Reverse Imaging for Wide-spectrum Generalization of Cardiac MRI Segmentation

Yidong Zhao[1], Peter Kellman[2], Hui Xue[3], Tongyun Yang[1], Yi Zhang[1], Yuchi Han[4], Orlando Simonetti[4], Qian Tao[1]

[1] Department of Imaging Physics, Delft University of Technology, Delft, The Netherlands
[2] National Heart Lung and Blood Institute, National Institutes of Health, Bethesda, Maryland, USA
[3] Microsoft Research
[4] Cardiovascular Division, The Ohio State University Wexner Medical Center, Columbus, Ohio, USA
q.tao@tudelft.nl

Abstract. Pretrained segmentation models for cardiac magnetic resonance imaging (MRI) struggle to generalize across different imaging sequences due to significant variations in image contrast. These variations arise from changes in imaging protocols, yet the same fundamental spin properties, including proton density, T_1, and T_2 values, govern all acquired images. With this core principle, we introduce *Reverse Imaging*, a novel physics-driven method for cardiac MRI data augmentation and domain adaptation to fundamentally solve the generalization problem. Our method reversely infers the underlying spin properties from observed cardiac MRI images, by solving ill-posed nonlinear inverse problems regularized by the prior distribution of spin properties. We acquire this "spin prior" by learning a generative diffusion model from the multiparametric SAturation-recovery single-SHot acquisition sequence (mSASHA) dataset, which offers joint cardiac T_1 and T_2 maps. Our method enables approximate but meaningful spin-property estimates from MR images, which provide an interpretable "latent variable" that lead to highly flexible image synthesis of arbitrary novel sequences. We show that Reverse Imaging enables highly accurate segmentation across vastly different image contrasts and imaging protocols, realizing wide-spectrum generalization of cardiac MRI segmentation.

Keywords: CMR Segmentation · Reverse Imaging · MR Physics.

1 Introduction

Segmentation is crucial for evaluating cardiac biomarkers from cardiac magnetic resonance imaging (MRI), for example, the ejection fraction from balanced steady-state free-precession (bSSFP) cine imaging [3,4]. Learning-based methods have become the standard practice for cardiac MRI segmentation [15,3,4,17],

mostly trained on widely accessible bSSFP cine images with manual annotations [3,4,17]. However, models trained on bSSFP cine are prone to fail in test images with heterogeneous contrasts [4,17]. This can be partially mitigated by data augmentation (e.g. blur, gamma correction, etc.) [15], which effectively generalizes pretrained models to bSSFP cine images cross vendors and centers [27,4,6].

However, a much stronger contrast change comes from different MRI sequences, which cannot be tackled by common image augmentation techniques. For example, an alternative to bSSFP is the gradient echo sequence (GRE), which is less affected by field inhomogeneities [23] and more suitable for patients with implanted devices [10]. However, GRE has much poorer contrasts and easily fails the models trained on bSSFP. An even stronger contrast change occurs in quantitative imaging, such as T_1 mapping by the modified Look-Locker Inversion Recovery Sequence (MOLLI) [18]. In MOLLI, the image readouts have substantial contrast variations during the spin relaxation process, alternating between bright and dark blood, with low myocardium-blood contrast. Consequently, segmentation models trained on bSSFP images (source domain) can fail completely on MOLLI images (target domain) even with extensive data augmentation.

To tackle the cross-sequence generalization challenge, previous works have focused on disentangling "*content*" and "*style*" [20,19]. BayeSeg learns to remove the style and extracts the boundary for segmentation [11,12]. Domain adaptation techniques strive to learn content and style embeddings [5,28]. The learned "content" is supposed to be shared between sequences, while cross-sequence translation is achieved by interchanging the learned "style" [5,28,9]. However, such strategies require data from new sequences in order to learn the disentangling, and, importantly, the extracted content or style is barely interpretable. In this study, we argue that the fundamental but overlooked "content" is the underlying *spin properties*, including the voxel-wise proton density (PD), T_1, and T_2 [13] of blood and tissue. Given this content, the style is fully governed by the MRI signal models of new pulse sequences. This motivates us to infer the spin properties from the observed image, for more interpretable and physics-grounded image translation to tackle the cross-sequence domain generalization challenge.

