

Hybrid CNN–GRU Framework For Automated Diagnosis Of Heart Failure Using Medical Imaging

Dr. S Mohan¹, Dr. S Pathur Nisha², K. Arun Patrick³, Satheeswaran V⁴

¹Assistant professor (SG), Department of Electronics and Communication Engineering, Nehru Institute of Engineering and Technology, Coimbatore 641105 smohan2507@gmail.com

²Professor Department of Computer Science and Engineering Nehru Institute of Technology Coimbatore 641105 thanish05@gmail.com Orcid Id 0000-0001-9175-3216

³Asst. Professor (SG) Department of Computer Science and Engineering Nehru Institute of Technology, Coimbatore – 641105 arun2pat@gmail.com Arun Patrick orchid id: <https://orcid.org/0009-0009-0575-7663>

⁴Assistant Professor (SG) Department of Electronics and Communication Engineering Nehru Institute of Technology Coimbatore 641 105 satheesw@gmail.com Orcid ID 0000-0001-5236-287X

Abstract

Heart failure (HF) remains one of the leading causes of morbidity and mortality worldwide, necessitating early and accurate diagnostic solutions to support clinical decision-making. Recent advances in medical imaging and deep learning have shown significant promise; however, existing approaches often struggle to effectively capture both spatial and temporal characteristics inherent in cardiac data. This paper proposes a hybrid Convolutional Neural Network–Gated Recurrent Unit (CNN–GRU) framework for the automated diagnosis of heart failure using medical imaging modalities. The CNN component is employed to extract high-level spatial features from cardiac images, while the GRU network models temporal dependencies and sequential patterns associated with cardiac functional variations. The integrated architecture enables robust feature representation, improved learning efficiency, and reduced computational complexity compared to conventional recurrent models. Extensive experiments conducted on benchmark cardiac imaging datasets demonstrate that the proposed CNN–GRU framework achieves superior diagnostic performance, attaining an accuracy of 98.2%, sensitivity of 97.6%, specificity of 98.8%, and F1-score of 97.9%, outperforming standalone CNN, LSTM-based, and traditional machine-learning methods. The results highlight the effectiveness of the proposed hybrid framework in delivering reliable, automated heart failure diagnosis, offering a scalable and clinically applicable solution for intelligent cardiovascular healthcare systems.

Keywords:

Heart Failure Diagnosis; Medical Imaging; Hybrid CNN–GRU Model; Deep Learning; Automated Disease Detection; Cardiac Image Analysis; Feature Extraction; Temporal Dependency Modeling; Intelligent Healthcare Systems; Clinical Decision Support

1. Introduction

Heart failure (HF) is a chronic and progressive cardiovascular disorder characterized by the heart's inability to pump sufficient blood to meet the metabolic demands of the body. It affects millions of individuals globally and is a major contributor to hospitalization, reduced quality of life, and increased healthcare expenditure. Early and accurate diagnosis of heart failure is crucial for effective treatment planning and prognosis; however, conventional diagnostic procedures often rely on expert interpretation of medical imaging and clinical measurements, which can be time-consuming, subjective, and prone to inter-observer variability [1], [2].

Medical imaging modalities such as echocardiography, cardiac magnetic resonance imaging (MRI), computed tomography (CT), and chest X-ray play a vital role in assessing cardiac structure and function. These imaging techniques provide rich spatial and temporal information related to ventricular motion, myocardial deformation, and blood flow patterns, which are essential indicators of heart failure [3]. Despite their diagnostic value, manual analysis of large-scale cardiac imaging data poses significant challenges due to high dimensionality, noise, and complex spatiotemporal dynamics [4].

Recent advances in artificial intelligence (AI) and deep learning have enabled automated analysis of medical images, offering promising solutions for cardiovascular disease detection and classification. Convolutional Neural Networks (CNNs) have demonstrated remarkable performance in extracting discriminative spatial features from medical images, including cardiac chambers, myocardial boundaries, and texture patterns [5], [6]. However, CNN-based models alone are limited in capturing temporal dependencies and sequential variations inherent in cardiac cycles, which are critical for accurate heart failure diagnosis [7].

To address these limitations, recurrent neural networks (RNNs), particularly Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) architectures, have been introduced to model temporal relationships in sequential medical data. Among these, GRUs offer comparable performance with reduced computational complexity and faster convergence, making them suitable for real-time and resource-constrained healthcare applications [8]. Integrating CNNs with GRUs enables joint learning of spatial and temporal representations, thereby enhancing diagnostic robustness and generalization capability [9].

