

Enhancing Cardiovascular Disease Prediction through Multi-Modal Self-Supervised Learning

Francesco Girlanda¹

fgirlanda@student.ethz.ch

Olga Demler^{1,2}

odemler@bwh.harvard.edu

Bjoern Menze³

bjoern.menze@uzh.ch

Neda Davoudi^{1,3,4}

neda.davoudi@ai.ethz.ch

¹ Department of Computer Science

ETH Zürich

Zürich, Switzerland

² Brigham and Women's Hospital

Harvard Medical School

Boston, Massachusetts, USA

³ Department of Quantitative Biomedicine

University of Zürich

Zürich, Switzerland

⁴ ETH AI Center

ETH Zürich

Zürich, Switzerland

Abstract

Accurate prediction of cardiovascular diseases remains imperative for early diagnosis and intervention, necessitating robust and precise predictive models. Recently, there has been a growing interest in multi-modal learning for uncovering novel insights not available through uni-modal datasets alone. By combining cardiac magnetic resonance images, electrocardiogram signals, and available medical information, our approach enables the capture of holistic status about individuals' cardiovascular health by leveraging shared information across modalities. Integrating information from multiple modalities and benefiting from self-supervised learning techniques, our model provides a comprehensive framework for enhancing cardiovascular disease prediction with limited annotated datasets.

We employ a masked autoencoder to pre-train the electrocardiogram ECG encoder, enabling it to extract relevant features from raw electrocardiogram data, and an image encoder to extract relevant features from cardiac magnetic resonance images. Subsequently, we utilize a multi-modal contrastive learning objective to transfer knowledge from expensive and complex modality, cardiac magnetic resonance image, to cheap and simple modalities such as electrocardiograms and medical information. Finally, we fine-tuned the pre-trained encoders on specific predictive tasks, such as myocardial infarction. Our proposed method enhanced the image information by leveraging different available modalities and outperformed the supervised approach by 7.6% in balanced accuracy.

1 Introduction

Cardiovascular diseases (CVDs) are the leading cause of death worldwide resulting in 20.5 million deaths in 2021 [26] which makes the early diagnosis and prevention of CVD a critical task in medical research. Patient data, including laboratory values, imaging data, medical histories from the electronic health record (EHR), and lifestyle contribute to a comprehensive patient profile and are essential for healthcare decisions. Integrating these diverse data sources in real time facilitates more effective prevention and treatment strategies. Several medical data are used for detecting CVD. Cardiac magnetic resonance (CMR) images are the gold standard for assessing cardiac structure, and electrocardiograms (ECGs) are commonly used for electrophysiological evaluation. Early algorithms focus on automated ECG interpretation [26, 30] or CMR image analysis [54]. A major limitation of current models is that they are limited to algorithms trained on a single modality and don't translate between different modalities, however, to capture the complexity of human biology, there is a need to go beyond traditional clinically-focused and expert-curated features and include critically important, but often neglected, data types on which doctors' evaluation rely on [17]. Machine learning models can leverage the complementary information present in different modalities to develop a joint characterization of physiological states similar to the conventional approach of diagnosis by doctors and further enhancing their effectiveness. Several studies have employed multimodal data to improve diagnostic capabilities. Khurshid *et al.* developed a deep learning-based model to estimate left ventricular (LV) mass from 12-lead ECGs, which showed a strong correlation with LV mass on cardiac MRI imaging and cardiovascular events [25]. Christensen *et al.* developed a model that learned to interpret cardiac ultrasound images by correlating them with expert cardiologists' annotations [11]. Borsos *et al.* proposed a multimodal deep learning architecture combining tabular data and imaging to predict dichotomized mRS scores three months post-Acute Ischemic Stroke [8]. Amal *et al.* provided a comprehensive review of multimodal approaches in healthcare for CVD analysis, emphasizing various fusion strategies [9]. Ensemble learning improves deep learning model performance through various approaches but it is challenging due to high training costs, the need for inducing diversity among models, complex model selection, and combining predictions, requiring further research to optimize these aspects. [15]

Recent models have proven to be fruitful for use in biomedical prediction tasks, but the risk of overfitting remains due to the limited size of annotated datasets for supervised learning. Results of generative models such as Autoencoders (AEs) [7] on multi-modal clinical measurements show that they perform well on aligning the embeddings from diverse modalities and constructing a holistic representation for characterizing physiological state [53]. AEs are employed to learn cross-modal representations from large multimodal datasets. Healthcare datasets such as UK Biobank [56] serve as an excellent resource for learning clinically relevant representations. Hager *et al.* attempted to combine images and tabular data for multi-modal pre-training of representation by optimizing a CLIP loss and predicting myocardial infarction (MI) for the downstream task [19]. They have shown that multi-modal prediction using CMR and tabular data outperforms uni-modal prediction. However, it relies solely on the image encoder for downstream tasks, ignoring the extensive information contained in the tabular data or other relevant modalities such as ECG.

