

Multi-depth rubidium PET graph neural network for major adverse cardiac event prediction

Arthur Chevalley^{1,2}, Maria M. Asiain², Eric Moulton⁵, Ran Klein⁴, Robert DeKemp³, John O. Prior¹, Mario Jreige¹, Christel H. Kamani^{1,6}, and Adrien Depeursinge^{1,2,*}

¹ Department of Nuclear Medicine and Molecular Imaging, Centre Hospitalier Universitaire Vaudois (CHUV), Lausanne, Switzerland

² Institute of Informatics, HES-SO Valais-Wallis, Sierre, Switzerland

³ Division of Cardiology, University of Ottawa Heart Institute, Ottawa, ON, Canada

⁴ Division of Nuclear Medicine and Molecular Imaging, University of Ottawa, Ottawa, ON, Canada

⁵ Jubilant DraxImage Inc., Kirkland, QC, Canada

⁶ Department of Cardiology, CHUV, Lausanne, Switzerland

* Corresponding autor adrien.depeursinge@chuv.ch

Abstract. Major adverse cardiac events (MACE) are among the leading causes of death worldwide. Quantitative myocardium perfusion imaging (MPI) based on rubidium (Rb) PET allows measuring myocardial blood flow (MBF) in the left ventricle (LV), which proved to be predictive for MACE. Clinical routine relies on an approximate 3D LV center-line (CL) flattened to a 2D polar map (PM). PM-based artificial intelligences (AI) were proposed to predict MACE outperforming traditional non-image-based risk scores but discarding the true LV geometry and epi/endocardium perfusion. Using our fully automated 3D Rb PET processing pipeline, we developed tissue-depth-specific graph neural networks (GNNs) operating in the initial 3D PET space for our 232 patients cohort, with 47 MACE events. We compared various models using a five-fold cross-validation and fold-wise bootstrapping confidence interval estimation. Because of the substantial class imbalance, evaluation focused on average precision (AP). As an initial baseline, we compared MBF intensity along the CL to train a logistic regression classifier. When extracting min, mean and max intensity on the CL, an AP of 45.9[44.0,48.0] is reached, which was outperformed by adding epi/endocardium values with AP of 54.3[52.0,56.5]. A CNN trained on the CL-based PM reached an AP of 50.7[48.6,52.9]. The proposed GNN leveraging multi-depth epicardium, CL and endocardium perfusion (MBF, uptake and activity) in the 3D PET space reached top AP of 60.3[58.3, 62.3], AUC 82.5[81.2, 83.8] and C-Index of 78.2[77.0, 79.4]. This approach could allow more accurate identification of patients at risk of MACE when compared to current clinical assessment based global LV CL MBF.

Keywords: MACE · PET MPI · Graph Neural Network.

1 Introduction

Cardiovascular diseases (CVD) are among the leading causes of death worldwide [36]. Clinical manifestations of CVD are referred to as major adverse cardiac events (MACE) including cardiac death, myocardial infarct (MI), stroke, delayed revascularisation (>6 months post-PET/CT), congestive heart failure and de novo stable angina. Early MACE prognosis is therefore a key element for personalized patient handling. While many risk factors have been identified, e.g. smoking, nutrition, cardiovascular conditions [24], studies showed that a significant proportion of the CVD population does not suffer from any of these established risk factors [18].

Myocardial Perfusion Imaging (MPI) of the left ventricle (LV) is a powerful non-invasive tool to detect and localise heart diseases. In nuclear medicine, quantitative MPI can be obtained using PET e.g. Rubidium (^{82}Rb), which allows computing the myocardial blood flow (MBF). While many works investigated the use of artificial intelligence (AI) for MACE prediction using non-quantitative SPECT MPI, most of them extracted imaging and clinical features to train machine learning (ML) models. Even though they achieve convincing results, intensity features discard locality information as well as all geometrical considerations. More recent works used computer vision models using the MPI as input with great results, highlighting the potential complex, subtle patterns in MPI. However, given the dimensional complexity underlying MPI, they may currently be underexploited and requiring more advanced processing tools. In addition, most works convert the 3D MPI to a flattened clinical representation, referred to as polar map (PM), thus altering the LV geometry. In this work, we proposed a new multi-depth approach for MACE prediction using ^{82}Rb PET, investing the importance of epi/endocardial perfusion. We further compared the classical features-based ML to graph neural network (GNN), suggesting the importance of subtle multi-depth perfusion patterns overlooked by classic AI and global quantitative MPI.

