Improved Myocardial Mapping by Cardiac MRI in Transthyretin Amyloidosis Cardiomyopathy Using Consistency-Parametrized AI Segmentation Model

Mateusz C. Florkow MTFW@NOVONORDISK.COM
Andreas Heindl EHIN@NOVONORDISK.COM
Novo Nordisk R&D Digital Hub, London, United Kingdom
Majd Zreik MXJZ@NOVONORDISK.COM
Novo Nordisk A/S, The Netherlands
Ahmed Emam AHWQ@NOVONORDISK.COM
Novo Nordisk A/S, United States of America

Editors: Under Review for MIDL 2026

Abstract

Disease monitoring relies on accurate, standardised and consistent assessment of longitudinal data acquired over different time points. However, most current image segmentation approaches operate on individual images in isolation and do not leverage available information from other time points. In this work, we introduce a simple, parameterized loss function that incorporates a guiding mask to constrain image segmentation and promote consistency across repeated acquisitions in the same patient. The constraint can be adjusted according to image quality characteristics (e.g., signal-to-noise ratio, contrast). We trained a U-Net–based network using this loss and compared its performance against a baseline U-Net and SAM2. The approach was applied on a clinical dataset of transthyretin amyloid cardiomyopathy patients. Results indicate that the proposed model improved T1 mapping and extracellular volume estimation—a key marker for transthyretin amyloid cardiomyopathy.

Keywords: AI, segmentation, T1 mapping, ECV mapping, deep learning, guiding mask, cardiac MRI, MOLLI, medical image analysis.

1. Introduction

Tissue mapping in ATTR Transthyretin amyloidosis (ATTR) cardiomyopathy is a cardiac disease in which misfolded transthyretin protein aggregates as amyloid fibrils that deposit in the myocardium. The deposited amyloid fibrils cause the myocardium to stiffen and lead to expansion of the extracellular matrix. Current efforts to develop a drug that can deplete the amyloid fibrils from the heart can be a potential new treatment option for ATTR patients. However, monitoring of myocardial amyloid burden and its response to treatment is technically challenging due to imaging and analysis limitations. Amyloid burden can be assessed by measuring the myocardial extracellular volume (ECV) using cardiac magnetic resonance imaging (MRI). ECV is derived from pre- and post-contrast T1 maps on MRI (Figure 1) which are typically reconstructed from 8–12 T1-weighted (T1w) images (also known as echoes) acquired at different time points (echo times) of the same anatomical cross-section. To reduce motion artifacts between the multiple T1w images and improve mapping quality, accurate alignment of the T1w images is a prerequisite for T1 and ECV mapping (Messroghli et al., 2016).

Segmentation and Alignment of T1w images Scaling-up ECV analysis in large clinical trials requires reliable and consistent analysis of myocardial T1 maps. Additionally, several T1w alignment methods depend on accurate and consistent segmentation of the myocardium in the T1w images (El-Rewaidy et al., 2018; Fahmy et al., 2019; Zhu et al., 2020). However, current automatic segmentation algorithms typically process each T1w image independently and therefore cannot leverage concomitant information from T1w images obtained during the same acquisition. This leads to two key limitations. First, the segmented myocardial borders can vary substantially in shape and size between images, even though they represent the same anatomical cross-section. Second, the myocardial delineation in echoes with poor image quality or contrast might fail when it could benefit from higher-quality echoes. In this study, we present a novel AI-based segmentation approach that exploits information from high-contrast T1w images to improve both the consistency of the myocardial shape and the segmentation accuracy in lower qualities echoes across the entire series. We also report its application for generating accurate T1 maps and ECV maps in a cohort of ATTR patients.