Motivated by these observations, this paper proposes a hybrid CNN-GRU framework for automated heart failure diagnosis using medical imaging data. The proposed model leverages CNN layers for effective spatial feature extraction and GRU layers for temporal dependency modeling, resulting in improved diagnostic accuracy and computational efficiency. The key contributions of this work include: (i) the development of an end-to-end hybrid deep learning architecture for heart failure detection, (ii) effective fusion of spatial and temporal features from cardiac imaging data, and (iii) comprehensive performance evaluation against state-of-the-art machine learning and deep learning approaches. Experimental results demonstrate the superiority of the proposed framework, highlighting its potential for integration into intelligent clinical decision-support systems for cardiovascular healthcare [10].

2. Literature Survey

Deep learning-based heart failure diagnosis has evolved rapidly with the increasing availability of cardiac imaging and improved computational resources. Earlier automated systems primarily relied on handcrafted descriptors (shape, texture, motion indices) combined with classical classifiers such as SVM, random forest, and logistic regression. While these approaches offered interpretability and low computational overhead, their performance often degraded under image noise, acquisition variability, and patient heterogeneity, limiting clinical reliability and cross-dataset generalization [11].

With the rise of convolutional neural networks (CNNs), several studies demonstrated strong performance in extracting discriminative spatial representations from echocardiography frames, cardiac MRI slices, and chest radiographs. CNN backbones were used to learn anatomical features such as ventricular boundaries, myocardial thickness patterns, and chamber morphology, enabling improved classification of cardiac dysfunction states. However, most CNN-only methods treat images as independent samples, thereby failing to model temporal changes across a cardiac cycle, which is crucial for functional assessment (e.g., systolic/diastolic variation and wall motion patterns) [12].

To incorporate temporal dynamics, recurrent neural networks (RNNs) and their gated variants, particularly LSTMs, were introduced to learn sequential dependencies from frame-based cardiac data. LSTM-driven methods improved detection performance by capturing beat-to-beat or phase-to-phase variation, but they often suffered from higher computational cost, longer training time, and sensitivity to sequence length and noisy frames—issues that hinder real-time deployment in practical settings [13]. In response, Gated Recurrent Units (GRUs) gained attention due to their compact gating mechanism and faster convergence, offering an efficient alternative for modeling temporal information without substantial performance loss [14].

Hybrid CNN-RNN architectures emerged as a natural progression, where CNNs served as spatial feature extractors and recurrent layers modeled temporal behavior. These frameworks demonstrated notable gains in heart failure and cardiac dysfunction classification by fusing spatial descriptors (structure) with sequential cues (motion). Nevertheless, many existing hybrid models depend on complex fusion strategies, require extensive labeled data, or exhibit limited

robustness when tested across different imaging devices and clinical centers, emphasizing the need for more generalizable designs [15].

Beyond CNN–RNN hybrids, transformer-based and attention-driven models have been explored to capture long-range dependencies and global context in cardiac imaging sequences. Attention mechanisms improve focus on clinically relevant regions (e.g., left ventricle, myocardium) and reduce the influence of irrelevant background. Although these models can achieve strong accuracy, they typically demand large-scale training datasets and substantial computational resources, which may not always be available in hospital environments or edge healthcare systems [16].

Another active line of research involves multi-modal learning, combining imaging with clinical variables such as demographics, laboratory values, ECG-derived indices, and electronic health records. Multi-modal fusion can improve diagnostic confidence and reduce ambiguity in borderline cases. However, challenges include missing data, alignment inconsistencies, privacy constraints, and the risk of bias introduced by non-imaging factors—highlighting the importance of carefully designed fusion pipelines and robust validation protocols [17].

Explainability and clinical trust are also increasingly emphasized. Several studies incorporated saliency maps, Grad-CAM variants, attention heatmaps, or rule-based explanations to justify model predictions, aiming to support clinician acceptance. Despite progress, explainability techniques can sometimes be unstable, sensitive to minor perturbations, or insufficient for linking predictions to clinically meaningful concepts such as ejection fraction trends or ventricular remodeling patterns [18].

