a single study-level embedding. This allows the model to weigh views by relevance while preserving information from the entire study.

3.3 Multimodal Contrastive Pre-Training

To transfer information from multi-view Echo studies to ECGs, we align the embeddings from both modalities within a shared 512-dim latent space using

Table 1. Internal pre-training dataset details from TUM Klinikum Rechts der Isar. Age is presented as Mean $\pm$ SD.

Split	#Pairs	#Patients	#ECGs	#Echo Studies	#Echo Views	Age	Sex (%Male)
Train	11,859	6,836	11,516	8,712	179,571	65.6 $\pm$ 16.1	63.1%
Val	1,248	565	1,213	934	20,633	65.3 $\pm$ 16.7	64.7%
Test	1,358	566	1,290	924	20,903	64.8 $\pm$ 16.6	64.1%

the CLIP [14] contrastive objective. After modality-specific projection layers, the ECG and Echo study embeddings are pulled together for matched pairs and pushed apart for unmatched pairs within a batch.

For a batch of size B , the ECG $\rightarrow$ Echo contrastive loss is defined as:

$$\mathcal{L}_{\text{ECG} \rightarrow \text{Echo}} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(\cos(z_{\text{ecg}}^i, z_{\text{echo}}^i)/\tau)}{\sum_{j=1}^B \exp(\cos(z_{\text{ecg}}^i, z_{\text{echo}}^j)/\tau)} \quad (1)$$

where τ denotes the temperature parameter, and z_{ecg} and z_{echo} are the ECG and Echo study embeddings after the projection layers respectively. An analogous Echo $\rightarrow$ ECG loss is computed by swapping the modalities. The final training objective is a weighted combination of the two directional losses.

4 Experimental Setup

4.1 Dataset and Pre-processing

We pre-train using paired ECG-Echo studies collected at our institution, TUM Klinikum Rechts der Isar. An ECG and an Echo study of a patient are a valid pair if they were acquired within 7 days. We split the data at the patient level (Table 1).

ECGs (12-lead, 10-second, 500 Hz) are processed following OTIS [19], including baseline drift removal using asymmetric least-squares smoothing and lead-group normalization. We use $T = 1008$ timesteps and apply the same augmentations as in OTIS.

Echo studies include up to 128 views. Following EchoPrime [21], we extract the Echo regions, remove annotations, and downsample the spatial resolution to 224. We sample 16 frames with stride 2, and zero-pad shorter clips. Each Echo view is normalized using EchoPrime train set statistics without augmentation.

4.2 Baselines

We compare our ECG model against two categories of baselines: (1) unimodal ECG models, such as OTIS (7.1M params) [19], xECG (57.0M) [11], and ECG-Founder (30.7M) [9]; and (2) multimodal ECG models, such as EchoingECG (126.6M) [5] and PTACL (19.3M) [16]. To ensure fair comparison, we follow the data pre-processing and training protocols in the original papers or codebases.

Table 2. Downstream datasets and label distributions. LVEF: 0/1/2 = reduced/mildly reduced/normal; SHD: 0/1 = absence/presence. These tasks assess how well the ECG representations capture morphological information.

Task	Dataset	Train (N)	Val (N)	Test (N)
LVEF clf (0/1/2)	Internal	294 (33/33/33%)	564 (11/16/73%)	562 (14/15/71%)
	EchoNext	300 (33/33/33%)	4,112 (17/7/76%)	4,827 (16/7/77%)
	UKB	255 (33/33/33%)	2,834 (1/4/95%)	2,827 (1/4/95%)
SHD clf (0/1)	EchoNext	72,475 (48/52%)	4,626 (57/43%)	5,442 (57/43%)

4.3 Evaluation Strategy

Structural Cardiac Phenotype Prediction We assess the quality of the extracted ECG representations on two downstream tasks which require morphological information: (1) left ventricular ejection fraction (LVEF) classification (reduced $\leq 40\%$, mildly reduced $40 - 50\%$, and normal $\geq 50\%$ [12]) and (2) structural heart disease (SHD) classification (absence vs. presence of moderate/greater disease [3]). For both tasks, we train a k -nearest neighbors classifier (kNN) with $k = 5$ on the ECG features. As kNN performance is deterministic, we report AUROC with 95% CI via 1,000 bootstrap resamples of the test set.

