

Opportunistic Cardiac Health Assessment: Estimating Phenotypes from Localizer MRI through Multi-Modal Representations

Busra Nur Zeybek¹, Özgün Turgut¹, Yundi Zhang¹, Jiazhen Pan^{1,4},
Robert Graf^{1,2}, Sophie Starck¹, Daniel Rueckert^{1,3,4}, and Sevgi Gokce
Kafali¹

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

² Department of Diagnostic and Interventional Neuroradiology, School of Medicine,
TUM University Hospital, Germany.

³ Department of Computing, Imperial College London, United Kingdom

⁴ Munich Center for Machine Learning (MCML), Germany
s.kafali@tum.de

Abstract. Cardiovascular diseases are the leading cause of death. Cardiac phenotypes (CPs), e.g., ejection fraction, are the gold standard for assessing cardiac health, but they are derived from cine cardiac magnetic resonance imaging (CMR), which is costly and requires high spatio-temporal resolution. Every magnetic resonance (MR) examination begins with rapid and coarse localizers for scan planning, which are discarded thereafter. Despite non-diagnostic image quality and lack of temporal information, localizers can provide valuable structural information rapidly. In addition to imaging, patient-level information, including demographics and lifestyle, influence the cardiac health assessment. Electrocardiograms (ECGs) are inexpensive, routinely ordered in clinical practice, and capture the temporal activity of the heart. Here, we introduce **C-TRIP** (Cardiac Tri-modal Representations for Imaging Phenotypes), a multi-modal framework that aligns localizer MRI, ECG signals, and tabular metadata to learn a robust latent space and predict CPs using localizer images as an opportunistic alternative to CMR. By combining these three modalities, we leverage cheap spatial and temporal information from localizers and ECG, respectively, while benefiting from patient-specific context provided by tabular data. Our pipeline consists of three stages. First, encoders are trained independently to learn uni-modal representations. The second stage fuses the pre-trained encoders to unify the latent space. The final stage uses the enriched representation space for CP prediction, with inference performed exclusively on localizer MRI. Proposed C-TRIP yields accurate functional CPs, and high correlations for structural CPs. Since localizers are inherently rapid and low-cost, our C-TRIP framework could enable better accessibility for CP estimation.⁵

Keywords: Contrastive Learning · MRI · Opportunistic Screening · Cardiac Biomarkers.

⁵ Code available at <https://github.com/basrazey/C-TRIP>.

1 Introduction

Cardiovascular diseases (CVDs) are the leading cause of death globally [1]. Several important functional cardiac phenotypes (FCPs), such as left and right ventricular ejection fraction (LVEF, RVEF), and structural cardiac phenotypes (SCPs), such as left ventricular mass (LVM) and right ventricular end-diastolic volume (RVEDV), are established clinical biomarkers for CVD detection [5]. Cardiac cine magnetic resonance imaging (CMR) is the reference standard for extracting SCPs (LVM, RVEDV) and FCPs (LVEF, RVEF) [18].

However, CP calculation from CMR is a complex process [18]. As a dynamic acquisition, CMR requires high spatio-temporal resolution and electrocardiogram (ECG) gated reconstruction to bin data from the same cardiac phases. Given the need for a dedicated full cardiac MRI protocol and additional ECG hardware for a CMR scan, there is an evident need for more efficient CP estimation. In routine clinical practice, ECG serves as a first-line cardiac assessment tool due to its low cost, rapid acquisition, and ability to capture the heart’s electrical activity [20]. Yet, ECG cannot directly quantify CPs. A recent study proposed an ECG-based cardiac screening tool by applying contrastive learning to align the representation spaces of ECG and CMR to enrich the learned features, while predicting CPs solely on ECG during inference [21].

Beyond imaging and ECG, clinical decision making also relies on patient-level metadata, such as sex, body-mass-index (BMI), and lifestyle, as cardiac structure and function are strongly influenced by these factors [24]. A previous study on multi-modal contrastive learning proved that aligning CMR and tabular data increased the CVD classification accuracy [9]. This is also reflected in clinical settings, where physicians integrate information across modalities rather than interpreting each modality in isolation.