2 Related work

Previous studies have shown that automated analysis of SPECT MPI allows better long-term risk stratification beyond expert visual interpretation, leveraging subtle or visually normal perfusion defects, as demonstrated in large cohorts such as the REFINE-SPECT [23], where several works investigated the use of feature-based ML for MACE. [30] demonstrated that ML models leveraging only a limited set of automatically derived imaging features combined with key clinical variables can better predict MACE than standard visual SPECT interpretation or traditional multivariable risk models. Further work reinforced this conclusion by adding coronary CT angiography features to provide highly accurate patient-level MACE predictions [2].

While these works relied on predefined clinical and imaging features, subsequent work have shown that deep learning (DL) can extract prognostic information directly from MPI. While [19] used raw MPI data, enabling fully data-driven

estimation of cardiovascular risk, most works convert the 3D MPI to PMs, from which convolutional neural networks (CNNs) were trained [33, 28].

These works relied on SPECT and few investigated PET for MACE prediction even though quantitative MPI was shown to be predictive using DL for all cause mortality [32]. [16] investigated MBF derived quantities (myocardial reserve and resistance) shown to be predictive of MACE without using DL. [4] used [^{82}Rb] PET PMs to train CNNs and features-based models. Comparing approaches, they showed that intensity features were the best performing in survival prediction while CNNs slightly outperformed them for the classification task.

While these works demonstrated first relevance of AI to predict MACE from MPI, they were not able to fully leverage quantitative PET MPI due to the limited datasets available for MACE prediction¹ or because of important geometric distortions of MPI such as the use of CNN with PM, discarding subtle patterns of perfusion anomalies. We thus propose a new approach, partially investigated by [34], using GNN to keep the signal as raw as possible.

3 Methods

This section presents the proposed Rb PET processing pipeline to obtain the multi-depth LV centerlines as well as the DL algorithms to process them. Training scheme and the selected evaluation metrics are also detailed.




Fig. 1. Illustration of the usual clinical pipeline (top) and proposed automatic pipeline (bottom) for graph-based characterization of multi-depth transmural LV perfusion.

3.1 Clinical [^{82}Rb] PET MPI

The current clinical MPI processing pipeline is generally as follows. First, the tracer is injected and a dynamic scan is conducted for 5-6 minutes before being reconstructed in 20-21 frames [^{82}Rb] activity volumes. An approximate LV

¹ Public release of large scale datasets such as REFINE-PET could alleviate this [29].

centerline (CL) is then semi-automatically computed using commercial software, e.g. FlowQuant or Corridor 4DM [8, 11]. This step is crucial for a proper MBF quantification but still depends on the software and users [14]. Using a one-tissue compartment model (1TCM), one can model the blood dynamics between the myocardium tissue $C_m(t)$ and arterial blood $C_a(t)$ activities through uptake K_1 and washout k_2 rates [20]. From the uptake rate K_1 , the MBF can be computed via the Renkin-Crone formula, Eq. (1) [20].

$$K_1(x) = MBF(x) \cdot (1 - \alpha e^{-\beta/MBF(x)}), \text{ with } \alpha=0.79 \text{ and } \beta=0.55. \quad (1)$$

Using the 1TCM and solving Eq. (1) on the CL points, a single PM with 3 or 17 territories [5] is often obtained by flattening the LV to the polar domain. While this approach is widely used, PMs distort the LV 3D geometry and do not allow multi-depth perfusion characterization (see Fig 1). In addition, [14] showed the impact of user and software on MBF values² due to CL fitting and $C_a(t)$. To alleviate this we propose an automatic multi-depth CL generation pipeline.

