Piecewise Dynamic Diffusion Regularization for Reconstruction of Cardiac Cine MRI

Florian Fürnrohr[1,2] and Reinhard Heckel[1,2]

[1] Technical University of Munich, Munich, Germany
[2] Munich Center for Machine Learning (MCML), Munich, Germany
reinhard.heckel@tum.de

Abstract. Cardiac cine MRI enables visualization of the beating heart during free breathing, but severe undersampling and motion make reconstruction highly challenging. A central challenge for reconstruction is incorporating powerful priors of cardiac anatomy while remaining computationally efficient. We propose *Piecewise Dynamic Diffusion Regularization (PDDR)*, a reconstruction method that integrates a spatiotemporal diffusion model as a generative prior within a variational reconstruction framework for cine MRI. The model employs dedicated spatial layers to encode anatomical structure and temporal layers to capture cardiac motion learned from gated cine data. PDDR leverages the dynamic prior in a piecewise manner, enabling the efficient use of spatiotemporal diffusion models for processing of long real-time sequences. Experiments on retrospectively accelerated and prospective real-time cine MRI demonstrate that PDDR outperforms classical, unsupervised, and diffusion-based methods, delivering high-quality reconstructions with substantially reduced computation time compared to state-of-the-art baselines. These results highlight PDDR as a practical and scalable solution for free-breathing, real-time cardiac MRI. Code is available at https://github.com/MLI-lab/pddr.

Keywords: Dynamic Imaging · Cardiac MRI · Diffusion Model Prior

1 Introduction

Cardiac cine MRI is an indispensable, non-invasive, clinical imaging technology for evaluation of cardiac function through a video of the beating heart.

Image reconstruction in cardiac MRI is particularly challenging due to rapid heart motion and limited data acquisition speed in magnetic resonance imaging (MRI). Conventional techniques therefore are based on combining information from multiple cardiac cycles, assuming consistent cycles. This requires ECG-gated and breathhold acquisitions for temporal binning of the data [19].

Real-time cine MRI uses continuously acquired measurements. It enables individual imaging of true physiology under free-breathing conditions. As a result, the acquisition is more efficient, comfortable, and robust [5]. In real-time cine, acquisition of fully sampled reference data of sufficient spatiotemporal resolution is practically impossible, making video reconstruction exceptionally challenging.

In real-time acquisitions, only very few measurements are available for each video frame. Reliable reconstruction therefore depends on exploiting the spatiotemporal cardiac structure [16,18]. Data-driven approaches can learn problem-specific structure from representative datasets and use this knowledge for improved reconstruction [11,8]. Consequently, strong spatiotemporal priors are required for accurate cardiac reconstruction.

Diffusion models provide excellent priors, capturing complex image distributions [9]. Diffusion-based reconstruction methods have shown strong performance in static MRI [10]. A spatiotemporal diffusion model could learn the distribution of cardiac motion and provide an effective generative prior for sparsely acquired dynamic measurements.

To date, high-dimensionality and computationally demanding sampling techniques have hindered widespread application of spatiotemporal diffusion priors for video reconstruction [6,27]. In cardiac MRI, diffusion models have primarily been used for reconstructing short gated acquisitions. Long real-time cine sequences are beyond the practical scope of existing diffusion-based methods [27].

In this work, we propose *Piecewise Dynamic Diffusion Regularization* for reconstruction of cardiac cine MRI videos. The customized dynamic diffusion prior uses spatial layers to introduce knowledge about cardiovascular anatomy, while temporal layers model the dynamics of cardiac motion and enable exchange of information across the temporal dimension. Using a variational approach, the diffusion model can be applied as piecewise regularizer in real-time cardiac MRI. The framework constitutes a flexible, robust, and computationally efficient reconstruction method that scales from gated cine to long free-breathing acquisitions with hundreds of video frames.

2 Background

Here we provide background on cardiac MRI reconstruction and diffusion models.

