

Machine Learning for Automated Mitral Regurgitation Detection from Cardiac Imaging

Ke Xiao¹[0000-0001-6453-0420], Erik Learned-Miller¹[0000-0002-3778-9135],
Evangelos Kalogerakis¹[0000-0002-5867-5735], James Priest²[0000-0002-8349-4784],
and Madalina Fiterau¹[0000-0003-4179-7274]

¹ University of Massachusetts Amherst, Amherst MA 01003, USA

{kexiao, elm, kalo, mfiterau}@cs.umass.edu

² Stanford University, Stanford CA 94305, USA

jpriest@stanford.edu

Abstract. Mitral regurgitation (MR) is a heart valve disease with potentially fatal consequences that can only be forestalled through timely diagnosis and treatment. Traditional diagnosis methods are expensive, labor-intensive and require clinical expertise, posing a barrier to screening for MR. To overcome this impediment, we propose a new semi-supervised model for MR classification called CUSSP. CUSSP operates on cardiac magnetic resonance (CMR) imaging slices of the 4-chamber view of the heart. It uses standard computer vision techniques and contrastive models to learn from large amounts of unlabeled data, in conjunction with specialized classifiers to establish *the first ever automated MR classification system using CMR imaging sequences*. Evaluated on a test set of 179 labeled – 154 non-MR and 25 MR – sequences, CUSSP attains an F1 score of 0.69 and a ROC-AUC score of 0.88, setting the first benchmark result for detecting MR from CMR imaging sequences.

1 Introduction

Mitral regurgitation. Mitral regurgitation (MR) [7] is a valvular heart disease in which the mitral valve does not close completely during systole when the left ventricle contracts, causing regurgitation – leaking of blood backwards – from the left ventricle (LV), through the mitral valve, into the left atrium (LA) – Figure 1. MR can be caused by either organic or functional mechanisms [6], with organic MR leading to atrial and annular enlargement and functional MR increasing atrial pressure. As MR progresses, it may cause arrhythmia, shortness of breath, heart palpitations and pulmonary hypertension [14]. Left undiagnosed and untreated, MR may cause significant hemodynamic instability and congestive heart failure which can lead to death [17], while acute MR usually necessitates immediate medical intervention [22]. Thus, *early detection and assessment of MR are crucial for optimal treatment outcomes*, with the best short-term and long-term results obtained in asymptomatic patients operated on in advanced repair centers with low operative mortality ($< 1\%$) and high repair rates ($\geq 80 - 90\%$) [7].




Fig. 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.

MR diagnosis. MR is often only detected following symptom onset. Among patients with asymptomatic MR, quantitative grading of mitral regurgitation is a powerful indicator for clinical treatment such as immediate cardiac surgery [8]. Clinically, MR is usually diagnosed with doppler echocardiography, with *cardiovascular magnetic resonance* (CMR) subsequently used to assess the MR severity and to accurately quantify the regurgitant volume, one of the indicators of severity [20]. Most studies that have evaluated CMR for assessing the mitral regurgitant volume use the difference between left ventricular stroke volume (LVSV) and forward stroke volume (FSV). LVSV is usually estimated with the short-axis (SA) view CMR – a 4-D tensor – while FSV is most commonly determined by aortic phase-contrast velocity-encoding images [20]. This diagnosis and assessment process requires continuous involvement from expert clinicians along with specific order and post-processing for the phase-contrast images of the proximal aorta or main pulmonary artery during the acquisition of the CMR data. The associated expense with this standard diagnostic procedure thus poses an obstacle to the large-scale screening for MR in the general population.

