



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Structure-Contrast Disentangled INRs for Accelerated Multi-contrast MRI Reconstruction

Zach Vavasour^{1,2}, Kyle Vavasour^{3,4}, and Mark Chiew^{1,2}

¹ Department of Medical Biophysics, University of Toronto, Toronto, Canada

² Physical Sciences, Sunnybrook Research Institute, Toronto, Canada

³ Department of Computer Science, University of Toronto, Toronto, Canada

⁴ Department of Statistical Sciences, University of Toronto, Toronto, Canada

`zach.vavasour@mail.utoronto.ca`

Abstract. Implicit neural representations (INRs) have demonstrated significant promise as a patient-specific alternative to data-driven deep learning (DL) for accelerated MRI reconstruction. However, without additional priors, they struggle to match state-of-the-art supervised reconstruction performance, particularly at higher acceleration factors. To address this, recent research has explored leveraging minimally biased information sharing between multiple contrasts within a subject’s own dataset. In this work, we demonstrate that straightforward extensions of INR methods to the multi-contrast setting do not efficiently capitalize on the complementary information present in the shared underlying structures between contrasts. We therefore propose a novel structure-contrast **DISentangled Implicit Neural Representation (DISINR)** framework for multi-contrast MRI reconstruction, specifically designed to disentangle the shared structural information from each contrast image during the training process. We demonstrate that our method out-performs previous patient-specific multi-contrast and single-contrast methods on the 0.3T M4RAW dataset, producing robust reconstructions at accelerations up to $8\times$. Additionally, we visualize the internal features of our model and demonstrate that it learns a disentangled structure representation, enabling it to better support mixed acceleration protocols. Our code is available at: <https://github.com/chiew-group/DISINR>.

Keywords: MRI · Reconstruction · Implicit Neural Representations · Multi-contrast.

1 Introduction

Magnetic Resonance Imaging (MRI) is a critical imaging modality in a wide variety of clinical settings [1, 9, 19], offering strong flexibility for both structural and functional imaging with greater soft tissue contrast than alternatives such as X-ray or CT [18]. However, due to the physical constraints of the sampling process, MRI acquisition is notably very slow, with standard clinical protocols requiring many minutes to acquire quality images of the target anatomy.

Over the past decade, supervised deep learning (DL) methods have demonstrated state-of-the-art reconstruction performance for accelerated (undersampled) imaging [2, 16]. However, their data-driven nature also introduces new challenges and risks. Firstly, obtaining raw MRI training datasets in sufficient quantity presents a significant hurdle. Moreover, the diversity of the chosen training data inevitably injects biases into the resulting model. When *domain-shift* occurs, (e.g., a model trained on high-field data was applied to out-of-distribution low-field), the quality of the resulting reconstructions can suffer [3].

Patient-specific reconstruction methods (also known as zero-shot) do not require any external training data and therefore avoid many of the challenges associated with supervised DL. By leveraging the architecture of DL without the data-driven training paradigm, these methods inherit the implicit signal priors of neural networks. Recently, Implicit Neural Representations (INRs) [15, 11, 13] in particular have demonstrated their strong capabilities for subject-specific MRI reconstruction [14, 5, 4]. INRs are typically light-weight multi-layer perceptrons (MLPs) that are trained to fit a single instance of a signal as a continuous function of spatial coordinates.

While INR-based MRI reconstruction has been shown to improve upon conventional methods, without additional prior information it struggles to compete with state-of-the-art supervised DL. Early methods, such as NeRP [14] explored incorporating priors from a scan of the same subject from a previous time-point. IMJENSE [5] incorporated joint estimation of coil sensitivities with INR-based reconstruction. More recently, Niessen et al. [12] extended INR methods to joint multi-contrast reconstruction, where information between multiple acquired contrasts of the same anatomy can be shared.

The extension of the reconstruction problem to multiple contrasts introduces a new powerful prior: consistency of the underlying structure between the contrasts. Previous research has demonstrated in data-driven DL settings that joint multi-contrast (MC) methods can surpass single contrast (SC) [21, 20]. However, an important limitation is that they become specialized to the specific set of contrasts used during training, further compounding risks of domain-shift. Because INRs are entirely patient specific, they can be applied to any arbitrary set of contrasts without these risks.