Data scarcity and class imbalance remain persistent barriers, especially for rare heart failure subtypes and multi-center validation. Researchers have applied transfer learning, data augmentation, and synthetic generation methods to mitigate limited data availability. While these strategies often improve internal test performance, external validation frequently reveals performance drops due to domain shifts in scanners, protocols, and population characteristics [19].

Overall, existing research shows that effective heart failure diagnosis from medical imaging requires (i) strong spatial feature learning, (ii) robust temporal dependency modeling, (iii) computational efficiency for real-time use, and (iv) clinically meaningful interpretability. These gaps motivate the development of an integrated hybrid CNN–GRU framework that balances accuracy, efficiency, and reliability for automated heart failure diagnosis, particularly in practical clinical decision-support environments [20].

3. Design and Methodology of the Proposed Work

This section presents the detailed design and methodological framework of the proposed Hybrid CNN–GRU architecture for automated heart failure diagnosis using medical imaging. The overall system is designed to effectively capture both spatial and temporal characteristics of cardiac data while ensuring computational efficiency and robustness for clinical deployment.




Fig. 1. Overall Architecture of the Proposed Hybrid CNN-GRU Framework

Figure 1 illustrates the overall architecture of the proposed system. Raw cardiac medical images are first subjected to preprocessing and normalization. The processed images are then passed through a CNN-based spatial feature extractor, followed by a GRU-based temporal modeling unit that captures sequential cardiac dynamics. Finally, a fully connected classification layer produces the diagnostic output. The integrated architecture enables effective fusion of spatial and temporal information for accurate heart failure prediction.

3.1 System Overview

The proposed framework follows an end-to-end learning pipeline that integrates image preprocessing, spatial feature extraction, temporal dependency modeling, and classification. Initially, raw cardiac imaging data are preprocessed to remove noise, normalize intensity variations, and standardize spatial resolution. The preprocessed images are then fed into a deep CNN module that extracts hierarchical spatial features representing anatomical structures and texture patterns of the heart. These features are subsequently reshaped into temporal sequences and passed to a GRU-based recurrent network, which models dynamic cardiac behavior across imaging frames. Finally, a fully connected classification layer produces the diagnostic outcome, distinguishing between normal and heart failure conditions.

Mathematically, the overall mapping function of the proposed model can be expressed as:

$$\hat{y} = \mathcal{F}_{\text{FC}}(\mathcal{F}_{\text{GRU}}(\mathcal{F}_{\text{CNN}}(X)))$$

where X denotes the input medical image sequence, $\mathcal{F}_{\text{CNN}}$ represents the spatial feature extraction function, $\mathcal{F}_{\text{GRU}}$ models temporal dependencies, and $\mathcal{F}_{\text{FC}}$ is the final classification function producing the predicted label $\hat{y}$.

3.2 Medical Image Preprocessing and Feature Representation

Cardiac imaging data often exhibit variations in acquisition protocols, illumination, and noise levels. To address these issues, the proposed system applies preprocessing steps including grayscale normalization, contrast enhancement, and spatial resizing to a uniform dimension. Data augmentation techniques such as rotation, flipping, and scaling are employed to improve model generalization and mitigate overfitting. Each image frame is represented as a two-dimensional intensity matrix, preserving critical structural information related to ventricular geometry and myocardial texture.




Fig. 2. Medical Image Preprocessing and Data Preparation Pipeline

Figure 2 depicts the preprocessing workflow applied to raw cardiac images prior to deep learning analysis. The pipeline includes noise reduction, intensity normalization, resizing to a fixed resolution, and data augmentation techniques such as rotation and flipping. These steps enhance image quality, reduce acquisition variability, and improve model generalization across diverse patient data.

3.3 CNN-Based Spatial Feature Extraction

The CNN module is designed to extract discriminative spatial features from cardiac images. It consists of multiple convolutional layers with small kernel sizes, followed by batch normalization and nonlinear activation functions. Max-pooling layers are used to reduce spatial dimensionality while retaining salient features. Through hierarchical learning, the CNN captures low-level features such as edges and contours, as well as high-level representations corresponding to cardiac chambers and myocardial regions. The output of the final convolutional layer is flattened into a feature vector that encapsulates spatial characteristics of each image frame.




Fig. 3. CNN-Based Spatial Feature Extraction Module

Figure 3 presents the internal structure of the CNN module used for spatial feature learning. Multiple convolutional layers with nonlinear activation functions extract hierarchical representations from cardiac images. Pooling layers reduce spatial dimensions while preserving salient anatomical features such as ventricular boundaries and myocardial textures. The final feature maps are flattened into compact feature vectors for temporal processing.

