End-to-end Adaptive *k*-space Sampling, Reconstruction and Registration for Dynamic MRI

George Yiasemis , Jan-Jakob Sonke , Jonas Teuwen

Netherlands Cancer Institute, Amsterdam, Netherlands
g.yiasemis@nki.nl

Abstract. Dynamic MRI supports a range of clinical applications, including cardiac function assessment, organ motion tracking, and radiotherapy guidance. However, fully sampling dynamic k-space is often infeasible due to acquisition time constraints and physiological motion, such as respiration and heartbeat. This necessitates undersampling, which degrades reconstructed image quality and can impair deformation field estimation for registering dynamic images to a static reference. Accurate registration is critical for motion correction, treatment planning, and quantitative analysis, particularly in cardiac imaging and MR-guided radiotherapy. Extending E2E-ADS-Recon, we add deformable registration and jointly optimize its adaptive sampling–reconstruction pipeline with registration. The framework first uses a DL-based adaptive sampling strategy to optimize dynamic k-space acquisition for each case. A reconstruction module then produces images from undersampled moving data that are optimized for deformation field estimation, followed by a registration module that aligns the reconstructed dynamic images with a static reference. The framework is independent of specific reconstruction and registration architectures, enabling plug-and-play integration of different modules. It is jointly trained using supervised and unsupervised losses, allowing end-to-end optimization across all components. Controlled experiments and ablations evaluate warped-image similarity on undrsampled cardiac cine MRI, a suitable dynamic setting, and externally on aortic MRI.

Keywords: Dynamic MRI · Cardiac MRI · Adaptive Sampling · Image Reconstruction and Registration · Deep Learning · End-to-end Learning

1 Introduction

MRI is an essential imaging modality owing to its non-invasive nature, absence of ionizing radiation, and excellent soft tissue contrast. Dynamic MRI further enables real-time cardiac function assessment, organ motion tracking, and adaptive radiotherapy, making it valuable for both diagnosis and treatment guidance. However, dynamic MRI is limited by slow acquisition, which is particularly challenging in real-time settings and is further complicated by physiological motion, such as respiration [6] and heartbeat [14].

Motion and acquisition-time constraints make full k-space sampling difficult in dynamic imaging, making undersampling the practical route to acceleration. Even when dynamic acquisitions are possible, such as ECG-triggered cardiac MRI [19], undersampling can reduce scan time, lower costs, and increase scanner

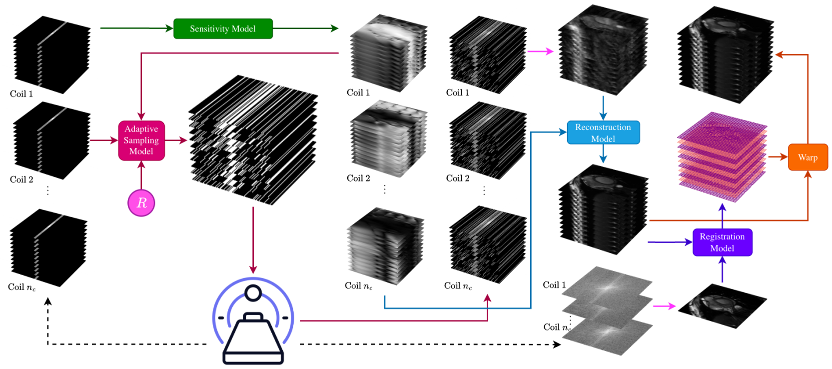

Fig. 1: End-to-end pipeline. Sensitivities are estimated from ACS data; the adaptive sampling model selects patterns from initial moving data and sensitivities; sampled data are reconstructed; the registration model estimates motion aligning the sequence to a reference. The warped output is used for loss computation.

throughput. Deep learning (DL) has substantially improved undersampled MRI reconstruction [8, 21], often outperforming compressed sensing (CS) [18] by learning data-driven reconstructions from limited k-space data.

Beyond reconstruction, DL-based adaptive sampling has shown improved performance over non-adaptive learned, fixed, or random sampling. However, most methods target static imaging; E2E-ADS-Recon [36], an end-to-end adaptive sampling and reconstruction framework, specifically addresses dynamic MRI but does not include registration. Despite these advances, undersampling can degrade image quality and affect downstream motion estimation and registration, which are central to dynamic MRI. In cardiac MRI, registration supports motion correction, segmentation, and functional assessment [15]; in adaptive radiotherapy, MRI-based registration enables tumor motion tracking for accurate dose delivery while limiting exposure to healthy tissue.