Unlike quantitative MRI that directly maps spin properties [18,16,7], retrieving the underlying spin properties from a qualitative image such as bSSFP is ill-posed because infinitely many solutions may explain the same image. This challenge parallels that of accelerated MRI, where partial measurements are insufficient for image reconstruction, yet the inverse problem can be solved by incorporating a regularizer that defines the prior distribution of plausible images [21]. Recent MRI reconstruction works use generative diffusion models [14] as *a priori* regularizers, allowing high-quality reconstruction from partial measurements [26,8]. Following the same spirit, we introduce a diffusion-based generative prior on spin properties, called "spin prior", to regularize the ill-posed inverse problem of spin property estimation. Specifically, we leverage the multiparametric SAturation-recovery single-SHot acquisition (mSASHA) datasets, which perform joint cardiac T_1 and T_2 mapping, to learn the spin prior. We term the process of inferring spin properties from observed images as *Reverse*

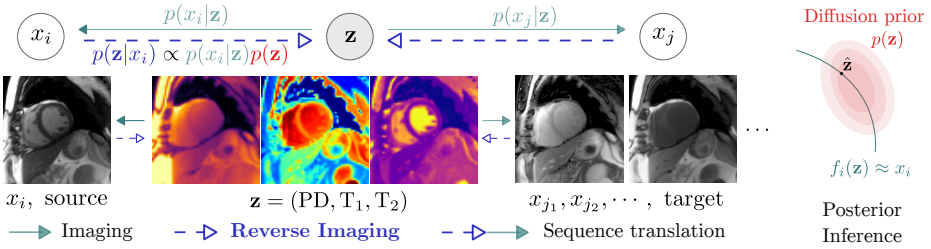

Fig. 1. Images acquired by sequences i and j from a subject share the same underlying spin properties $\mathbf{z}$. Reverse Imaging (dashed arrows) estimates $\hat{\mathbf{z}}$ from the prior $p(\mathbf{z})$ that explains the observed image x_i. Images of an arbitrary sequence can be generated by MR physics from $\hat{\mathbf{z}}$ (solid arrows). In this process, the only unknown is the prior $p(\mathbf{z})$, which we learn from the mSASHA data by a diffusion model (red).

Imaging. Reverse Imaging uniquely enables physics-grounded data augmentation and domain adaptation without the need for target domain data. We will show that Reverse Imaging leads to high-quality zero-shot generalization to a wide spectrum of MRI sequences. We make the following contributions:

- We propose an interpretable cross-sequence translation framework for cardiac MRI, which explicitly formulates the translation as applying MRI physics forward models to estimated tissue spin properties (i.e. PD, T_1 and T_2).
- We introduce Reverse Imaging, a novel approach that reversely infers spin properties from observed MR images using a physics-guided reverse diffusion process. The diffusion model serves as a generative prior that regularizes the estimation of the underlying spin properties.
- With Reverse Imaging, we achieve high-quality zero-shot generalization of cardiac MRI segmentation, from bSSFP to a wide spectrum of unseen contrasts and imaging protocols.

2 Method

2.1 Physics-based Cross-Sequence Translation

Physics-Prior Decomposition Let x_i and x_j be images acquired using sequences i and j, on which we aim to perform cross-sequence translation. The direct translation $p(x_j|x_i)$ from x_i to x_j is intractable. To solve this, we propose a decomposition of $p(x_j|x_i)$, by introducing an underlying physics $\mathbf{z} = (\mathrm{PD}, \mathrm{T}_1, \mathrm{T}_2)$ as a "latent variable" :

$$p(x_j|x_i) = \int_{\mathbf{z}} p(x_j, \mathbf{z}|x_i) d\mathbf{z} = \int_{\mathbf{z}} p(x_j|\mathbf{z}, x_i) p(\mathbf{z}|x_i) d\mathbf{z} = \int_{\mathbf{z}} p(x_j|\mathbf{z}) p(\mathbf{z}|x_i) d\mathbf{z}. \quad (1)$$

A probability graph model is shown in Fig. 1. Reverse Imaging can be formulated as posterior inference $p(\mathbf{z}|x_i)$, expressed as:

$$\log p(\mathbf{z}|x_i) \propto \log p(x_i|\mathbf{z}) + \log p(\mathbf{z}) + \mathrm{Const}. \quad (2)$$

As a result of (1) and (2), the inference $p(x_i|x_j)$ relies on two likelihoods $p(x_i|\mathbf{z})$ and $p(x_j|\mathbf{z})$, which are directly accessible from the imaging physics forward model. The only missing piece for inference, therefore, is the prior distribution on the spin properties $p(\mathbf{z})$. Moreover, we can interchange x_i and x_j in (1) once the prior $p(\mathbf{z})$ is known, allowing dual-direction translation of both $p(x_j|x_i)$ and $p(x_i|x_j)$.