3.2 Fully Automatic Multi-Depth MPI Processing Pipeline

To reduce the CL location variability and to allow multi-depth LV characterization, we developed an automated two-step pipeline depicted in Fig. 1. An nnU-Net processes 3D PET to generate a static segmentation mask, LV_{Seg} [3]. The mask is then used to compute the LV CL, LV_{CL} , to obtain mid-muscle perfusion.

Subtracting LV_{Seg} from its convex hull allows to retrieve the blood pool cavity, LV_{BP} . By using erosion, dilation and subtraction between LV_{Seg} and LV_{BP} , LV inner and outer borders, $LV_{inner/outer}$ are extracted and post-processed to remove small objects. A KNN classifier [13] is trained to separate outer and inner points, then used to predict the probability of each class for all the segmentation mask points, thus separating the LV in dense epi-/endocardium $Epi/Endo_{Dense}$. Finally, the intersection of Epi_{Dense} and $Endo_{Dense}$ defines the CL points. Doing this recursively allows to obtain different muscle-depth signals, e.g. $Epi/Endo_{CL}$ is computed by completing the same procedure on $Epi/Endo_{Dense}$ instead of LV_{Seg} , corresponding to the CL between LV_{CL} and $LV_{Outer/Inner}$.

As the resulting point cloud is defined in the scanner space, the Rb activity can be easily sampled. We use the model proposed in [22] based on the arterial input function from FlowQuant [8], to predict 3D MBF and K_1 volumes, which are further sampled according to the point clouds. FlowQuant's 3/17 segments PM parcellation is then mapped to the predicted CL.

Graph Construction A graph is constructed by connecting neighboring CL point clouds vertices closer than $\sqrt{3}$ voxels. After outlier filtering (due to basal irregularities), graphs have a median number of 3133 nodes (1637 to 6187), which is correlated to LV_{Seg} volume (Pearson's correlation of 0.94). Due to scanner

² Note that other works [27, 15] reported limited MBF variations across software.

resolution, the graphs are downsampled by a factor of ten to maintain the heart morphology differences. Minimum spanning trees allow drastically reducing the graph complexity, easing training and mitigating over-smoothing. CL point cloud normals are automatically estimated to obtain epi-endo signals across LV tissue depth as the intersection between LV_{CL} normals and Epi_{CL} & $Endo_{CL}$. Myocardial flow reserve (MFR, defined as $\frac{MBF_{Stress}}{MBF_{Rest}}$) was shown to be predictive for MACE [9]. As the graphs change between stress and rest, due to heart deformation, MFR cannot be computed. However, to help the models learn combined rest-stress features, we used RANSAC [12] to register the graphs in their domain.

3.3 Dataset

Participants with suspected myocardial ischemia were enrolled to undergo [^{82}Rb] cardiac Silicon PhotoMultiplier (SiPM) PET/CT under approved ethics study protocol *anonymous*. Following-up 232 patients revealed 47 who experienced MACE after Rb exam, with an average of 572 days (2 to 862 days) between the exam and a potential event. Quantitative PMs and segmental analysis were obtained with FlowQuant. In addition, quantitative 3D volumes and automatic CLs were obtained with our proposed pipeline. MACE subtypes included cardiac death, delayed revascularisation (>6 months post-PET/CT), MI, congestive heart failure, or de novo stable angina. We extracted min, mean and max of global MBF at stress and rest as [9] showed global MBF, MFR and capacity to be predictive of cardiovascular events. Important clinical features (CF) were collected including age, sex, height, weight, blood pressure and heart rate as potential non-MPI MACE predictors.

3.4 Machine learning models

In this section, we present the various ML graph and baseline models used to predict MACE. Following earlier work on PET [4] and SPECT [2, 30, 23] we investigated the performance of baseline feature-based and CNN models based on the PM representation [4, 28, 33], which include important geometric distortions of the perfusion along the LV.