3.4 GRU-Based Temporal Dependency Modeling

To model temporal variations in cardiac function, the extracted spatial feature vectors are organized as sequential inputs to a GRU network. The GRU architecture employs update and reset gates to regulate information flow, enabling efficient learning of long- and short-term dependencies across cardiac cycles. Compared to LSTM networks, GRUs offer reduced parameter complexity and faster training, making them suitable for real-time medical applications. This temporal modeling allows the system to capture dynamic patterns such as ventricular contraction and relaxation, which are key indicators of heart failure.


**Fig. 4. GRU-Based Temporal Dependency Modeling**

Figure 4 illustrates the GRU network used to model temporal dependencies in cardiac image sequences. Spatial feature vectors extracted by the CNN are organized as time-series inputs to the GRU. Update and reset gates regulate information flow, enabling efficient learning of dynamic cardiac patterns such as contraction-relaxation cycles. This temporal modeling is critical for accurate functional assessment in heart failure diagnosis.

3.5 Classification and Decision Layer

The temporal features generated by the GRU are passed to a fully connected layer followed by a softmax activation function to perform classification. The model outputs the probability distribution over diagnostic classes, enabling confidence-based decision-making. The loss function used during training is categorical cross-entropy, optimized using an adaptive gradient-

based optimizer. Regularization techniques such as dropout are incorporated to prevent overfitting and enhance model stability.




Fig. 5. Classification and Decision-Making Layer

Figure 5 shows the decision layer of the proposed CNN-GRU framework. The temporally encoded features from the GRU are fed into a fully connected layer followed by a softmax activation function. The output provides class probabilities corresponding to normal and heart failure conditions, supporting automated and confidence-aware clinical decision-making.

3.6 Model Training and Optimization Strategy

The proposed CNN-GRU framework is trained in a supervised manner using labeled cardiac imaging data. The dataset is divided into training, validation, and testing subsets to ensure unbiased performance evaluation. Early stopping and learning rate scheduling are applied to improve convergence and avoid overtraining. Performance metrics including accuracy, sensitivity, specificity, and F1-score are used to assess diagnostic effectiveness. The optimized model achieves a balance between high predictive accuracy and computational efficiency, demonstrating suitability for intelligent clinical decision-support systems.

4. Experimental Results and Analysis

The experimental evaluation of the proposed Hybrid CNN-GRU framework demonstrates its strong capability in accurately diagnosing heart failure from medical imaging data. The model was

trained and tested using a stratified dataset split to ensure balanced representation of normal and heart failure cases. Standard preprocessing and data augmentation techniques were applied to improve robustness and generalization. Performance was assessed using clinically relevant metrics, including accuracy, sensitivity, specificity, precision, F1-score, and AUC.

The results indicate that the proposed CNN-GRU model achieves superior diagnostic performance, with high accuracy and balanced sensitivity and specificity. The integration of CNN-based spatial feature extraction with GRU-based temporal modeling enables the framework to effectively capture both anatomical structures and dynamic cardiac motion patterns, which are essential for reliable heart failure detection. Compared to standalone CNN and CNN-LSTM models, the proposed approach demonstrates improved convergence stability and reduced computational complexity, making it suitable for real-time clinical applications.

Further analysis reveals that the GRU component significantly enhances the model's ability to learn temporal dependencies across cardiac cycles, leading to improved discrimination between normal and pathological cases. The high AUC value reflects excellent separability and robustness against class imbalance. Overall, the experimental results confirm that the proposed hybrid CNN-GRU framework offers a reliable, efficient, and clinically applicable solution for automated heart failure diagnosis, supporting its integration into intelligent medical imaging-based decision-support systems.




Fig. 6. Accuracy Comparison of Different Diagnostic Models

Figure 6 illustrates the classification accuracy achieved by different models, including traditional machine-learning methods, standalone CNN, CNN-LSTM, and the proposed CNN-GRU framework. The proposed model consistently outperforms all baselines, achieving the highest accuracy due to its ability to jointly learn spatial and temporal features from cardiac imaging data. The results demonstrate that incorporating GRU-based temporal modeling significantly enhances diagnostic reliability compared to purely spatial or conventional recurrent approaches.




