CME LIVE OPEN ACCESS

Cardiology & Vascular Research

Review Article

A Medical Imaging Approach for Recognising Mitral Regurgitation Through Machine Learning Methods in Cardiac Imaging

Vijay Singh[1], Anwar Ali Sathio[1], Saeed Anwar[1], Raja Vavekanand[2*] and R Danish[3]

[1]*Faculty of Computing Science and Information Technology, BBSUL.*

[2]*Irish Computing Society.*

[3]*Department of Public Health, UOS.*

Corresponding Author: Raja Vavekanand, 2Irish Computing Society.

Received: 2024 Jul 11 **Accepted:** 2024 Jul 30 **Published:** 2024 Aug 13

Abstract

Mitral regurgitation (MR) is a serious heart valve disease that can have devastating consequences if left untreated. Timely diagnosis and treatment are crucial to prevent further complications, but traditional diagnostic methods pose significant challenges. These methods are not only expensive but also labor-intensive, requiring specialized clinical expertise, which creates barriers to effective MR screening. To address these challenges, we propose a novel semi supervised model for MR classification called CUSSP. CUSSP is designed to process cardiac imaging slices from the 4- chamber view of the heart, utilizing standard computer vision techniques and contrastive models to learn from large amounts of unlabeled data. This approach enables the model to leverage the vast amounts of available imaging data, even when labeled data is scarce. By employing specialized classifiers, CUSSP creates the first automated MR classification system, revolutionizing the diagnosis and treatment of this critical heart condition. The significance of CUSSP lies in its ability to overcome the limitations of traditional diagnostic methods. Cardiac imaging is a complex and time-consuming process, requiring specialized expertise to interpret and diagnose MR. CUSSP automates this process, enabling healthcare professionals to focus on high- value tasks while ensuring accurate and timely diagnoses. The performance of CUSSP is impressive, achieving an F1 score of 0.69 and a ROC-AUC score of 0.88 on a test set of 179 labeled sequences, comprising 154 non-MR and 25 MR cases. These results establish the initial benchmark for this new task, demonstrating the potential of CUSSP to transform MR diagnosis. The CUSSP model is trained using a semi supervised approach, combining the strengths of both supervised and unsupervised learning. This approach enables the model to learn from large amounts of unlabeled data, leveraging the inherent patterns and relationships within the imaging data. By incorporating specialized classifiers, CUSSP achieves high accuracy and robustness, even in the presence of limited labeled data. The use of contrastive models in CUSSP is a key innovation, enabling the model to learn from unlabeled data and adapt to new patterns and variations. This approach allows CUSSP to generalize well to unseen data, ensuring accurate diagnoses even in cases with complex or rare presentations. The implications of CUSSP are far-reaching, with the potential to improve patient outcomes, reduce healthcare costs, and enhance the efficiency of cardiac care. By automating MR diagnosis, CUSSP can help address the growing demand for cardiac imaging services, enabling healthcare providers to focus on high-value tasks and improving patient care. CUSSP represents a significant breakthrough in MR diagnosis, offering a novel semisupervised approach to automated classification. By leveraging computer vision techniques, contrastive models, and specialized classifiers, CUSSP achieves high accuracy and robustness, establishing a new benchmark for MR diagnosis. With its potential to transform cardiac care, CUSSP is poised to make a meaningful impact on patient outcomes and healthcare efficiency.

Keywords: Machine Learning, Mitral Regurgitation (MR), Medical Imaging, Cardiac Imaging.

1. Introduction

1.1 Mitral Regurgitation

Mitral regurgitation (MR) is a valvular heart disease characterized by the incomplete closure of the mitral valve during systole, the phase when the left ventricle contracts, resulting in the backflow of blood from the left ventricle (LV) into the left atrium (LA) – as illustrated in Figure 1. MR can arise from organic or functional causes [1-6]. With organic MR leading to the enlargement of the atrium and annulus, and functional MR increasing atrial pressure [7]. As MR progresses, it can lead to arrhythmia, shortness of breath, heart palpitations, and pulmonary hypertension

[8-14]. If left undiagnosed and untreated, MR can cause significant hemodynamic instability and congestive heart failure, potentially leading to death [15-17]. Acute MR typically requires immediate medical intervention [18-22]. Therefore, early detection and assessment are crucial for optimal treatment outcomes. The best short-term and long-term results are achieved in asymptomatic patients who undergo surgery in advanced repair centers with low operative mortality (<1%) and high repair rates (≥80–90%) [7].

1.2 MR Diagnosis

MR is often detected only after symptoms appear. For asymptomatic patients, the quantitative grading of MR is a strong indicator for clinical treatment, such as immediate cardiac surgery [8]. Clinically, MR is typically diagnosed using Doppler echocardiography, followed by cardiovascular magnetic resonance (CMR) to assess MR severity and accurately quantify the regurgitant volume, a key severity indicator [20]. Studies evaluating CMR for MR assessment often use the difference between left ventricular stroke volume (LVSV) and forward stroke volume (FSV), with LVSV typically estimated from short-axis (SA) view CMR – a 4-D tensor – and FSV determined by aortic phase-contrast velocity- encoding images [20]. This diagnostic process requires continuous involvement from expert clinicians and specific processing of phase-contrast images of the proximal aorta or main pulmonary artery during CMR data acquisition. The expense associated with this standard diagnostic procedure presents a barrier to large-scale MR screening in the general population.