Every magnetic resonance imaging (MRI) examination starts with coarse but rapid localizer scans, which are acquired for planning purposes. Although localizers have lower image quality and are not targeted for diagnosis [22], they contain structural tissue information and require minimal scan and reconstruction time. Due to their simplicity and availability, they possess certain advantages over CMR (**Table 1**). Localizer MRIs, even by themselves, represent a promising yet underexplored opportunity for opportunistic screening, as reported by [4] for detection of osteoporotic vertebral fractures. Under a multi-modal framework, localizers have a bigger potential to be a screening tool. Thus, we introduce **C-TRIP** (Cardiac Tri-modal Representations for Imaging Phenotypes), a contrastive learning approach that anchors on localizer MRI and aligns its representation space with ECG-derived functional information and tabular clinical context. In contrast to prior studies [19,21,9,2], our contributions are:

- **Localizer-centric tri-modal alignment:** We use routinely acquired rapid and coarse localizers as the central modality for cheap spatial information, then align it with functional (ECG) and contextual (tabular) information.
- **Multi-modal training with single-modality inference:** Multi-modal contrastive alignment is performed during training, while CP prediction is conducted exclusively on localizer MRI at inference.

Table 1. Comparison of routinely acquired localizer MRI and CMR.

Aspect	Localizer MRI	Cardiac CINE MRI
Purpose	Scan planning; not for diagnosis	CP quantification and diagnosis
Scan time	Very short (seconds per stack)	Longer
Temporal info	None (static images)	Multi-frame across the cardiac cycle
Motion	Not ECG-gated; susceptible to motion	ECG-gated; captures cardiac motion
Image quality	Lower resolution; variable contrast	High spatial resolution and contrast
Standardization	Less standardized across sites	Established clinical protocols
Coverage	Coarse or incomplete coverage	Full short-axis + long-axis views
Clinical role	Not intended for diagnosis	Reference standard for CPs
Cost	Routine acquisition	Dedicated acquisition required
Availability	Near-universal in cardiac MRI exams	Only during full cardiac protocol




Fig. 1. Our multi-modal contrastive learning framework. First stage includes pre-training encoders for each modality. Second stage involves unification of the latent spaces from three modalities, in which the data from the same subjects were pulled, and different subjects were pushed away. The last stage uses a regression head to predict all 18 cardiac phenotypes solely from localizer MRIs.

- **Strong performance in low-data regimes:** C-TRIP yields robust performance for both FCPs and SCPs.
- **Opportunistic low-cost screening tool:** Rather than replacing CMR, C-TRIP acts as an opportunistic low-cost screening tool when full CMR is limited, enabling better accessibility to CP estimation.

2 Methods

We propose the C-TRIP framework shown in **Fig. 1**. Our framework consists of three sequential stages. In the first stage, we pre-train three separate encoders for localizer, ECG, and tabular data using self-supervised masked data modeling. The second stage fuses the three pre-trained encoders to minimize the contrastive loss and achieve patient-level alignment in a shared latent embedding space. In the third stage, we predict CPs solely from localizers.

Stage I: Single-modality Masked Data Modeling: To obtain a robust latent representation, we train uni-modal vision transformer (ViT) based masked autoencoders (MAEs)[6,10], for which the architecture is adapted for each individual modality. All encoders process the 25% visible input, while lightweight decoders reconstruct the randomly masked 75% of the data. We minimize mean squared error (MSE) to reconstruct the masked patches. Multi-channel ECG signals are divided into non-overlapping 1D patches similar to [21], and reconstructed to capture underlying temporal dynamics. For tabular data, the input is processed as a 1D sequence of discrete feature tokens. We project numerical and categorical features into a uniform embedding space, similar to [23]. Full sequence is reconstructed by optimizing MSE for numerical features and cross-entropy loss for categorical features.

Stage II: Multimodal Contrastive Learning: We align the pre-trained uni-modal representations into a joint latent space to enrich the anatomical features extracted from the localizers with the functional and temporal information from the ECG and the clinical context from the tabular data. We utilize a localizer-centric contrastive loss that aligns localizers (L), ECG (E) and tabular data (T), without aligning T and E .