Likelihoods from MR Imaging Physics The log-likelihood of observing image x_i from spin properties $\mathbf{z}$ under sequence i is:

$$\log p(x_i|\mathbf{z}) \propto \|f_i(\mathbf{z}) - x_i\|_2^2, \tag{3}$$

where the signal equation (i.e. forward model) f_i can be expressed by the Bloch equation of MR imaging [13]. Below, we list the signal equations for the sequences involved in this work, namely, bSSFP, MOLLI T_1 mapping, and GRE. We denote the flip angle (FA) as ω, the repetition and echo times as TR and TE. The bSSFP imaging equation [24] with short TR and TE is:

$$f_{\mathrm{SS}}(\mathbf{z}) = \frac{\mathrm{PD}\sin(\omega)}{1+\cos(\omega)+[1-\cos(\omega)]\mathrm{T}_1/\mathrm{T}_2}. \tag{4}$$

For T_1 mapping, the phase-insensitive readout [25] at inversion time t_{INV} is:

$$f_{\mathrm{MOLLI}}(\mathbf{z}) = \left| f_{\mathrm{SS}}(\mathbf{z}) \left(1 - \mathrm{INV}\cdot\exp(-\frac{t_{\mathrm{INV}}}{\mathrm{T}_1^*})\right)\right|, \tag{5}$$

where the apparent T_1 is $\mathrm{T}_1^* = \left(\frac{1}{\mathrm{T}_1}\cos^2(\frac{\omega}{2}) + \frac{1}{\mathrm{T}_2}\sin^2(\frac{\omega}{2})\right)^{-1}$ and the inversion factor $\mathrm{INV} = 1 + \frac{\sin(\frac{\omega}{2})}{\sin(\omega)}\left(1+\cos(\omega)+[1-\cos(\omega)]\frac{\mathrm{T}_1}{\mathrm{T}_2}\right)$. For GRE cine images from patients with implanted devices, let $\mathrm{E}_1 = \exp(-\frac{\mathrm{TR}}{\mathrm{T}_1})$ and $\mathrm{E}_2 = \exp(-\frac{\mathrm{TE}}{\mathrm{T}_2})$, the GRE signal model [13] is described as

$$f_{\mathrm{GRE}}(\mathbf{z}) = \mathrm{PD}\sin(\omega)\frac{1-\mathrm{E}_1}{1-\cos(\omega)\mathrm{E}_1}\mathrm{E}_2. \tag{6}$$

2.2 Generative Prior on Spin Properties from mSASHA

To learn the spin prior $p(\mathbf{z})$, we leverage an advanced quantitative MRI sequence mSASHA, which enables joint T_1 and T_2 mapping [7]. mSASHA consists of a train of saturation recovery pulses (SR) with SR time TS, followed by a T_2-prep pulse with duration TD and echo time TE. Its signal equation is:

$$f_{\mathrm{mSASHA}}(\mathbf{z}) = A(\mathbf{z})\left\{1 - \left[1 - \left(1 - e^{-\frac{\mathrm{TS}}{\mathrm{TE}}}\right)e^{-\frac{\mathrm{TE}}{\mathrm{T}_2}}\right]e^{-\frac{\mathrm{TD}}{\mathrm{T}_1}}\right\}, \tag{7}$$

where the unknowns $A(\mathbf{z}), \mathrm{T}_1, \mathrm{T}_2$ can be extracted from multiple acquisitions of f_{mSASHA} with varying TS, TD and TE by least-square fitting [7]. A encodes PD and sequence parameters, and the approximation $\mathrm{PD} \approx A$ can be used [1,2].