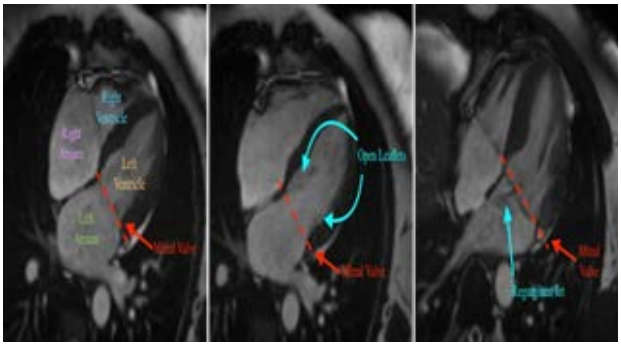

Figure 1: Three Cardiovascular Magnetic Resonance (CMR) Images Showing the Long-axis Four-chamber View of the heart.

Left: A Heart with Normal Mitral Valve. Middle: A Heart with Normal Mitral Valve When the Valve Leaflets are Open. Right: A Heart With Mitral Regurgitation. The Red Dotted Line Denotes the Mitral Valve.

diagnose MR, enabling wide-scale screening. We use long-axis 4CH CMR imaging data from the UK Biobank, including data from over 30,000 subjects, of which 704 were labeled by an expert cardiologist [1]. While the 4CH view can identify MR when the regurgitant jet is visible, the imaging lacks comprehensive annotations or diagnoses for individual patients. To address this challenge, we rely on weakly supervised and unsupervised methods. Weakly supervised deep learning has proven successful in detecting other heart pathologies. For instance, Fries et al used a weakly supervised deep learning method (CNN- LSTM) to classify aortic valve malformations from aortic valve cross-section CMR data in the UK Biobank, successfully extracting critical features of the aortic valve opening shape [9]. Similarly, Vimalesvaran et al. developed a deep learning- based pipeline to detect aortic valve pathology using 3CH CMR imaging from three hospitals, with fully annotated data sets including landmarks, stenotic jets, and regurgitant jets [21]. Unlike these studies, we faced the challenge of extracting complex mitral valve regurgitant features from 4CH CMR images with no annotations for landmarks, regurgitant jets, or easily extractable features, and only a small amount of binary MR labels. To our knowledge, this is the first study to identify MR using 4CH CMR imaging data in an automated pipeline.

1.3 Noval Approach

We propose an automated five-stage pipeline called Cardiovascular Magnetic Resonance U-Net Localized Self-Supervised Predictor (CUSSP). This approach integrates several existing neural network architectures to address the challenges of MR classification. Specifically, we use a U-Net for segmenting heart chambers, which helps localize the area around the mitral valve [18]. We enhance the appearance of the valve using histogram equalization, then use a Barlow Twins network to learn unsupervised representations of blood flow around the valve. A Siamese network is employed to learn differences between MR and non- MR instances. During training, CUSSP leverages a large amount of unlabeled CMR images with minimal supervision from a small set of MR labels annotated by a cardiologist. However, at test time, CUSSP operates fully automatically [23, 26].

2. Methodology

2.1 Segmentation of the Cardiac Magnetic Resonance Images

The CMR imaging data from the UK Biobank relevant to MR detection includes long-axis 2- chamber (2CH) and long-axis 4-chamber (4CH) views, as illustrated in Figure 2. Additionally, the short-axis view CMR offers precise descriptions of the left ventricle. Both long-axis and short-axis views are utilized to estimate heart measurements pertinent to MR detection, although only the long-axis 4CH view is employed for the deep learning models.

As part of the preprocessing steps, we performed semantic

Citation: Vavekanand, R., Singh, V., Sathio, A, A., Anwar, S., Danish, R. (2024). A Medical Imaging Approach for Recognising Mitral Regurgitation Through Machine Learning Methods in Cardiac Imaging. Cardiology & Vascular Research. 2 (1). 1-8.