Towards machine learning for MR diagnosis. Although quantitatively assessing mitral regurgitant volume requires specific CMR imaging sequences and expert analysis, four-chamber (4CH) CMR images provide a comprehensive view of all four heart chambers, including the mitral valve as it opens and closes, as shown in Figure 1. Thus, we propose to train a model that uses 4CH CMR to automatically diagnose MR, making wide screening possible. As training data, we use the long axis 4CH CMR imaging data from the UK Biobank [1], from over 30,000 subjects, out of which $N=704$ were labeled by an expert cardiologist. While the 4CH view has the potential to identify MR when the regurgitant jet is visible, the imaging is not accompanied by comprehensive annotations or diagnoses of diseases/conditions for individual patients. This is in contrast to Zhang et al. [27], where tens of thousands of annotated color doppler echocardiography images are available for MR assessments. To overcome this difficulty, we rely on weakly supervised and unsupervised methods. Weakly supervised deep learning




Fig. 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.

has proved successful in detecting other heart pathologies. Specifically, Fries et al. [9] proposed a weakly supervised deep learning method (CNN-LSTM) to classify aortic valve malformation from the aortic valve cross section CMR present in the UK Biobank, wherein the critical feature of the aortic valve opening shape was easily extracted from the aortic valve cross section CMR imaging data. Meanwhile, Vimala et al. [21] proposed a deep learning based pipeline for detecting aortic valve pathology using 3CH CMR imaging from three hospitals. The data set was fully annotated with landmarks, stenotic jets and regurgitant jets. Unlike these prior two 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 the best of our knowledge, this is *the first study on identifying MR using the 4CH CMR imaging data in an automated pipeline*.

Our approach. We propose an automated five stage pipeline named Cardiovascular magnetic resonance **U**-Net localized **S**elf-**S**upervised **P**redictor (CUSSP). Our approach incorporates several different preexisting neural network architectures in the pipeline, discussed in Section 2.3, to address the challenges inherent to the MR classification task. Specifically, we use a U-Net [18] to perform segmentation of the heart chambers, which we then use to localize the area around the mitral valve. We apply histogram equalization to enhance the appearance of the valve. We then use a Barlow Twins [26] network to learn, without supervision, representations of the blood flow around the valve, and a Siamese network [25] to learn differences between instances of MR and non-MR. During training, CUSSP leverages a large amount of unlabeled CMR images, and *minimal supervision*, in the form of a comparatively small set of MR labels manually annotated by cardiologist. However, *at test time CUSSP is fully automated*.

Contribution. Our work is the *first study on automated detection of mitral regurgitation (MR) from CMR*, providing a benchmark for the classification of MR in an automated pipeline from long axis 4CH CMR images. As a screening tool, it has the potential to support hospital diagnostics and improve patient care.

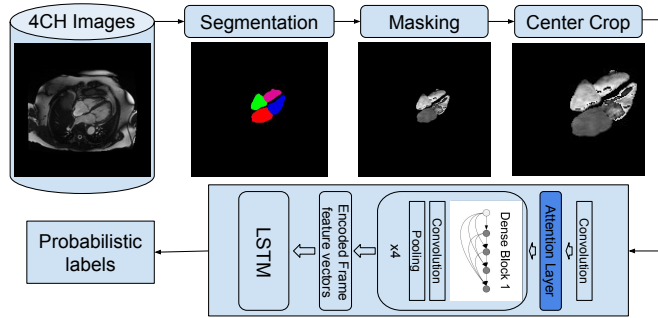


Fig. 3. Overview of the CNN-LSTM method pipeline for MR classification.

2 Methods

2.1 Segmentation of the cardiac magnetic resonance images

The CMR imaging data from the UK Biobank that is relevant to MR detection includes long-axis 2-chamber (2CH) view and long-axis 4-chamber (4CH) view, which are all shown in Figure 2. In addition, the short-axis view CMR provides accurate description of the left ventricle. Both long-axis views and short-axis view are used to estimate heart measurements relevant to the MR detection task, while only the long-axis 4CH view is used for the deep learning models.