In healthcare, ECG signals and basic medical information, are readily available or easily measured. However, CMR imaging, which provide detailed anatomical and physiological information about the heart, is costly to obtain and may not always be accessible. Turgut *et al.* tried a self-supervised contrastive learning [10] approach that transfers domain-specific

information from CMR images to ECG embeddings. They predict the subject-specific risk of various CVDs and determine distinct cardiac phenotypes solely from ECG data and demonstrate that learned ECG embeddings incorporate information from CMR image with self-supervised training on both modalities [47]. However, integrating tabular data, which includes demographics, lifestyle, and lab tests, with imaging data in multi-modal datasets is essential for informed clinical decision-making in healthcare [4]. This project aims to improve the existing classification of CVD using only ECG signals and patient medical information for diagnosing MI by incorporating the embedding that is constructed with multiple modalities including CMR images. This study demonstrates – to the best of our knowledge, for the first time – how to successfully integrate all available modalities in different formats as in realistic medical scenarios for predicting different downstream tasks. Furthermore, it improves the alignment of different modalities in latent space by leveraging the huge amount of unannotated data with self-supervised learning. Importantly, we increase the balanced accuracy by 7.6% on a well-characterized clinical dataset for MI diagnosis. Our code is publicly available at <https://github.com/FraGirila/MMSSL-for-CVD-Pred>.

2 Data

Our analyses were performed on UK Biobank which is a prospective cohort study of over 500,000 individuals including their clinical data from across the United Kingdom [24, 36, 39].

Cardiac MR images: The whole CMR image dataset is composed of 60,623 participants with dimensions of $208 \times 208 \times Z \times 50$ with Z as a variable number of slices dependent on the size of the heart through 50 time points. To reduce data dimensionality, we selected the subset of CMR information that best characterizes a heartbeat. We used a pre-trained segmentation model from [5, 6] to label the left ventricle and calculate heart expansion based on the number of labeled pixels. This allowed us to determine the end-diastolic phase (maximum left ventricle expansion), the end-systolic phase (minimum expansion), and the mid-phase, defined as the midpoint between the two, as outlined in [47]. Additionally, we reduced the 3D volume to 2D images by selecting the middle slice on the z -axis. This preprocessing approach extracts essential heartbeat information, reducing model complexity and improving speed.

ECG: We used 12-lead ECG data recorded for 10 seconds at rest with a sampling frequency of 500Hz and removed the measurement noise with 0.5 Hz high-pass bandwidth filter with order of 5, followed by powerline filtering [29]. The whole ECG dataset is composed of 50,531 participants each of them with 5000 time points and 12 leads.

Tabular data: As the aim of this research is to take advantage of data that are easily available to make predictions on CVD, we integrate medical information as tabular features in the signal encoder that will be used to predict the downstream task. The list of 33 important clinical features included in this study are evaluated by medical experts and summarized a table in tabular data section in supplementary materials. Tabular dataset is composed of 45,257 participants with 33 features. Categorical features were encoded using ordinal encoding and continuous features were normalized between 0 and 1.

We paired the ECG, CMRI, tabular features, and MI labels for multi-modal training. Some participants didn't have all three modalities available resulted in 45112 final dataset which we divided into 80%, 10%, and 10% for training, validation, and test sets respectively.

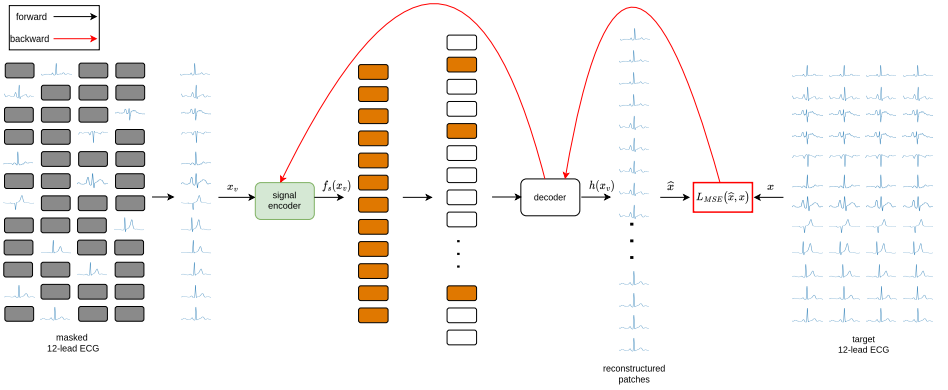


Figure 1: Visualization of the pre-training pipeline for ECG signals. The signal is divided into patches and a fraction of them is masked. The ECG encoder takes the visible patches that will be reconstructed via decoder. MAE loss measures the distance between the reconstruction and the original signal. A trained ECG encoder will be used in the next steps.