2.1 Cardiac MRI reconstruction problem

We consider the reconstruction of a complex video consisting of N frames $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_N] \in \mathbb{C}^{N \times H \times W}$ from undersampled multi-coil MRI measurements $\mathbf{y} = [\mathbf{y}_1, \dots, \mathbf{y}_N] \in \mathbb{C}^{N \times C \times L}$ of the beating heart. The linear forward model for time frame τ and receiver coil c is given as

$$\mathbf{y}_{\tau,c} = \mathbf{M}_\tau \mathbf{F} \mathbf{S}_c \mathbf{x}_\tau + \mathbf{n}_{\tau,c},$$

where $\mathbf{S}_c$ are coil sensitivity maps, $\mathbf{F}$ is the two-dimensional discrete Fourier transform, $\mathbf{n}_{\tau,c}$ is additive noise, and $\mathbf{M}_\tau$ is a masking operator encoding the sampling pattern. We stack in the coil dimension C and define $\mathbf{A}_\tau = \mathbf{M}_\tau \mathbf{F} \mathbf{S}$, which lets us write the measurement model as $\mathbf{y}_\tau = \mathbf{A}_\tau \mathbf{x}_\tau + \mathbf{n}_\tau$.

Cardiac MRI reconstruction is inherently challenging, as slow data acquisition results in only very few measurements corresponding to a single frame $\mathbf{y}_\tau$, and cardiac and respiratory motion induce changes across frames.

To alleviate this problem, it is common to perform acquisitions while an electrocardiogram (ECG) is recorded and patients hold their breath, allowing the data to be binned. However, that can lead to binning artifacts and elimination of dynamic variability in the reconstructions [19].

In this work, we consider real-time MRI reconstruction without using binning or gating. Real-time reconstruction enables visualization of true cardiac physiology in free-breathing acquisition [5].

Unsupervised machine learning methods are widely used for real-time cine reconstruction [29,24,13,7]. While these methods flexibly adapt to various acquisition parameters and physiological dynamics, they require long reconstruction times and do not incorporate learned prior information about cardiac images.

Standard supervised techniques are challenging due to the lack of ground truth training data. Existing methods are trained using binned reference acquisitions and a fixed synthetic measurement model, making generalization to real-time acquisitions difficult [20,25].

2.2 Diffusion models

Denoising diffusion models [9,15] define a forward diffusion process, transforming a clean data sample $\mathbf{x}_0 \sim p_{data}$ to standard Gaussian noise $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$. Each intermediate sample in the diffusion process $t \in [0, T]$ can be described as $\mathbf{x}_t = \sqrt{1-\sigma_t^2}\mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}$, with $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$ and a fixed variance schedule σ_t. The reverse process, parameterized by a neural network $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)$ intended to predict the noise $\boldsymbol{\epsilon}$, learns to reverse the diffusion process. Training of the network is done by minimizing the objective $\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{x}_0 \sim p_{data}, t \sim \mathcal{U}(0,T), \boldsymbol{\epsilon} \sim \mathcal{N}(0,\mathbf{I})} \left[\|\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) - \boldsymbol{\epsilon}\|_2^2 \right]$. Image generation starts from a random Gaussian vector and applies iterative denoising to sample from the data distribution p_{data}.

Diffusion models have been applied successfully in reconstruction of accelerated static MRI [10,4,17], where they serve as robust generative image priors. Solving inverse problems with diffusion models have predominantly relied on posterior sampling techniques [21,3]. In contrast, variational inference offers a fast and flexible alternative by treating sampling with stochastic optimization. Adopting the diffusion prior as regularizer and enforcing consistency with the measurement data, the regularization by denoising diffusion methodology [14] formulates image reconstruction as a variational optimization problem. Ozturkler et al. [17] have applied this to reconstruction of static MRI and report enhanced robustness to distribution shifts and faster sampling rates compared to traditional Langevin-based diffusion sampling.

3 Piecewise Dynamic Diffusion Regularization

We propose a diffusion model-based reconstruction method for cardiac MRI that uses an efficient spatiotemporal diffusion prior as a piecewise regularizer within a variational reconstruction framework. The method, called Piecewise Dynamic Diffusion Regularization (PDDR), allows flexible adaptation to various acquisition lengths, hardware restrictions, and runtime requirements.