Feature-based models Following [4], we trained logistic regression (LR) models using ElasticNet [38] based on two kinds of intensity features: min-mean-max (MMM) intensities or 18 intensity radiomics [37] features (Rad_{int}). Intensity was extracted from three signals: flow (MBF), uptake rate (K_1) and raw normalised myocardium tracer activity (C_m) and at each depth to investigate the importance of transmural perfusion (i.e. epi, CL and endo).

Graph Neural Network To account for the complex and variable heart morphology, we use graph neural networks (GNN) to remove geometric distortions coming from PM representations and avoid interpolation. While standard neural

networks assume grid-like data structures, GNNs are uniquely suited for non-Euclidean domains represented as graphs $G = (V, E)$, where V represents nodes and E represents edges. GNNs can capture the L -hop neighborhood of a node through L successive layers, allowing the model to learn relevant connectivity patterns. The fundamental operation of a GNN layer involves updating the representation of a node v based on its neighbours u in the set $N(v)$. We used a L -layer and K -head TransformerConv [31] following the node update as

$$\mathbf{h}_v^{l,k} = \text{ReLU}(\mathbf{W}_1^{l,k} h_v^{l-1} + \sum_{u \in N(v)} \text{softmax}\left(\frac{(\mathbf{W}_3^{l,k} h_v^{l-1})^T (\mathbf{W}_4^{l,k} h_u^{l-1})}{\sqrt{d}}\right) \mathbf{W}_2^{l,k} h_u^{l-1}) \quad (2)$$

where $h_v^{l,k}$ is the hidden state of node v at layer l for head k , d the node degree and $\mathbf{W}_i^{l,k}$ the weight matrix of head k at layer l . Depending on the architecture, $\mathbf{h}_v^l$ is computed by concatenating or averaging $h_v^{l,k}$ over the K heads.

3.5 Model training and evaluation

Given the unbalanced proportion of MACE events and the high temporal variability between the exam and the potential event, the MACE prediction task is formulated both as classification and survival. All models were trained using a stratified five-fold cross-validation. Optuna [1] was used to select a unique set of parameters performing well on all folds, using the same splits as the final models to avoid data leakage. Fold-wise bootstrapping is used to compute performance metrics, whose variability is estimated via a second bootstrap on the pooled metrics to determine 95% confidence intervals (CI).

ElasticNet LR models were implemented with scikit-learn [26] and using Gridsearch for $L1_{ratio} \in [0.1, 1.0]$. The Cox model [7] was used for the survival task. The GNNs were implemented in PyTorch 2.1.1 [25] and Torch Geometric 2.7 [10] and trained jointly for classification and survival tasks. Our final GNN has four layers with skip connections and 32 hidden channels with 6 heads concatenated. The resulting node-features were averaged using graph-wise mean across nodes, further fed to two distinct fully-connected layer, one predicting event probability and the other risk hazard. The resulting model was trained using AdamW [21] with a weight decay of 0.2, learning rate of 0.01 and dropout probability of 40%, optimising an addition of cross-entropy loss and Cox partial negative log likelihood loss proposed in deepSurv [17], weighted by a factor 0.15, for 100 epochs³. To encode spatial information, an approximated 17-segment label is added as a feature by mapping FlowQuant PM to the scanner space. Note that the feature-based approach do not consider the segments' information as some segments are not well represented and detrimental to the models. Future works will investigate 17-segment delimitation in 3D volume. Code is available upon request. The dataset cannot be shared due to legal restrictions.

³ Number of epochs was set to 100 as the losses reached a plateau. All models were trained on a NVIDIA A100 with five-fold training time of 2 to 5 minutes.

We report two probability-based metrics: the area under the receiver operating characteristic curve (AUC) and the average precision (AP) summarising the precision–recall curve. While AUC is suboptimal for imbalanced datasets, AP focuses on the model’s ability to detect the positive minority class across predicted probabilities, particularly appropriate for less represented MACE events. We also report two thresholded metrics. The F1-score combines precision and recall into a single harmonic-mean measure, emphasising correct detection of the minority class. The Matthews correlation coefficient (MCC) provides a correlation-like summary of the confusion matrix. MCC is more reliable for overall performances in imbalanced datasets [6] but less focused on positive class detection. To evaluate the survival task, we used the Concordance index (C-Index) [35] measuring concordance between predicted hazard score and time-to-event.