Several traditional and DL-based registration methods have been developed for medical imaging [4, 7, 16, 25]. In accelerated MRI, some approaches jointly optimize reconstruction and registration [10, 26, 27], while others use learned motion models [33]. Motivated by the need for accurate motion estimation in undersampled dynamic MRI, we introduce the first end-to-end DL-based framework that integrates adaptive undersampling, reconstruction, and registration for dynamic MRI. An overview is shown in Fig. 1. Our contributions are:

- Extending E2E-ADS-Recon, we add deformable registration and jointly optimize the complete sampling–reconstruction–registration pipeline end-to-end.
- We provide a detailed and reproducible methodology covering all components of the proposed framework; code is available at github.com/NKI-AI/direct.
- We evaluate the framework under multiple configurations, validating on cardiac cine data and demonstrating generalization to out-of-distribution data.

2 Background and Related work

2.1 Dynamic MRI Reconstruction

Dynamic MRI reconstruction aims to recover a moving image $\mathbf{x} \in \mathbb{C}^{n \times n_t}$ from acquired dynamic k-space data, where $n = n_x \times n_y$ denotes the spatial dimensions, n_c represents the number of scanner coils, and n_t the temporal dimension of the sequence. Given fully sampled data $\mathbf{y} \in \mathbb{C}^{n \times n_c \times n_t}$, the underlying image sequence can be obtained by applying the inverse Fast Fourier Transform $\mathcal{F}^{-1}$, followed by the root-sum-of-squares (RSS) as follows for $\tau = 1, \cdots, n_t$,:

$$\mathbf{x}_{\cdot,\tau} := \mathrm{RSS} \circ \mathcal{F}^{-1}(\mathbf{y}_{\cdot,\cdot,\tau}),\ \mathrm{RSS}(\mathbf{w}) = \big(\sum_{k=1}^{n_c} |\mathbf{w}_{\cdot,k}|^2\big)^{1/2} \in \mathbb{R}^n, \quad \mathbf{w} \in \mathbb{C}^{n \times n_c}. \quad (1)$$

To reduce the acquisition time, undersampling is applied to the k-space: $\tilde{\mathbf{y}}^{\mathbf{M}} := \mathbf{M}(\mathbf{y})$. Undersampling is characterized by a binary dynamic mask operator $\mathbf{M} \in \{0,1\}^{n \times n_t}$, which zeros-out non-acquired k-space samples. The resulting forward model for each frame τ is described by the forward operator $\mathcal{T}_{\mathbf{M}_\tau, \mathbf{S}_\tau}$:

$$\tilde{\mathbf{y}}^{\mathbf{M}}_{\cdot,\cdot,\tau} \ = \ \mathcal{T}_{\mathbf{M}_\tau, \mathbf{S}_\tau}(\mathbf{x}_\tau) \ := \ \mathbf{M}_\tau \circ \mathcal{F} \circ \mathcal{C}_{\mathbf{S}_\tau}(\mathbf{x}_{\cdot,\tau}), \quad (2)$$

where $\mathcal{C}_{\mathbf{S}_\tau}$ denotes the coil sensitivity-encoding operator, which decomposes an image into individual coil images using coil sensitivity profiles $\mathbf{S}_\tau \in \mathbb{C}^{n^2 \times n_c}$. A reconstruction for (2) can be formulated as a solution to an optimization problem.

2.2 Deep Learning-based MRI Reconstruction

DL approaches to accelerated MRI include direct mapping and unrolled optimization [1, 2, 13, 24, 35, 38]. Open-access k-space datasets [31, 32] have accelerated research, predominantly on static [11,24,35] and recently, dynamic MRI [30,37,40].

2.3 Adaptive MRI Acquisition

DL methods have sought to replace standard equidistant, random uniform, or Gaussian sampling [34] with learned strategies, including training-dataset-based sampling optimization [23, 42] and adaptive, case-specific patterns jointly learned with reconstruction networks [3, 12]. Most work focuses on static MRI; [36] extended adaptive sampling to dynamic (2D + t) imaging with E2E-ADS-Recon, an end-to-end adaptive dynamic undersampling and reconstruction framework.

Unlike earlier single-factor methods, E2E-ADS-Recon supports multiple acceleration factors (ARs) using phase-specific or unified sampling patterns.