3 Method

We divided the training into four steps. First, we pre-trained the ECG encoder with an MAE that learns the latent space by reconstructing the ECG [47]. For this part, we used all available ECGs. We employed the MAE pretraining for ECG as masked signal reconstruction is shown to be the most efficient self-supervised learning (SSL) method for ECG data [35, 40] and we think that the reconstruction task applied on all the leads is more effective in distinguishing between different ECG compared to the comparison between different augmentations made by SimCLR method. The image encoder with ResNet50 [40] backbone is separately pre-trained on CMR images using SimCLR [9] loss to map two different views of the same image closer in the latent space. The selection of SimCLR is due to its demonstrated superior performance in pretraining on medical data compared to other self-supervised learning methods [12, 19]. Training the image encoder in advance helps transfer the knowledge from the image encoder to the signal encoder. In the next step, we use paired CMRI, ECG, and tabular data to improve the signal encoder of ECG. This is done using the CLIP loss [60] to transfer the learning from CMR images to ECG and tabular data. Eventually, we fine-tuned ECG and tabular encoder using a balanced dataset composed of all data from individuals having the disease and the same amount of randomly chosen healthy people.

3.1 ECG masked autoencoder

Masked Autoencoder is a widely recognized self-supervised technique renowned for its ability to learn robust representations through its masking and reconstruction mechanism [22]. This approach aids in the acquisition of a robust ECG encoder capable of mitigating the redundancy inherent in ECG data, thereby enhancing sample discrimination. Additionally, this process facilitates the transfer of learning from CMRI to ECG and tabular encoders by leveraging an existing encoder. The pipeline for this step is summarized in Figure 1.

Masking: The first step consists of dividing the whole 12-leads ECG signal into rectangular non-overlapping patches, masking a specified percentage of them using random sampling. While a mask ratio of 75% is typically recommended [22], we have increased it to

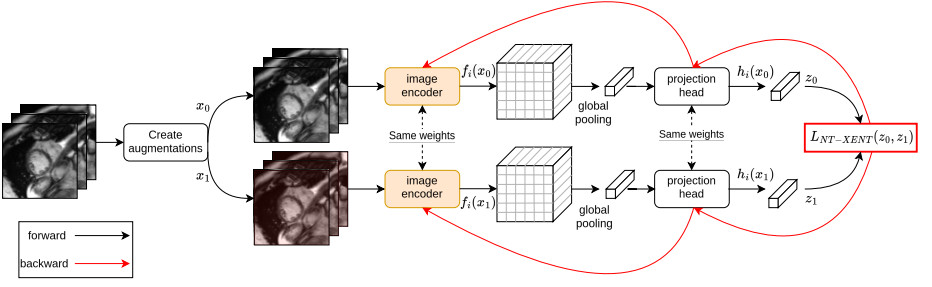


Figure 2: Visualization of the pre-training pipeline for CMRI. The image stack is augmented and given to image encoders. The resulting embeddings are then given to projection heads for optimizing the NT-Xent loss. The resulting image encoder will be later used for multi-modal self-supervised and uni-modal supervised training.

80% resulting in a larger effective receptive field that considers the ECG’s redundancy such that the reconstruction task cannot be easily solved by interpolation.

We have as input a multi-channel time series $x \in \mathbb{R}^{C \times T}$, we divide the input in a flattened list of N patches with size D reshaping the original input to $x \in \mathbb{R}^{N \times D}$. Afterward, we randomly draw a binary mask on the patches for obtaining the visible patches $x_v = x[m] \in \mathbb{R}^{N_v \times D}$ and masked ones $x_m = x[1 - m] \in \mathbb{R}^{N_m \times D}$ such that $N = N_v + N_m$.

The backbone encoder is a ViT. We chose this model because we treat the ECG signal and its patches as an image that will be reconstructed thanks to its redundancy. Furthermore, it is a transformer-based technique that showed great reconstruction capabilities thanks to its ability to focus on relevant parts of the visible patches [13]. We haven’t used a ResNet encoder as in the image encoding because we think that a reconstruction of the signal together with the attention mechanism are more beneficial for taking advantage of the periodicity and redundancy of the signal to learn quality ECGs embeddings. We apply this model to the visible patches. The encoder embeds patches by a linear projection with added positional embeddings and then processes the result via a series of Transformer blocks.

MAE decoder is given a set of encoded visible patches and mask tokens (introduced after the encoder) that represent a learned vector that indicates the presence of a missing patch to be predicted. Its task is to perform the prediction of pixels for each of the missing patches. The decoder’s output $\hat{x} \in \mathbb{R}^{N \times D}$ is then reshaped to form a reconstructed image $\hat{x} \in \mathbb{R}^{C \times T}$.

MAE $h(\cdot)$ consists of the composition of the signal encoder $f_s(\cdot)$ and decoder $g(\cdot)$ which, when provided with visible patches x_v , yielding the reconstructed signal \hat{x} .

$$\hat{x} = h(x_v) = g(f_s(x_v)) \quad (1)$$

Further information about the augmentations and the training setting can be found in experimental setup section in supplementary materials.

3.2 Image encoder pre-training

We pre-trained the image encoder (Figure 2) with SimCLR framework to learn representations by maximizing agreement between differently augmented views of the image via a contrastive loss in the latent space [9]. The first step of this self-supervised technique consists of an input example x and generating two image views \tilde{x}_i and \tilde{x}_j which we consider as positive pairs.

