

Cross-Modal Contrastive Learning of ECG and Angiography Representations for Severe Stenosis Classification

Nikola Cenikj^{1,4,*} , Özgün Turgut¹ , Alexander Müller⁴, Alexander Steger⁴, Jan Kehr⁴, Marcus Brugger⁴, Daniel Rueckert^{1,2,3} , Eimo Martens⁴ , and Philip Müller¹ 

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

² Department of Computing, Imperial College London, UK

³ Munich Center for Machine Learning (MCML), Munich, Germany

⁴ Department of Internal Medicine, TUM University Hospital, Munich, Germany
`nikola.cenikj@tum.de`

Abstract. Coronary artery stenosis is a common cardiovascular disease, with severe, untreated cases posing significant risks of heart attack. Although coronary (X-ray) angiograms remain the standard for stenosis diagnosis, they are invasive, time- and resource-intensive, and therefore only performed on patients with a high probability of disease based on symptoms and prior clinical tests. However, a subset of patients, especially those without symptoms, may remain undiagnosed. Detecting indications of stenosis from ECGs, which are fast, cheap, non-invasive, and thus routinely acquired even in asymptomatic patients, would support early diagnosis. However, as no reliable stenosis-specific signal has been identified in ECGs, they can not currently be used for stenosis risk stratification. To address this, we introduce *StenCE*, a pretraining framework, allowing stratification of patients based on features derived directly from ECGs. Evaluations across varying stenosis severity thresholds and additional ECG disease classification tasks demonstrate consistent performance improvements across different ECG encoders, outperforming previous work. The obtained models successfully detect signals for stenosis diagnosis in ECGs and are the first to achieve high performance in severe stenosis classification. The source code is available at <https://github.com/NikolaCenic/ecg-stenosis-cls>.

Keywords: Deep-Learning · Multi-modal Pretraining · Coronary Artery Stenosis · Coronary X-ray Angiography · Electrocardiogram.

1 Introduction

Coronary artery stenosis is a common cardiovascular disease that gets progressively more severe over time. As severe cases often lead to heart failure, early




Fig. 1: Overview of the proposed approach: Multi-view angiography and 12-lead ECG from the same patient are encoded with transformer-based modality-specific encoders. The ECG encoder is pretrained to extract features aligned with the frozen angiography encoder, trained for coronary stenosis classification. The pretrained ECG encoder is then fine-tuned for coronary stenosis and additional cardiac abnormalities diagnosed from ECGs.

diagnosis is critical for improving the survival rate. Diagnosis of severe stenosis is typically performed through coronary (X-ray) angiograms, where multiple angiography views capture different segments of the coronary arteries. However, X-ray angiography is an invasive procedure with a small mortality risk and is typically reserved for patients with a high likelihood of stenosis based on symptoms and clinical tests. Consequently, asymptomatic patients may remain undiagnosed while the disease progresses. Identifying stenosis indicators from a fast, non-invasive, and routinely-acquired modality such as ECG would enable early diagnosis, even for asymptomatic patients. However, despite being used for diagnosing different cardiovascular diseases, ECGs provide limited information for coronary stenosis, as patients with severe stenosis often have normal ECGs [5].

In this work, we aim to detect signals of stenosis in ECGs and thus enable the early identification of stenosis risks. To achieve this, we develop a deep learning based stenosis classifier working only on ECG inputs. To train this model, we rely on information distillation from an X-ray angiogram encoder. More precisely, as shown in Fig. 1, we employ multi-modal contrastive learning between an ECG-encoder and an angiography encoder trained for stenosis classification, followed by fine-tuning of the pretrained ECG model. Our contributions are as follows:

1. We propose *StenCE*, a contrastive pretraining framework that aligns ECG representations with features from a multi-view angiography stenosis classification model, thereby enabling an ECG encoder to detect stenosis signals from ECGs alone.
2. Our evaluation on clinical stenosis classification demonstrates an AUC of 0.822 for the most severe cases, showcasing that severe stenosis can be identified solely from ECGs, enabling early stenosis detection.
3. Evaluations among multiple stenosis severities as well as on additional cardiac abnormalities on the EchoNext dataset, further demonstrate the utility of our pretraining framework.

Table 1: Distribution of patients, number of pairs, and stenosis samples based on different severity thresholds across train, validation, and test splits.