4 Results

Comparison of the predictive value of the multi-depth perfusion is listed in Table 1. All models are trained using both rest and stress MPI. The first row reports the results when training a CNN with FlowQuant’s PMs. The next two double rows compare the results of models using either the CL solely or concatenated epi-, CL and endocardium (*ECE*) perfusion with MMM or *Rad_{int}* features. The next multi-row reports the performances using GNN for the different included signals using nodes with epi-, CL and endocardium perfusion. The last two GNN rows report the performances when combining K_1 and tracer activity values referred to as *Multi* and the addition of clinical features.

Table 2 reports the results of the two best performing models of Table 1 in the survival task. It is worth recalling that the GNN is jointly trained for survival and classification. LR with *multi* is omitted due to worst performances.

Model	Depth	AUC	AP	F1	MCC
CNN _{PM}	MBF _{CLPM}	74.2[72.7,75.6]	50.8[48.6,52.9]	47.5[45.7,49.5]	33.74[31.31,36.12]
LR _{MMM}	MBF _{CL}	69.8[68.4,71.1]	45.9[44.0,48.0]	35.3[33.4,37.3]	15.0[12.4,17.6]
	MBF _{ECE}	79.3[78.0,80.5]	54.3[52.0,56.5]	48.2[46.3,50.0]	34.8[32.5,37.1]
LR _{Rad_{Int}}	MBF _{CL}	74.3[73.0,75.6]	47.7[45.8,49.7]	44.7[42.9,46.4]	28.7[26.5,30.8]
	MBF _{ECE}	76.5[75.3,77.7]	50.4[48.4,52.5]	47.2[45.5,48.9]	32.7[30.5,34.9]
GNN	MBF _{CL}	76.1[74.9, 77.2]	47.8[46.1, 49.6]	48.4[46.8, 49.9]	34.3[32.2, 36.2]
	MBF _{ECE}	80.1[78.8, 81.4]	55.8[53.8, 58.0]	<i>53.7[52.2, 55.2]</i>	43.1[41.4, 44.9]
	Multi _{ECE}	<i>80.2[78.7, 81.6]</i>	<i>57.4[54.9, 59.8]</i>	51.6[49.9, 53.3]	39.3[37.2, 41.5]
	Multi _{ECE} +CF	82.5[81.2, 83.8]	60.3[58.3, 62.3]	53.7[51.9, 55.6]	<i>42.2[39.7, 44.6]</i>

Table 1. MACE classification performance comparing of the relevance of transmural perfusion (i.e. epi-, CL and endocardial depths). *Depth* column details the signal considered (i.e. MBF or *Multi* combining MBF, K_1 and normalized activity) sampled at a given depth (ECE for all epicardium, CL and endocardium values). Rest and stress values were concatenated for all models. Bold indicates best & italic second best.

Model	Depth	C-Index
Cox _{MMM}	MBF _{CL}	65.4[63.9,66.8]
	MBF _{ECE}	72.5[71.3,73.6]
GNN	MBF _{CL}	75.8[74.6, 76.9]
	MBF _{ECE}	76.8[75.7, 78.0]
	Multi _{ECE}	79.6[78.3, 80.8]
	Multi _{ECE} +CL	78.2[76.9, 79.4]

Table 2. MACE prognosis prediction (i.e. survival), following notation from Table 1. Rest and stress values were concatenated. Bold indicates best and italic second best.

5 Discussions and Conclusions

In this work we addressed two research hypotheses for MACE prediction. First, the prognostic value of multi-depth perfusion was investigated. Multiple ML approaches, previously reported to be predictive [4], were trained using either CL perfusion only or the concatenation of epicardium, CL and endocardium perfusion (Table 1). To assess transmural perfusion at these three depths, we proposed an automated Rb PET MPI processing pipeline which segments the LV and further derives epi-, CL and endocardium perfusion. All multi-depth models outperformed the CL-only models for both classification and survival tasks, often with non-overlapping CIs. A CNN using rest-stress PMs slightly outperformed non-graph based CL-only models, though not significantly.