After the augmentations, we use neural network $f(\cdot)$ that takes the two positive pairs and encodes them into representation embeddings. We use a ResNet50 as our backbone encoder since it has been proved that self-supervised techniques benefit more from bigger models compared to supervised techniques [9]. We perform a forward pass for each of the two views and then make a global pooling. After the global pooling layer, we obtain $h_i = f(\tilde{x}_i) = \text{ResNet}(\tilde{x}_i)$ where $h_i \in \mathbb{R}^d$ with $d = 2048$.

We pass the resulting tensors to a projection head, denoted as $g(\cdot)$, responsible for mapping representations to the space where contrastive loss is applied. This projection head is removed during inference and replaced with a linear probing layer. Despite its seemingly counterintuitive nature, it has been theorized that this approach enhances the generalization of the learned embeddings, even in the absence of knowledge about the downstream task [18]. With the resulting projected embeddings, we calculate the NT-Xent loss [9] [20] such that the image encoder maps two augmentations closer in the latent space.

Our projection head consists of a multilayer perceptron (MLP) with a single hidden layer to obtain $z_i = g(h_i) = W^{(2)}\sigma(W^{(1)}h_i)$ where σ is a ReLU non-linearity, z_i represents the projected representation of the input, $W^{(1)} \in \mathbb{R}^{H \times I}$ and $W^{(2)} \in \mathbb{R}^{O \times H}$ are weight matrices associated with the hidden layer and projection layer respectively [9]. In our setting the input of the projection head is an embedding of size $I = 2048$, we used a hidden layer size of $H = 2048$ and a projection size of $O = 2048$.

Given a set $\{\tilde{x}_k\}$ including a positive pair of examples \tilde{x}_i and \tilde{x}_j , the contrastive prediction task aims to identify \tilde{x}_j in $\{\tilde{x}_k\}_{k \neq i}$ for a given \tilde{x}_i .

We sample a mini-batch of N samples, we augment them resulting in $2N$ data points. Then, given a positive pair we treat the other $2(N - 1)$ within a mini-batch as negatives similar to Chen *et al.* [9]. The NT-Xent (the normalized temperature-scaled cross-entropy loss) loss is defined as:

$$\ell_{i,j} = -\log \frac{\exp(\text{sim}(z_i, z_j) / \tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(\text{sim}(z_i, z_k) / \tau)} \quad (2)$$

where $\mathbb{1}_{[k \neq i]} \in \{0, 1\}$ is an indicator function evaluating to 1 if $k \neq i$, $\text{sim}(u, v) = u^\top v / \|u\| \|v\|$ denote the dot product between ℓ_2 normalized u and v (i.e. cosine similarity) and τ denotes a temperature parameter [9]. Further information about the image augmentations and parameters used can be found in experimental setup section in supplementary materials.

3.3 Multimodal Contrastive Learning

We leverage the previously pre-trained encoders and introduce a third encoder specifically tailored to encode tabular features. By incorporating this additional encoder, we aim to capture the rich information embedded within the tabular data, complementing the representations learned from CMRI and ECG modalities as illustrated in Figure 3.

Data augmentations: We have the paired time series data x_s , tabular data x_t , and imaging data x_i that are augmented and then processed by a signal encoder $f_s(\cdot)$, a tabular encoder f_t and an image encoder $f_i(\cdot)$ to generate the embeddings for the downstream task. As in the previous step, we take advantage of the same ECG and image augmentations. We augment the tabular features by corrupting a certain amount of randomly selected subjects' data following [9]. We used a corruption rate of 0.3. In the corruption process, the value of each feature is deliberately altered by sampling it with replacement from all observed values for that particular feature within the dataset. This technique is commonly referred

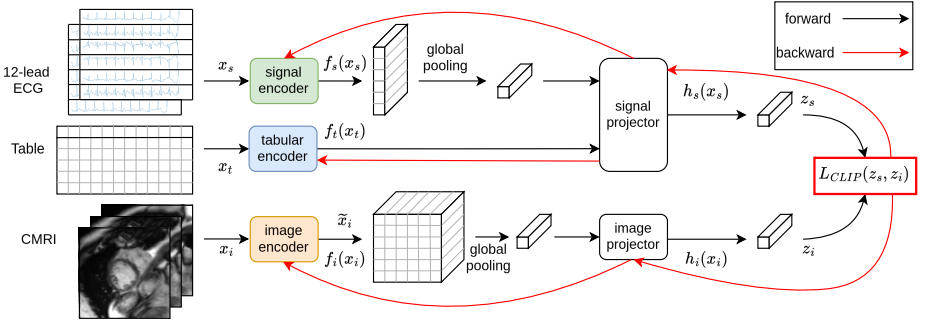


Figure 3: Visualization of the multi-modal pre-training pipeline. The three modalities are encoded taking advantage of the previous pretraining steps. The embeddings of ECG and tabular data are concatenated and optimized together. The projection head maps the two embeddings into the shared latent space where the CLIP loss is optimized.

to as sampling from the empirical marginal distribution. By introducing this corruption mechanism, the model is exposed to a range of possible variations in the data, enabling it to learn more robust and generalized representations.