To model the prior of spin properties $\mathbf{z} = (\mathrm{PD}, \mathrm{T}_1, \mathrm{T}_2)$, we propose to train a denoising diffusion probabilistic model (DDPM) [14] in the $\mathbf{z}$ space. DDPM adds noise to the real $\mathbf{z}$'s from mSASHA in the forward process $\mathbf{z}_t = \sqrt{1-\beta_t}\mathbf{z}_{t-1} + \sqrt{\beta_t}\epsilon$ for T steps, where β_t is the noise scheduler at t and $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. The learned model can generate data by starting with a Gaussian distribution $\mathbf{z}_T \sim \mathcal{N}(0, \mathbf{I})$ and reversing the diffusion process:

$$\mathbf{z}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{z}_t + (1-\alpha_t) \nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t) \right) + \sigma_t \epsilon, \tag{8}$$

where $\alpha_t = 1 - \beta_t$ and $\sigma_t^2 = \frac{1-\overline{\alpha}_{t-1}}{1-\overline{\alpha}_t}\beta_t$. The diffusion model learns a denoising backbone parameterized by θ, $\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t) \approx -\frac{1}{\sqrt{1-\overline{\alpha}_t}}\epsilon_\theta(\mathbf{z}_t, t)$, $\overline{\alpha}_t = \prod_{r=1}^{t} \alpha_r$, as the generative prior for $\mathbf{z}$.

2.3 Reverse Imaging and Cross-Sequence Synthesis

Reverse Imaging searches for $\mathbf{z}$ in the prior space that best explains the observation x_i. Following conditional generation principles [8], we replace $\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t)$ in (8) with $\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t | x_i)$:

$$\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t | x_i) \propto \nabla_{\mathbf{z}_t} \log p(x_i | \mathbf{z}_t) + \nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t). \tag{9}$$

We approximate the time-step-dependent likelihood $\log p(x_i|\mathbf{z}_t)$ by: $\log p(x_i|\mathbf{z}_t) \approx \log p(x_i|\tilde{\mathbf{z}}_0(\mathbf{z}_t))$, where $\tilde{\mathbf{z}}_0(\mathbf{z}_t) \approx (\mathbf{z}_t - \sqrt{1-\overline{\alpha}_t}\epsilon_\theta(\mathbf{z}_t))/\sqrt{\overline{\alpha}_t}$ [8]. Starting from $\mathbf{z}_T \sim \mathcal{N}(0, \mathbf{I})$, the recursive process for reverse the imaging sequence x_i is

$$\mathbf{z}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{z}_t - \frac{1-\alpha_t}{\sqrt{1-\overline{\alpha}_t}} \epsilon_\theta(\mathbf{z}_t, t) \right) + \sigma_t \epsilon - \xi \frac{\partial}{\partial \mathbf{z}_t} \left\| f_i\left(\tilde{\mathbf{z}}_0(\mathbf{z}_t)\right) - x_i \right\|_2^2, \tag{10}$$

where ξ is a step size parameter. The guided reverse diffusion process in (10) yields the estimated spin properties $\hat{\mathbf{z}} = \mathbf{z}_0$ in the prior distribution that reconstructs x_i. After Reverse Imaging, we can synthesize images of novel sequences using the imaging physics equations f_j described in (4)-(6). Synthesizing other imaging sequences is also possible by sequence-specific Bloch equations. In addition, we can also use the image physics f_j of x_j to guide generation by replacing f_i, x_i with f_j, x_j in (10) and translating from target to source (T2S).

3 Experiments

Datasets We used the ACDC dataset [3] training split with bSSFP cine images of 100 subjects as the source domain dataset. The mSASHA dataset for learning the spin prior consists of 63 subjects with 186 short-axis slices. For evaluation of cross-sequence generalizability, we included two datasets imaged with different sequences: 1) MOLLI T_1 mapping data: 49 subjects with 131 short-axis slices in total. Each slice contains 11 baseline images with different contrasts. 2) cine MRI of 25 subjects with implanted devices: 16 subjects with the bSSFP sequence (Dev-bSSPF) and 9 with GRE sequence (Dev-GRE). Both the MOLLI and the Device datasets have significantly different contrasts than those of the ACDC.