Downstream dataset details are summarized in Table 2. To assess generalization, models are evaluated across three cohorts on LVEF classification: our internal dataset, EchoNext [3], and UK Biobank (UKB) [17]. For SHD classification on EchoNext, we utilize the provided splits and evaluate the strength of the ECG embeddings by scaling training data from 0.1% to 100%.

Retrieval of Echo Studies with Similar Phenotypes using ECG Queries

We assess the learned joint embedding space via phenotype-aware cross-modal retrieval. Given a query ECG, Echo study embeddings are ranked by cosine similarity, and the top- k Echo studies are retrieved. A retrieved Echo study is a match if its phenotype value lies within $\pm 0.5\sigma$ of the query ECG’s phenotype, where σ denotes the standard deviation of that phenotype in the test set. For models without an Echo view aggregator, the Echo study embedding is calculated by averaging the embeddings of the Echo views belonging to the study. We evaluate four cardiac phenotypes: LV end-diastolic and end-systolic volumes (EDV, ESV), stroke volume (SV), and ejection fraction (EF); and report Precision@1 (i.e., $k = 1$) and mean rank (MnR) across the test set.

5 Results and Discussion

5.1 Echo2ECG Extracts Powerful Out-of-the-Box Features

To evaluate diagnostic utility, we assess Echo2ECG on (i) LVEF (Table 3) and (ii) SHD classification (Figure 2) against unimodal and multimodal baselines. Across both tasks, Echo2ECG consistently achieves superior performance.

Table 3. Best and second-best AUROC [95% CI] ($\uparrow$) for ECG-based LVEF classification. Echo2ECG provides high-quality, out-of-the box ECG features, achieving the highest AUROC on our internal dataset and on EchoNext. On UKB, it ranks second behind PTACL, which was pre-trained on UKB by pairing ECG with CMR—a superior modality to Echo for LVEF quantification.

Paired Modality	Method		Internal	External	
			Internal	EchoNext	UKB
-	OTIS	(7.1M)	0.688 [0.650, 0.727]	0.697 [0.680, 0.713]	0.637 [0.574, 0.694]
-	xECG	(57.0M)	0.732 [0.696, 0.769]	0.695 [0.677, 0.711]	0.531 [0.474, 0.591]
-	ECGFounder	(30.7M)	0.746 [0.710, 0.781]	0.686 [0.670, 0.702]	0.653 [0.592, 0.708]
Text+Echo	EchoingECG	(126.6M)	0.740 [0.703, 0.776]	0.676 [0.661, 0.692]	0.604 [0.541, 0.662]
CMR	PTACL	(19.3M)	0.723 [0.686, 0.759]	0.704 [0.689, 0.719]	0.729 [0.677, 0.770]
Echo	Echo2ECG	(7.1M)	0.785 [0.750, 0.819]	0.723 [0.708, 0.738]	0.692 [0.642, 0.736]

For LVEF classification, a kNN classifier trained on Echo2ECG’s ECG features achieves the highest AUROC on our internal dataset and on EchoNext, improving AUROC by 5.2% and 2.7% respectively over the second-best model on each dataset. On UKB, our model ranks second behind PTACL, which was pre-trained on UKB by pairing ECG with CMR—a superior modality to Echo for LVEF quantification.

For SHD classification, a kNN classifier trained on just 1% of the ECG features extracted by Echo2ECG outperforms nearly all other models trained on the full 100% dataset (Figure 2(a)), showing that Echo2ECG provides robust ECG features that work well under low-data regimes. Remarkably, a kNN trained on just 0.1% of the ECG features extracted by Echo2ECG outperforms all competing models trained on the same amount of data—including EchoingECG—while Echo2ECG remains 18 $\times$ smaller (Figure 2(b)).