Feature encoding: Our tabular encoder is an MLP with one hidden layer of size 384 that generates embedding of size 384. The weights are initialized using Xavier uniform distribution [16]. As both the image encoder and ECG encoder, we retained the architecture from the preceding step and utilized the pre-trained model weights. As the MAE encoder yielded embeddings for individual patches, we incorporated a global pooling layer to collapse the additional dimension leading to 1D embeddings. As discussed in [8], we use a joint fusion strategy to merge ECG and tabular encodings. We will refer to the concatenated embeddings of the ECG encoder and tabular encoder as signal embedding. They represent the encoding of the modalities that are cheap and easily available and that will be contrasted to the CMRI modality which is expensive and more complex to acquire. We remark that only the signal embedding will be used in the prediction of the downstream task.

Signal projection: Both signal and CMRI embeddings are given to two SimCLR projection heads $g_s(\cdot)$ and $g_i(\cdot)$ to generate projections $z_s = g_s(f_s(x_s))$ and $z_i = g_i(f_i(x_i))$, respectively. The two projection heads have a hidden size of 256 and a final size of 128. The resulting projections are ℓ_2 -normalized and mapped into a shared latent space.

CLIP loss: Furthermore, to facilitate the alignment of representations across all modalities and enhance the model’s ability to learn informative features, we employ the CLIP loss function due to its effectiveness in encouraging the model to learn representations that can match input samples from different modalities [12].

We obtain one loss for each modality:

$$\mathcal{L}^{\text{sig}} = \mathbb{E}_{p(x_s, x_i)} \mathbb{E}_{p(x_i^j)} \left[-\log \frac{\exp(z_s^\top z_i / \tau)}{\sum_{j \in B} \exp(z_s^\top z_i^j / \tau)} \right] \quad (3)$$

$$\mathcal{L}^{\text{img}} = \mathbb{E}_{p(x_i, x_s)} \mathbb{E}_{p(x_s^j)} \left[-\log \frac{\exp(z_i^\top z_s / \tau)}{\sum_{j \in B} \exp(z_i^\top z_s^j / \tau)} \right] \quad (4)$$




Figure 4: Pre-trained signal and tabular encoders are fine-tuned in a supervised manner including the labels for downstream task prediction.

then we combine it with a weighted average to obtain the total loss

$$\mathcal{L}_{\text{CLIP}} = (1 - \lambda) \mathcal{L}^{\text{sig}} + \lambda \mathcal{L}^{\text{img}} \quad (5)$$

where τ is the temperature parameter and λ is the parameter to balance the loss of the signal and imaging modalities. Further information about the multimodal SSL step and parameters used can be found in experimental setup section in supplementary materials.

3.4 Signal encoder finetuning

After enriching our embedding with multimodal contrastive learning, we fine-tune the signal encoder in a supervised fashion by training on a balanced dataset by optimizing the cross-entropy loss to predict better the downstream task as shown in Figure 4. This balanced dataset comprises all patients with diagnosed disease and randomly chosen individuals from the healthy cohort. We evaluate the final model on the same validation and test set. As in [67], we replace the global pooling with a Multi-Head Attention [68] to map the encoder output to 1D embedding. Further information about the finetuning step and parameters used can be found in experimental setup section in supplementary materials.

To address the risk of overfitting in the fine-tuning step, our approach leverages a multi-modal framework that has demonstrated strong performance even in low-data regimes, as evidenced by [19] where they tested the performance of learned encoders in a low-data regime by sampling the fine-tuning dataset down to 10% and 1% of its original size. The results showed improved representation even with fewer samples compared to uni-modality methods which mitigates the concerns about overfitting with small fine-tuning dataset.

4 Results

In this section, we undertake a thorough comparison between different self-supervised techniques as well as a comparison between self-supervised and supervised learning methodologies for cardiovascular disease prediction. To address our imbalanced dataset issue, we employ evaluation metrics such as the receiver operating characteristic area under the curve (ROC AUC) and balanced accuracy which is a more reliable evaluation metric for imbalanced datasets. Our prediction task specifically focuses on MI which is a critical medical condition that requires timely detection and intervention. Our task results in a binary classification for predicting the event based on the labels acquired from UK Biobank and shown

Table 1: Comparison of different diagnostic modalities and training strategy for CVD prediction. Columns indicate which pre-training (pretr.)/training strategy is used. Best scores are in **BOLD** font. The second best is underlined. Our approach outperforms all baseline models with regard to AUC and balanced accuracy (Bal. Acc) metrics.

SSL Modality	MAE pretr.	Image pretr.	ECG Train	Tab. Train	CMRI Train	AUC [%]	Bal. Acc [%]
Tabular	✗	✗	✗	✓	✗	0.62	0.5
ECG	✗	✗	✓	✗	✗	0.68	0.6
ECG	✓	✗	✓	✗	✗	0.76	0.69
CMRI	✗	✗	✗	✗	✓	0.74	0.67
ECG (MMCL [5])	✓	✓	✓	✗	✓	<u>0.77</u>	<u>0.7</u>
ECG + Tabular (Ours)	✓	✓	✓	✓	✓	0.80	0.71

as vascular disease in tabular features table in Supplementary materials. Since annotated data is only used during the fine-tuning step, the encoder obtained at the end of multi-modal training remains flexible and can be repurposed for other cardiovascular disease predictions such as stroke. We have extended the method for predicting stroke and showed the results in supplementary materials.