Experimental Settings The DDPM is pretrained on ACDC images for $36,000$ steps and then fine-tuned with the real spin properties of from mSASHA for $6,000$ steps on the resolution of 128×128. We employ the hugging-face `diffuser` for the implementation of DDPM [22]. The estimated spin properties were resized to their original resolution after Reverse Imaging. We perform $T = 1000$ forward steps and 200 reverse steps of diffusion in (10) for Reverse Imaging, with the step size $\xi = 400$. As the exact FA for ACDC scans is unknown, we approximate it with $\omega = 45^\circ$ as typically used. We open source the estimated spin properties for the ACDC dataset and code[1]. For evaluating the generalization performance, we include the following methods:

- **Baseline** We train an nnUNet[15] with the ACDC data as the baseline;
- **BayeSeg** We compare our model against BayeSeg [11,12] which extracts the boundary and contour information from images for generalization;
- **RI-T2S** We translate the MOLLI and GRE data into bSSFP cine by Reverse Imaging (RI) and perform segmentation on the translated (T2S) images with the baseline model;
- **RI-Aug.** We perform Reverse Imaging on the ACDC cine images and extend the nnUNet framework by synthesizing MOLLI and GRE images through the estimated spin properties for physics-based augmentation in (5) and (6).

4 Results and Discussion

4.1 Reverse Imaging

We first show an example of Reverse Imaging a bSSFP cine in Fig. 2-(a). The estimated spin properties $\hat{\mathbf{z}}$ accurately reconstruct the bSSFP image, indicating a high likelihood. The predicted PD has low contrast between tissues, while in T_1 and T_2 estimation, blood has a high T_1 and T_2, myocardium has high T_1 but low T_2, and fat has low T_1 but a high T_2. These match the reference relative magnitude of the spin properties [29]. Based on estimated spin properties, images of various imaging protocols can be generated with large contrast variations, as shown in Fig. 2-(b). Reverse Imaging does not provide precise spin property estimation, but we note that they are not used for quantitative evaluation purposes w.r.t. T_1 or T_2, instead, they serve to synthesize novel contrasts that are essential for wide-spectrum generalization. Reverse Imaging also enables the T2S translation from MOLLI to bSSFP cine (Fig. 2-(c)). Practically, our reverse imaging results of cine images can be used for the prior training although the mSASHA dataset is not widely available.

4.2 Segmentation Performance

We now demonstrate that Reverse Imaging can significantly improve the generalization of segmentation models. In Table 1, we list the segmentation accuracy

[1] `https://github.com/Ido-zh/cmr_reverse.git`

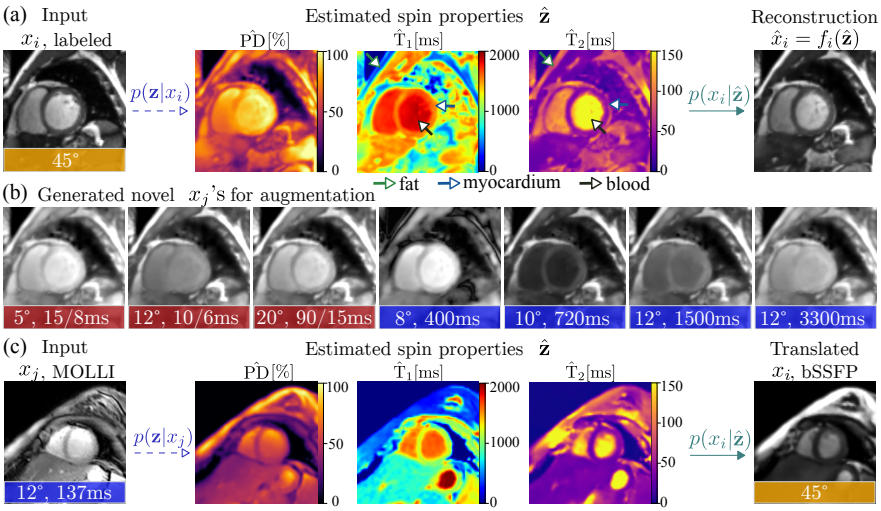

Fig. 2. (a) Reverse Imaging a bSSFP image x_i (yellow). The estimated $\hat{\mathbf{z}}$ accurately reconstructs $\hat{x}_i \approx x_i$. (b) Generating GRE (red) and MOLLI (blue) images with $\hat{\mathbf{z}}$. The FA in degrees and TR/TE or t_{INV} in milliseconds are given for each image. (c) Reverse Imaging a MOLLI readout for target-to-source (T2S) translation.

measured by the Dice score of the left ventricle (LV), myocardium (MYO) and right ventricle (RV) in the challenging MOLLI and Device datasets. To further evaluate the robustness against contrast variation, we also report the Dice for each baseline image of MOLLI in Fig. 3, as MOLLI readouts show strong contrast variation during inversion recovery. Qualitative examples of segmentation results are shown in Fig. 4.