2. Methods

Model Development The proposed model was derived from a UNet (Ronneberger et al., 2015; Isensee et al., 2020) and was optimized to segment the myocardium from cardiac MR images. It was trained and validated using a publicly available native T1 mapping dataset (Fahmy et al., 2019), utilizing data from a modified MOLLI sequence (210 patients; 5 slices; 11 echo-times). The dataset was split patient-wise into training (n=142), validation (n=26), and testing (n=42). The model had two 2D slices as inputs: (1) a 200x200 T1w grayscale image to be segmented; and (2) a binary mask that is used to guide the segmentation. This guiding mask is meant to be derived from another echo, from the same acquisition or patient, and should be accurate as it will influence the segmentation of the remaining echoes. In this work, we used the mask of the first T1w image in the MOLLI sequence, where the myocardium-to-blood image contrast is typically at its maximum. Manual quality check of the guiding mask was used to ensure accurate delineation of the myocardial borders. To integrate the information of the guiding mask into the segmentation, we formulated a customized loss function that includes a parameterized term $\alpha \in [0, 1]$, controlling how the guiding mask influences the segmentation output. This parametrization formulation facilitates controlling the trade-off between model fit with underlying image (low α) versus consistency with the guiding mask (high α). The guiding parameter, α, was applied to both the input guiding mask and the loss function as follows:

$$L = (1 - \alpha) * f(pred, y) + \alpha * f(pred, T_{1-\alpha}(\hat{y} * \alpha)) \quad (1)$$

Where $f(.)$ is a standard objective function, y is the ground truth target, $\hat{y}$ is the guidance mask and $T_{(1-\alpha)}$ is a partial translation operation of the guiding mask towards the ground truth mask (full translation if $\alpha = 0$ and no translation if $\alpha = 1$). In this work, f was defined as the average of binary cross-entropy loss and dice loss. A UNet model with the same architecture but with a standard loss and applied independently on each echo was used as a reference base-model for evaluating the performance of the proposed model. A SAM2 model (Ravi et al., 2024) fined-tuned on the open dataset using low-rank adaptation

(Hu et al., 2021) was used as a state-of-the-art for comparisons. The agreement between the manual and automatic segmentations was measured using subjective qualitative scoring and quantitative Dice score (DSC).

Application to ATTR cardiomyopathy The proposed model was applied to automatically delineate the myocardium and estimate the pre- and post-contrast T1 mapping for 272 unseen scans from 83 ATTR patients. The dataset was collected as a part of a multi-center clinical study and included pre- and post-contrast T1 mapping MOLLI sequence, and the collection of a blood sample. A previously validated segmentation-based T1w image alignment and mapping method (Fahmy et al., 2019; Zhu et al., 2020) was used to generate the pre- and post-contrast T1 maps of the myocardium as explained in Figure 1. The T1 value of the blood was also computed by fitting an exponential to the average intensity of the blood pool in each echo, in the pre- and post-contrast images. An expert image analyst blinded to the underlying automatic T1 maps reconstruction pipeline performed a quality check of all segmented images and confirmed the accuracy of the T1 measurements. The agreement between the manual and automatic segmentations was evaluated using the DSC measurement. Additionally, pixel-wise coefficient of determination for T1 estimation was used to assess the mapping accuracy. From the pre- and post-contrast T1 maps, ECV maps could be computed by applying the following formula pixel-wise:

$$ECV = (1 - Hct) * \frac{\Delta R1_{myocardium}}{\Delta R1_{blood}} \tag{2}$$

Where Hct is the hematocrit value, i.e. the proportion of red blood cells in the blood, $R1 = \frac{1}{T1}$ and Δ is the difference operator (post - pre). To estimate the clinical impact of the proposed method, the difference in average ECV was computed between the two methods.

3. Results

The developed model effectively improved segmentation consistency without compromising individual image accuracy. On the public dataset, an average $\pm$ standard deviation DSC of 0.87 ± 0.11 have been obtained for the base-model and of 0.86 ± 0.09 for the consistency-regularized model ($\alpha = 0$), in line with a previous study which reported a DSC of 0.86 [3]. As the consistency-weighting, α, was increased to 0.25 and 1, the DSC decreased to 0.81 ± 0.12 and 0.75 ± 0.10, respectively. This decrease in DSC is expected because the public dataset was segmented without having cross-image consistency as an objective. Nevertheless, qualitative scoring indicated that the consistency-parametrized model was deemed better in 64.3% of cases when trading off between accuracy and consistency.