Split	Patients	Pairs	Samples per Severity Category					
			$=0 \geq 100$	$=0 \geq 90$	$=0 \geq 70$	$\leq 50 \geq 99$	$\leq 25 \geq 70$	$< 70 \geq 70$
Train	3001	3587	762 533	762 1338	762 1699	1888 856	1427 1699	1888 1699
Val	204	234	41 27	41 81	41 112	122 46	83 112	122 112
Test	209	246	40 32	40 98	40 127	119 57	81 127	119 127

2 Related Work

Coronary Artery Stenosis. Most existing stenosis detection work focuses on single angiographic views, without considering different views of a patient [1, 6, 10, 12, 17, 20]. While [2] evaluates at patient-level, [3] is the only study training on the patient-level, using all available views. ECGs, however, have been part of general time-series pretraining [19] as well as ECG-specific pretraining [13, 14], but coronary stenosis diagnosis from ECGs remains challenging [5]. In [21], 12-lead ECGs are used to classify stenosis (≤ 50 | ≥ 99), yielding an AUC of 0.654. Performance increases to 0.717 when using only clinical risk factors and further to 0.847 when combining the two modalities. The work in [22] trains a model for ECG stenosis classification (< 70 | ≥ 70), and achieves an AUC of 0.57. Both [21] and [22] show that ECGs alone are not sufficient for reliable stenosis diagnosis. However, neither of them leverages cross-modal pretraining for the extraction of stenosis-relevant features out of the ECG signal.

Cross-modal knowledge transfer. Cross-modal knowledge transfer closes the gap between two modalities by aligning their feature representations. Initially designed for contrastive language-image pairing (CLIP) [15], it has been widely adopted in the medical domain, for tasks including radiology report generation [4], anomaly detection [11], and biomedical foundation models [23]. In cardiology, cross-modal pretraining has been used to align ECG representations with modalities such as echocardiograms [9] and cardiac MRIs [16, 18]. However, no prior work has explored contrastive pretraining between angiography and ECG pairs, despite them being suitable for cross-modal alignment.

3 Materials and Methods

3.1 Datasets

ECG-Angio Stenosis Dataset. We use a hospital dataset of 3,414 patients (70% male, with a mean BMI of 27), who underwent a X-ray angiogram and have a paired 10-second 12-lead ECGs, captured up to 2 days prior to the angiogram. As some patients have multiple such ECGs, the total number of matched pairs is 4,067. For each patient, we have annotations for the blockage in the segments of the coronary arteries [8], extracted from structured reports done by clinicians,




Fig. 2: Given paired ECG and coronary angiography data from a single patient, both modalities are encoded and projected into a shared embedding space using transformer-based encoders and linear projections. Using contrastive loss, the ECG encoder learns to extract stenosis-relevant features, which are then utilized in fine-tuning for detecting severe stenosis using only a 12-lead ECG signal.

used for clinical decision making. The blockage annotations are categorized into: 0%, 25%, 50%, 70%, 90%, 99%, and 100%, and a patient-level label is obtained as the maximum segment-level severity. We split the patients into disjoint train, validation, and test sets, with the distribution within each split shown in Table 1. **EchoNext Dataset.** To further assess the ECG encoders, we also evaluate them on EchoNext [7], a public dataset of ECGs with labels for different cardiac abnormalities derived from echocardiograms. We evaluate on the following binary tasks: left ventricular ejection fraction, pulmonary artery systolic pressure, moderate aortic stenosis, moderate mitral regurgitation, and moderate right ventricular systolic dysfunction.

3.2 Method

We propose *StenCE*, a stenosis-driven contrastive pretraining framework for a 12-lead ECG encoder. Given a 12-lead ECG paired with coronary angiography data for patient i , each modality is processed with a dedicated encoder. During pretraining, the ECG encoder is trained to extract features aligned with those produced by the angiography encoder using a contrastive loss, after which the ECG encoder is fine-tuned for downstream tasks, including stenosis classification.

ECG Encoders: We explore two different options for the ECG encoder: (i) *OTIS* [19], a 7.1M-parameter transformer pretrained on time-series data including ECGs, i.e. without any cross-modal knowledge transfer, and (ii) *EchoingECG* [9], a 128M-parameter 1D-ResNet pretrained on triplets of echocardiogram-text-ECG, i.e. with knowledge transfer from echocardiograms.