As a pre-processing step, we performed semantic segmentation on the CMR imaging data, using masks (Figure 2) generated by a U-Net [18] segmentation model to highlight regions of interest to MR classification. U-Net is currently the leading model architecture for medical imaging segmentation, with various U-Net variants developed for different applications. TerausNet [12] is a U-Net variant that reshapes the U-Net encoder to match the VGG11 architecture, allowing it to use pre-trained VGG11 [19] model weights for faster convergence and improved segmentation results. While most medical imaging segmentation models are trained using supervised learning, weakly supervised segmentation methods such as VoxelMorph augmented segmentation [28], ACNN [16], CCNN [13], graph-based unsupervised segmentation [15], and GAN-based unsupervised segmentation [23,24] also produce comparable segmentation results. For the segmentation of the 4CH, 2CH, SA, and aorta view CMR imaging dataset from the UK Biobank, Bai et al. [2] offer a supervised segmentation model.

We manually labeled 100 CMR images for each view and trained a supervised segmentation model with the TerausNet [12] architecture. Then, segmentation outputs, shown in Figure 2, are used to compute measurements of cardiac structure and function for the four chambers of the heart, as summarized in Table 1. The short-axis view CMR segmentation output is used to estimate the left ventricle and right ventricle measurements, while the long-axis 4CH view and 2CH view outputs are used to estimate the left atrium and right atrium measurements. Specifically, the left atrial volume is estimated using the biplane method

with segmentation of both the 2CH and 4CH view, while the right atrial volume is estimated using single plane method with segmentation of the 4CH view.

2.2 Baseline models

We first considered a random forest (RF) classifier [3] trained for MR classification on the tabular heart measurements in Table 1. We divided the 18 features by body surface area (BSA) prior to training the RF.

Next, we developed a deep learning baseline model for MR classification, a weakly supervised CNN-LSTM, following the principles in Fries et al. [9] and operating on the 4CH CMR imaging data. Fries et al. [9] used CMR imaging sequences from the UK Biobank, however, the objective of their work was the identification of aortic valve malformations. Their proposed deep learning architecture – CNN-LSTM – used DenseNet [11] as the CNN of choice to encode CMR imaging frames and the LSTM to encode embeddings of all frames within each sequence for a final classification of aortic valves into tricuspid (normal) and bicuspid (pathological). We point out that *our MR classification problem is considerably more challenging* due to the lack of direct view of the mitral valve in the CMR imaging data. Moreover, the flow information provided from the 4CH view CMR imaging data is difficult to learn and encode in the model, an issue which we alleviated in the CUSSP framework.

The CNN-LSTM pipeline, shown in Figure 3, includes an image segmentation model, and an image classification model. It uses the 4CH CMR from the UK Biobank. The CMR data is center-cropped using the center of mass of the CMR imaging frames. The resulting sequence provided to the CNN-LSTM, which generates probabilistic labels of MR for the sample.

In the CNN-LSTM model architecture, the CNN serves as the frame encoder, which encodes each frame of each sequence into a representation vector. The model uses DenseNet-121 pre-trained on ImageNet as the CNN. To better learn the attention span of the frame encoder, we added an attention layer to the DenseNet-121 after the first convolutional layer. After the bi-directional LSTM, a multi-layer perceptron (MLP) performs the final classification.

Left Atrium	Right Atrium	Left Ventricle	Right Ventricle
Vol Max (mL)	Vol Max (mL)	End Systolic Vol (mL)	End Systolic Vol (mL)
Vol Min (mL)	Vol Min (mL)	End Diastolic Vol (mL)	End Diastolic Vol (mL)
Stroke Vol (mL)	Stroke Vol (mL)	Stroke Vol (mL)	Stroke Vol (mL)
Eject. fraction (%)	Eject. fraction (%)	Eject. fraction (%)	Eject. fraction (%)
		Cardiac Output (L/min)	
		Mass (g)	