3.1 Piecewise regularization

In variational reconstruction of a dynamic video $\mathbf{x} \in \mathbb{C}^{N\times H\times W}$, we optimize for consistency with the measurements and consistency with a prior, in our case in form of a diffusion model $\boldsymbol{\epsilon}_\theta$. In real-time cine MRI the acquisition duration is not restricted, and the total number of video frames N can be large. Therefore, computational restrictions inhibit the application of the diffusion model to the full video sequence.

We propose providing a stochastic estimate of the full regularization signal, by applying the dynamic diffusion model $\boldsymbol{\epsilon}_\theta$ to a random Q-sized block of consecutive frames $\mathbf{x_q} = [\mathbf{x}_q, \dots, \mathbf{x}_{q+Q-1}] \in \mathbb{C}^{Q\times H\times W}$ in each optimization step of the reconstruction. The variational objective of our method therefore is

$$\min_{\mathbf{x}} \sum_{\tau=1}^{N} \|\mathbf{A}_\tau \mathbf{x}_\tau - \mathbf{y}_\tau\|_2^2 + \mathbb{E}_{q,t,\boldsymbol{\epsilon}} \left[w_t \left\| \boldsymbol{\epsilon}_\theta \left(\mathbf{x_{q,t}}, t \right) - \boldsymbol{\epsilon} \right\|_2^2 \right], \tag{1}$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \boldsymbol{I})$ is Gaussian noise, $t \sim \mathcal{U}(\{0, \dots, T\})$ denotes the diffusion step, and $q \sim \mathcal{U}(\{1, \dots, N-Q\})$ defines the position of the regularization block $\mathbf{x_q}$ in the video. The diffusion model $\boldsymbol{\epsilon}_\theta$ is applied to the noisy input block $\mathbf{x_{q,t}} = \sqrt{1-\sigma_t^2}\mathbf{x_q} + \sigma_t \boldsymbol{\epsilon}$, and the estimated noise residual is weighted with a time-dependent weight $w_t = \lambda \frac{\sigma_t}{\sqrt{1-\sigma_t^2}}$, where λ is a hyperparameter that balances data consistency and regularization strength.

The objective (1) is minimized by gradient-based optimization with total number of K steps. In iteration k of the reconstruction we choose the diffusion step according to $t = \lfloor T' \cdot \frac{K-k}{K} \rfloor$, with $0 < T' < T$ and T being the number of diffusion steps the model was trained on [12,14]. The size Q of the regularization block is a hyperparameter that balances the memory consumption, the runtime, and regularizing effect during inference.

3.2 Dynamic diffusion prior

As mentioned, for real-time cine MRI reconstruction temporal and spatial relations of the frames have to be leveraged since each frame has only very few measurements associated with it. Our reconstruction approach is therefore based on an architecture that efficiently learns spatiotemporal relations between frames.

We propose a separable spatiotemporal U-Net with skippable temporal integration [1]. As main building block, we introduce a separable spatiotemporal residual block, consisting of a 2D spatial layer and a 1D temporal layer, making it computationally more efficient compared to a naive 3D residual block while still capturing temporal correlations. Spatial and temporal layers are convolutional with residual connection, producing outputs $\hat{\boldsymbol{x}}_s$ and $\hat{\boldsymbol{x}}_t$, respectively. The output of the spatiotemporal residual block is then computed as weighted sum $\hat{\boldsymbol{x}} = \sigma(\alpha)\hat{\boldsymbol{x}}_s + (1-\sigma(\alpha))\hat{\boldsymbol{x}}_t$, based on the learnable weight parameter $\alpha \in \mathbb{R}$ and the sigmoid function $\sigma : \mathbb{R} \to [0, 1]$. Optionally, the temporal layer can be bypassed, which enables pretraining of the spatial modules only $\hat{\boldsymbol{x}} = \hat{\boldsymbol{x}}_s$.

The separable spatiotemporal block is then used to build a U-Net architecture with four levels of spatial downsampling, each using four spatiotemporal residual blocks, and one attention block at the bottleneck.

We compared this architecture to using naive 3D residual blocks, and found that our architecture requires significantly lower memory. This enables using bigger block sizes Q for a given GPU, which in turn leads to better performance.