Angiogram Encoder: The angiogram encoder is based on the SegmentMIL [3] model. It consists of an angio view encoder, encoding individual views of coronary

X-ray angiogram studies, and a views aggregator component to compute study-level predictions. The view aggregator is based on a transformer decoder with query tokens for study-level, artery-level, and segment-level stenosis predictions. It was pretrained jointly on all these three tasks.

Knowledge Transfer through Contrastive Alignment: To transfer the knowledge from the angiography encoder to the ECG encoder, we follow the approach in [18] on pairs of individual ECGs and the associated angiography studies. More precisely, we apply the CLIP loss [15] on pairs of ECG and angiography features projected into a shared space using linear projections. To compute study-level angiography features, we introduce a new **CLS** query token into the views aggregator (extending the existing stenosis query tokens) and take its output features. We keep the angio view encoder frozen, while training the views aggregator, the ECG encoder, and the projections.

Supplementary Stenosis Loss: During contrastive alignment, we also keep training study-level, artery-level, and segment-level stenosis classification on the respective query tokens of the angiography views aggregator, following the procedure from the original SegmentMIL [3]. We balance this loss with the CLIP loss using the coefficient $W_{\text{Sten}} = 0.3$.

Fine-Tuning: After pretraining, we only keep the ECG encoder and fine-tune it on stenosis classification. Here we use only the ECG samples but derive their stenosis targets from their associated angiography studies. To address label imbalance, we use the class-weighted binary cross-entropy loss.

3.3 Experiments

Baselines. We group the baselines into single- and cross-modal pretrained ECG encoders. As single-modality approaches, we consider *OTIS* [19], *xECG* [14], and *ECGFounder* [13]. OTIS is pretrained on general time-series data, including ECGs, while xECG and ECGFounder are specifically pretrained using ECGs. As cross-modal pretrained models, we include *EchoingECG* [9] and *PTACL* [16], both pretrained using contrastive learning. EchoingECG uses pairs of echocardiograms, text and ECGs, while PTACL focuses on pairs of ECGs and CMRs.

Evaluation Setup. We evaluate the learned ECG encoders using linear probing and full fine-tuning on stenosis classification across multiple severity thresholds. Additionally, we evaluate the encoders on EchoNext’s tasks, formulated as multi-label classification. We report the mean AUC and the 95% confidence interval computed over three independent random seeds. We analyze the statistical significance of the performance differences using Welch’s t-test ($p < 0.05$).

4 Results

In Table 2, we report the results under both full fine-tuning and linear probing for EchoingECG and OTIS backbones, pretrained with StenCE, referred to as EchoingECG-StenCE and OTIS-StenCE. We compare them against OTIS, EchoingECG, as well as the other ECG baselines.

Table 2: Comparison of OTIS-StenCE and EchoingECG-StenCE encoders against OTIS, EchoingECG, and other ECG baselines under both full fine-tuning and linear probing. We report mean AUC with 95% confidence intervals over three seeds for stenosis classification at different severity thresholds, and mean AUC across five EchoNext tasks. The best results for OTIS and EchoingECG encoders are **bold**. The results indicate that our models have strong performance for stenosis classification, particularly at the highest severity threshold ($= 0 \mid = 100$). Furthermore, StenCE pretraining consistently boosts performance for both OTIS and EchoingECG across linear probing and full fine-tuning.

Model	Training	Stenosis AUCs		EchoNext
		$= 0 \mid = 100$	$= 0 \mid \geq 90$	Mean AUC
xECG	ECG	0.672	0.646	0.774
ECGFounder		0.664, 0.683	0.625, 0.663	0.772, 0.777
PTACL	ECG+CMR	0.623	0.682	0.771
		0.588, 0.673	0.661, 0.693	0.768, 0.773
EchoingECG	ECG+Echo+Text	0.530	0.489	0.766
		0.520, 0.536	0.486, 0.492	0.763, 0.769
EchoingECG	ECG+Echo+Text	0.719	0.704	0.765
		0.705, 0.734	0.696, 0.717	0.760, 0.768
EchoingECG - StenCE	ECG+Angio	0.822	0.668	0.770
		0.817, 0.828	0.640, 0.696	0.764, 0.774
OTIS	ECG	0.650	0.661	0.780
		0.624, 0.672	0.651, 0.673	0.778, 0.782
OTIS - StenCE	ECG+Angio	0.692	0.667	0.783
		0.678, 0.702	0.650, 0.680	0.780, 0.784