segmentation on the CMR imaging data, using masks (Figure 2) generated by a U-Net segmentation model to emphasize regions of interest for MR classification. U-Net is currently the leading architecture for medical imaging segmentation, with various U-Net variants developed for specific applications [18]. Ternaus Net, a U-Net variant, reshapes the U-Net encoder to match the VGG11 architecture, allowing it to leverage pre-trained VGG11 model weights for faster convergence and improved segmentation outcomes [12, 19]. While most medical imaging segmentation models are trained using supervised learning, weakly supervised segmentation methods such as Voxel Morph augmented segmentation, ACNN, CCNN, graph-based unsupervised segmentation, and GAN-based unsupervised segmentation also achieve comparable results. For the segmentation of the 4CH, 2CH, SA, and aorta view CMR imaging dataset from the UK Biobank, Bai et al. provide a supervised segmentation model [2, 13, 15, 16, 23, 24, 27]. In Figure 2, we present examples of the segmentation outputs for the long- axis 4CH (left), 2CH (middle), and short-axis (right) CMR imaging data. We manually labeled 100 CMR images for each view and trained a supervised segmentation model using the Ternaus Net architecture [2]. The segmentation outputs, shown in Figure 2, are used to calculate measurements of cardiac structure and function for the four heart chambers, summarized in Table 1. The short-axis view CMR segmentation output is used to estimate measurements for the left and right ventricles, while the long-axis 4CH and 2CH views are used to estimate measurements for the left and right atria. Specifically, the left atrial volume is estimated using the biplane method with segmentation of both the 2CH and 4CH views, while the right atrial volume is estimated using the single plane method with segmentation of the 4CH view.

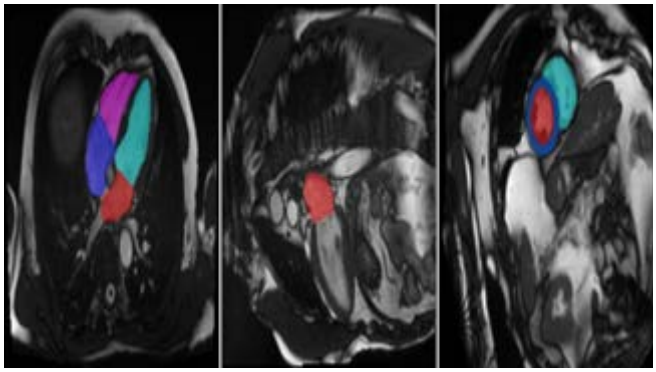

Figure 2: Example of the Segmentation Outputs of the Long Axis 4ch (left), 2ch (Middle) CMR View Imaging Data and the Short Axis (Right) CMR Imaging Data.

2.2 Three MR Classification Models

We consider two baseline models, random forests (Section 2.2.1) and a CNNLSTM (Section 2.2.2). We then present our CUSSP model in section 2.2.3.

2.3 Random Forest

baseline We first considered a random forest (RF) classifier trained for MR classification on the tabular heart measurements derived from the semantic segmentation masks, as described in Section 2.1. We divided the 18 features by body surface area (BSA) prior to training the RF [3].

Table 1: Cardiac measurements derived from the semantic segmentation of the CMR.

Left Atrium	Right Atrium	Left Ventricle	Right Ventricle
Vol Max (mL)	*Vol Max (mL)*	*End-Systolic Vol (mL)*	*End- Systolic Vol (mL)*
Volume Min (mL)	*Vol Mi (mL)*	*End- Diastolic Vol (mL)*	*End- Diastolic Vol (mL)*
Stroke Vol (mL)	*Stroke Vol (mL)*	*Stroke Vol (mL)*	*Stroke Vol (mL)*
Ejection Fraction (%)	*Ejection Fraction (%)*	*Ejection Fraction (%) Cardiac Output (L/min) Mass (g)*	*Ejection Fraction (%)*

2.4 Weakly Supervised CNN-LSTM Baseline

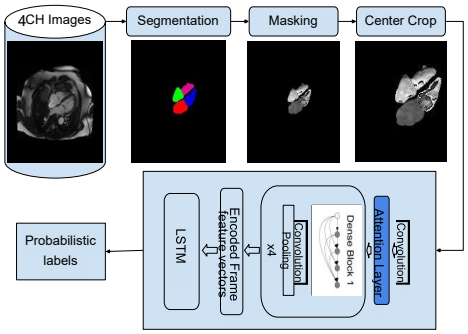

Figure 3: Overview of the CNN-LSTM Method Pipeline for MR Classification.

2.5 The CUSSP Framework

encoder is fine-tuned in the Siamese network, it is combined with a 3-layer multi-layer perceptron (MLP) network to form a classifier, which is trained on the same labeled dataset. To enhance computation efficiency and training accuracy, we also tested the framework using a smaller window of 25 frames, since MR occurs between diastole and systole.

2.6 Conceptualization

To enhance the encoding of blood flow information pertinent to MR classification from the 4CH CMR view, we explored self-supervised representation learning methods that can harness all the unlabeled CMR sequences available in the UK Biobank. Typically, self-supervised representation learning for visual data involves maximizing the similarity between representations of diverse distorted versions of a sample. Among various self- supervised architectures, including Sim CLR , SWAV , and BYOL , we opted for Barlow Twins due to its batch-size independence [4, 5, 10, 26]. With labeled data, our siamese network compares the representation differences between classes by sampling two inputs from different classes, as demonstrated in [25]. Thus, our CUSSP MR classification pipeline leverages both self-supervised and supervised representation learning.