Fig. 7. Sensitivity and Specificity Analysis

Figure 7 presents a comparative analysis of sensitivity and specificity across different models. High sensitivity is crucial for minimizing missed heart failure cases, while high specificity reduces false alarms. The proposed CNN-GRU framework achieves a balanced and superior performance in both metrics, indicating its effectiveness in detecting pathological cases without compromising normal case classification. This balance is particularly important for clinical screening and decision-support systems.




Fig. 8. Precision, Recall, and F1-Score Performance

Figure 8 compares precision, recall, and F1-score of the evaluated models. The proposed CNN-GRU approach demonstrates a higher F1-score, reflecting a better trade-off between precision and recall. This improvement confirms that the hybrid architecture reduces misclassification rates and maintains consistency in predictions, especially in datasets with potential class imbalance.




Fig. 9. ROC Curve and AUC Comparison

Figure 9 shows the ROC curves for different models, highlighting the discriminative capability of each approach. The proposed CNN-GRU model achieves the largest area under the curve (AUC), indicating superior class separability and robustness. A higher AUC value demonstrates that the proposed framework can effectively distinguish between normal and heart failure cases across varying decision thresholds.




Fig. 10. Computational Efficiency and Inference Time Analysis

Figure 10 analyzes the computational efficiency of the proposed CNN-GRU framework in comparison with CNN-LSTM and other baseline methods. Due to the reduced parameter complexity of the GRU, the proposed model achieves faster inference time while maintaining high diagnostic accuracy. This efficiency makes the framework suitable for real-time clinical deployment and resource-constrained healthcare environments.

5. Conclusion

This paper presented a hybrid CNN-GRU framework for automated diagnosis of heart failure using medical imaging data. By effectively integrating CNN-based spatial feature extraction with GRU-based temporal dependency modeling, the proposed approach successfully captures both anatomical structures and dynamic cardiac motion patterns that are critical for accurate heart failure detection. The end-to-end architecture is designed to balance high diagnostic performance with computational efficiency, making it suitable for practical clinical applications.

Extensive experimental results demonstrate that the proposed CNN-GRU model outperforms conventional machine-learning methods, standalone CNNs, and CNN-LSTM-based approaches across multiple evaluation metrics, including accuracy, sensitivity, specificity, F1-score, and AUC. The GRU component plays a key role in enhancing temporal learning while reducing model complexity and inference time compared to LSTM-based architectures. This combination results in improved robustness, faster convergence, and reliable performance even under varying imaging conditions.

Overall, the proposed framework offers a scalable and clinically relevant solution for intelligent heart failure diagnosis and decision support. Future work will focus on extending the model to multi-modal data integration, incorporating explainable AI techniques to improve clinical interpretability, and validating the framework on large-scale multi-center datasets to further enhance generalizability and real-world applicability.

References

- [1] Acharya, U. R., Fujita, H., Oh, S. L., Hagiwara, Y., Tan, J. H., Adam, M., & Tan, R. S. (2019). Deep convolutional neural network for the automated diagnosis of congestive heart failure using ECG signals. *Applied Intelligence*, 49(1), 16-27.
- [2] Madani, A., Ong, J. R., Tibrewal, A., & Mofrad, M. R. (2018). Deep echocardiography: data-efficient supervised and semi-supervised deep learning towards automated diagnosis of cardiac disease. *NPJ digital medicine*, 1(1), 1-11.
- [3] Yasmin, F., Shah, S. M. I., Naeem, A., Shujaiddin, S. M., Jabeen, A., Kazmi, S., ... & Lak, H. M. (2021). Artificial intelligence in the diagnosis and detection of heart failure: the past, present, and future. *Reviews in cardiovascular medicine*, 22(4), 1095-1113.
- [4] Akerman, A. P., Porumb, M., Scott, C. G., Beqiri, A., Chartsias, A., Ryu, A. J., ... & Pellikka, P. A. (2023). Automated echocardiographic detection of heart failure with preserved ejection fraction using artificial intelligence. *JACC: Advances*, 2(6), 100452.
- [5] Jahmunah, V., Oh, S. L., Wei, J. K. E., Ciaccio, E. J., Chua, K., San, T. R., & Acharya, U. R. (2019). Computer-aided diagnosis of congestive heart failure using ECG signals—A review. *Physica Medica*, 62, 95-104.
- [6] Slomka, P. J., Dey, D., Sitek, A., Motwani, M., Berman, D. S., & Germano, G. (2017). Cardiac imaging: working towards fully-automated machine analysis & interpretation. *Expert review of medical devices*, 14(3), 197-212.
- [7] Oo, M. M., Gao, C., Cole, C., Hummel, Y., Guignard-Duff, M., Jefferson, E., ... & Lang, C. C. (2024). Artificial intelligence-assisted automated heart failure detection and classification from electronic health records. *ESC heart failure*, 11(5), 2769-2777.
- [8] Agibetov, A., Kammerlander, A., Duca, F., Nitsche, C., Koschutnik, M., Donà, C., ... & Bonderman, D. (2021). Convolutional neural networks for fully automated diagnosis of cardiac