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

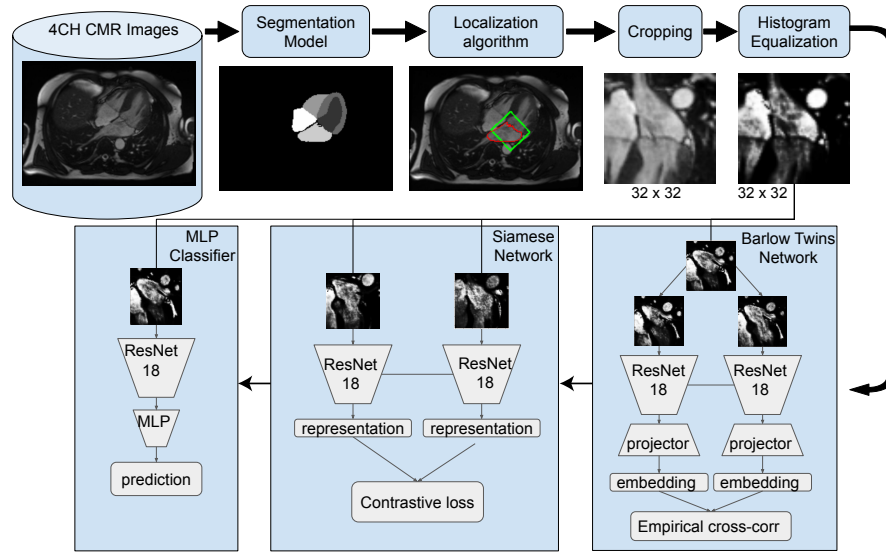


Fig. 4. Overview of the CUSSP pipeline for MR classification, with its 5 steps: (1) segmentation, (2) localization, (3) cropping, (4) equalization, and (5) prediction. In particular, the prediction stage of the CUSSP method contains three stages: (i) the feature encoder is trained in the Barlow-Twins 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.

2.3 The CUSSP framework

Conceptualization. To better encode the blood flow information relevant to MR classification from the 4CH CMR view, we investigated self-supervised representation learning methods which can leverage all the unlabeled CMR sequences present in the UK Biobank. Typically, self-supervised representation learning for visual data involves maximizing the similarity between representations of various distorted versions of a sample. Among the many self-supervised architectures, SimCLR [5], SwAV [4], and BYOL [10], we chose Barlow Twins [26], since it does not require large batches. With the labeled data, our siamese network compares the representation differences between classes by sampling two inputs from different classes as performed in [25]. Thus, our CUSSP MR classification pipeline takes advantage of both self-supervised and supervised representation learning.

Test-time pipeline. Our CUSSP method consists of five main steps, shown in Figure 4, the first four for pre-processing, and the last one for prediction, with three stages using network components trained for MR classification. The pre-processing of the CMR imaging sequence is shown in Figure 8 in the Appendix. We used the segmentation model in 2.1 to locate the mitral valve and determine the orientation of the left ventricle using the contours and centers of the heart

chambers derived from the segmentation output. We then cropped 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 used by the downstream networks.

Training process. The first step involves training a representation encoder in a Barlow Twins network using over 30,000 unlabeled pre-processed sequences. We chose Barlow Twins for its versatility and robustness to distortions such as blurring, different sizes of the relevant areas and intensity variations, which are common in CMR images. We used a ResNet-18 model pretrained on ImageNet as an encoder. Its output vector is 512-dimensional. We selected ResNet-18 because we found that more complex encoders would easily memorize the limited labeled dataset, reducing the effectiveness of the embeddings. The projector network has three fully connected layers, all with 2048 output units. After training the encoder with the unlabeled dataset, it is fine-tuned in a siamese network using a comparatively smaller labeled set, as indicated in Section 3. 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. Contrastive learning helps because it uses the very limited labels in our possession to obtain representations focused on encoding differences between classes. 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 improve computation efficiency and training accuracy, we also tested the framework using a smaller window of 25 frames, since MR occurs between diastole and systole. The code is available at https://github.com/Information-Fusion-Lab-Umass/CUSSP_UKB_MR.

3 Experiments and Discussion

Experimental setup. 4CH CMR images were used to conduct experiments with both the CNN-LSTM method and the CUSSP method. We used a total of 704 labeled sequences, with 525 sequences selected for the training set, including 452 labeled as non-MR and 73 labeled as MR. The remaining 179 sequences were used for testing, with 154 labeled as non-MR and 25 labeled as MR. Considering the substantial class imbalance, we opted to use F1 score as our primary evaluation metric, along with precision and recall.