Localizer-Centric Contrastive Loss: We independently align the localizer embeddings with the auxiliary modalities by optimizing an InfoNCE-based contrastive loss [11,16]. Generically, let M denote an auxiliary modality representing ECG (E) or tabular data (T). The loss is defined symmetrically in both directions between localizer MRI (L) and modality M . The bidirectional alignment ($\mathcal{L}_{L \leftrightarrow M}$) pulls the embeddings of the same patient closer while pushing the embeddings of different patients in the batch.

$$\mathcal{L}_{L \leftrightarrow M} = \frac{1}{2} (\mathcal{L}_{L \rightarrow M} + \mathcal{L}_{M \rightarrow L}) \quad (1)$$

$$\mathcal{L}_{L \rightarrow M} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\text{sim}(z_L^{(i)}, z_M^{(i)})/\tau_{LM})}{\sum_{j=1}^N \exp(\text{sim}(z_L^{(i)}, z_M^{(j)})/\tau_{LM})} \quad (2)$$

In the above formulation, N is the batch size, $\text{sim}(\cdot, \cdot)$ denotes cosine similarity, z_L and z_M are the extracted representations, and τ_{LM} is a learnable temperature parameter. The overall loss objective is then $\mathcal{L}_{Total} = \frac{1}{2} (\mathcal{L}_{L \leftrightarrow E} + \mathcal{L}_{L \leftrightarrow T})$. By utilizing the localizer MRI as the central modality, ECG and tabular embeddings are aligned indirectly. We intentionally avoided a direct $E + T$ alignment to prevent the model from exploiting non-imaging shortcuts. This, additionally, reduced computational overhead during loss computation.

Stage III: Inference to Predict CPs from Localizers: We attach a lightweight MLP regression head to predict CPs from localizers only during inference.

2.1 Baselines and Ablation Studies

To evaluate the proposed **C-TRIP** framework, we compared our approach against several baselines and conducted targeted ablation studies: **Baseline networks** included supervised schemes to predict the CPs as a regression task, with different types of inputs: CMR (the reference, CMR_{sup}), localizers (L_{sup}), tabular data (T_{sup}), and ECG (E_{sup}). **The first ablation study** analyzed a tri-modal scheme on two different conditions: tabular data injected with CPs explicitly ($L + E + T^p$) inspired by [9], versus tabular data without any CPs ($L + E + T$, i.e., proposed **C-TRIP**). This way, we could understand if no explicit supervision (i.e., no CP injection) could still generalize well in a wide test population. **The second ablation study** aimed to choose the most optimum modality combination, comparing bimodal and trimodal contrastive pre-training strategies: $L + E$, $L + T$, $L + E + T$ (i.e., our proposed **C-TRIP**).

2.2 Dataset and Preprocessing

We utilized paired multi-modal data from the UK Biobank [17], with imaging parameters detailed in [15]. To prevent data leakage, we used subject-level splits with demographically balanced (age, sex, LVEF) training, validation, and test sets ($N_{train} : 14577, N_{val} : 3177, N_{test} : 3123$). Downstream experiments evaluated only subjects with complete multi-modal pairings and 18 ground truth CPs [3]. For CMR_{sup} , we extracted 3 short-axis time frames from the middle baso-apical slice (end-diastole, end-systole, mid-phase) per subject, identical to [21]. For localizer MRI, we used 3 oblique-transverse mid-slices from the 18-slice "AA_Heart_Scout" scan. The heart was segmented using [8] and cropped from thorax. Images were normalized and resized to 224×224 . Standard 10-second, 12-lead resting ECGs (500Hz) were corrected for baseline drift to eliminate respiratory and motion artifacts. Tabular data consists of numerical (e.g., BMI, metabolic rates, QRS duration) and categorical features (demographics, medical history, smoking status). Recorded during the same assessment, all modalities provide concurrent representation of the heart.