In the ATTR dataset, average DSC were consistently improved reaching an average improvement of 0.34 for $\alpha = 0.25$ and 0.39 for $\alpha = 1$ as shown in Figure 2-A, B, C. Improvements were higher for the post-contrast T1w images, not seen during training (increase in DSC by 0.37 for post vs 0.30 for pre at $\alpha = 0.25$), indicating a better generalizability when a mask is provided. The consistency-weighting parameter, α, successfully modulated the model's response, showing higher consistency towards the guiding mask with increasing α.

Although the relationship between consistency and α was not linear (Figure 2-D), an $\alpha = 0.25$ was found to be a good trade-off between concordance with the guiding mask

and adaptation to image intensities and was used in the remaining of this study. Figure 3 provides a visualization of the segmentations. Improvements can be seen for echoes with low contrast between the myocardium and blood pool (Figure 3 – echoes 6) or when high contrast structures are visible in the myocardium (Figure 3 – echoes 3).

T1 maps generated by the consistency-parametrized model were fitted with an average coefficient of determination above 0.9 in 98.7% of the cases (vs 77.8% with the UNet base-model) and 26.4% (vs 98.9%) of cases had more than 10% of pixels with a coefficient of determination under 0.95. That is, the developed model effectively reduced the number of cases with inconsistent segmentations from 99% to only 26% of all cases. This significantly improved the reproducibility and consistency of T1, and hence ECV, mapping analysis due to the reduced need for subjective manual refinement of the segmentation. Example T1 maps are given in Figure 4. We further computed the average ECV from the T1 maps derived from the masks obtained by the base and proposed models. The average percentage difference between the two methods was -2.8% $\pm$ 6.7%. Since ECV is a direct function of the T1 maps, improved T1 mapping is expected to directly improve the ECV mapping. Defining a change in ECV by 10% as clinically significant, 8.3% of scans processed by the consistency model have clinically significant improvement in ECV measurements.

4. Conclusion

Our results demonstrate that incorporating a guiding mask into a parameterized deep-learning framework produces robust, consistent segmentations and enables fully automated T1 mapping with accurate fits—and therefore reliable ECV mapping. By standardizing the processing pipeline, this AI-enabled approach improves reproducibility and the quality of anatomical analysis, allowing for more detailed assessment of disease progression in ATTR cardiomyopathy patients. Fewer required manual corrections make ECV studies more scalable and cost-effective, while more accurate ECV measurements can accelerate clinical research by better characterizing disease and exploring links between ECV, biomarkers, and clinical outcomes. Although we showed application of the proposed framework on images collected at different time points of the same MRI scan, the framework can be extended to images collected during different acquisitions to improve the consistency when monitoring the disease changes. Future work will validate this method across a wider range of anatomical structures and imaging modalities.