As a second research question, we investigated the use of GNNs to best leverage multi-depth perfusion signals across the diverse heart morphologies. While CNN and GNN reached similar results when using CL only (overlapping CIs), multi-depth GNN outperforms CNN with non-overlapping CIs (bottom rows of Table 1). When comparing *ECE* models, feature-based and GNN reached similar results for AUC ($\sim 80\%$) and AP ($\sim 55\%$), GNN significantly outperformed all models in terms of thresholded metrics (F1 +5.1% and MCC +8.3%). When comparing survival performances (Table 2), GNNs systematically outperformed all Cox models, even solely based on MBFs reaching 76.8% C-Index (+4.3%).

Finally, the GNN approach was enriched by adding *Multi* inputs (K_1 and normalized activity) on top of MBF. While the classification performances are similar to MBF only models, the survival performance is significantly increased (+2.8% C-Index). Finally, adding the CF as a late fusion before the final fully-connected layers did not significantly impacted the performance, motivating further investigations to build optimal cardiac GNNs.

To conclude, we demonstrated the relevance of multi-depth perfusion characterization for MACE prediction in Rb PET MPI. To best leverage the latter, we proposed a fully automated pipeline and graph-based ML, which could allow more reproducible and accurate identification of patients at risk of MACE when compared to current clinical assessment based on global LV CL MBF.

Acknowledgments. This work was partially funded by the Swiss National Science Foundation (SNSF) with the project 219430.

Disclosure of Interests. R.D.K. & R.K. receive royalty revenues from Rb technologies licensed to Jubilant DraxImage. E.M. is an employee at Jubilant DraxImage.

References

1. Akiba, T., et al.: Optuna: A next-generation hyperparameter optimization framework. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 2623–2631 (2019)
2. Alahdab, F., et al.: Patient-level explainable machine learning to predict major adverse cardiovascular events from SPECT MPI and CCTA imaging. *PLoS One* **18**(11), e0291451 (2023)
3. Amirian, M., et al.: Left Ventricle Segmentation in Dynamic 82Rb PET/CT Using Deep Convolutional Neural Networks. In: Annual Congress of the European Association of Nuclear Medicine, October 4–8, 2025 | Barcelona, Spain. p. 67 (2025)
4. Bors, S., et al.: Predicting MACE from [82Rb] PET: Can AI outperform more traditional quantitative assessment of the myocardial perfusion? (2023)
5. American Heart Association Writing Group on Myocardial Segmentation Registration for Cardiac Imaging, ., et al.: Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart: a statement for healthcare professionals from the Cardiac Imaging Committee of the Council on Clinical Cardiology of the American Heart Association. *Circulation* **105**(4), 539–542 (2002)
6. Chicco, D., Jurman, G.: The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC genomics* **21**(1), 6 (2020)
7. Cox, D.R.: Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)* **34**(2), 187–202 (1972)
8. Dekemp, R.A., et al.: Multisoftware reproducibility study of stress and rest myocardial blood flow assessed with 3D dynamic PET/CT and a 1-tissue-compartment model of 82Rb kinetics. *Journal of Nuclear Medicine* **54**(4), 571–577 (2013)
9. Dietz, M., et al.: Comparison of the prognostic value of impaired stress myocardial blood flow, myocardial flow reserve, and myocardial flow capacity on low-dose Rubidium-82 SiPM PET/CT. *Journal of Nuclear Cardiology* **30**(4), 1385–1395 (2023)
10. Fey, M., Lenssen, J.E.: Fast graph representation learning with PyTorch Geometric. arXiv preprint arXiv:1903.02428 (2019)
11. Ficaro, E.P., et al.: Corridor4DM: the Michigan method for quantitative nuclear cardiology. *Journal of nuclear cardiology* **14**(4), 455–465 (2007)
12. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* **24**(6), 381–395 (1981)
13. Fix, E.: Discriminatory analysis: nonparametric discrimination, consistency properties, vol. 1. USAF school of Aviation Medicine (1985)
14. Gould, K.L., et al.: Reliability and reproducibility of absolute myocardial blood flow: does it depend on the PET/CT technology, the vasodilator, and/or the software? *Current Cardiology Reports* **23**(3), 12 (2021)
15. Jreige, M., et al.: Quantification of Myocardial Blood Flow Using High-Resolution, High-Sensitivity and 214-ps SiPM82Rb-PET/CT: A comparison of three software packages. (2019)
16. Kadoya, Y., et al.: Prognostic utility of quantitative perfusion PET in patients with prior CABG: incremental value of myocardial flow reserve and coronary vascular resistance. *Circulation: Cardiovascular Imaging* **18**(12), e018204 (2025)
17. Katzman, J.L., et al.: DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network. *BMC medical research methodology* **18**(1), 24 (2018)