2.7 Test-time Pipeline

Our CUSSP method consists of five main steps, depicted in Figure 4, with the first two steps representing data preprocessing and the latter three steps utilizing network components trained for MR classification, as described in the next section. The preprocessing of the CMR imaging sequence is detailed in Figure 8 in the Appendix. We employed the segmentation model in 2.1 to identify the mitral valve and the orientation of the left ventricle. We then extracted a square patch with the mitral valve at its center, positioned horizontally. After cropping, we applied histogram equalization to the patch with the pixel intensity range of the left atrium. The resulting patches are utilized by the downstream networks.

2.8 Training Process

The first step involves training a representation encoder in a Barlow Twins network using over 30,000 unlabeled preprocessed sequences. We utilized ResNet-18 with an output dimension of 512 as the encoder, with hidden dimension and projector output dimensions of 2048. After training the encoder with the unlabeled dataset, it is fine-tuned in a siamese network using a relatively smaller labeled set, as described in 3.1. During training, two sequences are sampled from the labeled dataset, with the first being non-MR and the second being either MR or non-MR. The two sequences are passed through the representation encoder to obtain embeddings, which are then used to calculate the contrastive loss. The model is trained to maximize contrastive loss when the two samples are non-MR and MR and to minimize it when both are non-MR. Once the representation encoder is fine-tuned in the siamese network, it is combined with a 3-layer multi-layer perceptron (MLP) network to form a classifier, which is trained on the same labeled dataset. To enhance computation efficiency and training accuracy, we also tested the framework using a smaller window of 25 frames, since MR occurs between diastole and systole.

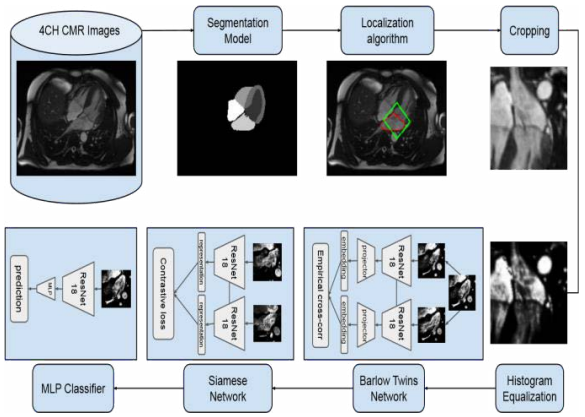

Figure 4: Overview of the CUSSP Pipeline for MR Classification, With its 5 Steps: (1) Segmentation, (2) Localization, (3) Cropping, (4) Equalization, and (5) Prediction.

Citation: Vavekanand, R., Singh, V., Sathio, A, A., Anwar, S., Danish, R. (2024). A Medical Imaging Approach for Recognising Mitral Regurgitation Through Machine Learning Methods in Cardiac Imaging. Cardiology & Vascular Research. 2 (1). 1-8.

3. Experiments

3.1 Experimental Setup

We used 4CH CMR images to conduct experiments with both the CNN-LSTM and CUSSP methods. Our dataset comprised 704 labeled sequences, divided into a training set and a test set. For the training set, we selected 525 sequences, of which 452 were labeled as non-MR and 73 as MR. The remaining 179 sequences were reserved for testing, with 154 labeled as non-MR and 25 as MR. Given the significant class imbalance, we chose the F1 score as our primary evaluation metric to balance precision and recall. Additionally, we reported precision and recall to provide a comprehensive evaluation of the model performance.

3.2 Random Forest Classification Results

The random forest model is trained with 10- fold cross validation, with a random search over a parameter grid of n_estimators (10–100), max_depth (2– 16), max_features (sqrt,log2), min_samples_leaf (2–8). The optimal hyperparameter setting found is: n_estimators = 20, max_features = log2, max_depth = 8, min_samples_leaf = 2. The best results obtained are presented in Table 2.

Table 2: Experimental results show that CUSSP-1, 2, 3, and 25 are trained using the Barlow Twins-MLP model without fine-tuning with the Siamese network, while USSP-SIAM and CUSSPSIAM-25 are trained using the same model.

Model	Pos. Acc	Neg. Acc	Precision	Recall	F1	AUC
RF	0.09	0.99	0.43	0.09	0.14	0.58
CNN-LSTM	0.53	0.86	0.45	0.53	0.44	0.72
CUSSP-1	0.38	0.87	0.29	0.38	0.32	0.65
CUSSP-2	0.29	0.87	0.25	0.29	0.27	0.63
CUSSP-3	0.38	0.90	0.35	0.38	0.36	0.66
CUSSP-SIAM	0.55	0.96	0.66	0.55	0.60	0.80
CUSSP-SIAM-25	0.62	0.96	0.8	0.62	0.69	0.88

3.3 CNN-LSTM Classification Results

We conducted experiments on the Dense Net-LSTM classification model using various input image sizes, attention layer configurations, and masks. The best CNN-LSTM model attains a F1-score of 0.44, shown in Table 2.