In this work we show that a straightforward extension of existing INR based frameworks to the MC setting is not sufficient to capitalize on the shared information between contrasts. Drawing inspiration from neural style transfer [6], we hypothesize that each contrast can be generated by modulating a shared structure representation. To this end, we propose a novel structure-contrast **DIS**entangled **INR** (DISINR) architecture designed to maximize information sharing between contrasts. Specifically, we introduce an architectural bottle-neck where structure and contrast encodings are combined, and demonstrate that by randomly masking structure features during training, we promote the learning of a shared structural representation. The main contributions of this work are as follows: 1) proposing a novel encoder-decoder style INR architecture for MC MRI reconstruction 2) quantitatively evaluating our method on the 0.3T M4RAW dataset

against both SC and MC patient-specific methods 3) exploring our architecture’s effectiveness at sharing complementary structure information.




Fig. 1. Overview of DISINR. The SIREN-based structure encoder (top left) takes spatial coordinates and produces a structure latent vector for each pixel. A contrast embedding (bottom left) is then concatenated onto the structure latent and passed through a small decoder (bottom right) to produce a contrast image.

2 Methods

2.1 Multi-contrast MRI Reconstruction

In the multi-contrast setting, the goal is to jointly reconstruct a set of images $X_1, X_2, \dots, X_c$, where each image $X_c \in \mathbb{C}^{h \times w}$ is of the same anatomy, but acquired with a different contrast. Each image is related to the acquired k-space as defined by the forward model

$$k_c = M_c F S X_c + \epsilon_c = A_c X_c + \epsilon_c \quad (1)$$

where M_c are the binary under-sampling masks, F is the Fourier transform, S are the coil sensitivities, and ϵ_c is noise. We define A_c as the overall forward model for each contrast.

2.2 Proposed Framework

Structure-Contrast Implicit Network. We model the set of contrast images with an encoder-decoder style INR network, parametrized by θ , which learns to map discrete spatial coordinates $(x, y) \in \mathbb{R}^2$ and a contrast index $c = 1, 2, \dots, N$ to pixel intensities. The network is composed of two encoders. The structure encoder $f_{str}(x, y)$ is a SIREN [15], which takes in spatial coordinates and produces a structural latent vector for each pixel. The contrast encoder $f_{con}(c)$ is a

single linear layer that takes in a contrast index and produces a contrast latent. The structure and contrast latent representations are then concatenated at the bottle-neck and passed through a small MLP decoder. The overall framework is presented in Fig. 1 Each contrast can be queried from the network as follows

$$\hat{X}_c(x, y) = f_{dec}(f_{str}(x, y), f_{con}(c)) \quad (2)$$

To train the network, we apply the forward model to compute a joint loss L across all contrasts, which is composed of a data consistency term and a total-variation (TV) based regularization term:

$$L = \sum_{c=1}^N |k_c - A_c \hat{X}_c|_1 + \lambda (|\nabla_x(\hat{X})|_1 + |\nabla_y(\hat{X})|_1) \quad (3)$$

Joint Total Variation Loss. The TV loss penalty is enforced jointly across contrasts to further emphasize consistency of underlying structure. Specifically, the gradient operator from equation (3) is defined as $\nabla_d(X) = \sqrt{\sum_c^N (\partial_d X_c)^2}$. In doing so, we incentivize alignment of edges and tissue boundaries.

Structure Latent Masking. To prevent "memorizing" of contrast information through different paths within the network, we additionally randomly mask features outputted by the structure encoder during the training process. By doing so, we encourage redundancy and impede the structure module from encoding specific contrast information.

Coil Sensitivity Estimation. In this work we jointly estimate coil sensitivities during the reconstruction process using the polynomial parametrization module proposed by Feng et al. in IMJENSE [5]. In our MC setting, we assume each contrast X_c shares a single set of sensitivity maps S .