(a) Fine-tuning Performance

Model	Training	Stenosis AUCs		EchoNext
		$= 0 \mid = 100$	$= 0 \mid \geq 90$	Mean AUC
EchoingECG	ECG+Echo+Text	0.643	0.625	0.667
		0.636, 0.646	0.609, 0.634	0.662, 0.67
EchoingECG - StenCE	ECG+Angio	0.816	0.637	0.739
		0.806, 0.829	0.636, 0.638	0.738, 0.740
OTIS	ECG	0.592	0.503	0.678
		0.585, 0.599	0.493, 0.510	0.672, 0.683
OTIS - StenCE	ECG+Angio	0.676	0.625	0.695
		0.674, 0.677	0.624, 0.628	0.685, 0.710

(b) Linear probing Performance

Our models are capable of detecting clear diagnostic signals of severe stenosis from ECGs alone. The full fine-tuning results in Table 2a demon-




Fig. 3: Comparison of the fine-tuned performance on different stenosis thresholds. While, as expected, the detection performance drops when trying to identify less severe cases, our StenCE framework still enables the separation of severe cases from both healthy (0% blockage) and mild (50% blockage) cases.

strate strong performance at the most severe stenosis threshold, achieving an AUC of 0.822 for the $(= 0 \mid = 100)$ setting. Our model is the first to achieve such high performance on ECG-based stenosis classification, confirming its ability to extract a strong diagnostic signal directly from ECGs, thereby supporting early severe stenosis detection.

StenCE pretraining yields performance improvement in the majority of tasks for both OTIS and EchoingECG. In full fine-tuning setting, OTIS-StenCE significantly outperforms OTIS by 4% AUC on the severity threshold $(= 0 \mid = 100)$, and insignificantly on $(= 0 \mid \geq 90)$. For EchoingECG, fine-tuning EchoingECG-StenCE leads to an 11% AUC improvement on the $(= 0 \mid = 100)$ threshold, achieving the best overall performance among all models. However, on the $(= 0 \mid \geq 90)$ threshold, EchoingECG outperforms EchoingECG-StenCE by 4% AUC. On the EchoNext tasks, all models achieve comparable fine-tuning results due to the large scale of the EchoNext dataset. Nevertheless, for both OTIS and EchoingECG, the StenCE variants retain an advantage over their base versions. The benefit of StenCE pretraining is even more pronounced in the linear probing setting (Table 2b), where for both OTIS and EchoingECG, the StenCE variants significantly outperform the base models on each task. The largest gain is observed for EchoingECG-StenCE, with a 17% AUC improvement over EchoingECG on the $(= 0 \mid = 100)$ threshold. Similarly, OTIS-StenCE improves over OTIS by 8% and 12% AUC on the $(= 0 \mid = 100)$ and $(= 0 \mid \geq 90)$ thresholds, respectively. Similarly, on EchoNext, the StenCE models in linear probing, outperform OTIS by 2% and EchoingECG by 7% AUC.

Detection performance drops for less severe cases. In Fig. 3, we compare the fine-tuned performance on different stenosis severity thresholds. Starting from an AUC of 0.822 for the most severe cases, performance decreases to 0.704 for the less severe $(= 0 \mid \geq 90)$ threshold, and approaches random performance for the hardest to diagnose stenosis with severity of $(< 70 \mid \geq 70)$. The performance at this threshold indicates that the model is not yet suitable for clinical use. This trend is consistent with prior findings in [21] and [22]. Although their models and datasets are not publicly available, which prevents direct comparison, our results indicate consistent improvements. In [22], an AUC of 0.57 is

Table 3: Ablation study on architectural design and stenosis supervision in StenCE. We investigate using a frozen angiography encoder and no CLS token, as well as different values of W_{Sten} . Experiments are conducted in linear probing using OTIS- and EchoingECG-based ECG encoders. We report the AUC for severe stenosis and the EchoNext tasks. The best performing model in each category is **bold**. Results show that unfrozen angiography encoder with a CLS token with a moderate stenosis supervision ($W_{\text{Sten}} = 0.3$) yields the best performance.