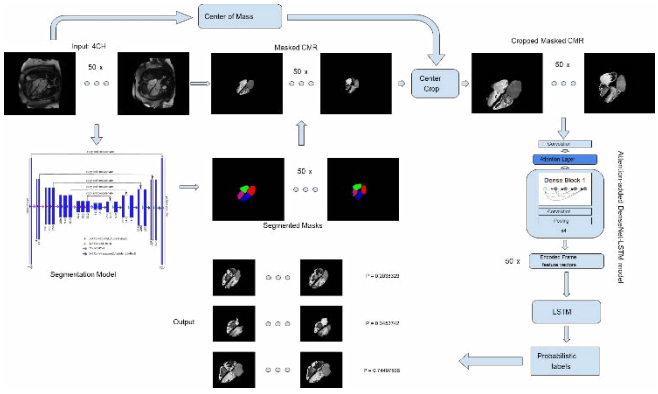

Figure 5: Detailed Overview of the CNN-LSTM Method Pipeline for MR Classification

3.4 CUSSP Classification Results

We evaluated various configurations of the CUSSP model to determine the relative benefits of different components. In the first configuration, we combined the ResNet18 model with a 3-layer MLP to train a classifier using the labeled training set after pre-training it in the Barlow Twins network with the unlabeled dataset. During classifier training, we weighted the cross-correlation loss from the Barlow Twins network and the cross-entropy loss from the binary classification using three different configurations. For CUSSP-1, the cross-correlation loss had a weight of 0.9, while the cross-entropy loss had a weight of 0.1. For CUSSP-2, both losses were weighted equally at 0.5. For CUSSP-3, the weights were reversed to 0.1 for the cross-correlation loss and 0.9 for the cross-entropy loss. Both CUSSP-1 and CUSSP-3 outperformed CUSSP-2, although their performance remained low, highlighting the necessity for fine-tuning, which we discuss below.

In the second scenario, we fine-tuned the encoder with a Siamese network to enhance the quality of the encoded representations post-training in the Barlow Twins network. To prevent overfitting and limit the model's capacity, we froze the parameters of all layers except the last block of the ResNet18 encoder during the training of the Siamese network and the classifier. This fine-tuned model, named CUSSP- SIAM, exhibited significant performance improvement. In the final configuration, CUSSP-SIAM-25, we

truncated the number of frames in the training sequences from 50 to 25 frames, focusing on the interval during which mitral regurgitation occurs. This approach further refined the model's performance. The results are summarized in Table 2, and the ROC-AUC curve for CUSSP-SIAM-25 is shown in Figure 5.

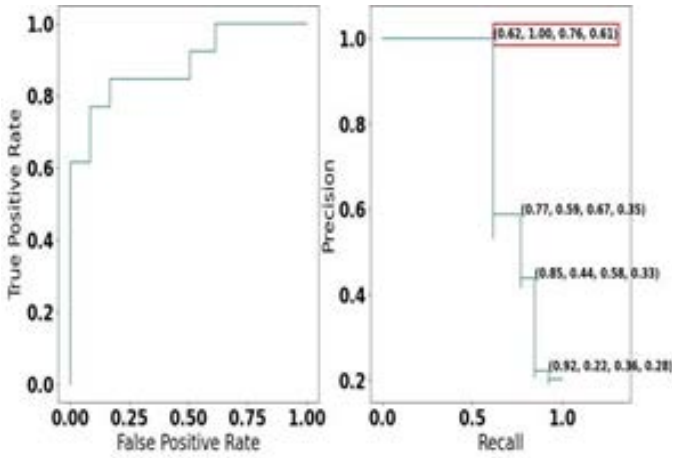

Figure 6: The ROC AUC Curve and the Precision-recall Curve of CUSSP. The Annotated Coordinates on the Precision-recall Curve Plot are (Recall, Precision, f1-score, Threshold).