3 Experiments and Results

3.1 Data and Sampling

We evaluated our method on 183 brains scans from the the publicly available M4RAW dataset [8], each containing three paired contrasts: T1-weighted, T2-weighted, and FLAIR. Each scan was acquired fully-sampled at 256×192 with 4 slices and 4 coil channels. This dataset was acquired on a custom low-field 0.3T scanner, resulting in notably lower SNR than standard clinical scans. The M4RAW dataset provides multiple repetitions (2-3); therefore, to increase SNR the average was used. Phase-ramp correction was applied before averaging. To simulate the accelerated MRI setting, we retrospectively under-sampled using a set of binary sampling masks, generated independently for each contrast. The masks contained a fully-sampled 16×16 center region followed by 2D random sampling for the remaining k-space points at acceleration factors $R=4,6,8$.

3.2 Comparisons

Comparison Methods. We evaluated our method against four existing patient specific MRI reconstruction methods. Firstly, we evaluate it against **IMJENSE** [5], the current state-of-the-art SC INR method. Secondly, we evaluate it against a simple extension of IMJENSE to the MC setting (**MC-IMJENSE**), where additional outputs for each contrast are added to the SIRENs. Third, we evaluate it against the hash-encoder [11] based MC INR method proposed by Niessen et al. [12] (**MC-INR**). We additionally introduce a TV-loss for MC-INR, as the baseline model exhibited noisy reconstructions with the low-SNR M4RAW dataset. Finally, we evaluate it against SC L1-wavelet compressed sensing.

Comparison Metrics. To quantify reconstruction quality, we report two metrics: peak-signal-to-noise ratio (PSNR) and structural similarity index (SSIM). ESPIRiT-derived brain masks were applied to both the predicted and reference images before computing all metrics.




Fig. 2. Example reconstructions for a patient from the M4RAW dataset at R=4.

3.3 Implementation Details

For all experiments the structure encoder SIREN has 7 hidden layers, each with 256 hidden features. The SIREN is initialized following the strategy proposed by Sitzmann et al. [15], with $w_0 = 10.0$. The structure encoder is then followed by a linear projection layer down to a length 64 "structure latent vector". During training, 25% of structure features were randomly masked at each iteration (implemented with dropout). The contrast encoder is a single linear layer outputting a length 64 embedding. The decoder MLP has two hidden layers with 64 features using ReLU activations, followed by a linear layer with two outputs corresponding to the real and imaginary components of the image. For all experiments the TV regularization weight $\lambda = 5.4$ was used. For the MC-INR method, we use the implementation provided by Niessen et al. [12], with an additional TV-loss term. MC-INR does not jointly predict coil sensitivities; therefore, they were pre-estimated with ESPIRiT [17]. Our method was trained with the Adam [7] optimizer, with a learning rate of 10^{-4} for the INR and 10^{-1} for the coil sensitivity estimator. Optimization was performed for 3000 iterations with the learning rate decayed by 0.8 every 500 steps. Our method is implemented in PyTorch and all experiments were performed on a single NVIDIA H100 GPU.

3.4 Results

Table 1 shows PSNR and SSIM scores for DISINR across acceleration factors on the M4RAW dataset. DISINR consistently outperforms existing methods with an average improvement over MC-IMJENSE of +0.2dB at R=4, increasing to +0.3dB at R=8. A paired t -test between DISINR and MC-IMJENSE was performed, showing statistical significance across all contrasts and acceleration factors ($p < 10^{-4}$ in all cases). Fig. 2 presents an example reconstruction for one dataset, demonstrating higher quality images and better reconstruction of high frequency details. We note that MC-IMJENSE performs only marginally better than IMJENSE, indicating that a straight-forward extension to the MC setting without explicitly disentangling structure and contrast leads to sub-optimal information sharing across contrasts.

Fig. 3 shows how performance varies when only one of the three contrasts is increasingly under-sampled. For one example subject, We fixed the T1 and T2 acceleration factors at $4\times$ and progressively increased the acceleration for the FLAIR. MC-IMJENSE exhibits a steeper decrease in reconstruction quality when compared to DISINR as the FLAIR acceleration increases, particularly at higher acceleration factors above $8\times$. This indicates that the shared structural information learned by DISINR from the combination of all three contrasts can counter-act the significantly higher acceleration of the FLAIR.