To evaluate our model, we compare the performance across different modalities, namely CMR image, ECG signal, and tabular data from medical information. We assess our model’s predictive accuracy when trained on individual modalities and combinations thereof. In examining the results of self-supervised training from table 1, it becomes apparent that approaches solely utilizing tabular data or ECG signal lack the necessary granularity to effectively discriminate between the two classes, likely due to the inherent generality of the information provided by these modalities. However, when incorporating (MAE) pretraining, the performance of ECG-based models notably improves. This enhancement suggests that MAE pretraining contributes to bolstering the robustness and diversifying the training of the self-supervised method applied to ECG encoders, yielding more meaningful embeddings. Notably, utilizing only CMR images yields respectable performance, underscoring its status as the most informative modality for this predictive task. Moreover, the MMCL technique successfully transfers knowledge from CMR image to ECG, leading to a performance boost compared to uni-modal settings. In our method, we further extend this by adding the tabular modality, ultimately outperforming the previous settings. In summary, our approach leverages the strengths of all modalities, showcasing the effectiveness of multi-modal learning in enhancing predictive accuracy for CVD prediction.

We comprehensively compare our self-supervised method with supervised learning with Neural Network (NN) for cardiovascular disease prediction. We evaluate the performance of our model under both paradigms to assess the efficacy of self-supervised pre-training in enhancing predictive accuracy when the available annotated data is limited. As a baseline, we train the NN [4] in a supervised manner on tabular data, ECG signals, and a combination of them. We compared our method with simple NN supervised approaches due to the limited availability of annotated data, which restricted our ability to explore more advanced supervised methods. We extracted the features from ECG signal [2, 8] for concatenating them

with tabular features. As seen in table 2, the performance of supervised learning on tabular data only and ECG signals only results in balanced accuracy of 54% and 66% respectively. Training the model on a combination of ECG signals and tabular information results in 65% balanced accuracy which is likely due to the overfitting on the small task-specific data using only 1524 samples with balanced diseased and healthy individuals. However, our approach with pre-training the model on multiple modalities including images and fine-tuning on labeled data outperforms the supervised method by 7.6% increase in balanced accuracy.

Table 2: Comparison between self-supervised and supervised techniques

Modality	ECG	Tabular	AUC [%]	Balanced Acc [%]
Supervised NN	✗	✓	0.55	0.54
Supervised NN	✓	✗	0.71	0.66
Supervised NN	✓	✓	0.70	0.65
Ours	✓	✓	0.80	0.71

5 Conclusions

This work presents a novel approach for enhancing the predictive performance for CVD diagnosis with limited data availability by leveraging transfer learning from CMR images to ECG signals and tabular clinical data. Our method, which incorporates self-supervised contrastive learning, offers a comprehensive framework for extracting meaningful embedding from diverse medical modalities and enhancing predictive capabilities by 7.6% in balanced accuracy compared to supervised methods when using only ECG and tabular data for prediction.

Through extensive experimentation on UK Biobank dataset, we have demonstrated the effectiveness of our approach with significant improvement in predictive performance compared to supervised models. Moreover, we demonstrated that the steps in our pipeline progressively make the model learn richer embeddings and succeed in effectively capturing and aligning relevant features across modalities. This enables the model to learn informative representations that are robust and generalizable, ultimately leading to enhanced predictive accuracy for several downstream tasks.

Furthermore, the utilization of self-supervised learning techniques mitigates the need for large annotated datasets, making our approach scalable and applicable to real-world clinical settings where labeled data may be limited. We have demonstrated the potential for benefiting from other patient information such as text data with large language models. In summary, our proposed method represents a promising avenue for advancing the field of personalized healthcare by learning holistic representations of the physiological state.

Our work is limited by the reliance on the UK Biobank, which may not fully capture the diversity of broader populations due to its specific demographic characteristics. This could affect the model’s generalizability and clinical effectiveness. Future studies should assess our approach using external data from more diverse populations to overcome these limitations and improve its global applicability. Additionally, exploring other SSL techniques might enhance patient discrimination and lead to more meaningful embeddings. Finally, assessing the limits of transfer learning from expensive to more affordable modalities could provide insights into which improvements could be made to improve these results.

Acknowledgement

This research has been conducted using the UK Biobank resource under application number 81959. We also thank Prof. Samia Mora from Harvard Medical School and Brigham and Women’s Hospital for her invaluable medical insight and guidance, which greatly contributed to the success of our research. This project was supported by the grant #2023-N-306 of the 1st Joint Call of the Swiss Data Science Center (SDSC) and the Strategic Focus Area “Personalized Health and Related Technologies (PHRT)” of the ETH Domain (Swiss Federal Institutes of Technology). N.D. is partially supported by the ETH AI Center postdoctoral fellowship. B.M. and N.D. acknowledge support of the Helmut-Horten-Foundation.