Table 1. Segmentation accuracy on novel sequences measured by Dice score [%]. Statistical significance ($p < 0.05$) is indicated by *.

Methods	**MOLLI**			**Device**		
	RV	MYO	LV	RV	MYO	LV
Baseline [15]	24.0 ± 12.9	47.4 ± 19.5	39.9 ± 22.0	68.9 ± 25.2	86.9 ± 8.6	91.9 ± 7.6
BayeSeg [12]	51.5 ± 26.9	41.4 ± 19.9	57.0 ± 24.1	53.6 ± 21.0	62.6 ± 12.6	72.6 ± 19.1
RI-T2S	63.6 ± 36.5	69.8 ± 20.1	82.7 ± 21.0	66.4 ± 24.2	82.3 ± 15.2	88.6 ± 12.1
RI-Aug.	$\mathbf{87.0 \pm 17.8}^*$	$\mathbf{86.5 \pm 9.8}^*$	$\mathbf{91.6 \pm 9.3}^*$	$\mathbf{87.4 \pm 9.5}^*$	$\mathbf{93.1 \pm 4.0}^*$	$\mathbf{96.0 \pm 3.5}^*$

The baseline nnUNet generalizes poorly in the MOLLI images (Table 1), especially when the blood-myocardium contrast is inverted (see #4-#7 in Fig. 3 and LL-#2 in Fig. 4). BayeSeg can improve performance on dark blood images, yet it can still fail on low-contrast images (e.g. Fig. 4 LL-#3-#5). The device

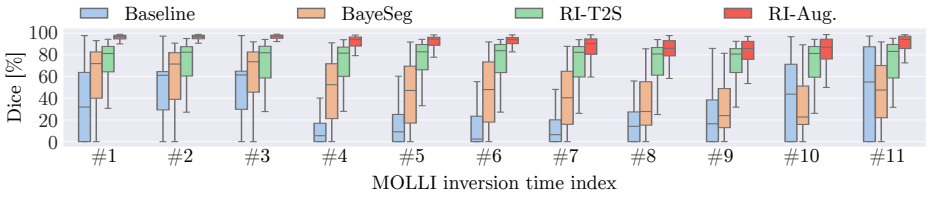

Fig. 3. Segmentation accuracy on each MOLLI readout during inversion recovery.

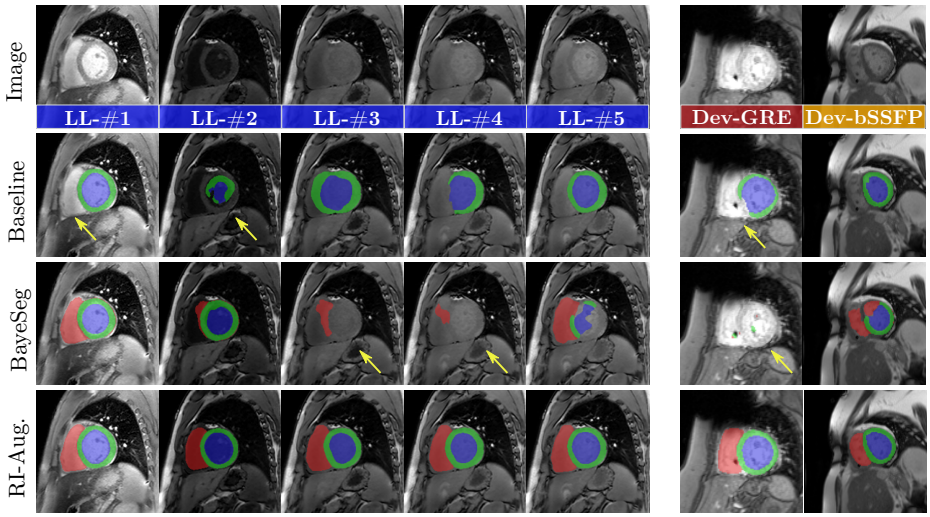

Fig. 4. Segmentation of 5 MOLLI baseline images (LL-#1-#5) and 2 device cine images, all of which have different contrast than bSSFP cine images used for training.

images have bright blood but lower contrast, and the baseline nnUNet works more robustly than BayeSeg with higher Dice scores (Table 1). Performing the T2S transform (RI-T2S) significantly improved the segmentation accuracy for MOLLI, but a performance drop is observed in the device dataset. This is because the short TE ($\approx$ 1.5 ms) in GRE minimizes T_2 weighting (cf. (6)). In this case, Reverse Imaging is limited by the little amount of information carried by the observed image. However, this can be solved by using Reverse Imaging for augmentation instead (RI-Aug.), which achieves the highest segmentation accuracy in both datasets, significantly better than all other methods (Table 1, $p < 0.05$). With RI-Aug., the trained model consistently achieves highly accurate segmentation on images with various contrasts, including low-contrast, bright and dark blood, and GRE images (Fig. 4).