Ablations. We additionally conducted an ablation study to evaluate the contribution of each component of our proposed approach. Table 2 quantifies the effect of structure feature masking, the encoder-decoder architecture, and joint TV. We observe that the structure feature masking provides the greatest benefit,

Table 1. Quantitative results for DISINR across different acceleration factors on the M4RAW dataset. Values report the average across T1, T2, and FLAIR (mean $\pm$ std). * indicates a statistically significant improvement over MC-IMJENSE ($p < 0.05$).

Method	Metric	R=4	R=6	R=8
L1-Wavelet	PSNR	29.57 ± 0.33	27.94 ± 0.30	26.87 ± 0.36
	SSIM	0.879 ± 0.02	0.849 ± 0.02	0.829 ± 0.02
IMJENSE	PSNR	32.64 ± 0.20	31.35 ± 0.22	30.69 ± 0.21
	SSIM	0.921 ± 0.01	0.899 ± 0.01	0.888 ± 0.01
MC-IMJENSE	PSNR	32.69 ± 0.20	31.41 ± 0.22	30.76 ± 0.20
	SSIM	0.921 ± 0.01	0.901 ± 0.01	0.889 ± 0.01
MC-INR	PSNR	31.49 ± 0.23	29.80 ± 0.23	28.76 ± 0.22
	SSIM	0.910 ± 0.01	0.882 ± 0.02	0.863 ± 0.02
DISINR	PSNR	$32.89 \pm 0.20^*$	$31.69 \pm 0.22^*$	$31.06 \pm 0.20^*$
	SSIM	$0.926 \pm 0.01^*$	$0.907 \pm 0.01^*$	$0.896 \pm 0.01^*$




Fig. 3. Mean PSNR of fixed R=4 T1 and T2 with different FLAIR acceleration.

indicating its importance in enforcing shared structure. Notably, MC-IMJENSE with masking applied in the the penultimate hidden layer results in lower performance, highlighting its specific effect of encouraging redundancy of structural features when coupled with our DISINR architecture.

4 Discussion and Conclusion

Our results demonstrate that by explicitly disentangling contrast and structural information, DISINR outperforms previous patient-specific INR reconstruction methods. Within a subject, DISINR learns a shared structural representation, which can then modulated through the decoder to produce a given contrast. To support this interpretation, Fig. 4 demonstrates that the contrast embedding space encodes semantically meaningful contrast information, enabling modulation of the structure latent even with unseen synthetic contrast embeddings.

Table 2. Ablation study for our DISINIR approach. Average PSNR across all three contrasts is reported.

Method	M4RAW R=6
DISINR	31.69
DISINR - Masking	31.45
DISINR - Masking - jTV	31.35
MC-IMJENSE	31.41
MC-IMJENSE + Masking	31.03
MC-INR - TV	26.98


**Fig. 4. Visualization of contrast embedding space and interpolated pseudo-contrasts.** a) For one subject at R=4, UMAP of learned contrast embeddings for the 3 originally available contrasts (black circle) along with synthetic contrast embeddings (red cross) generated by interpolating between original embeddings. b) Original DISINR outputs and generated pseudo-contrasts corresponding to interpolated contrast embeddings (highlighted in red).

Fig. 4a presents the UMAP projection [10] of the learned contrast embeddings for one subject, along with three new synthetic contrast embeddings obtained by linearly interpolating between the originals. Fig. 4b demonstrates that the predicted images resulting from these synthetic contrast embeddings appear as new, yet semantically plausible, pseudo-contrast images that exhibit a blend of features from the original embeddings. For example, the T2-FLAIR image has similar grey-white matter contrast to the FLAIR, yet has bright CSF.

The hash-encoder based MC-INR method appears to suffer significantly in our experiments with the M4RAW data, especially without the added TV-penalty term, as demonstrated in Table 2. As the multi-resolution hash encoder is designed to fit high frequency features much faster than SIRENs, in combination with the weighted loss function employed by Niessen et al. we hypothesize that the MC-INR model’s emphasis on high frequency information was counter-productive when applied to the low-SNR M4RAW data.

Additionally, we note that data used in this work contains sets of contrasts with matched dimensions and without significant motion. We hypothesize that when there is structural mis-alignment between the contrasts, the model will struggle to learn a disentangled structure representation and therefore less efficiently leverage the shared structural information during the reconstruction process. Future work will explore extending the DISINR method to handle mismatched and misaligned datasets.