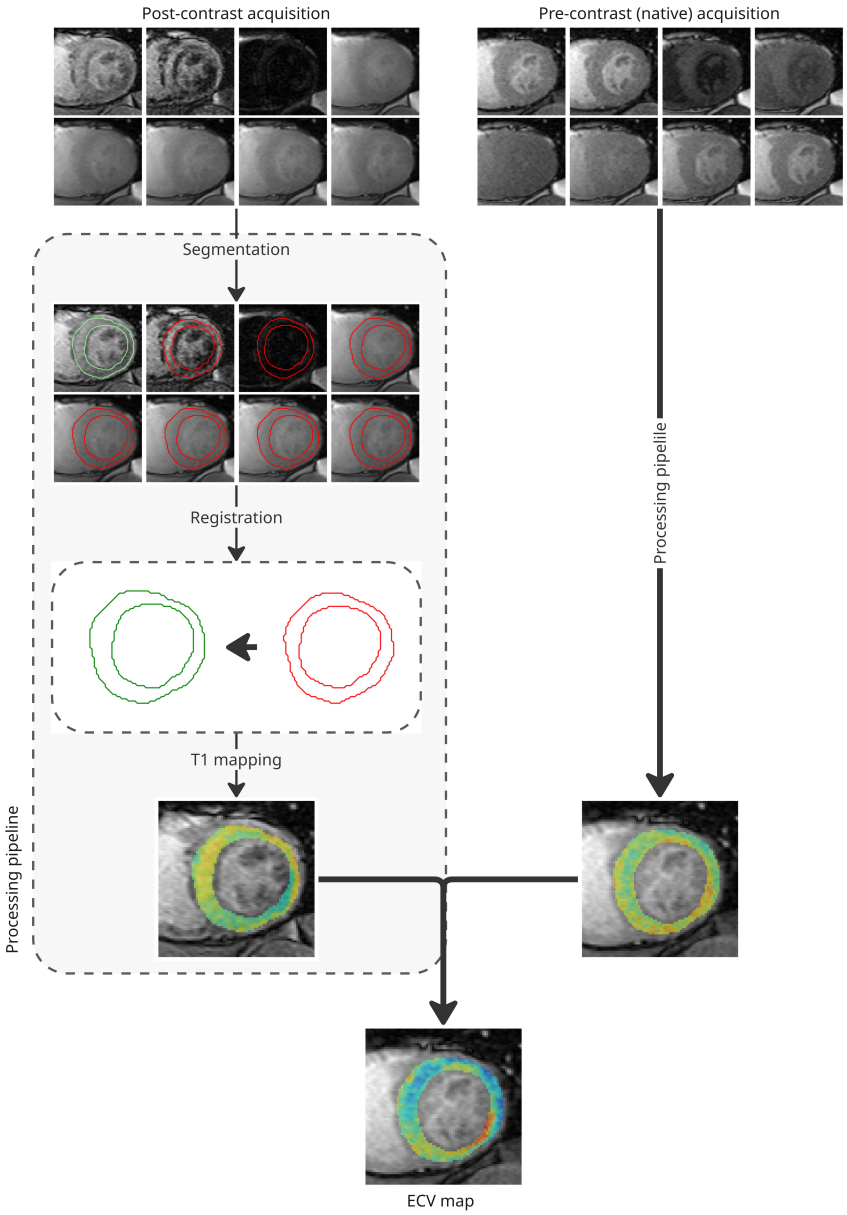

Figure 1: Proposed workflow for the computation of T1 maps and ultimately ECV maps from T1-weighted MOLLI sequences. The myocardium is segmented on all echoes of the acquisition. The segmentation is then used to perform an affine registration (polar transform [3]) to the first echo. An exponential fitting is then applied to extract the T1 maps. This process is performed on the native (pre-contrast) and post-contrast images. The native and post-contrast T1 maps are combined pixelwise to obtain the ECV map