2.3 Implementation

All experiments were implemented in PyTorch/Lightning on a single NVIDIA A40 GPU optimized via AdamW [13] with cosine annealing and early stopping. In Stage I, uni-modal MAEs (embedding dimension (dim): $L : 768, E : 384, T : 384$) were pre-trained for 400 epochs (batch size (bs) 256, learning rate (lr) 5×10^{-4}). Random augmentations (cropping, scaling, rotation) were applied only to the localizer training set. In Stage II, fully-trainable encoders projected representations into a shared 256-dim space. Contrastive alignment was optimized for 150 epochs (bs 256) with pair-specific temperature set to $\tau_{LE} = 0.1$ and $\tau_{LT} = 0.25$ after hyperparameter tuning. In Stage III, downstream regression head was trained for 100 epochs solely on Stage II localizer embeddings for C-TRIP. To preserve Stage II representations, differential lr's were used: 1×10^{-3}

Table 2. Performance comparison across modalities. T/ T^p : tabular data without/with CPs; E: electrocardiogram; L: localizer MRI. CMR_{sup} : provided as reference. Bold text and gray shading denote best MD and LoA.

	LET ^p T	LVEF [%]		LVM [g]		RVEF [%]		RVEDV [mL]	
		MD	LoA	MD	LoA	MD	LoA	MD	LoA
CMR_{sup}	✓✓✓✓	0.3	[−8.0, 8.5]	−0.3	[−24.5, 24.0]	−0.2	[−8.7, 8.4]	2.8	[−39.2, 44.8]
L_{sup}	✓✓✓✓	0.1	[−10.8, 11.1]	0.7	[−23.4, 24.8]	0.1	[−10.3, 10.4]	1.8	[−39.8, 43.4]
E_{sup}	✓✓✓✓	0.4	[−10.9, 11.6]	−0.8	[−35.9, 34.3]	0.2	[−10.8, 11.2]	−0.6	[−60.6, 59.4]
T_{sup}	✓✓✓✓	0.1	[−11.3, 11.5]	2.3	[−22.5, 27.0]	−0.5	[−11.5, 10.4]	5.0	[−38.9, 48.9]
L+T	✓✓✓✓	0.1	[−10.6, 10.8]	0.3	[−23.1, 23.8]	0.0	[−10.2, 10.2]	1.1	[−39.8, 41.9]
L+E	✓✓✓✓	0.3	[−10.6, 11.1]	0.9	[−23.5, 25.2]	0.2	[−10.1, 10.5]	1.7	[−39.9, 43.3]
L+E+T ^p	✓✓✓✓	0.2	[−10.5, 11.0]	0.9	[−22.1, 24.0]	0.2	[−10.0, 10.5]	1.2	[−39.5, 41.9]
L+E+T (C-TRIP)	✓✓✓✓	0.3	[−10.4, 11.1]	0.2	[−23.2, 23.7]	0.2	[−10.0, 10.4]	0.2	[−40.8, 41.1]

(regression head), and 1×10^{-4} (pre-trained encoder). Baselines used the identical head trained from scratch.

2.4 Evaluation

The reference CPs were extracted from CMR. Accuracy was evaluated using mean differences (MD), limits of agreement (LoA), and Pearson’s R [14]. We examined the effects of the number of fine-tuning samples during downstream analysis at stage III to mimic a realistic low-data regime (e.g., 1% of the samples).

3 Results

Table 2 demonstrates the CP estimation in all techniques. The CMR_{sup} achieved the best performance in FCPs, as expected, while localizer-based techniques are competitive for SCPs. L_{sup} already produces strong results (i.e., small MDs) for FCPs. ECG_{sup} yields the worst overall performance. The inclusion of explicit CPs in the tabular metadata, $L + E + T^p$, does not improve the performance, in fact, leads to an increase in MDs in SCPs. Our proposed C-TRIP yields the smallest MD, with the second tightest LoA overall, showing its robust performance in SCPs, while still being competitive in FCPs.

Fig. 2A shows that the attention maps for C-TRIP are more focal and localized than the baseline L_{sup} . In slice 1 (short-axis view), C-TRIP emphasizes the *donut* of the LV and the *crescent* of the RV. C-TRIP shifts its attention to the specific chambers in slice 3 (a more basal/outflow view).