2.4 Deformable Image Registration

Deformable, or non-rigid, registration aligns a moving image $\mathbf{z}_{\mathrm{mov}} \in \mathbb{R}^{n \times n_t}$ to a fixed reference image $\mathbf{z}_{\mathrm{ref}} \in \mathbb{R}^n$ by estimating spatial deformation fields

$\boldsymbol{\phi} = (\boldsymbol{\phi}_1, \cdots, \boldsymbol{\phi}_{n_t}) \in \mathbb{R}^{2\times n\times n_t}$. Each displacement field $\boldsymbol{\phi}_\tau$ maps coordinates between the reference image and temporal phase $(\mathbf{z}_{\text{mov}})_{\cdot,\tau}$ through a warping operation $\mathcal{W} : \mathbb{R}^n \times \mathbb{R}^{2\times n} \to \mathbb{R}^n$ [4], yielding

$$\mathbf{z}_{\text{reg}} = \{\mathcal{W}\left((\mathbf{z}_{\text{mov}})_{\cdot,\tau}, \boldsymbol{\phi}_\tau\right)\}_{\tau=1}^{n_t} \in \mathbb{R}^{n\times n_t}. \tag{3}$$

The goal is to estimate an optimal deformation field $\boldsymbol{\phi}^*$ that aligns the moving image sequence with the reference, commonly formulated as an unsupervised minimization problem using an image similarity loss $\mathcal{L}_{\text{sim}}$:

$$\boldsymbol{\phi}^* := \arg\min_{\boldsymbol{\phi}} \mathcal{L}_{\text{sim}}\left(\mathbf{z}_{\text{reg}}, \{\mathbf{z}_{\text{ref}}\}_{\tau=1}^{n_t}\right). \tag{4}$$

3 Methods

3.1 Deep Learning Framework

Sensitivity Estimation Coil sensitivity maps $\mathbf{S}$ are estimated from the fully sampled central k-space region $\tilde{\mathbf{y}}_{\text{mov}}^{\mathbf{M}^{\text{acs}}}$ and refined with a 2D U-Net $\mathcal{S}_{\boldsymbol{\sigma}}$ [22,24].

Adaptive Sampling Model The adaptive dynamic sampling (ADS) module from E2E-ADS-Recon [36], predicts a dynamic binary sampling mask from initially undersampled moving data $\tilde{\mathbf{y}}_{\text{mov}}^{\mathbf{M}^0}$, sensitivities $\mathbf{S}$, and a target AR R:

$$\mathbf{M} := \mathcal{A}_{\boldsymbol{\omega}}\left(\tilde{\mathbf{y}}_{\text{mov}}^{\mathbf{M}^0}; \mathbf{S}, R\right), \qquad \mathbf{M} = (\mathbf{M}_1, \ldots, \mathbf{M}_{n_t}). \tag{5}$$

The ADS module uses cascaded 3D U-Net-like encoders and MLPs to generate sampling probabilities, which are binarized using a straight-through estimator [5,39] to satisfy the sampling budget. Unless otherwise specified, the ACS data are used as the initial acquisition, $\tilde{\mathbf{y}}_{\text{mov}}^{\mathbf{M}^0} = \tilde{\mathbf{y}}_{\text{mov}}^{\mathbf{M}^{\text{acs}}}$.

Reconstruction Model The framework is agnostic to the reconstruction network. In the main experiments, we use vSHARP [38], due to its performance in the CMRxRecon 2023 [31]. vSHARP unrolls $T = 10$ half-quadratic variable-splitting ADMM iterations with learned denoising (four-scale 3D U-Net; 16/32/64/128 filters), data consistency ($T_x = 6$ gradient steps), and Lagrange-multiplier updates; other settings follow [36]. Given the undersampled dynamic image $\tilde{\mathbf{x}}_{\text{mov}}^{\mathbf{M}}$ and sensitivity maps $\mathbf{S}$, vSHARP produces

$$\hat{\mathbf{x}}_{\text{mov}} = \mathcal{V}_{\boldsymbol{\theta}}\left(\tilde{\mathbf{x}}_{\text{mov}}^{\mathbf{M}}; \mathbf{S}\right), \qquad \left(\tilde{\mathbf{x}}_{\text{mov}}^{\mathbf{M}}\right)_{\cdot,\tau} = \mathcal{R}_{\mathbf{S}_\tau}\left[\mathcal{F}^{-1}\left(\tilde{\mathbf{y}}_{\text{mov}}^{\mathbf{M}}\right)_{\cdot,\cdot,\tau}\right], \tag{6}$$

Here, $\mathcal{R}_{\mathbf{S}_\tau}$ denotes coil combination: $\mathcal{R}_{\mathbf{S}_\tau}(\mathbf{v}) := \sum_{k=1}^{n_c}\left(\mathbf{S}_\tau^k\right)^* \mathbf{v}_{\cdot,k}$, $\mathbf{v} \in \mathbb{C}^{n\times n_c}$.