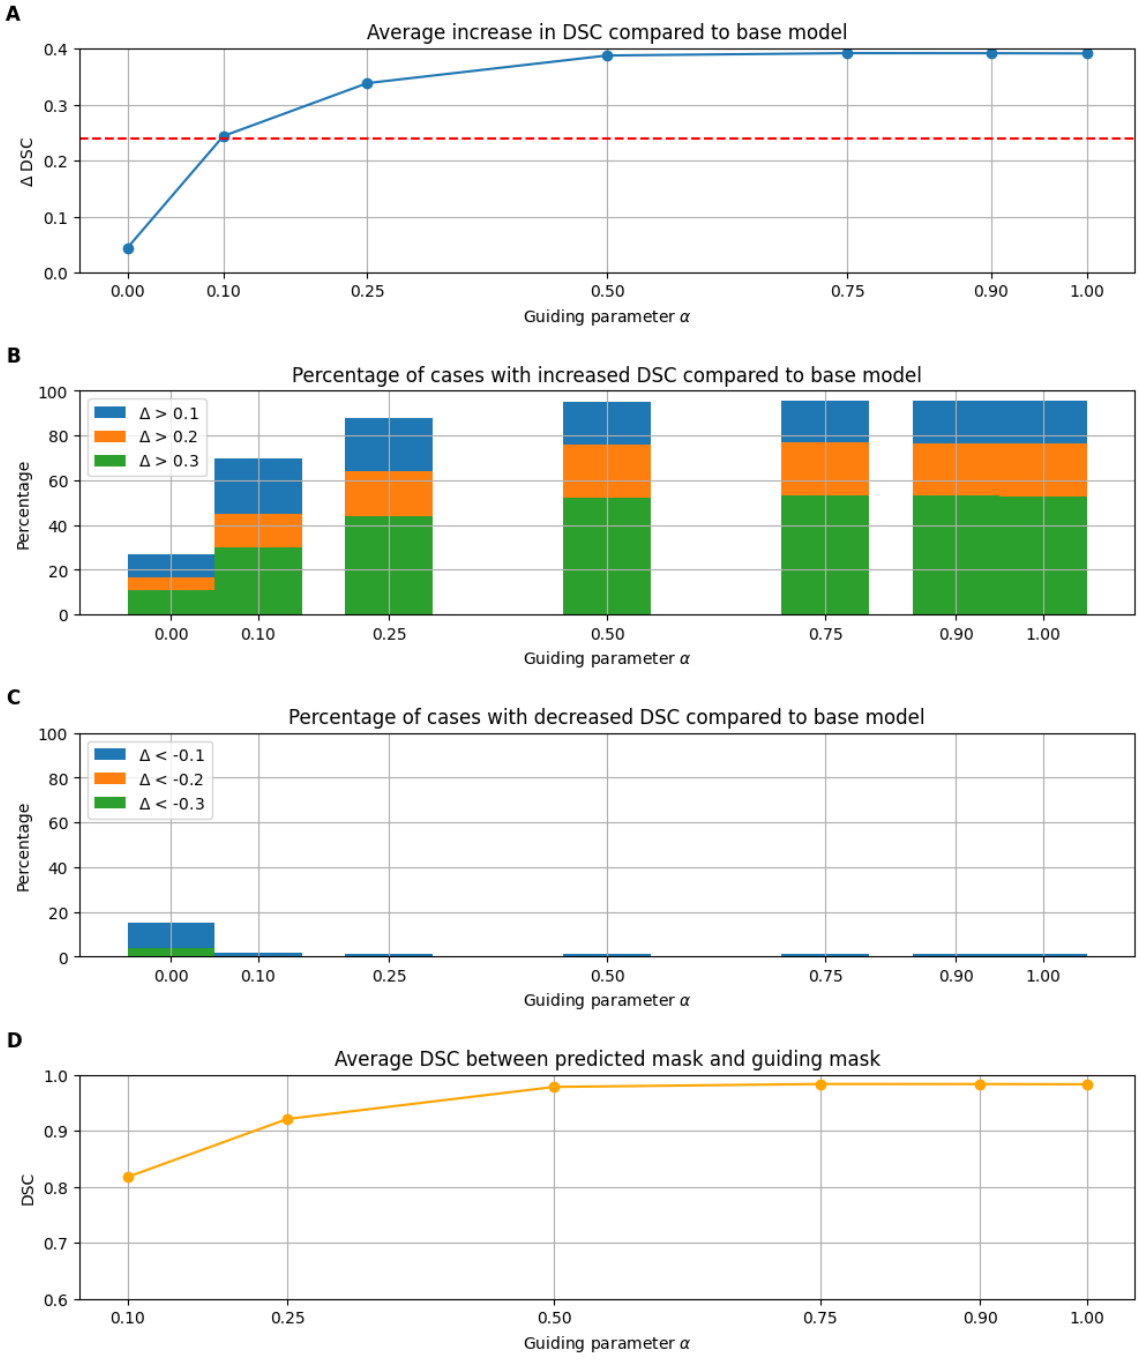

Figure 2: Performance of the proposed method on native and post-contrast images when using the ground truth mask of the first echo of the acquisition as the guiding mask. (A) Increase in dice score (DSC) between the conventional UNet and the proposed method depending on the guiding parameter α. The dashed red line corresponds to the improvement in DSC of SAM2 with regards to the base mod-el. (B-C) Percentage of cases that have DSC increased/decreased by more than 0.1, 0.2, 0.3 with respect to the conventional UNet. (D) DSC between the prediction and the guiding mask.