The multi-modal latent space (e.g., UMAPs) before and after localizer-centric alignment is shown in **Fig. 2B** color-coded by LVM and sex. L and T representations become deeply intertwined after alignment due to high structural and functional correlation (e.g., males have larger LVMs, LVEF is influenced by sex). The ECG cluster remains distinct due to a more complex relationship between




Fig. 2. A) Attention maps. (head-averaged [CLS] token weights from the final ViT block) While L_{sup} , exhibits leaked attention to non-cardiac structures (e.g., lungs and chest wall), C-TRIP was able to focus on biologically relevant structures. Images are cropped to the heart region for visibility. **B) UMAPs** (cosine distance, 15 neighbors, min_dist=0.1) L_2 -normalized joint embeddings show representation space before and after alignment. Our localizer-centric design reflects itself in UMAPs. Due to strong anatomical correlation, L and T merge easily. The relationship between L and E is more complex, so E remains more distinct, but is pulled toward the same phenotypic gradient/trend from low to high LVM, or male/female distinction.

L and E since there is an inherent domain gap between electrophysiology (ECG) and anatomy (localizers). Yet, in the overall latent space, we observe a gradient from lower to higher LVM values, and male/female distinction. This particular trend indicates that the C-TRIP learned a robust mapping for CPs.

Fig. 3 shows CMR_{sup} achieves the highest correlation for FCPs, even at low-data regimes. L_{sup} shows moderate performance across all data regimes. E_{sup} yields the worst performance in both FCPs and SCPs. T_{sup} leads to robust performance in SCPs, but not in FCPs. C-TRIP achieves high correlations in SCPs, even at 1% data regime (Pearson $R > 0.70$), while still performing the second-best after CMR for FCPs.

4 Discussion and Conclusion

Proposed C-TRIP estimated CPs solely from routinely acquired localizers, to act as an **opportunistic assessment tool for cardiac health at low-cost**. Here, we proved; tabular data provides a strong population-level prior, ECG is the primary source of calibration, reducing MD. Yet, by itself, ECG performance still suffers. By aligning three modalities, C-TRIP uses the ECG to correct for generalized guesses from the tabular data to consider subject-specific information, while exploiting the spatial prior from the localizers. The impact of tri-modal alignment is most pronounced in SCPs, particularly RV, indicating that the complex geometry of the RV requires the full integration of three modalities. Our high accuracy in LVM estimation could also be useful in detecting LV hypertrophy, a result of coronary artery disease (CAD). Other real-world applications include preliminary assessment in resource-limited settings and real-time decisions on the necessity of a full CMR according to the estimations by C-TRIP.

C-TRIP learned a robust anatomical prior; only the heart-specific pixels are relevant for CP estimation, while mitigating the non-cardiac noise in L_{sup}




Fig. 3. Performance and scaling comparison of our proposed multimodal approach against unimodal supervised baselines across four CPs. (Top Row) Scatter plots comparing model predictions to ground truth values using 100% of the available fine-tuning data for LVEF, LVM, RVEF, RVEDV. Pearson R is reported in the legends. The plots show 500 subjects for better visibility. **(Bottom Row)** Fine-tuning data scaling behavior. The plots demonstrate model performance (Pearson R) as a function of the fine-tuning subset size (1%, 10%, and 100%, displayed on a logarithmic scale). The shaded regions represent the 95% confidence intervals. Our proposed method C-TRIP (in green) is compared against single modality baselines.

as shown in attention maps. The UMAPs showed no *one-blob* for C-TRIP, partly due to our localizer-centric design, partly by the natural modality gap [12]. While tabular data and MRI share a direct structural correlation (e.g., body scale and sex affect heart size), the relationship between ECG and MRI is more nuanced. The residual separation of the ECG cluster represents the unique electrophysiological information provided by ECG, which adds to the already learned structural information of the localizer. Our localizer-centric design also specifically prevents shortcut correlations between non-imaging features (i.e., older subjects having certain ECGs). This is crucial since CPs are imaging-derived biomarkers.