18. Khot, U.N., et al.: Prevalence of conventional risk factors in patients with coronary heart disease. *Jama* **290**(7), 898–904 (2003)
19. Li, Y.L., et al.: Predicting long-term time to cardiovascular incidents using myocardial perfusion imaging and deep convolutional neural networks. *Scientific reports* **14**(1), 3802 (2024)
20. Lortie, M., et al.: Quantification of myocardial blood flow with ^{82}Rb dynamic PET imaging. *European journal of nuclear medicine and molecular imaging* **34**(11), 1765–1774 (2007)
21. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017)
22. Moulton, E., et al.: Artificial Intelligence-Accelerated Parametric Imaging of Myocardial Blood Flow with Rubidium-82 PET (2024)
23. Otaki, Y., et al.: 5-year prognostic value of quantitative versus visual mpi in subtle perfusion defects: results from refine spect. *Cardiovascular Imaging* **13**(3), 774–785 (2020)
24. Palaniappan, L.P., et al.: 2026 Heart Disease and Stroke Statistics: A Report of US and Global Data From the American Heart Association. *Circulation* (2026)
25. Paszke, A., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
26. Pedregosa, F., et al.: Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
27. Piccinelli, M., et al.: Direct comparison of two clinical software packages for the quantification of absolute myocardial blood flow and flow reserve from ^{82}Rb PET/CT images. (2023)
28. Pieszko, K., et al.: Time and event-specific deep learning for personalized risk assessment after cardiac perfusion imaging. *npj Digital Medicine* **6**(1), 78 (2023)
29. Ramirez, G., et al.: The REgistry of flow and perfusion imaging for artificial Intelligence with PET (REFINE PET): rationale and design. *Journal of Nuclear Cardiology* p. 102449 (2025)
30. Rios, R., et al.: Determining a minimum set of variables for machine learning cardiovascular event prediction: results from REFINE SPECT registry. *Cardiovascular research* **118**(9), 2152–2164 (2022)
31. Shi, Y., et al.: Masked label prediction: Unified message passing model for semi-supervised classification. *arXiv preprint arXiv:2009.03509* (2020)
32. Singh, A., et al.: Deep learning for explainable estimation of mortality risk from myocardial positron emission tomography images. *Circulation: Cardiovascular Imaging* **15**(9), e014526 (2022)
33. Singh, A., et al.: Direct risk assessment from myocardial perfusion imaging using explainable deep learning. *Cardiovascular Imaging* **16**(2), 209–220 (2023)
34. Spier, N., et al.: Classification of polar maps from cardiac perfusion imaging with graph-convolutional neural networks. *Scientific reports* **9**(1), 7569 (2019)
35. Uno, H., et al.: On the C-statistics for evaluating overall adequacy of risk prediction procedures with censored survival data. *Statistics in medicine* **30**(10), 1105–1117 (2011)
36. Vaduganathan, M., et al.: The global burden of cardiovascular diseases and risk: a compass for future health (2022)
37. Van Griethuysen, J.J., et al.: Computational radiomics system to decode the radiographic phenotype. *Cancer research* **77**(21), e104–e107 (2017)
38. Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **67**(2), 301–320 (2005)