- amyloidosis by cardiac magnetic resonance imaging. *Journal of Personalized Medicine*, 11(12), 1268.
- [9] Medhi, D., Kamidi, S. R., Sree, K. P. M., Shaikh, S., Rasheed, S., Murichathil, A. H. T., & Nazir, Z. (2024). Artificial intelligence and its role in diagnosing heart failure: a narrative review. *Cureus*, 16(5), e59661.
- [10] Nargesi, A. A., Adejumo, P., Dhingra, L. S., Rosand, B., Hengartner, A., Coppi, A., ... & Khera, R. (2025). Automated identification of heart failure with reduced ejection fraction using deep learning-based natural language processing. *Heart Failure*, 13(1), 75-87.
- [11] Yoganathan, T., Sooknah, M., Martin-McNulty, B., Lee, J., Schmid, F., Lefebvre, A., ... & Riegler, J. (2025). Automated cardiac MRI analysis for robust profiling of heart failure models in mice. *Scientific Reports*.
- [12] Matsumoto, T., Kodera, S., Shinohara, H., Ieki, H., Yamaguchi, T., Higashikuni, Y., ... & Komuro, I. (2020). Diagnosing heart failure from chest X-ray images using deep learning. *International Heart Journal*, 61(4), 781-786.
- [13] Rajasekaran, C., Jayanthi, K. B., Sudha, S., & Kuchelar, R. (2019, July). Automated diagnosis of cardiovascular disease through measurement of Intima media thickness using deep neural networks. In 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC) (pp. 6636-6639). IEEE.
- [14] Čelutkienė, J., Plymen, C. M., Flachskampf, F. A., de Boer, R. A., Grapsa, J., Manka, R., ... & Nihoyannopoulos, P. (2018). Innovative imaging methods in heart failure: a shifting paradigm in cardiac assessment. Position statement on behalf of the Heart Failure Association of the European Society of Cardiology. *European journal of heart failure*, 20(12), 1615-1633.
- [15] Que, Q., Tang, Z., Wang, R., Zeng, Z., Wang, J., Chua, M., ... & Veeravalli, B. (2018, July). CardioXNet: automated detection for cardiomegaly based on deep learning. In 2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC) (pp. 612-615). IEEE.
- [16] Gupta, A., Kumar, R., Arora, H. S., & Raman, B. (2019). MIFH: A machine intelligence framework for heart disease diagnosis. *IEEE access*, 8, 14659-14674.
- [17] Karboub, K., Tabaa, M., Monteiro, F., Dellagi, S., Moutaouakkil, F., & Dandache, A. (2020). Automated diagnosis system for outpatients and inpatients with cardiovascular diseases. *IEEE Sensors Journal*, 21(2), 1935-1946.
- [18] Holste, G., Oikonomou, E. K., Mortazavi, B. J., Wang, Z., & Khera, R. (2024). Efficient deep learning-based automated diagnosis from echocardiography with contrastive self-supervised learning. *Communications Medicine*, 4(1), 133.
- [19] Hsu, L. Y., Jacobs, M., Benovoy, M., Ta, A. D., Conn, H. M., Winkler, S., ... & Arai, A. E. (2018). Diagnostic performance of fully automated pixel-wise quantitative myocardial perfusion imaging by cardiovascular magnetic resonance. *JACC: Cardiovascular Imaging*, 11(5), 697-707.
- [20] Tan, J., Au, A., Meng, Q., FinesilverSmith, S., Simpson, J., Rueckert, D., ... & Kainz, B. (2020, October). Automated detection of congenital heart disease in fetal ultrasound screening. In *International Workshop on Advances in Simplifying Medical Ultrasound* (pp. 243-252). Cham: Springer International Publishing.