SCP estimation was more accurate than FCPs in L -based techniques (higher Pearson R). CMR maintains the highest R in FCPs, as it captures temporal cardiac motion. In low-data regimes (1–10% fine-tuning), L_{sup} and E_{sup} yielded low performance. The relatively strong performance of T_{sup} in SCPs is expected, as T_{sup} can extract accurate population statistics (e.g., males have higher LVM than females). Yet, T_{sup} could not produce subject-specific functional information. C-TRIP integrates functional, contextual, and spatial information, providing **better low-data efficiency**. Although C-TRIP yielded reduced performance in FCPs, the model learns temporal relations implicitly via multi-modal training with ECG. ECG addition at inference may improve FCP estimation, but the lack of temporal alignment with MRI in UK Biobank may limit its added value.

Despite UK Biobank’s healthy volunteer bias [7], 30% of subjects had LVEF $< 55\%$, enabling robust associations. C-TRIP reduces regression-to-the-mean in T_{sup} , showing it does not rely solely on population averages. This aligns with opportunistic assessment, seeking incidental findings from healthy individuals.

Our study had limitations. Although UK Biobank was conducted in four sites, it offers a homogeneous dataset under a single vendor, single scanner type [17]. Additionally, conventional localizers are not acquired in a standardized manner, as opposed to UK Biobank. Thus, future work includes validating C-TRIP under less homogeneous settings. Furthermore, CMR_{sup} does not consider full stack of cine, which was chosen to ensure architectural comparability across L -based models. Hence, it does not reflect the full diagnostic capacity of clinical CMR. Although we did not focus on disease classification, future work will include risk classification based on EF, as well as disease detection (e.g., CAD).

Our results support **the potential clinical utility of C-TRIP**, not as a replacement for CMR, but as a low-cost opportunistic screening tool. By enhancing rapid localizer MRI with routinely collected data (ECG and tabular), we enable reliable CP estimation when comprehensive CMR cannot be performed.

Acknowledgments. SGK has been sponsored by the Alexander von Humboldt Foundation. 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. Cardiovascular diseases (CVDs). [https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-\(cvds\)](https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds))
2. Alvarez-Florez, L., Bujalance-Gomez, A., Raijmakers, F., Ruiperez-Campillo, S., Kolk, M.Z.H., Wiers, J., Vogt, J., Bekkers, E.J., Išgum, I., Tjong, F.V.Y.: Dual-Phase Cross-Modal Contrastive Learning for CMR-Guided ECG Representations for Cardiovascular Disease Assessment (Feb 2026). <https://doi.org/10.48550/arXiv.2602.12883>
3. Bai, W., Suzuki, H., Huang, J., Francis, C., Wang, S., Tarroni, G., Guitton, F., Aung, N., Fung, K., Petersen, S.E., et al.: A population-based phenome-wide association study of cardiac and aortic structure and function. *Nature medicine* **26**(10), 1654–1662 (2020)
4. Bazzocchi, A., Garzillo, G., Fuzzi, F., Diano, D., Albinini, U., Salizzoni, E., Battista, G., Guglielmi, G.: Localizer sequences of magnetic resonance imaging accurately identify osteoporotic vertebral fractures. *Bone* **61**, 158–163 (2014)
5. Chen, Y., Xie, W., Zhang, J., Qiu, H., Zeng, D., Shi, Y., Yuan, H., Zhuang, J., Jia, Q., Zhang, Y., Dong, Y., Huang, M., Xu, X.: Myocardial Segmentation of Cardiac MRI Sequences With Temporal Consistency for Coronary Artery Disease Diagnosis. *Frontiers in Cardiovascular Medicine* **9**, 804442 (Feb 2022). <https://doi.org/10.3389/fcvm.2022.804442>

6. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale (Jun 2021). <https://doi.org/10.48550/arXiv.2010.11929>
7. Fry, A., Littlejohns, T.J., Sudlow, C., Doherty, N., Adamska, L., Sprosen, T., Collins, R., Allen, N.E.: Comparison of sociodemographic and health-related characteristics of uk biobank participants with those of the general population. *American journal of epidemiology* **186**(9), 1026–1034 (2017)
8. Graf, R., Platzek, P., Riedel, E.O., Ramschütz, C., Starck, S., Möller, H.K., Atad, M., Völzke, H., Bülow, R., Schmidt, C.O., et al.: Vibesegmentator: full body mri segmentation for the nako and uk biobank. *European Radiology* pp. 1–15 (2025)
9. Hager, P., Menten, M.J., Rueckert, D.: Best of both worlds: Multimodal contrastive learning with tabular and imaging data. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 23924–23935 (2023)
10. He, K., Chen, X., Xie, S., Li, Y., Dollar, P., Girshick, R.: Masked Autoencoders Are Scalable Vision Learners. In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 15979–15988. IEEE, New Orleans, LA, USA (Jun 2022). <https://doi.org/10.1109/CVPR52688.2022.01553>
11. Li, Y., Fan, H., Hu, R., Feichtenhofer, C., He, K.: Scaling language-image pre-training via masking. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 23390–23400 (2023)
12. Liang, V.W., Zhang, Y., Kwon, Y., Yeung, S., Zou, J.Y.: Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *Advances in Neural Information Processing Systems* **35**, 17612–17625 (2022)
13. Loshchilov, I., Hutter, F.: Decoupled Weight Decay Regularization (Jan 2019). <https://doi.org/10.48550/arXiv.1711.05101>
14. Obuchowski, N.A., Reeves, A.P., Huang, E.P., Wang, X.F., Buckler, A.J., Kim, H.J., Barnhart, H.X., Jackson, E.F., Giger, M.L., Pennello, G., et al.: Quantitative imaging biomarkers: a review of statistical methods for computer algorithm comparisons. *Statistical methods in medical research* **24**(1), 68–106 (2015)
15. Petersen, S.E., Matthews, P.M., Francis, J.M., Robson, M.D., Zemrak, F., Boubertakh, R., Young, A.A., Hudson, S., Weale, P., Garratt, S., et al.: Uk biobank’s cardiovascular magnetic resonance protocol. *Journal of cardiovascular magnetic resonance* **18**(1), 8 (2016)
16. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PmLR (2021)
17. Raisi-Estabragh, Z., Harvey, N.C., Neubauer, S., Petersen, S.E.: Cardiovascular magnetic resonance imaging in the uk biobank: a major international health research resource. *European Heart Journal-Cardiovascular Imaging* **22**(3), 251–258 (2021)
18. Rajiah, P.S., François, C.J., Leiner, T.: Cardiac MRI: State of the art. *Radiology* **307**(3), e223008 (2023)
19. 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: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 217–227. Springer (2025)
20. Siontis, K.C., Noseworthy, P.A., Attia, Z.I., Friedman, P.A.: Artificial intelligence-enhanced electrocardiography in cardiovascular disease management. *Na-*

- ture Reviews. *Cardiology* **18**(7), 465–478 (2021). <https://doi.org/10.1038/s41569-020-00503-2>
21. Turgut, Ö., 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)
 22. You, J.Y., Lee, J.W., Seo, J., Chai, J.W., Chae, H.D., Kang, H.S.: Readability of extraspinal organs on scout images of lumbar spine MRI according to different protocols. *PLoS ONE* **16**(5), e0251310 (May 2021). <https://doi.org/10.1371/journal.pone.0251310>
 23. Zhang, Y., Hager, P., Liu, C., Shit, S., Chen, C., Rueckert, D., Pan, J.: Towards cardiac mri foundation models: Comprehensive visual-tabular representations for whole-heart assessment and beyond. *Medical Image Analysis* p. 103756 (2025)
 24. Zhu, Z., Li, F.R., Jia, Y., Li, Y., Guo, D., Chen, J., Tian, H., Yang, J., Yang, H.H., Chen, L.H., Zhang, K., Yang, P., Sun, L., Shi, M., Zhang, Y., Qin, L.Q., Chen, G.C.: Association of Lifestyle With Incidence of Heart Failure According to Metabolic and Genetic Risk Status: A Population-Based Prospective Study. *Circulation: Heart Failure* **15**(9) (Sep 2022). <https://doi.org/10.1161/CIRCHEARTFAILURE.122.009592>