Registration–Motion Estimation Model The framework can incorporate any deformable registration method. In the main experiments, we use a lightweight two-dimensional U-Net $\mathcal{R}_{\boldsymbol{\psi}}$ [22], which takes the reconstructed moving sequence and reference image $\mathbf{x}_{\text{ref}} \in \mathbb{R}^n$ as channel-wise input and predicts deformation fields for all temporal phases:

$$\boldsymbol{\phi} = \mathcal{R}_{\boldsymbol{\psi}}\left([|\hat{\mathbf{x}}_{\text{mov}}|, \mathbf{x}_{\text{ref}}]\right) \in \mathbb{R}^{2\times n\times n_t}. \tag{7}$$

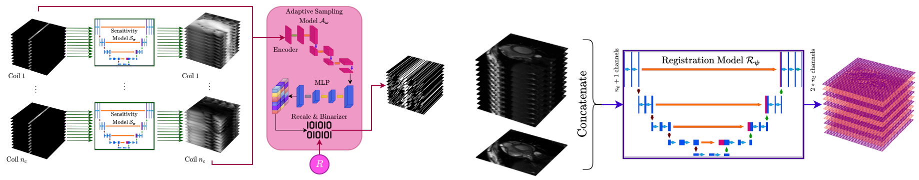

(a) Adaptive dynamic sampling module. (b) Registration module.

Fig. 2: Overview of the adaptive sampling and registration components.

3.2 End-to-End Framework

The proposed framework combines sensitivity estimation, adaptive sampling, reconstruction, and registration in a sequential pipeline. Sensitivity maps $\mathbf{S}$ are estimated using $\mathcal{S}_{\boldsymbol{\sigma}}$ and, together with the initial undersampled data $\tilde{\mathbf{y}}_{\text{mov}}^{\mathbf{M}^0}$, are used by $\mathcal{A}_{\boldsymbol{\omega}}$ to generate dynamic sampling masks $\mathbf{M}$ for a target acceleration factor R. The sampled data $\tilde{\mathbf{y}}_{\text{mov}}^{\mathbf{M}}$ are reconstructed by $\mathcal{V}_{\boldsymbol{\theta}}$, after which $\mathcal{R}_{\boldsymbol{\psi}}$ estimates deformation fields that align the reconstructed sequence with a reference frame, as illustrated in Fig. 1; module architectures are shown in Fig. 2.

3.3 Training Loss Function

The framework is trained primarily to improve registration, while reconstruction provides high-fidelity inputs for motion estimation. Let $\tilde{\mathbf{y}}_{\text{mov}}^{\mathbf{M}} \in \mathbb{C}^{n \times n_c \times n_t}$ denote the undersampled moving data obtained using the mask $\mathbf{M}$ predicted by $\mathcal{A}_{\boldsymbol{\omega}}$, and let $\mathbf{y}_{\text{mov}} \in \mathbb{C}^{n \times n_c \times n_t}$ and $\mathbf{y}_{\text{ref}} \in \mathbb{C}^{n \times n_c}$ denote the fully sampled moving and reference k-space measurements. The total loss is

$$\mathcal{L} := \alpha \mathcal{L}_{\text{rec}} + \beta \mathcal{L}_{\text{reg}}, \qquad \alpha, \beta > 0. \tag{8}$$

The reconstruction loss is

$$\mathcal{L}_{\text{rec}} = \mathcal{L}_{\text{ssim}} \left(|\hat{\mathbf{x}}_{\text{mov}}| , \mathbf{x}_{\text{mov}} \right) + \mathcal{L}_1 \left(|\hat{\mathbf{x}}_{\text{mov}}| , \mathbf{x}_{\text{mov}} \right), \tag{9}$$

where $\hat{\mathbf{x}}_{\text{mov}} = \mathcal{V}_{\boldsymbol{\theta}}(\tilde{\mathbf{y}}_{\text{mov}}^{\mathbf{M}}; \mathbf{S})$, $\mathbf{x}_{\text{mov}} = \text{RSS} \circ \mathcal{F}(\mathbf{y}_{\text{mov}})$, and $\mathcal{L}_{\text{ssim}} := \mathcal{L}_{\text{ssim2D}} + \mathcal{L}_{\text{ssim3D}}$ combines 2D and 3D SSIM losses [38].
The registration loss combines image similarity and deformation smoothness:

$$\mathcal{L}_{\text{reg}} = \mathcal{L}_{\text{sim}} \left(\mathbf{x}_{\text{reg}}, \{\mathbf{x}_{\text{ref}}\}_{\tau=1}^{n_t} \right) + \mathcal{L}_{\text{smooth}} \left(\boldsymbol{\phi} \right), \quad \text{with} \tag{10}$$

$$\mathbf{x}_{\text{reg}} = \left\{ \mathcal{W} \left(|\hat{\mathbf{x}}_{\text{mov}}|_{\cdot,\tau} , \boldsymbol{\phi}_\tau \right) \right\}_{\tau=1}^{n_t} \in \mathbb{R}^{n \times n_t}, \qquad \mathbf{x}_{\text{ref}} = \text{RSS} \circ \mathcal{F} \left(\mathbf{y}_{\text{ref}} \right) \in \mathbb{R}^n. \tag{11}$$

The reference image is repeated across temporal frames as $\{\mathbf{x}_{\text{ref}}\}_{\tau=1}^{n_t}$, $\mathcal{L}_{\text{sim}}$ uses the same losses as (9), and $\boldsymbol{\phi} = \mathcal{R}_{\boldsymbol{\psi}}([|\hat{\mathbf{x}}_{\text{mov}}|, \mathbf{x}_{\text{ref}}])$. The smoothness term $\mathcal{L}_{\text{smooth}}(\boldsymbol{\phi})$ penalizes spatial gradients of $\boldsymbol{\phi}$ [4,9].

4 Experiments

4.1 Datasets and Setup

We used the CMRxRecon 2023 cardiac cine dataset [31] (471 scans, 3,185 slices, 12 frames, $n_c = 10$), split into 251/100/100 scans for training, validation, and testing, and the CMRxRecon 2024 aorta dataset [32] for external validation (111 scans, 883 slices). All experiments used retrospective 1D line undersampling with a 4% ACS center fraction and $R \in \{4, 6, 8\}$ during training and inference. Each sequence used the 6^{th} cardiac phase as the registration reference (end-systole [20]); the remaining $n_t = 11$ phases formed the undersampled moving data.

4.2 Evaluation

We assess motion estimation indirectly using SSIM, PSNR, and NMSE between the reference and registered moving images (reconstructions warped by predicted fields), averaged across temporal phases. Because the moving sequences were acquired rather than simulated, no ground-truth deformation fields exist.

4.3 Ablation and Component Analysis

Our proposed pipeline is the first to integrate adaptive sampling, reconstruction, and registration into an end-to-end framework for dynamic MRI, making direct comparisons with existing baselines impossible. We therefore evaluate each component through controlled comparisons and ablations.

Baseline Setup We use phase-specific ADS [36], where the sampler predicts a distinct undersampling pattern for each temporal phase rather than a single shared pattern, initialized from ACS data. The pipeline is trained end-to-end with α=β=1 using Adam without weight decay, batch size 1, a 2k-iteration linear warm-up to learning rate 3×10^{-3}, and step decay every 10k iterations (factor 0.8) for 52k iterations; we select the checkpoint with the best validation SSIM.

Ablation Studies We modify one factor at a time while keeping all other settings fixed: (i) registration module vs. VoxelMorph [4], TransMorph [7], ILK [17], TV-L_1 [41], and DEMONS [28]; (ii) vSHARP vs. 3D VarNet [24]; (iii) learned vs. fixed non-adaptive sampling schemes [29, 34, 42]; and (iv) loss weighting and joint vs. decoupled training.

4.4 Experimental Results

We report quantitative results on in-distribution cine and out-of-distribution aorta data. Representative qualitative results are shown in Fig. 7.

Registration Component Fig. 3 compares registration modules across both datasets. The proposed lightweight registration model consistently outperforms traditional methods (OptFlow ILK, TV-L_1, DEMONS) and DL-based alternatives (VoxelMorph, TransMorph) across all acceleration factors and metrics.