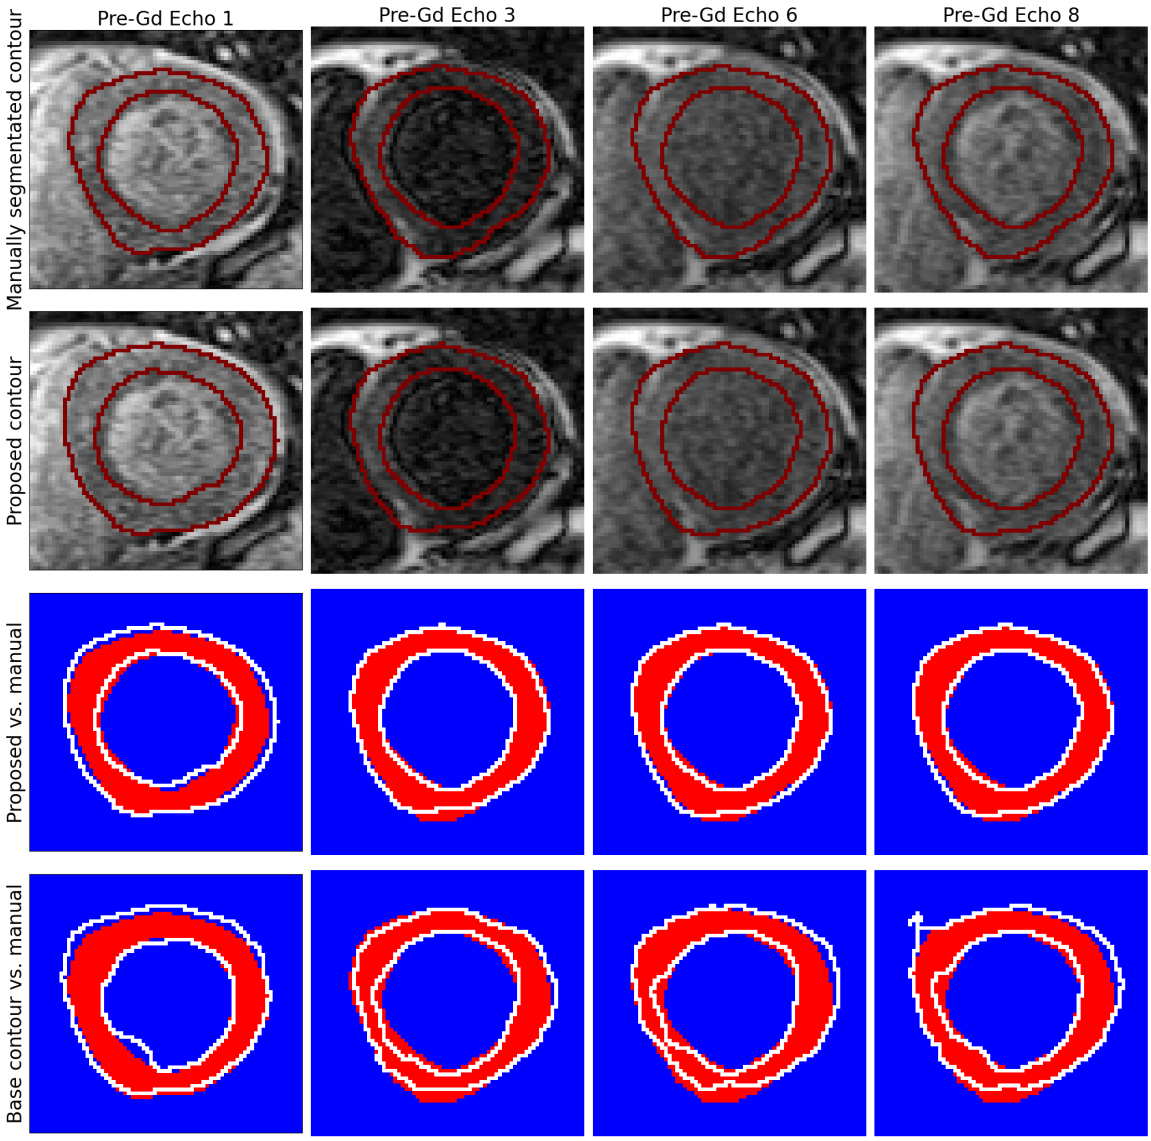

Figure 3: Qualitative analysis of the consistency model. The first row shows the ground truth (GT) masks over-laid on each echo of the acquisition, while the second row shows the output of the proposed meth-od. For display purposes, the underlying MR image contrast has been enhanced using histogram equalization. The last two rows display the proposed (row 3) and baseline (row 4) contour overlaid on top of the GT mask (in red).