References

- [1] Julián N Acosta, Guido J Falcone, Pranav Rajpurkar, and Eric J Topol. Multimodal biomedical ai. *Nature Medicine*, 28(9):1773–1784, 2022.
- [2] Guillaume Alain and Yoshua Bengio. What regularized auto-encoders learn from the data-generating distribution. *The Journal of Machine Learning Research*, 15(1):3563–3593, 2014.
- [3] Saeed Amal, Lida Safarnejad, Jesutofunmi A Omiye, Ilies Ghanzouri, John Hanson Cabot, and Elsie Gyang Ross. Use of multi-modal data and machine learning to improve cardiovascular disease care. *Frontiers in cardiovascular medicine*, 9:840262, 2022.
- [4] Dara Bahri, Heinrich Jiang, Yi Tay, and Donald Metzler. Scarf: Self-supervised contrastive learning using random feature corruption. *arXiv preprint arXiv:2106.15147*, 2021.
- [5] Wenjia Bai, Matthew Sinclair, Giacomo Tarroni, Ozan Oktay, Martin Rajchl, Ghislain Vaillant, Aaron M. Lee, Nay Aung, Elena Lukaschuk, Mihir M. Sanghvi, Filip Zemrak, Kenneth Fung, Jose Miguel Paiva, Valentina Carapella, Young Jin Kim, Hideaki Suzuki, Bernhard Kainz, Paul M. Matthews, Steffen E. Petersen, Stefan K. Piechnik, Stefan Neubauer, Ben Glocker, and Daniel Rueckert. Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. *Journal of Cardiovascular Magnetic Resonance*, 20(1):65, 2018. ISSN 1097-6647. doi: <https://doi.org/10.1186/s12968-018-0471-x>. URL <https://www.sciencedirect.com/science/article/pii/S1097664723006026>.
- [6] Wenjia Bai, Hideaki Suzuki, Chen Qin, Giacomo Tarroni, Ozan Oktay, Paul M. Matthews, and Daniel Rueckert. Recurrent neural networks for aortic image sequence segmentation with sparse annotations. In Alejandro F. Frangi, Julia A. Schnabel, Christos Davatzikos, Carlos Alberola-López, and Gabor Fichtinger, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, pages 586–594, Cham, 2018. Springer International Publishing. ISBN 978-3-030-00937-3.
- [7] Balázs Borsos, Corinne G Allaart, and Aart van Halteren. Predicting stroke outcome: a case for multimodal deep learning methods with tabular and ct perfusion data. *Artificial Intelligence in Medicine*, 147:102719, 2024.

- [8] Ting Chen, Yizhou Sun, Yue Shi, and Liangjie Hong. On sampling strategies for neural network-based collaborative filtering. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 767–776, 2017.
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [11] Matthew Christensen, Milos Vukadinovic, Neal Yuan, and David Ouyang. Vision–language foundation model for echocardiogram interpretation. *Nature Medicine*, pages 1–8, 2024.
- [12] Zi Yu Diong, Wen Yuan Lim, and Ching Pang Goh. Self-supervised learning in medical diagnostics: An examination of simclr and byol in image classification. In *2024 3rd International Conference on Digital Transformation and Applications (ICDXA)*, pages 210–214. IEEE, 2024.
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [14] Nick Erickson, Jonas Mueller, Alexander Shirkov, Hang Zhang, Pedro Larroy, Mu Li, and Alexander Smola. Autogluon-tabular: Robust and accurate automl for structured data. *arXiv preprint arXiv:2003.06505*, 2020.
- [15] Mudasir A Ganaie, Minghui Hu, Ashwani Kumar Malik, Muhammad Tanveer, and Ponnuthurai N Suganthan. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence*, 115:105151, 2022.
- [16] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feed-forward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.
- [17] Weikang Gong, Song Bai, Ying-Qiu Zheng, Stephen M Smith, and Christian F Beckmann. Supervised phenotype discovery from multimodal brain imaging. *IEEE Transactions on Medical Imaging*, 42(3):834–849, 2022.
- [18] Kartik Gupta, Thalaisyasingam Ajanthan, Anton van den Hengel, and Stephen Gould. Understanding and improving the role of projection head in self-supervised learning. *arXiv preprint arXiv:2212.11491*, 2022.
- [19] Paul Hager, Martin J Menten, and Daniel Rueckert. 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*, pages 23924–23935, 2023.

- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020.
- [22] 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*, pages 16000–16009, 2022.
- [23] Cun Ji, Mingsen Du, Yupeng Hu, Shijun Liu, Li Pan, and Xiangwei Zheng. Time series classification based on temporal features. *Applied Soft Computing*, 128:109494, 2022.
- [24] N Kappelhof, LA Ramos, M Kappelhof, Hendrikus JA van Os, Vicky Chalos, KR van Kranendonk, ND Kruij, YBWEM Roos, Wim H van Zwam, Irene C van der Schaaf, et al. Evolutionary algorithms and decision trees for predicting poor outcome after endovascular treatment for acute ischemic stroke. *Computers in Biology and Medicine*, 133:104414, 2021.
- [25] Shaan Khurshid, Samuel Friedman, James P Pirruccello, Paolo Di Achille, Nathaniel Diamant, Christopher D Anderson, Patrick T Ellinor, Puneet Batra, Jennifer E Ho, Anthony A Philippakis, et al. Deep learning to predict cardiac magnetic resonance–derived left ventricular mass and hypertrophy from 12-lead ecgs. *Circulation: Cardiovascular Imaging*, 14(6):e012281, 2021.
- [26] Shaan Khurshid, Samuel Friedman, Christopher Reeder, Paolo Di Achille, Nathaniel Diamant, Pulkit Singh, Lia X Harrington, Xin Wang, Mostafa A Al-Alusi, Gopal Sarma, et al. Ecg-based deep learning and clinical risk factors to predict atrial fibrillation. *Circulation*, 145(2):122–133, 2022.
- [27] Megan Lindstrom, Nicole DeCleene, Henry Dorsey, Valentin Fuster, Catherine O Johnson, Kate E LeGrand, George A Mensah, Christian Razo, Benjamin Stark, Justine Varieur Turco, et al. Global burden of cardiovascular diseases and risks collaboration, 1990–2021. *Journal of the American College of Cardiology*, 80(25):2372–2425, 2022.
- [28] Carl H Lubba, Sarab S Sethi, Philip Knaute, Simon R Schultz, Ben D Fulcher, and Nick S Jones. catch22: Canonical time-series characteristics: Selected through highly comparative time-series analysis. *Data Mining and Knowledge Discovery*, 33(6):1821–1852, 2019.
- [29] Dominique Makowski, Tam Pham, Zen J. Lau, Jan C. Brammer, François Lespinasse, Hung Pham, Christopher Schölzel, and S. H. Annabel Chen. NeuroKit2: A python toolbox for neurophysiological signal processing. *Behavior Research Methods*, 53(4): 1689–1696, feb 2021. doi: 10.3758/s13428-020-01516-y. URL <https://doi.org/10.3758%2Fs13428-020-01516-y>.
- [30] Peter A Noseworthy, Zachi I Attia, Emma M Behnken, Rachel E Giblon, Katherine A Bews, Sijia Liu, Tara A Gosse, Zachery D Linn, Yihong Deng, Jun Yin, et al. Artificial intelligence-guided screening for atrial fibrillation using electrocardiogram during sinus rhythm: a prospective non-randomised interventional trial. *The Lancet*, 400(10359):1206–1212, 2022.

- [31] 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*, pages 8748–8763. PMLR, 2021.
- [32] 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*, pages 8748–8763. PMLR, 2021.
- [33] Adityanarayanan Radhakrishnan, Sam F Friedman, Shaan Khurshid, Kenney Ng, Puneet Batra, Steven A Lubitz, Anthony A Philippakis, and Caroline Uhler. Cross-modal autoencoder framework learns holistic representations of cardiovascular state. *Nature Communications*, 14(1):2436, 2023.
- [34] Zahra Raisi-Estabragh, Ahmed Salih, Polyxeni Gkontra, Angélica Atehortúa, Petia Radeva, Ilaria Boscolo Galazzo, Gloria Menegaz, Nicholas C Harvey, Karim Lekadir, and Steffen E Petersen. Estimation of biological heart age using cardiovascular magnetic resonance radiomics. *Scientific Reports*, 12(1):12805, 2022.
- [35] Shinnosuke Sawano, Satoshi Kodera, Naoto Setoguchi, Kengo Tanabe, Shunichi Kushida, Junji Kanda, Mike Saji, Mamoru Nanasato, Hisataka Maki, Hideo Fujita, et al. Applying masked autoencoder-based self-supervised learning for high-capability vision transformers of electrocardiographies. *PloS one*, 19(8):e0307978, 2024.
- [36] Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, et al. Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine*, 12(3):e1001779, 2015.
- [37] Özgün Turgut, Philip Müller, Paul Hager, Suprosanna Shit, Sophie Starck, Martin J Menten, Eimo Martens, and Daniel Rueckert. Unlocking the diagnostic potential of ecg through knowledge transfer from cardiac mri. *arXiv preprint arXiv:2308.05764*, 2023.
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [39] Lohit Velagapudi, Nikolaos Mouchtouris, Richard F Schmidt, David Vuong, Omaditya Khanna, Ahmad Sweid, Bryan Sadler, Fadi Al Saiegh, M Reid Gooch, Pascal Jabbour, et al. A machine learning approach to first pass reperfusion in mechanical thrombectomy: prediction and feature analysis. *Journal of Stroke and Cerebrovascular Diseases*, 30(7):105796, 2021.
- [40] Donghwan Yun, Hyun-Lim Yang, Soonil Kwon, So-Ryoung Lee, Kyungju Kim, Kwangsoo Kim, Hyung-Chul Lee, Chul-Woo Jung, Yon Su Kim, and Seung Seok Han. Automatic segmentation of atrial fibrillation and flutter in single-lead electrocardiograms by self-supervised learning and transformer architecture. *Journal of the American Medical Informatics Association*, 31(1):79–88, 2024.