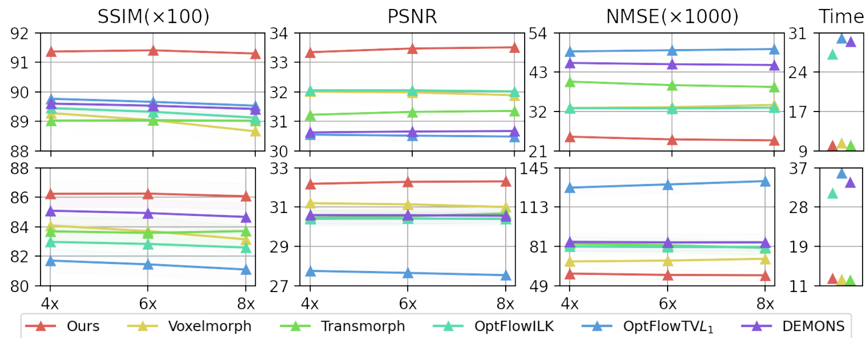

Fig. 3: Comparison of registration performance across different methods. Top: Cine. Bottom: Aorta (not seen during training).

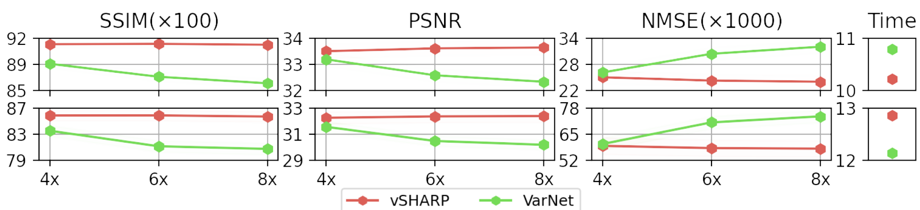

Fig. 4: Impact of reconstruction model on registration results. Top: Cine. Bottom: Aorta (not seen during training).

Reconstruction Component Fig. 4 shows that replacing vSHARP with VarNet reduces registration performance, suggesting that vSHARP better supports downstream motion estimation.

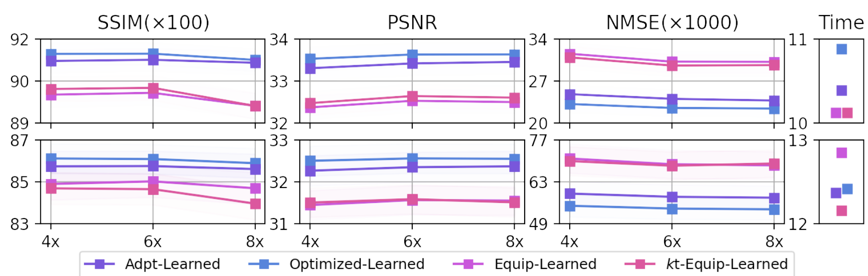

Fig. 5: Impact of learned vs fixed non-adaptive sampling schemes on registration results. Top: Cine. Bottom: Aorta (not seen during training).

Adaptive Sampling Component Fig. 5 shows that learned sampling strategies outperform fixed non-adaptive schemes across metrics on both datasets.

Loss Weighting and Training Strategy Fig. 6 shows that prioritizing registration ($\alpha < \beta$) improves registration but degrades reconstruction, partly by approximating the reference; prioritizing reconstruction ($\alpha > \beta$) has the opposite effect. Equal weighting gives the best balance. Joint optimization substantially improves registration over decoupled training with only a minor reconstruction loss.

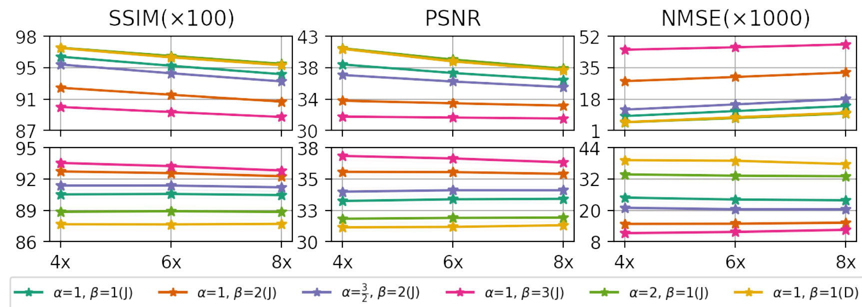

Fig. 6: Impact of joint (J) vs decoupled (D) training, and varying loss weights (cine). Top: Reconstruction performance. Bottom: Registration performance.