5.2 Multi-View Alignment Enables Accurate Morphological Cross-Modal Retrieval

In phenotype-aware ECG-to-Echo retrieval (Table 4), Echo2ECG achieves the highest Prec@1 and lowest MnR. Our study-level, multi-view alignment outperforms single-view alignment, highlighting the importance of aligning global electrical signals with comprehensive anatomy to encode morphological structure. Interestingly, EchoingECG performs worse than random retrieval, likely because its contrastive objective prioritizes ECG-Text over ECG-Echo alignment [5], failing to directly link the heart’s electrical activity with its morphological structure.




Fig. 2. AUROC (↑) for ECG-based SHD classification. (a) A kNN classifier trained on just 1% of data using ECG representations extracted by Echo2ECG outperforms most models trained on 100% of the data, meaning that Echo2ECG provides robust ECG features that work well under low-data regimes. (b) Echo2ECG outperforms EchoingECG at 0.1% training data, despite being 18× smaller.

Table 4. Cross-modal retrieval performance for structural cardiac phenotypes. Multi-view alignment allows Echo2ECG to accurately retrieve Echo studies from ECG queries.

Method	LVEDV		LVESV		LVSV		LVEF	
	P@1↑	MnR↓	P@1↑	MnR↓	P@1↑	MnR↓	P@1↑	MnR↓
Random retrieval	0.296	5.915	0.366	6.081	0.323	4.570	0.295	7.230
EchoingECG	0.241	10.111	0.337	10.998	0.325	5.320	0.232	13.035
Echo2ECG (<i>single-view</i> Echo)	0.332	5.376	0.517	4.366	0.334	5.125	0.415	5.346
Echo2ECG (<i>multi-view</i> Echo)	0.387	4.358	0.517	3.841	0.379	4.193	0.404	5.101

Table 5. Ablation of Echo2ECG components. Mean AUROC [95% CI] (↑) is reported. Multi-view alignment and attention pooling improve the quality of ECG features.

Alignment	View Agg.	LVEF clf (Internal)	SHD clf (EchoNext 1%)
ECG ↔ <i>single-view</i> Echo	-	0.747	0.743
ECG ↔ <i>multi-view</i> Echo	mean pool	0.775	0.752
ECG ↔ <i>multi-view</i> Echo	CLS pool	0.756	0.754
ECG ↔ <i>multi-view</i> Echo	attention pool	0.785	0.757

5.3 Learnable Pooling of Multi-View Echo Benefits Alignment

Our ablation study (Table 5) shows that multi-view alignment consistently outperforms single-view alignment, with attention pooling proving most effective at distilling key morphological features into ECG representations.

6 Conclusion

We present Echo2ECG, a self-supervised framework that enriches ECG representations with cardiac morphology from multi-view Echo studies. Particularly, we

resolve the representational mismatch of prior single-view methods and eliminate reliance on paired text data. Across three datasets, our lightweight ECG model serves as a powerful feature extractor, outperforming state-of-the-art baselines in structural phenotyping and morphology-aware cross-modal Echo retrieval. Remarkably, it achieves this while being $18\times$ smaller than the largest baseline.

Our work has several limitations. First, we do not account for beat- or phase-level synchronization between ECGs and Echos, which may limit the model’s ability to capture temporal relationships between electrical activity in ECGs and mechanical activity in Echos. Second, although we do not rely on text reports, our method requires paired ECG–Echo studies, which may be scarce at scale.