5 Conclusion

We introduced Reverse Imaging, a novel physics-grounded approach for cross-sequence generalization in cardiac MRI segmentation. Combining imaging physics and a generative diffusion prior, our method estimates underlying spin properties from observed images. Reverse Imaging provides an interpretable and physically grounded solution for domain adaptation, addressing the challenge of contrast variations across imaging protocols. Our experiments demonstrate that Reverse Imaging significantly boosts segmentation robustness, enabling zero-shot generalization to unseen sequences without requiring target-domain data. This approach offers a new paradigm for improving cardiac MRI segmentation generalizability, with potential applications to broader medical imaging tasks where domain shifts pose a significant challenge.

Acknowledgments. We acknowledge the financial support from the TU Delft AI Initiative and the Dutch Research Council (NWO).

Disclosure of Interests. We the authors hereby declare no competing interests in this paper.

References

1. Akçakaya, M., Basha, T.A., Weingärtner, S., Roujol, S., Berg, S., Nezafat, R.: Improved quantitative myocardial T2 mapping: Impact of the fitting model. Magnetic resonance in medicine **74**(1), 93–105 (2015)
2. Akçakaya, M., Weingärtner, S., Basha, T.A., Roujol, S., Bellm, S., Nezafat, R.: Joint myocardial T1 and T2 mapping using a combination of saturation recovery and t2-preparation. Magnetic resonance in medicine **76**(3), 888–896 (2016)
3. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Ballester, M.A.G., et al.: Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: is the problem solved? IEEE transactions on medical imaging **37**(11), 2514–2525 (2018)
4. Campello, V.M., Gkontra, P., Izquierdo, C., Martin-Isla, C., Sojoudi, A., Full, P.M., Maier-Hein, K., Zhang, Y., He, Z., Ma, J., et al.: Multi-centre, multi-vendor and multi-disease cardiac segmentation: the M&Ms challenge. IEEE Transactions on Medical Imaging **40**(12), 3543–3554 (2021)
5. Chen, C., Ouyang, C., Tarroni, G., Schlemper, J., Qiu, H., Bai, W., Rueckert, D.: Unsupervised multi-modal style transfer for cardiac MR segmentation. In: Statistical Atlases and Computational Models of the Heart. Multi-Sequence CMR Segmentation, CRT-EPiggy and LV Full Quantification Challenges: 10th International Workshop, STACOM 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 13, 2019, Revised Selected Papers 10. pp. 209–219. Springer (2020)
6. Chen, C., Qin, C., Ouyang, C., Li, Z., Wang, S., Qiu, H., Chen, L., Tarroni, G., Bai, W., Rueckert, D.: Enhancing MR image segmentation with realistic adversarial data augmentation. Medical Image Analysis **82**, 102597 (2022)