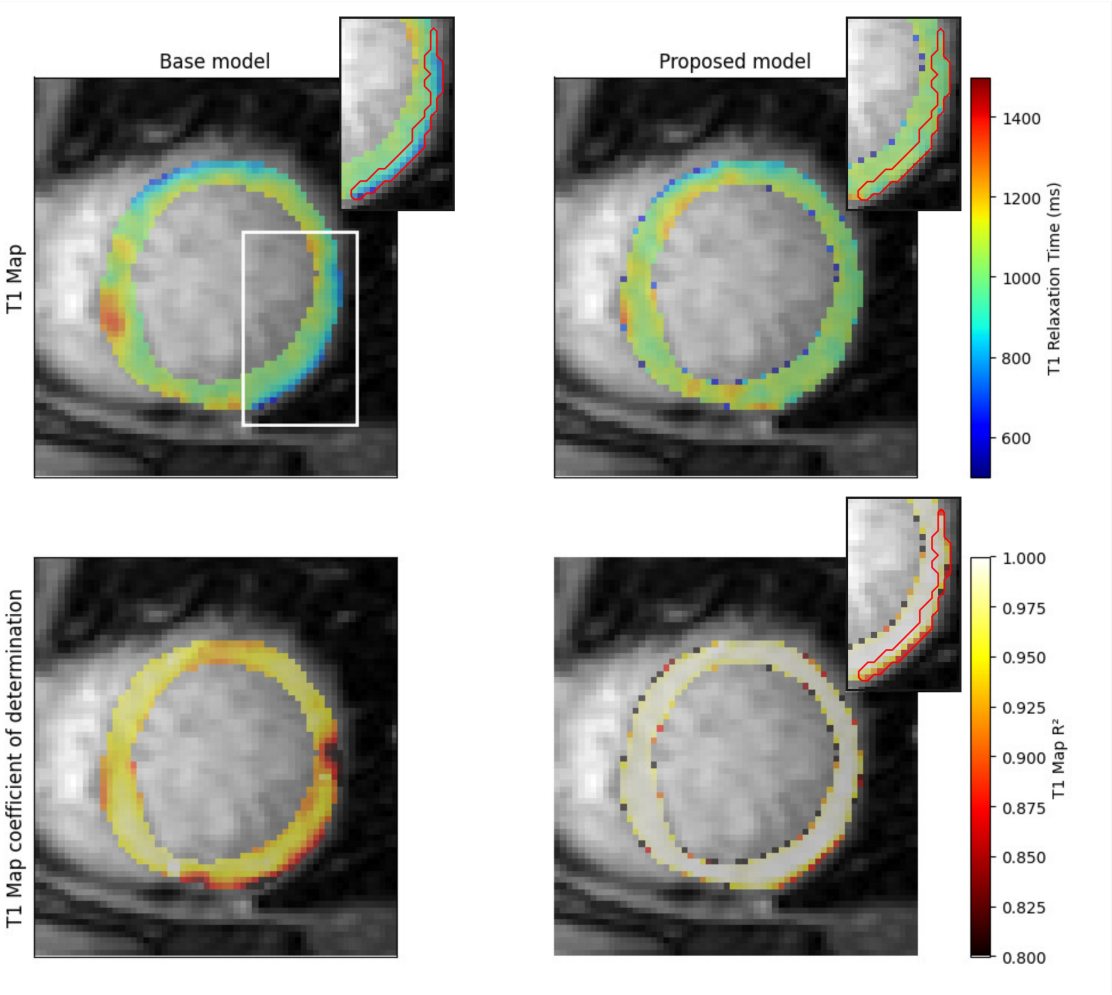

Figure 4: T1 maps and mapping coefficient of determination obtained for the UNet base-model and proposed methods. The white rectangle is zoomed in to highlight a hotspot present in the base model (average T1 = 848ms/average coefficient of determination = 0.86) but not in the proposed model (1010ms, 0.96). The lower coefficient of determination in the base model suggests that the hotspot is an artefact, potentially from a mis-segmentation of the blood pool. The higher errors at the interface of the myocardium probably indicate registration errors.

References

Hossam El-Rewaidy, Maryam Nezafat, Jihye Jang, Shiro Nakamori, Ahmed S. Fahmy, and Reza Nezafat. Nonrigid active shape model–based registration framework for motion correction of cardiac t1 mapping. *Magnetic Resonance in Medicine*, 80(2):780–791, 2018. doi: https://doi.org/10.1002/mrm.27068. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.27068.

Ahmed S. Fahmy, Hossam El-Rewaidy, Maryam Nezafat, Shiro Nakamori, and Reza Nezafat. Automated analysis of cardiovascular magnetic resonance myocardial native t1 mapping images using fully convolutional neural networks. *Journal of Cardiovascular Magnetic Resonance*, 21(1):7, 2019. ISSN 1097-6647. doi: https://doi.org/10.1186/s12968-018-0516-1. URL https://www.sciencedirect.com/science/article/pii/S1097664723001801.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021. URL https://arxiv.org/abs/2106.09685.