Random Forest Classification Results. The random forest model is trained with 10-fold cross validation, with a random search over a parameter grid of 10-100 estimators, 2-16 depth, 2-8 min samples. The optimal hyper-parameter setting found is: 20 estimators, log2 max features, max depth of 9 and a minimum of 2 samples per leaf node. The best results obtained are presented in Table 2.

CNN-LSTM Classification Results. We conducted experiments on the DenseNet-LSTM classification model using various input image sizes, attention layer

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

Table 2. Experimental results for Random Forest (RF) baseline, CNN-LSTM and CUSSP. CUSSP-1, CUSSP-2 and CUSSP-3 are trained with the BarlowTwins-MLP model without fine-tuning with the Siamese network. CUSSP-SIAM and CUSSP-SIAM-25 are trained with the BarlowTwins-Siamese-MLP model.

configurations, and masks. The best CNN-LSTM model attains a F1-score of 0.44, shown in Table 2, with further information on the performance under other settings summarized in the Appendix.

CUSSP Classification Results




Fig. 5. The ROC AUC curve and the precision-recall curve of CUSSP-SIAM-25 from Table 2. The annotated coordinates on the precision-recall curve plot are (recall, precision, F1-score, threshold).

We evaluated various configurations of the CUSSP model, to determine the relative benefits of different components. In the first configuration, the ResNet18 model was combined with a 3-layer MLP to train a classifier using the labeled training set after being trained in the Barlow-Twins network with the unlabeled dataset. During the classifier training, the cross-correlation loss from the Barlow-Twins network and the cross-entropy loss from the binary classification were weighted using three different

configurations. For CUSSP-1 the cross-correlation loss has a weight of 0.9, while the cross-entropy loss has a weight of 0.1. For CUSSP-2, the weights are 0.5 and 0.5, while for CUSSP-3 they are 0.1 and 0.9, respectively. Both CUSSP-1 and CUSSP-3 outperform CUSSP-2, though the performance is low, indicating the importance of fine-tuning, described below.

In the second scenario, we fine-tuned the encoder with a siamese network to enhance the quality of the encoded representations after training the Barlow Twins network. To prevent overfitting of the model and to limit its capacity,

we froze the parameters of all layers except the last block of the ResNet18 encoder when training the siamese network and the classifier. The resulting model, CUSSP-SIAM, showed a significant improvement in performance. In the final configuration CUSSP-SIAM-25, the number of frames in the training sequences was truncated from 50 frames to the 25 frames that correspond to the interval when mitral regurgitation occurs. The results are summarized in Table 2, while the ROC-AUC curve for CUSSP-SIAM-25 are shown in Figure 5.

CUSSP attains an F1 score of 0.69, and an ROC AUC of 0.88, outperforming the CNN-LSTM approach. This is due to CUSSP’s focus on the area around the valve to capture the blood flow through the valve, combining the advantages of Barlow Twins and contrastive learning. Meanwhile, the CNN-LSTM relies on attention, which does not seem to work as well. Additionally, using only the frames relevant to the task reduces the number of parameters and makes the model sample-efficient. This is essential for attaining high performance in the low label setting. In the future, we aim to use the pipeline on the large unlabeled dataset to scan for and adjudicate more MR cases.

In conclusion, we present the first automated mitral regurgitation classification system using CMR imaging sequences. The CUSSP model we developed, trained with limited supervision, operates on 4CH CMR imaging sequences and attains an F1 score of 0.69 and an ROC AUC of 0.88, opening up the opportunity for large-scale screening for MR.