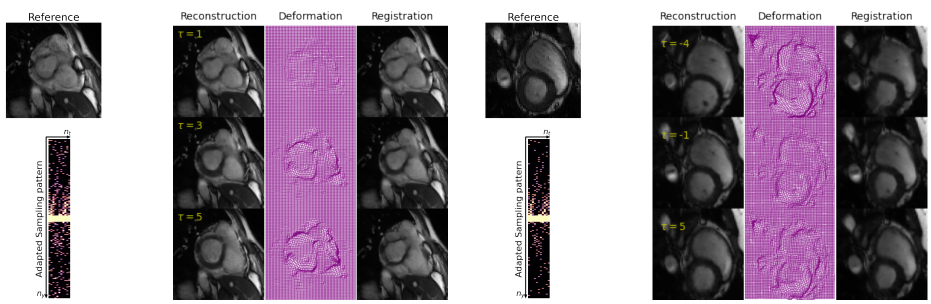

Fig. 7: Example results for two cases, shown at different temporal frames τ relative to the reference image. Left: Case I, $R = 6$. Right: Case II, $R = 8$.

5 Discussion and Conclusion

We presented a modular end-to-end deep learning framework that jointly optimizes adaptive undersampling, reconstruction, and registration to improve motion estimation from undersampled dynamic MRI while allowing different reconstruction and registration networks to be used. Across in-distribution cardiac cine and out-of-distribution aortic data, the framework achieves strong registration performance, shows benefits from learned sampling over fixed non-adaptive schemes, and improves motion estimation through joint rather than decoupled training. Together, these findings show that unified optimization of acquisition, reconstruction, and registration can improve deformation-field estimation in accelerated dynamic MRI while providing a flexible basis for self-supervised adaptive sampling, three-dimensional motion estimation, and direct motion estimation from undersampled measurements. Limitations include slice-wise two-dimensional motion estimation, evaluation metrics computed on registered reconstructions rather than ground-truth deformation fields, and reliance on fully sampled k-space during training. Since the moving sequences were acquired rather than simulated, no ground-truth deformation fields exist. Accordingly, warped-image similarity may reflect both reconstruction and deformation quality. Conclusions regarding deformation-field accuracy should therefore be interpreted cautiously.

References

1. Adler, J., Öktem, O.: Learned primal-dual reconstruction. IEEE transactions on medical imaging **37**(6), 1322–1332 (2018)
2. Aggarwal, H.K., et al.: Modl: Model-based deep learning architecture for inverse problems. IEEE transactions on medical imaging **38**(2), 394–405 (2018)
3. Bakker, T., et al.: On learning adaptive acquisition policies for undersampled multi-coil MRI reconstruction. In: Medical Imaging with Deep Learning (2022)
4. Balakrishnan, G., et al.: Voxelmorph: a learning framework for deformable medical image registration. IEEE transactions on medical imaging **38**(8), 1788–1800 (2019)
5. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv preprint (2013)
6. Bertelsen, A., Bernchou, U. et al.: The effect of respiration-induced target motion on 3d magnetic resonance images used to guide radiotherapy. Physics and Imaging in Radiation Oncology **24**, 167–172 (2022)
7. Chen, J., Frey, E.C. et al.: Transmorph: Transformer for unsupervised medical image registration. Medical image analysis **82**, 102615 (2022)
8. Chen, Y., Schönlieb, C.B. et al.: Ai-based reconstruction for fast mri—a systematic review and meta-analysis. Proceedings of the IEEE **110**(2), 224–245 (2022)
9. Commowick, O., et al.: An efficient locally affine framework for the smooth registration of anatomical structures. Medical Image Analysis **12**(4) (2008)
10. Corona, V., et al.: Variational multi-task mri reconstruction: Joint reconstruction, registration and super-resolution. Medical Image Analysis **68**, 101941 (2021)
11. Eo, T., Jun, Y. et al.: Kiki-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images. Magnetic resonance in medicine **80**(5), 2188–2201 (2018)
12. Gautam, S., Li, A., Ravishankar, S.: Patient-adaptive and learned mri data undersampling using neighborhood clustering (2024)
13. Hammernik, K., Klatzer, T. et al.: Learning a variational network for reconstruction of accelerated mri data. Magnetic resonance in medicine **79**(6), 3055–3071 (2018)
14. Ismail, T.F., Strugnell, W. et al.: Cardiac mr: from theory to practice. Frontiers in cardiovascular medicine **9**, 826283 (2022)
15. Khalil, A., et al.: An overview on image registration techniques for cardiac diagnosis and treatment. Cardiology Research and Practice **2018**, 1–15 (2018)
16. Kostelec, P.J., Periaswamy, S.: Image registration for mri. Modern signal processing **46**, 161–184 (2003)
17. Le Besnerais, G., Champagnat, F.: Dense optical flow by iterative local window registration. In: IEEE International Conference on Image Processing 2005. vol. 1, pp. I–137 (2005)
18. Lustig, M., Donoho, D.L. et al.: Compressed sensing mri. IEEE signal processing magazine **25**(2), 72–82 (2008)
19. Lyu, J., Qin, C. et al.: The state-of-the-art in cardiac mri reconstruction: Results of the cmrxrecon challenge in miccai 2023 (2024)
20. Mada, R.O., et al.: How to define end-diastole and end-systole?: Impact of timing on strain measurements. JACC: Cardiovascular Imaging **8**(2), 148–157 (2015)
21. Pal, A., Rathi, Y.: A review and experimental evaluation of deep learning methods for mri reconstruction. The journal of machine learning for biomedical imaging **1** (2022)

22. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI 2015. pp. 234–241. Springer (2015)
23. Shor, T.: Multi PILOT: Feasible learned multiple acquisition trajectories for dynamic MRI. In: Medical Imaging with Deep Learning (2023)
24. Sriram, A., Zbontar, J. et al.: End-to-end variational networks for accelerated mri reconstruction. In: MICCAI 2020. pp. 64–73. Springer (2020)
25. Sun, S., Han, K. et al.: Medical image registration via neural fields. Medical Image Analysis **97**, 103249 (2024)
26. Terpstra, M.L., Maspero, M. et al.: Real-time 3d motion estimation from undersampled mri using multi-resolution neural networks. Medical physics **48**(11), 6597–6613 (2021)
27. Terpstra, M.L., Maspero, M. et al.: ⊥-loss: A symmetric loss function for magnetic resonance imaging reconstruction and image registration with deep learning. Medical Image Analysis **80**, 102509 (2022)
28. Thirion, J.P.: Image matching as a diffusion process: an analogy with maxwell's demons. Medical image analysis **2**(3), 243–260 (1998)
29. Tsao, J., Boesiger, P., Pruessmann, K.P.: k-t blast and k-t sense: dynamic mri with high frame rate exploiting spatiotemporal correlations. Magnetic Resonance in Medicine **50**(5), 1031–1042 (2003)
30. Waddington, D.E., Hindley, N. et al.: On real-time image reconstruction with neural networks for mri-guided radiotherapy. arXiv preprint (2022)
31. Wang, C., et al.: Cmrxrecon: An open cardiac mri dataset for the competition of accelerated image reconstruction (2023)
32. Wang, Z., Wang, F. et al.: Cmrxrecon2024: A multi-modality, multi-view k-space dataset boosting universal machine learning for accelerated cardiac mri (2024)
33. Yang, F., Xue, Z. et al.: Robust fast inter-bin image registration for undersampled coronary mri based on a learned motion prior. IEEE Transactions on Biomedical Engineering (2024)
34. Yiasemis, G., Sánchez, C.I. et al.: On retrospective k-space subsampling schemes for deep mri reconstruction. Magnetic Resonance Imaging **107**, 33–46 (2024)
35. Yiasemis, G., Sonke, J.J. et al.: Recurrent variational network: a deep learning inverse problem solver applied to the task of accelerated mri reconstruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 732–741 (2022)
36. Yiasemis, G., Sonke, J.J., Teuwen, J.: End-to-end co-optimization of adaptive k-space sampling and reconstruction for dynamic mri (2026)
37. Yiasemis, G., et al.: Deep multi-contrast cardiac mri reconstruction via vsharp with auxiliary refinement network. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 183–192. Springer (2024)
38. Yiasemis, G., et al.: vsharp: variable splitting half-quadratic admm algorithm for reconstruction of inverse-problems. Magnetic Resonance Imaging p. 110266 (2024)
39. Yin, T., Wu, Z. et al.: End-to-end sequential sampling and reconstruction for mri. arXiv preprint arXiv:2105.06460 (2021)
40. Yoo, J., Jin, K.H. et al.: Time-dependent deep image prior for dynamic mri. IEEE Transactions on Medical Imaging **40**(12), 3337–3348 (2021)
41. Zach, C., Pock, T., Bischof, H.: A duality based approach for realtime tv-l1 optical flow. In: Hamprecht, F.A., Schnörr, C., Jähne, B. (eds.) Pattern Recognition. pp. 214–223. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)
42. Zhang, J., Zhang, H. et al.: Extending loupe for k-space under-sampling pattern optimization in multi-coil mri (2020)