Conclusion. This work introduces a structure contrast disentangled INR framework for MC MRI reconstruction. As our method is patient specific, it requires no training data and can be applied to any arbitrary set of contrasts. By explicitly modeling a shared structural representation between contrasts our method is better able to leverage shared structural information, achieving greater reconstruction performance than existing MC and SC methods and advancing the competitiveness of patient-specific reconstruction.

Acknowledgments. This research was partially supported by the Canadian Institutes of Health Research (PJT-197818) and the Canadian Foundation for Innovation. ZV is supported by the Queen Elizabeth II Graduate Scholarship in Science and Technology. MC is supported by the Canada Research Chairs program.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bohr, I.J., Kenny, E., Blamire, A., O'Brien, J.T., Thomas, A.J., Richardson, J., Kaiser, M.: Resting-State Functional Connectivity in Late-Life Depression: Higher Global Connectivity and More Long Distance Connections. *Frontiers in Psychiatry* **3** (2013). <https://doi.org/10.3389/fpsyt.2012.00116>, <http://journal.frontiersin.org/article/10.3389/fpsyt.2012.00116/abstract>
2. Chung, H., Ye, J.C.: Score-based diffusion models for accelerated MRI. *Medical Image Analysis* **80**, 102479 (Aug 2022). <https://doi.org/10.1016/j.media.2022.102479>, <https://linkinghub.elsevier.com/retrieve/pii/S1361841522001268>
3. Darestani, M.Z., Chaudhari, A.S., Heckel, R.: Measuring Robustness in Deep Learning Based Compressive Sensing. In: Meila, M., Zhang, T. (eds.) *Proceedings of the 38th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 139, pp. 2433–2444. PMLR (Jul 2021), <https://proceedings.mlr.press/v139/darestani21a.html>
4. Feng, J., Feng, R., Wu, Q., Shen, X., Chen, L., Li, X., Feng, L., Chen, J., Zhang, Z., Liu, C., Zhang, Y., Wei, H.: Spatiotemporal Implicit Neural Representation for Unsupervised Dynamic MRI Reconstruction. *IEEE Transactions on Medical Imaging* **44**(5), 2143–2156 (May 2025). <https://doi.org/10.1109/TMI.2025.3526452>, <https://ieeexplore.ieee.org/document/10829649/>
5. Feng, R., Wu, Q., Feng, J., She, H., Liu, C., Zhang, Y., Wei, H.: IM-JENSE: Scan-Specific Implicit Representation for Joint Coil Sensitivity and