4 Experiments and results

We demonstrate the reconstruction performance of the proposed dynamic diffusion prior on retrospectively accelerated cardiac MRI, where we achieve significantly better results compared to unsupervised and spatial baselines. Furthermore, we analyze the efficiency of piecewise regularization on long sequences, and show high-quality reconstructions of prospective real-time cine MRI.

4.1 Retrospective reconstruction experiment

We quantitatively study the reconstruction performance of our method on retrospectively undersampled gated data and compare to baselines.

We utilize public cine data from the CMRxRecon challenges 2023 and 2024 [26,28] consisting of 2D+time cardiac k-space measurements from 630 healthy subjects, containing multi-slice short-axis (SAX) and single-slice long-axis views (LAX). Data was acquired on a 3T MAGNETOM Vida scanner (Siemens Healthineers, Germany) in a breath-hold, ECG-gated approach and retrospectively segmented into 12 cardiac phases with a temporal resolution of 50 ms. The spatial resolution was 1.5×1.5 mm^2 with a slice thickness of 8.0 mm. Sensitivity maps were estimated by ESPIRiT [22] using the BART toolbox [23].

We train a dynamic diffusion model $\boldsymbol{\epsilon}_\theta$ with $T = 1000$ steps, using 5840 cine sequences from 450 individual subjects of CMRxRecon. Training followed the noise-matching methodology described in Section 2.2.

For evaluation, we generate variable density kt-Gaussian undersampling masks with acceleration factors ranging from 4 to 16 and no explicit fully sampled autocalibration signal (ACS). The masks are applied to 160 unseen test sequences from CMRxRecon and reconstructions are acquired for all tested methods. For testing reproducibility of the results, we report $mean \pm std$ over 5 runs with different random masks. Furthermore, we ran paired t-tests for statistical analysis, indicating significant differences to PDDR with $p < 0.05^*$ and $p < 0.01^{**}$.

We compare to the low-rank plus sparse (L+S) reconstruction approach [16]. We also compare to two untrained methods for cardiac MRI, Fourier-feature multi-layer perceptrons (FMLP) [13] and the time-dependent deep image prior (T-DIP) [29]. Moreover, we compare to diffusion posterior sampling (DPS) [3], an established diffusion-based inverse problems solver, using the proposed dynamic diffusion prior for reconstruction. Lastly, we propose spatial diffusion regularization (SDR), a variant of our method, using only a frame-based 2D image prior.

Table 1. Reconstruction performance for retrospective 12-fold acceleration of gated data. Comparison in terms of image metrics PSNR, SSIM, NMSE and computational requirements as GPU memory (VRAM) and reconstruction time (Time).

Method	PSNR [dB]	SSIM [%]	NMSE	VRAM [GB]	Time [s]
Zero-filled	20.17 ±0.04**	55.50 ±0.17**	1.220 ±0.014**	0.60 ±0.0**	0.015 ±0.00**
L+S	31.26 ±0.10*	85.44 ±0.18**	0.111 ±0.002	**1.16** ±0.0**	**3.98** ±0.00**
FMLP	24.60 ±0.65**	45.03 ±3.15**	1.375 ±0.483**	6.67 ±0.0**	1124 ±26.7**
T-DIP	29.73 ±0.22**	70.70 ±0.84**	0.346 ±0.256*	2.53 ±0.0**	445.5 ±8.12**
DPS	32.78 ±0.06	87.78 ±0.17	**0.069** ±0.002*	16.27 ±0.0**	419.0 ±0.11**
SDR	28.21 ±0.11**	79.13 ±0.20**	0.240 ±0.006**	4.66 ±0.0**	14.77 ±0.00**
PDDR	**32.84** ±0.08	**88.06** ±0.14	0.097 ±0.002	5.42 ±0.0	43.18 ±0.01

Hyperparameters were determined for each method individually, using a grid search over an validation set and fixed for evaluation on unseen test sets.

We find that PDDR achieves high quality reconstructions, outperforming the baselines, as the results reported in Table 1 show. We only report results for 12-fold acceleration, matching undersampling patterns required for real-time imaging. Experiments with accelerations 4, 8, and 16 show the same qualitative trend. PDDR robustly provides highest quality metrics, achieving up to 40 dB PSNR at 4-fold accelerated reconstruction.