Fabian Isensee, Paul F. Jaeger, Simon A. A. Kohl, Jens Petersen, and Klaus H. Maier-Hein. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, December 2020. ISSN 1548-7091, 1548-7105. doi: 10/ghns3w. URL http://www.nature.com/articles/s41592-020-01008-z.

Daniel R. Messroghli, James C. Moon, Vanessa M. Ferreira, Lars Grosse-Wortmann, Taigang He, Peter Kellman, Julia Mascherbauer, Reza Nezafat, Michael Salerno, Erik B. Schelbert, Andrew J. Taylor, Richard Thompson, Martin Ugander, Ruud B. van Heeswijk, and Matthias G. Friedrich. Clinical recommendations for cardiovascular magnetic resonance mapping of t1, t2, t2* and extracellular volume: A consensus statement by the society for cardiovascular magnetic resonance (scmr) endorsed by the european association for cardiovascular imaging (eacvi). *Journal of Cardiovascular Magnetic Resonance*, 19(1):75, 2016. ISSN 1097-6647. doi: https://doi.org/10.1186/s12968-017-0389-8. URL https://www.sciencedirect.com/science/article/pii/S1097664723011092.

Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos, 2024. URL https://arxiv.org/abs/2408.00714.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing.

Yanjie Zhu, Ahmed S. Fahmy, Chong Duan, Shiro Nakamori, and Reza Nezafat. Automated myocardial t2 and extracellular volume quantification in cardiac mri using transfer learning–based myocardium segmentation. *Radiology: Artificial Intelligence*, 2(1):

e190034, 2020. doi: 10.1148/ryai.2019190034. URL https://doi.org/10.1148/ryai.2019190034. PMID: 32076664.