7. Chow, K., Hayes, G., Flewitt, J.A., Feuchter, P., Lydell, C., Howarth, A., Pagano, J.J., Thompson, R.B., Kellman, P., White, J.A.: Improved accuracy and precision with three-parameter simultaneous myocardial T1 and T2 mapping using multi-parametric SASHA. Magnetic Resonance in Medicine **87**(6), 2775–2791 (2022)
8. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. arXiv preprint arXiv:2209.14687 (2022)
9. Cui, Z., Li, C., Du, Z., Chen, N., Wei, G., Chen, R., Yang, L., Shen, D., Wang, W.: Structure-driven unsupervised domain adaptation for cross-modality cardiac segmentation. IEEE Transactions on Medical Imaging **40**(12), 3604–3616 (2021)
10. Ferreira, P.F., Gatehouse, P.D., Mohiaddin, R.H., Firmin, D.N.: Cardiovascular magnetic resonance artefacts. Journal of Cardiovascular Magnetic Resonance **15**(1), 41 (2013)
11. Gao, S., Zhou, H., Gao, Y., Zhuang, X.: Joint modeling of image and label statistics for enhancing model generalizability of medical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 360–369. Springer (2022)
12. Gao, S., Zhou, H., Gao, Y., Zhuang, X.: BayeSeg: Bayesian modeling for medical image segmentation with interpretable generalizability. Medical Image Analysis **89**, 102889 (2023)
13. Hashemi, R.H., Bradley, W.G., Lisanti, C.J.: MRI: the basics: The Basics. Lippincott Williams & Wilkins (2012)
14. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
15. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods **18**(2), 203–211 (2021)
16. Kellman, P., Hansen, M.S.: T1-mapping in the heart: accuracy and precision. Journal of cardiovascular magnetic resonance **16**(1), 2 (2014)
17. Martín-Isla, C., Campello, V.M., Izquierdo, C., Kushibar, K., Sendra-Balcells, C., Gkontra, P., Sojoudi, A., Fulton, M.J., Arega, T.W., Punithakumar, K., et al.: Deep learning segmentation of the right ventricle in cardiac MRI: the M&Ms challenge. IEEE Journal of Biomedical and Health Informatics **27**(7), 3302–3313 (2023)
18. Messroghli, D.R., Radjenovic, A., Kozerke, S., Higgins, D.M., Sivananthan, M.U., Ridgway, J.P.: Modified look-locker inversion recovery (MOLLI) for high-resolution t1 mapping of the heart. Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine **52**(1), 141–146 (2004)
19. Ouyang, C., Chen, C., Li, S., Li, Z., Qin, C., Bai, W., Rueckert, D.: Causality-inspired single-source domain generalization for medical image segmentation. IEEE Transactions on Medical Imaging **42**(4), 1095–1106 (2022)
20. Ouyang, C., Kamnitsas, K., Biffi, C., Duan, J., Rueckert, D.: Data efficient unsupervised domain adaptation for cross-modality image segmentation. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22. pp. 669–677. Springer (2019)
21. Pauly, J.M.: Compressed sensing MRI. Signal Processing Magazine, IEEE **25**(2), 72–82 (2008)
22. von Platen, P., Patil, S., Lozhkov, A., Cuenca, P., Lambert, N., Rasul, K., Davaadorj, M., Nair, D., Paul, S., Berman, W., Xu, Y., Liu, S., Wolf, T.: Diffusers: State-of-the-art diffusion models. `https://github.com/huggingface/diffusers` (2022)

23. Ridgway, J.P.: Cardiovascular magnetic resonance physics for clinicians: part I. Journal of cardiovascular magnetic resonance **12**(1), 71 (2010)
24. Scheffler, K., Lehnhardt, S.: Principles and applications of balanced SSFP techniques. European radiology **13**, 2409–2418 (2003)
25. Schmitt, P., Griswold, M.A., Jakob, P.M., Kotas, M., Gulani, V., Flentje, M., Haase, A.: Inversion recovery TrueFISP: quantification of t1, t2, and spin density. Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine **51**(4), 661–667 (2004)
26. Song, Y., Shen, L., Xing, L., Ermon, S.: Solving inverse problems in medical imaging with score-based generative models. arXiv preprint arXiv:2111.08005 (2021)
27. Tao, Q., Yan, W., Wang, Y., Paiman, E.H., Shamonin, D.P., Garg, P., Plein, S., Huang, L., Xia, L., Sramko, M., et al.: Deep learning–based method for fully automatic quantification of left ventricle function from cine MR images: a multivendor, multicenter study. Radiology **290**(1), 81–88 (2019)
28. Xie, Q., Li, Y., He, N., Ning, M., Ma, K., Wang, G., Lian, Y., Zheng, Y.: Unsupervised domain adaptation for medical image segmentation by disentanglement learning and self-training. IEEE Transactions on Medical Imaging **43**(1), 4–14 (2022)
29. Xu, Z., Li, W., Wang, J., Wang, F., Sun, B., Xiang, S., Luo, X., Meng, Y., Wang, X., Wang, X., et al.: Reference ranges of myocardial t1 and t2 mapping in healthy chinese adults: a multicenter 3t cardiovascular magnetic resonance study. Journal of Cardiovascular Magnetic Resonance **25**(1), 64 (2023)