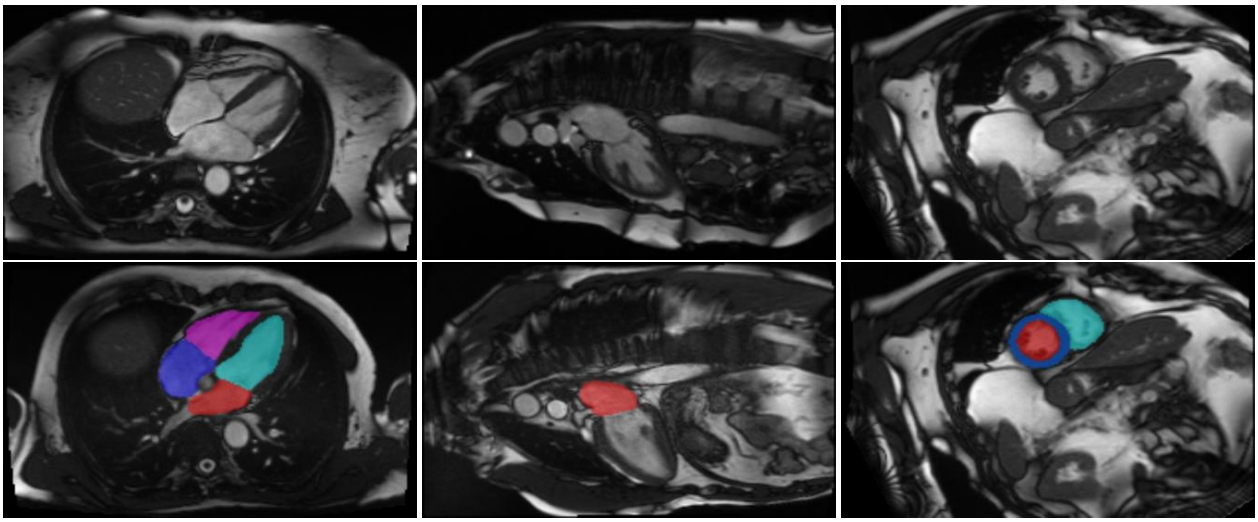

Figure 7: Overview of the dataset. Top: Example of the long axis 4CH (left), 2CH (middle) CMR view imaging data and the short axis (right) CMR imaging data. Bottom: Example of the segmentation outputs of the long axis 4CH (left), 2CH (middle) CMR view imaging data and the short axis (right) CMR imaging data.

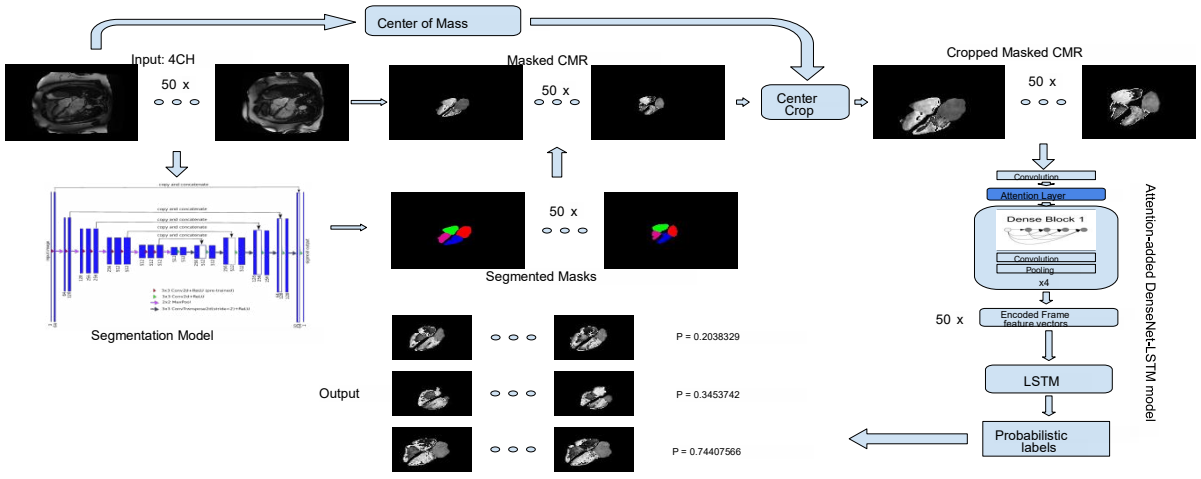

Figure 8: Detailed overview of the CNN-LSTM method pipeline for MR classification.

Citation: Vavekanand, R., Singh, V., Sathio, A, A., Anwar, S., Danish, R. (2024). A Medical Imaging Approach for Recognising Mitral Regurgitation Through Machine Learning Methods in Cardiac Imaging. Cardiology & Vascular Research. 2 (1). 1-8.

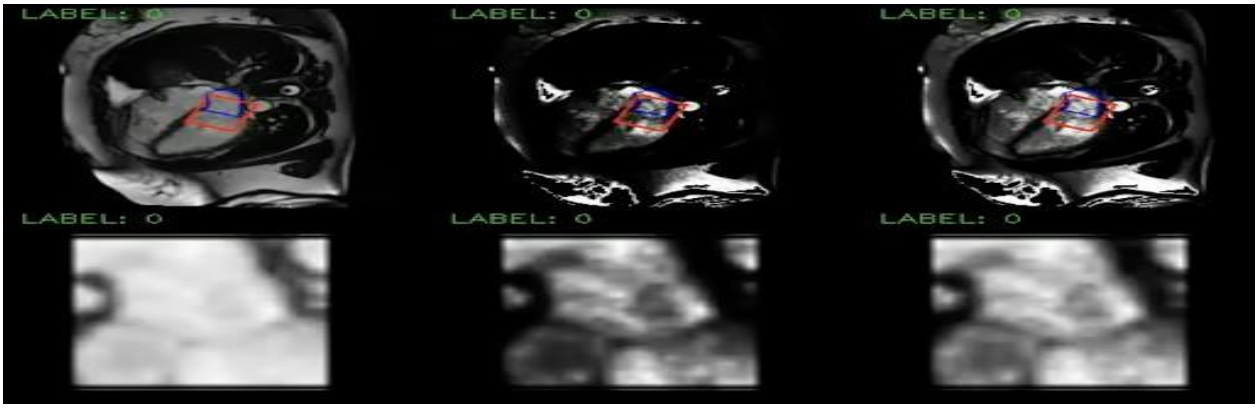

Figure 9: Detailed overview of the pre-processing steps for CUSSP. Top: Example of the 4CH CMR images in the original contrast (left), the left atrium histogram equalized contrast (middle), and the cropped patch histogram equalized contrast (right), with blue contours outline the left atrium, and the red square boxes outline the patch to crop. Bottom: Example of the cropped mitral valve patch as outlined in the red square boxes in the top row.