Ablation	OTIS		EchoingECG	
	= 0 = 100	EchoNext	= 0 = 100	EchoNext
Frozen Angio Encoder & no CLS token	0.686	0.691	<u>0.766</u>	0.727
$W_{\text{Sten}} = 0$	0.665	<u>0.692</u>	0.725	0.736
$W_{\text{Sten}} = 0.3$	<u>0.676</u>	0.695	0.816	0.739
$W_{\text{Sten}} = 1$	0.659	0.686	0.706	<u>0.738</u>

reported for the (< 70 | ≥ 70) threshold, whereas our model surpasses this by more than 5%. Similarly, in [21], the ECG-only model achieves an AUC of 0.654 for the (≤ 50 | ≥ 99) threshold, while our model attains 0.687 on the same threshold using our dataset.

Introducing stenosis supervision during StenCE pretraining improves performance. Table 3 shows an ablation analyzing the impact of architectural choices and the stenosis supervision in StenCE. We evaluate using OTIS and EchoingECG as ECG encoder backbones, assessing performance on stenosis classification at the ($= 0$ | $= 100$) severity threshold and on EchoNext. For architecture, we consider a setting in which the angiography encoder is frozen and, instead of introducing a CLS token, we use the study-level representation. To examine the effect of stenosis supervision, we vary the weight W_{Sten} (0, 0.3, and 1). Overall, incorporating stenosis supervision with a moderate weight ($W_{\text{Sten}} = 0.3$), while still keeping the pretraining primarily driven by the CLIP loss, consistently improves performance. This configuration achieves the best results in three out of four evaluations, and only ranks second, with a 1% difference, in stenosis classification using OTIS.

5 Discussion and Conclusion

Limitations. This study has three main limitations. First, each patient in our dataset underwent an X-ray angiogram, performed due to stenosis symptoms. While this is a selection bias that should be controlled for, such clinical studies might require applying X-ray angiograms on healthy patients, exposing them to severe risks. Second, although we use a large internal cohort collected across multiple devices, we do not have access to datasets from other institutions, which limits our ability to perform external validation. Last, our models rely only on the ECG signal, and unlike prior work that uses clinical risk factors [21], we do

not integrate complementary data sources. Incorporating such data may further improve the performance.

Conclusion. Current models for stenosis classification rely on using backbones pretrained on ECG-specific tasks, limiting their ability to capture stenosis-specific information. We overcome this by introducing cross-model pretraining with paired ECG–angiography data, enabling the ECG encoder to encode stenosis-relevant features seen in angiography. Although the obtained performances are still not sufficient for usage in clinical practice, our results demonstrate the benefit of cross-modal pretraining and enable future developments towards preliminary assessment of coronary artery stenosis, supporting early diagnosis.

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).

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

References

1. H. Lin, T. Liu, A. Katsaggelos, and A. Kline : StenUNet: Automatic Stenosis Detection from X-ray Coronary Angiography (2023). <https://doi.org/10.48550/arXiv.2310.14961>
2. C. Cong, Y. Kato, H. D. Vasconcellos, M. Ostovaneh, J. Lima, and B. Ambale-Venkatesh: Deep learning-based end-to-end automated stenosis classification and localization on catheter coronary angiography. *Frontiers in Cardiovascular Medicine* **10**, 944135 (02 2023). <https://doi.org/10.3389/fcvm.2023.944135>
3. Cenikj, N., Özgün Turgut, Müller, A., Steger, A., Kehrner, J., Brugger, M., Rueckert, D., Martens, E., Müller, P.: Multi-view stenosis classification leveraging transformer-based multiple-instance learning using real-world clinical data (2026). <https://doi.org/10.48550/arXiv.2602.02067>
4. Chen, Y., Du, C., Li, C., Hu, J., Shi, Y., Xiong, S., Zhu, X.X., Mou, L.: Uni-crossadapter: Multimodal adaptation of clip for radiology report generation (2025). <https://doi.org/10.48550/arXiv.2503.15940>
5. Crawford, M.H., Mendoza, C.A., O’Rourke, R.A., White, D.H., Boucher, C.A., , Gorwit, J.: Limitations of continuous ambulatory electrocardiogram monitoring for detecting coronary artery disease. *Annals of Internal Medicine* **89**(1), 1–5 (1978). <https://doi.org/10.7326/0003-4819-89-1-1>, PMID: 666154
6. D. L. Rodrigues, M. N. Menezes, F. J. Pinto, and A. L. Oliveira: Automated Detection of Coronary Artery Stenosis in X-ray Angiography using Deep Neural Networks (2021). <https://doi.org/10.48550/arXiv.2103.02969>
7. 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>, version 1.1.0
8. G. Sianos et. al: The SYNTAX Score: an angiographic tool grading the complexity of coronary artery disease. *EuroIntervention* **1**(2), 219–227 (2005)