Acknowledgments. This study was approved by the Ethics Committee of TUM Klinikum Rechts der Isar (reference number 2025-395-S-CB, application dated July 13, 2025). This research has been conducted using the UK Biobank Resource under Application Number 87802.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Brant, L.C., Ribeiro, A.H., Eromosele, O.B., Pinto-Filho, M.M., Barreto, S.M., Duncan, B.B., Larson, M.G., Benjamin, E.J., Ribeiro, A.L., Lin, H.: Prediction of atrial fibrillation from the ecg in the community using deep learning: A multinational study. *Circulation: Arrhythmia and Electrophysiology* **18**(10), e013734 (2025). <https://doi.org/10.1161/CIRCEP.125.013734>, <https://www.ahajournals.org/doi/abs/10.1161/CIRCEP.125.013734>
2. Cohen-Shelly, M., Attia, Z.I., Friedman, P.A., Ito, S., Essayagh, B.A., Ko, W.Y., Murphree, D.H., Michelena, H.I., Enriquez-Sarano, M., Carter, R.E., Johnson, P.W., Noseworthy, P.A., Lopez-Jimenez, F., Oh, J.K.: Electrocardiogram screening for aortic valve stenosis using artificial intelligence. *European Heart Journal* **42**(30), 2885–2896 (03 2021). <https://doi.org/10.1093/eurheartj/ehab153>, <https://doi.org/10.1093/eurheartj/ehab153>
3. Elias, P., Finer, J.: EchoNext: A Dataset for Detecting Echocardiogram-Confirmed Structural Heart Disease from ECGs. *PhysioNet* (Sep 2025). <https://doi.org/10.13026/3ykd-bf14>, <https://doi.org/10.13026/3ykd-bf14>, version 1.1.0
4. Elias, P., Poterucha, T.J., Rajaram, V., et al.: Deep learning electrocardiographic analysis for detection of left-sided valvular heart disease. *JACC* **80**(6), 613–626 (2022). <https://doi.org/10.1016/j.jacc.2022.05.029>, <https://www.jacc.org/doi/abs/10.1016/j.jacc.2022.05.029>
5. Gao, Y., Kim, S., McIntosh, C.: EchoingECG: An Electrocardiogram Cross-Modal Model for Echocardiogram Tasks . In: proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. vol. LNCS 15964. Springer Nature Switzerland (September 2025)
6. Gow, B., Pollard, T., Nathanson, L.A., Johnson, A., Moody, B., Fernandes, C., Greenbaum, N., Waks, J.W., Eslami, P., Carbonati, T., Chaudhari, A., Herbst, E., Moukheiber, D., Berkowitz, S., Mark, R., Horng, S.: MIMIC-IV-ECG: Diagnostic Electrocardiogram Matched Subset. *PhysioNet* (Sep 2023). <https://doi.org/10.13026/4nqg-sb35>, <https://doi.org/10.13026/4nqg-sb35>, version 1.0