References

1. Allen, N.E., Sudlow, C., Peakman, T., Collins, R., UK Biobank: UK Biobank data: Come and get it. *Science Translational Medicine* **6**(224) (2014)
2. Bai, W., Sinclair, M., Tarroni, G., Oktay, O., Rajchl, M., Vaillant, G., Lee, A.M., Aung, N., Lukaschuk, E., Sanghvi, M.M., et al.: Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. *Journal of Cardiovascular Magnetic Resonance* **20**(1) (2018)
3. Breiman, L.: Random forests. *Machine Learning* **45** (2001)
4. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. In: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020* (2020)
5. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. In: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event. Proceedings of Machine Learning Research*, vol. 119 (2020)
6. Dziadzko, V., Dziadzko, M., Medina-Inojosa, J.R., Benfari, G., Michelena, H.I., Crestanello, J.A., Maalouf, J., Thapa, P., Enriquez-Sarano, M.: Causes and mechanisms of isolated mitral regurgitation in the community: clinical context and outcome. *European Heart Journal* **40**(27) (2019)
7. Enriquez-Sarano, M., Akins, C.W., Vahanian, A.: Mitral regurgitation. *The Lancet* **373**(9672) (2009)
8. Enriquez-Sarano, M., Avierinos, J.F., Messika-Zeitoun, D., Detaint, D., Capps, M., Nkomo, V., Scott, C., Schaff, H.V., Tajik, A.J.: Quantitative determinants of the

- outcome of asymptomatic mitral regurgitation. *New England Journal of Medicine* **352**(9) (2005)
9. Fries, J.A., Varma, P., Chen, V.S., Xiao, K., Tejada, H., Saha, P., Dunmon, J., Chubb, H., Maskatia, S., Fiterau, M., et al.: Weakly supervised classification of aortic valve malformations using unlabeled cardiac MRI sequences. *Nature Communications* **10**(1) (2019)
 10. Grill, J., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Dörsch, C., Pires, B.Á., Guo, Z., Azar, M.G., Piot, B., Kavukcuoglu, K., Munos, R., Valko, M.: Bootstrap your own latent - A new approach to self-supervised learning. In: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020* (2020)
 11. Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017* (2017)
 12. Iglovikov, V., Shvets, A.: TerausNet: U-Net with VGG11 encoder pre-trained on ImageNet for image segmentation. *CoRR* **abs/1801.05746** (2018)
 13. Kervadec, H., Dolz, J., Tang, M., Granger, E., Boykov, Y., Ayed, I.B.: Constrained-CNN losses for weakly supervised segmentation. *Medical Image Anal.* **54** (2019)
 14. Mirabel, M., Jung, B., Baron, G., Messika-Zeitoun, D., Détaint, D., Vanoverschelde, J.L., Butchart, E.G., Ravaut, P., Vahanian, A.: What are the characteristics of patients with severe, symptomatic, mitral regurgitation who are denied surgery? *European Heart Journal* **28**(11) (2007)
 15. Nian, Y., Li, M., Cui, H., Hu, X., Xie, B., Li, K., Xiong, X., Xiao, J., Chen, W.: Graph-based unsupervised segmentation for lung tumor CT images. In: *2017 3rd IEEE International Conference on Computer and Communications (ICCC)*. IEEE (2017)
 16. Oktay, O., Ferrante, E., Kamnitsas, K., Heinrich, M.P., Bai, W., Caballero, J., Cook, S.A., de Marvao, A., Dawes, T., O'Regan, D.P., Kainz, B., Glocker, B., Rueckert, D.: Anatomically constrained neural networks (ACNNs): Application to cardiac image enhancement and segmentation. *IEEE Trans. Medical Imaging* **37**(2) (2018)
 17. Parcha, V., Patel, N., Kalra, R., Suri, S.S., Arora, G., Arora, P.: Mortality due to mitral regurgitation among adults in the United States: 1999-2018. In: *Mayo Clinic Proceedings*. vol. 95. Elsevier (2020)
 18. Ronneberger, O., Fischer, P., Brox, T.: 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*. Lecture Notes in Computer Science, vol. 9351 (2015)
 19. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: *3rd International Conference on Learning Representations, ICLR 2015, Conference Track Proceedings* (2015)
 20. Uretsky, S., Argulian, E., Narula, J., Wolff, S.D.: Use of cardiac magnetic resonance imaging in assessing mitral regurgitation: Current evidence. *Journal of the American College of Cardiology* **71**(5) (2018)
 21. Vimala, K., Uslu, F., Zaman, S., Galazis, C., Howard, J., Cole, G., Bharath, A.A.: Detecting aortic valve pathology from the 3-chamber cine cardiac MRI view. In: *Medical Image Computing and Computer Assisted Intervention - MICCAI 2022 - 25th International Conference, Singapore, September 18-22, 2022, Proceedings, Part I*. Lecture Notes in Computer Science, vol. 13431 (2022)