9. Gao, Y., Kim, S., McIntosh, C.: EchoingECG: An Electrocardiogram Cross-Modal Model for Echocardiogram Tasks, p. 175–185. Springer Nature Switzerland (Sep 2025). https://doi.org/10.1007/978-3-032-04971-1_17
10. H. Zhang, D. Zhang, Z. Gao, and H. Zhang: Joint Segmentation and Quantification of Main Coronary Vessels Using Dual-Branch Multi-scale Attention Network. pp. 369–378. Springer-Verlag, Strasbourg, France (2021). https://doi.org/10.1007/978-3-030-87193-2_35
11. Hua, L., Luo, Y., Qi, Q., Long, J.: Medicalclip: Anomaly-detection domain generalization with asymmetric constraints. *Biomolecules* **14**, 590 (05 2024). <https://doi.org/10.3390/biom14050590>
12. I. K. Lee, J. Shin, Y. Lee, J. Ku, and H. Kim: SSASS: Semi-Supervised Approach for Stenosis Segmentation (2023). <https://doi.org/10.48550/arXiv.2311.10281>
13. Li, J., Aguirre, A., Moura, J., Liu, C., Zhong, L., Sun, C., Clifford, G., Westover, B., Hong, S.: An electrocardiogram foundation model built on over 10 million recordings with external evaluation across multiple domains (2025). <https://doi.org/10.48550/arXiv.2410.04133>
14. Lunelli, R., Nicolson, A., Pröll, S.M., Reinstadler, S.J., Bauer, A., Dlaska, C.: Benchecg and xecg: a benchmark and baseline for ecg foundation models (2025). <https://doi.org/10.48550/arXiv.2509.10151>
15. 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 (2021). <https://doi.org/10.48550/arXiv.2103.00020>
16. Selivanov, A., Müller, P., Özgün Turgut, Stolt-Ansó, N., Rückert, D.: Global and local contrastive learning for joint representations from cardiac mri and ecg (2025). <https://doi.org/10.48550/arXiv.2506.20683>
17. T. Han et. al: Coronary artery stenosis detection via proposal-shifted spatial-temporal transformer in X-ray angiography. *Computers in Biology and Medicine* **153**, 106546 (2023). <https://doi.org/10.1016/j.combiomed.2023.106546>
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 (2025). <https://doi.org/10.48550/arXiv.2308.05764>
19. Özgün Turgut, Müller, P., Menten, M.J., Rueckert, D.: Towards generalisable time series understanding across domains (2025). <https://doi.org/10.48550/arXiv.2410.07299>
20. V. V. Danilov et. al: Real-time coronary artery stenosis detection based on modern neural networks. *Scientific Reports* **11**(1), 7582 (2021). <https://doi.org/10.1038/s41598-021-87174-2>
21. Xue, Z., Geng, S., Guo, S., Mu, G., Yu, B., Wang, P., Hu, S., Zhang, D., Xu, W., Liu, Y., Yang, L., Tao, H., Hong, S., Chen, K.: Screening for severe coronary stenosis in patients with apparently normal electrocardiograms based on deep learning. *BMC Medical Informatics and Decision Making* **24** (11 2024). <https://doi.org/10.1186/s12911-024-02764-0>
22. Yeh, C.H., Tsai, T.H., Chen, C.H., Chou, Y.J., Mao, C.T., Su, T.P., Yang, N.I., Lai, C.C., Chen, C.T., Sytwu, H.K., Tsai, T.F.: Artificial intelligence-enhanced electrocardiography improves the detection of coronary artery disease. *Computational and Structural Biotechnology Journal* **27**, 278–286 (2025). <https://doi.org/10.1016/j.csbj.2024.12.032>

23. Zhang, S., Xu, Y., Usuyama, N., Xu, H., Bagga, J., Tinn, R., Preston, S., Rao, R., Wei, M., Valluri, N., Wong, C., Tupini, A., Wang, Y., Mazzola, M., Shukla, S., Liden, L., Gao, J., Crabtree, A., Piening, B., Bifulco, C., Lungren, M.P., Naumann, T., Wang, S., Poon, H.: Biomedclip: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs (2025). <https://doi.org/10.48550/arXiv.2303.00915>