7. Johnson, A., Bulgarelli, L., Pollard, T., Gow, B., Moody, B., Horng, S., Celi, L.A., Mark, R.: MIMIC-IV. PhysioNet (Oct 2024). <https://doi.org/10.13026/kpb9-mt58>, <https://doi.org/10.13026/kpb9-mt58>, version 3.1
8. Ko, W.Y., Siontis, K.C., Attia, Z.I., Carter, R.E., Kapa, S., Ommen, S.R., Demuth, S.J., Ackerman, M.J., Gersh, B.J., Arruda-Olson, A.M., Geske, J.B., Asirvatham, S.J., Lopez-Jimenez, F., Nishimura, R.A., Friedman, P.A., Noseworthy, P.A.: Detection of hypertrophic cardiomyopathy using a convolutional neural network-enabled electrocardiogram. *JACC* **75**(7), 722–733 (2020). <https://doi.org/10.1016/j.jacc.2019.12.030>, <https://www.jacc.org/doi/abs/10.1016/j.jacc.2019.12.030>
9. Li, J., Aguirre, A.D., Junior, V.M., Jin, J., Liu, C., Zhong, L., Sun, C., Clifford, G., Brandon Westover, M., Hong, S.: An electrocardiogram foundation model built on over 10 million recordings. *NEJM AI* **2**(7), A10a2401033 (2025)
10. Li, Y., Wu, C.Y., Fan, H., Mangalam, K., Xiong, B., Malik, J., Feichtenhofer, C.: Mvitv2: Improved multiscale vision transformers for classification and detection. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4794–4804 (2022). <https://doi.org/10.1109/CVPR52688.2022.00476>
11. Lunelli, R., Nicolson, A., Pröll, S.M., Reinstadler, S.J., Bauer, A., Dlaske, C.: Bencheeg and xecg: a benchmark and baseline for ecg foundation models (2025), <https://arxiv.org/abs/2509.10151>
12. McDonagh, T.A., Metra, M., Adamo, M., et al.: 2021 esc guidelines for the diagnosis and treatment of acute and chronic heart failure: Developed by the task force for the diagnosis and treatment of acute and chronic heart failure of the european society of cardiology (esc) with the special contribution of the heart failure association (hfa) of the esc. *European Heart Journal* **42**(36), 3599–3726 (08 2021). <https://doi.org/10.1093/eurheartj/ehab368>, <https://doi.org/10.1093/eurheartj/ehab368>
13. Poterucha, T.J., Jing, L., Ricart, R.P., et al.: Detecting structural heart disease from electrocardiograms using ai. *Nature* **644**(8075), 221–230 (Aug 2025). <https://doi.org/10.1038/s41586-025-09227-0>, <https://doi.org/10.1038/s41586-025-09227-0>
14. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: Meila, M., Zhang, T. (eds.) *Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 139, pp. 8748–8763. PMLR (18–24 Jul 2021), <https://proceedings.mlr.press/v139/radford21a.html>
15. Ribeiro, A.L.P., Paixão, G.M.M., Gomes, P.R., Ribeiro, M.H., Ribeiro, A.H., Canazart, J.A., Oliveira, D.M., Ferreira, M.P., Lima, E.M., Moraes, J.L.d., Castro, N., Ribeiro, L.B., Macfarlane, P.W.: Tele-electrocardiography and bigdata: The CODE (clinical outcomes in digital electrocardiography) study. *J. Electrocardiol.* **57S**, S75–S78 (Nov 2019)
16. Selivanov, A., Müller, P., Turgut, Ö., Stolt-Ansó, N., Rueckert, D.: Global and Local Contrastive Learning for Joint Representations from Cardiac MRI and ECG . In: *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2025*. vol. LNCS 15960. Springer Nature Switzerland (September 2025)
17. Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., Liu, B., Matthews, P., Ong, G., Pell, J., Silman, A., Young, A., Sprosen, T., Peakman, T., Collins, R.: UK biobank: an open access

- resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Med.* **12**(3), e1001779 (Mar 2015)
18. Özgün Turgut, Müller, P., Hager, P., Shit, S., Starck, S., Menten, M.J., Martens, E., Rueckert, D.: Unlocking the diagnostic potential of electrocardiograms through information transfer from cardiac magnetic resonance imaging. *Medical Image Analysis* **101**, 103451 (2025). <https://doi.org/https://doi.org/10.1016/j.media.2024.103451>, <https://www.sciencedirect.com/science/article/pii/S1361841524003785>
 19. Özgün Turgut, Müller, P., Menten, M.J., Rueckert, D.: Towards generalisable time series understanding across domains (2025), <https://arxiv.org/abs/2410.07299>
 20. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 30. Curran Associates, Inc. (2017), https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
 21. Vukadinovic, M., Chiu, I.M., Tang, X., Yuan, N., Chen, T.Y., Cheng, P., Li, D., Cheng, S., He, B., Ouyang, D.: Comprehensive echocardiogram evaluation with view primed vision language ai. *Nature* (Nov 2025). <https://doi.org/10.1038/s41586-025-09850-x>, <https://doi.org/10.1038/s41586-025-09850-x>