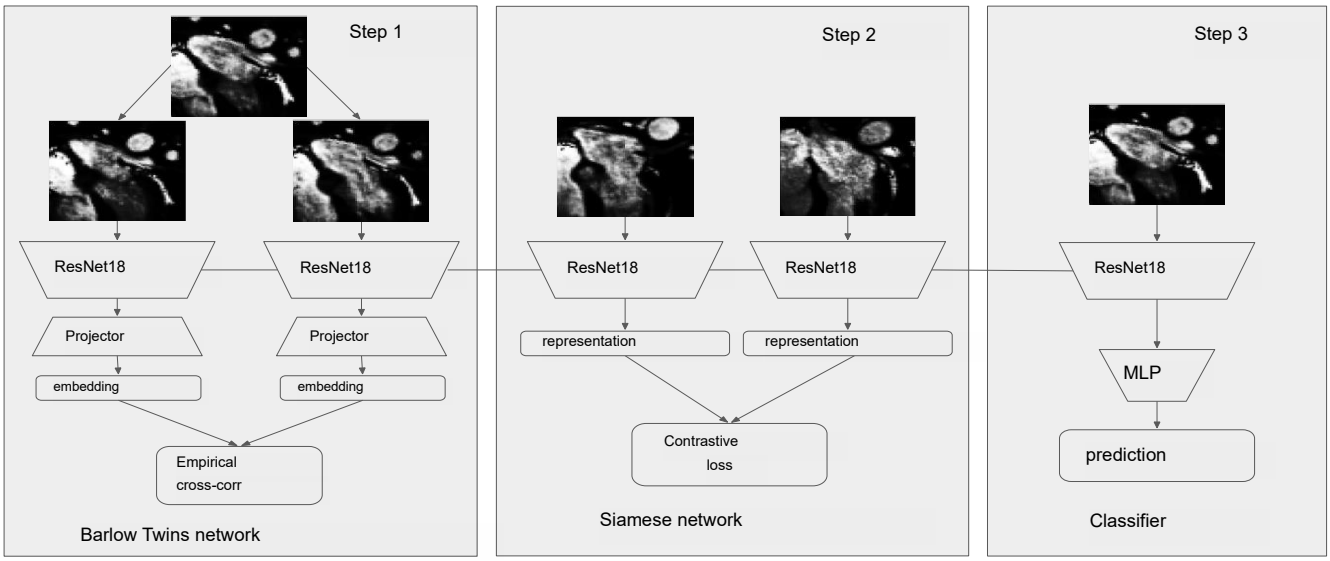

Figure 10: The model training stage of the CUSSP method contains three steps: (i) the feature encoder is trained in the Barlow-Twins network with unlabeled imaging data set, (ii) the feature encoder is fine-tuned in a siamese network with labeled imaging data set, and (iii) the feature encoder is assembled with a MLP, then trained with labeled imaging data set for the classification task of MR.

4. Conclusion

The world's first automated mitral regurgitation (MR) classification system is now available, a significant achievement in medical imaging. The CUSSP model, trained with limited supervision, has demonstrated exceptional performance in classifying 4CH Cardiac Magnetic Resonance (CMR) imaging sequences, paving the way for large-scale screening for MR. MR is a common cardiovascular condition where the mitral valve fails to close properly, allowing blood to flow backward into the heart. Early detection and diagnosis are crucial for effective treatment and improved patient outcomes. The automated CUSSP model has achieved remarkable results, with an F1 score of 0.69 and an ROC AUC of 0.88, indicating a high degree of accuracy in detecting MR, even with limited training data. The system's ability to operate on 4CH CMR imaging sequences allows for a more comprehensive evaluation of the mitral valve and surrounding cardiac structures, enabling the detection of subtle changes and patterns that may indicate MR, leading to more accurate diagnoses. The system's limited supervision training paradigm allows it to learn from limited labeled data, making it an attractive solution for applications where high-quality labeled data is scarce.