- Image Estimation in Parallel MRI. *IEEE Transactions on Medical Imaging* **43**(4), 1539–1553 (Apr 2024). <https://doi.org/10.1109/TMI.2023.3342156>, <https://ieeexplore.ieee.org/document/10356136/>
6. Gatys, L., Ecker, A., Bethge, M.: A Neural Algorithm of Artistic Style. *Journal of Vision* **16**(12), 326 (Sep 2016). <https://doi.org/10.1167/16.12.326>, <http://jov.arvojournals.org/article.aspx?doi=10.1167/16.12.326>
 7. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization (Jan 2017). <https://doi.org/10.48550/arXiv.1412.6980>, <http://arxiv.org/abs/1412.6980>, arXiv:1412.6980 [cs]
 8. Lyu, M., Mei, L., Huang, S., Liu, S., Li, Y., Yang, K., Liu, Y., Dong, Y., Dong, L., Wu, E.X.: M4Raw: A multi-contrast, multi-repetition, multi-channel MRI k-space dataset for low-field MRI research. *Scientific Data* **10**(1), 264 (May 2023). <https://doi.org/10.1038/s41597-023-02181-4>, <https://www.nature.com/articles/s41597-023-02181-4>
 9. Mann, R.M., Kuhl, C.K., Moy, L.: Contrast-enhanced MRI for breast cancer screening. *Journal of Magnetic Resonance Imaging* **50**(2), 377–390 (Aug 2019). <https://doi.org/10.1002/jmri.26654>, <https://onlinelibrary.wiley.com/doi/10.1002/jmri.26654>
 10. McInnes, L., Healy, J., Saul, N., Grokberger, L.: UMAP: Uniform Manifold Approximation and Projection. *Journal of Open Source Software* **3**(29), 861 (Sep 2018). <https://doi.org/10.21105/joss.00861>, <http://joss.theoj.org/papers/10.21105/joss.00861>
 11. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics* **41**(4), 1–15 (Jul 2022). <https://doi.org/10.1145/3528223.3530127>, <https://dl.acm.org/doi/10.1145/3528223.3530127>
 12. Niessen, N., Pirkel, C.M., Solana, A.B., Eichhorn, H., Spieker, V., Huang, W., Sprenger, T., Menzel, M.I., Schnabel, J.A.: INR Meets Multi-contrast MRI Reconstruction. In: Felsner, L., Küstner, T., Maier, A., Qin, C., Ahmadi, S.A., Kazi, A., Hu, X. (eds.) *Reconstruction and Imaging Motion Estimation, and Graphs in Biomedical Image Analysis*. pp. 23–33. Springer Nature Switzerland, Cham (2026)
 13. Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., Baraniuk, R.G.: WIRE: Wavelet Implicit Neural Representations. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 18507–18516 (Jun 2023)
 14. Shen, L., Pauly, J., Xing, L.: NeRP: Implicit Neural Representation Learning With Prior Embedding for Sparsely Sampled Image Reconstruction. *IEEE Transactions on Neural Networks and Learning Systems* **35**(1), 770–782 (Jan 2024). <https://doi.org/10.1109/TNNLS.2022.3177134>, <https://ieeexplore.ieee.org/document/9788018/>
 15. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit Neural Representations with Periodic Activation Functions. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 7462–7473. Curran Associates, Inc. (2020)
 16. Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C.L., Yakubova, N., Knoll, F., Johnson, P.: End-to-End Variational Networks for Accelerated MRI Reconstruction. In: Martel, A.L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga, M.A., Zhou, S.K., Racocanu, D., Joskowicz, L. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. pp. 64–73. Springer International Publishing, Cham (2020)

17. Uecker, M., Lai, P., Murphy, M.J., Virtue, P., Elad, M., Pauly, J.M., Vasanawala, S.S., Lustig, M.: ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA. *Magnetic Resonance in Medicine* **71**(3), 990–1001 (Mar 2014). <https://doi.org/10.1002/mrm.24751>, <https://onlinelibrary.wiley.com/doi/10.1002/mrm.24751>
18. Van Beek, E.J., Kuhl, C., Anzai, Y., Desmond, P., Ehman, R.L., Gong, Q., Gold, G., Gulani, V., Hall-Craggs, M., Leiner, T., Lim, C.T., Pipe, J.G., Reeder, S., Reinhold, C., Smits, M., Sodickson, D.K., Tempany, C., Vargas, H.A., Wang, M.: Value of MRI in medicine: More than just another test? *Journal of Magnetic Resonance Imaging* **49**(7) (Jun 2019). <https://doi.org/10.1002/jmri.26211>, <https://onlinelibrary.wiley.com/doi/10.1002/jmri.26211>
19. Wattjes, M.P., Steenwijk, M.D., Stangel, M.: MRI in the Diagnosis and Monitoring of Multiple Sclerosis: An Update. *Clinical Neuroradiology* **25**(S2), 157–165 (Oct 2015). <https://doi.org/10.1007/s00062-015-0430-y>, <http://link.springer.com/10.1007/s00062-015-0430-y>
20. Xin, B., Ye, M., Axel, L., Metaxas, D.N.: Fill the K-Space and Refine the Image: Prompting for Dynamic and Multi-Contrast MRI Reconstruction. In: Camara, O., Puyol-Antón, E., Sermesant, M., Suinesiaputra, A., Tao, Q., Wang, C., Young, A. (eds.) *Statistical Atlases and Computational Models of the Heart. Regular and CMRxRecon Challenge Papers*. pp. 261–273. Springer Nature Switzerland, Cham (2024)
21. Zhou, B., Dey, N., Schlemper, J., Salehi, S.S.M., Liu, C., Duncan, J.S., Sofka, M.: DSFormer: A Dual-Domain Self-Supervised Transformer for Accelerated Multi-Contrast MRI Reconstruction. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 4966–4975 (Jan 2023)