22. Watanabe, N.: Acute mitral regurgitation. *Heart* **105**(9) (2019)
23. Wu, X., Bi, L., Fulham, M.J., Feng, D.D., Zhou, L., Kim, J.: Unsupervised brain tumor segmentation using a symmetric-driven adversarial network. *Neurocomputing* **455** (2021)
24. Wu, X., Bi, L., Fulham, M.J., Kim, J.: Unsupervised positron emission tomography tumor segmentation via GAN based adversarial auto-encoder. In: 16th International Conference on Control, Automation, Robotics and Vision, ICARCV 2020, Shenzhen, China, December 13-15, 2020 (2020)
25. Xing, Z.J., Yin, F., Wu, Y.C., Liu, C.L.: Offline signature verification using convolution siamese network. In: Ninth International Conference on Graphic and Image Processing (ICGIP 2017). vol. 10615. SPIE (2018)
26. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. In: Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event. Proceedings of Machine Learning Research, vol. 139 (2021)
27. Zhang, Q., Liu, Y., Mi, J., Wang, X., Liu, X., Zhao, F., Xie, C., Cui, P., Zhang, Q., Zhu, X.: Automatic assessment of mitral regurgitation severity using the mask R-CNN algorithm with color Doppler echocardiography images. *Computational and Mathematical Methods in Medicine* **2021** (2021)
28. Zhao, A., Balakrishnan, G., Durand, F., Gutttag, J.V., Dalca, A.V.: Data augmentation using learned transformations for one-shot medical image segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019 (2019)

Appendix




Fig. 6. 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.

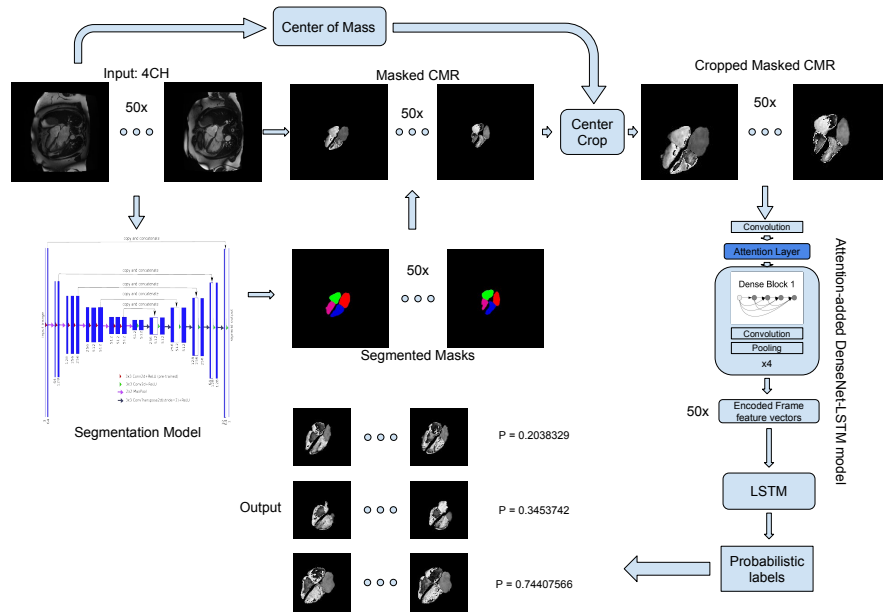


Fig. 7. Detailed overview of the CNN-LSTM method pipeline for MR classification.




Fig. 8. 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.