Reference

1. Allen, N. E., Sudlow, C., Peakman, T., Collins, R., & Uk biobank. (2014). UK biobank data: come and get it. Science translational medicine, 6(224), 224ed4-224ed4.
2. Bai, W., Sinclair, M., Tarroni, G., Oktay, O., Rajchl, M., et al. (2018). Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. Journal of cardiovascular magnetic resonance, 20(1), 65.
3. Breiman, L. (2001). Random forests. Machine learning, 45, 5-32.
4. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., et al. (2020). Unsupervised learning of visual features

by contrasting cluster assignments. Advances in neural information processing systems, 33, 9912-9924.
5. Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020, November). A simple framework for contrastive learning of visual representations. In International conference on machine learning (pp. 1597-1607). PMLR.
6. Dziadzko, V., Dziadzko, M., Medina-Inojosa, J. R., Benfari, G., Michelena, H. I., et al. (2019). Causes and mechanisms of isolated mitral regurgitation in the community: clinical context and outcome. European heart journal, 40(27), 2194-2202.
7. Enriquez-Sarano, M., Akins, C. W., & Vahanian, A. (2009). Mitral regurgitation. The Lancet, 373(9672), 1382-1394.
8. Enriquez-Sarano, M., Avierinos, J. F., Messika-Zeitoun, D., Detaint, D., Capps, M., et al. (2005). Quantitative determinants of the outcome of asymptomatic mitral regurgitation. New England Journal of Medicine, 352(9), 875-883.
9. Fries, J. A., Varma, P., Chen, V. S., Xiao, K., Tejeda, H., et al. (2019). Weakly supervised classification of aortic valve malformations using unlabeled cardiac MRI sequences. Nature communications, 10(1), 3111.
10. Grill, J. B., Strub, F., Altché, F., Tallec, C., Richemond, P., et al. (2020). Bootstrap your own latent-a new approach to self-supervised learning. Advances in neural information processing systems, 33, 21271-21284.
11. Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 4700-4708).
12. Iglovikov, V., & Shvets, A. (2018). Ternausnet: U-net with vgg11 encoder pre-trained on imagenet for image segmentation. arXiv preprint arXiv:1801.05746.
13. Kervadec, H., Dolz, J., Tang, M., Granger, E., Boykov, Y., et al. (2019). Constrained-CNN losses for weakly supervised segmentation. Medical image analysis, 54, 88-99.
14. Mirabel, M., Iung, B., Baron, G., Messika-Zeitoun, D., Détaint, D., et al. (2007). What are the characteristics of patients with severe, symptomatic, mitral regurgitation who are denied surgery?. European heart journal, 28(11), 1358-1365.
15. Nian, Y., Li, M., Cui, H., Hu, X., Xie, B., et al. (2017, December). Graph-based unsupervised segmentation for lung tumor CT images. In 2017 3rd IEEE International Conference on Computer and Communications (ICCC) (pp. 1884-1888). IEEE.
16. Oktay, O., Ferrante, E., Kamnitsas, K., Heinrich, M., Bai, W., et al. (2017). Anatomically constrained neural networks (ACNNs): application to cardiac image enhancement and segmentation. IEEE transactions on medical imaging, 37(2), 384-395.
17. Parcha, V., Patel, N., Kalra, R., Suri, S. S., Arora, G., et al. (2020, December). Mortality due to mitral regurgitation among adults in the United States: 1999-2018. In Mayo Clinic Proceedings (Vol. 95, No. 12, pp. 2633-2643). Elsevier.
18. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18 (pp. 234-241). Springer International Publishing.
19. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.
20. Uretsky, S., Argulian, E., Narula, J., & Wolff, S. D. (2018). Use of cardiac magnetic resonance imaging in assessing mitral regurgitation: current evidence. Journal of the American College of Cardiology, 71(5), 547-563.
21. Vimalesvaran, K., Uslu, F., Zaman, S., Galazis, C., Howard, J., et al. (2022, September). Detecting aortic valve pathology from the 3-chamber cine cardiac mri view. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 571-580). Cham: Springer Nature Switzerland.
22. Watanabe, N. (2019). Acute mitral regurgitation. Heart, 105(9), 671-677.
23. Wu, X., Bi, L., Fulham, M., Feng, D. D., Zhou, L., et al. (2021). Unsupervised brain tumor segmentation using a symmetric-driven adversarial network. Neurocomputing, 455, 242-254.
24. Wu, X., Bi, L., Fulham, M., & Kim, J. (2020, December). Unsupervised positron emission tomography tumor segmentation via GAN based adversarial auto-encoder. In 2020 16th International Conference on Control, Automation, Robotics and Vision (ICARCV) (pp. 448-453). IEEE.
25. Xing, Z. J., Yin, F., Wu, Y. C., & Liu, C. L. (2018, April). Offline signature verification using convolution siamese network. In Ninth International Conference on Graphic and Image Processing (ICGIP 2017) (Vol. 10615, pp. 415-423). SPIE.
26. Zbontar, J., Jing, L., Misra, I., LeCun, Y., & Deny, S. (2021, July). Barlow twins: Self-supervised learning via redundancy reduction. In International conference on machine learning (pp. 12310-12320). PMLR.
27. Zhao, A., Balakrishnan, G., Durand, F., Guttag, J. V., & Dalca, A. V. (2019). Data augmentation using learned transformations for one-shot medical image segmentation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 8543-8553).