DPS, using the same diffusion model, achieves similar video quality to PDDR but uses a costly sampling approach that requires 3 times more VRAM and about 10 times more reconstruction time compared to PDDR.

The unsupervised baselines FMLP and T-DIP are designed for scan-specific reconstruction of long prospective measurements and require model fitting during reconstruction, which leads to very high reconstruction times. FMLP fails to robustly reconstruct high quality videos in this retrospective setting.

Our hypothesis is that for high-performant dynamic MRI, it is important to exploit spatiotemporal correlations in the measurements for reconstruction. Consistent with this hypothesis, the spatial prior does not outperform the classical L+S method which leverages the temporal correlations between frames. This is supported by the reconstruction results using the dynamic diffusion prior in DPS and PDDR, which outperform the spatial prior of SDR in all image metrics by a large margin. This underscores the capabilities of the spatiotemporal diffusion model to capture and utilize strong priors of dynamic cardiac videos.

4.2 Piecewise regularization experiment

Next, we investigate PDDR's performance with respect to changes in the sampling of the regularization block based on simulated data.

Based on CMRxRecon data we simulate video sequences with up to $N = 120$ frames corresponding to 6 s of simulated measurement acquisition. As gated data is acquired with the assumption of periodically identical cardiac cycles, simulation of long acquisitions by repetitive sampling from the gated cycle is a valid

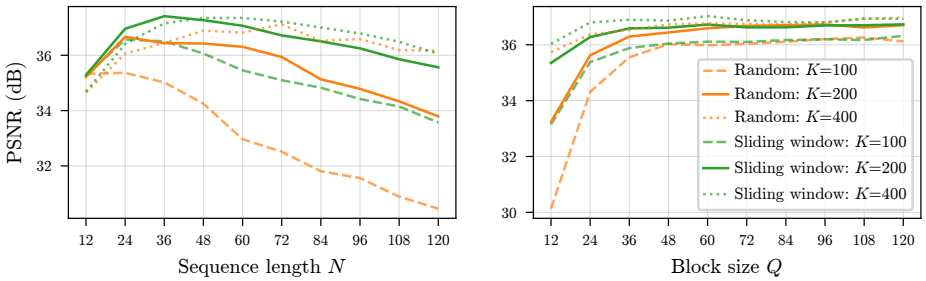

Fig. 1. The block sampling strategy. The sliding window approach outperforms random positions of the regularization block. *Left:* For a fixed block size of $Q = 12$, we vary the frame sequence length N. *Right:* For a fixed sequence length $N = 120$, we vary the size of the regularization block Q.

setup for evaluation of piecewise model performance. We applied retrospective undersampling masks with acceleration factor 8.

Figure 1 investigates the block sampling strategy and block size Q with respect to overall sequence length N and number of optimization steps K.

Sampling regularization blocks in a sliding window manner outperforms random sampling for the same computation budget, as determined by number of optimization steps K and block size Q. The smaller the ratio between Q and N, the worse random block sampling compares to the sliding window approach. The performance difference vanishes with more optimization steps K or larger block sizes Q, as both leads to more regularization coverage and therefore better approximation of the full regularization signal, but also leads to higher computational cost.

When using sliding window with a moderate number of iterations K, the experiment shows high reconstruction quality even with small block sizes Q relative to the number of video frames N. This enables the efficient application of otherwise expensive spatiotemporal diffusion models in prospective real-time acquisitions, where regularization of the full sequences is prohibited by memory and runtime constraints. The method can flexibly adopt to user-specific preferences, by trading memory, runtime, and reconstruction performance through selection of block size Q and optimization steps K.

4.3 Prospective real-time experiment

We apply PDDR to prospective free-breathing data and analyze reconstructions qualitatively and quantitatively. We use the pretrained model $\boldsymbol{\epsilon}_\theta$ from Section 4.1 with $K = 200$ optimization steps and block size $Q = 36$.

For our experiment, we use 10 prospective acquisitions of at least 5 s duration from the public OCMR dataset [2]. SAX and LAX data of 1.5T MAGNETOM Avanto and Sola (Siemens Healthineers, Germany) scanners is provided in spatial resolution of about 2.2×2.2 mm^2, slice thickness 8.0 mm, and temporal resolution of 50 ms.

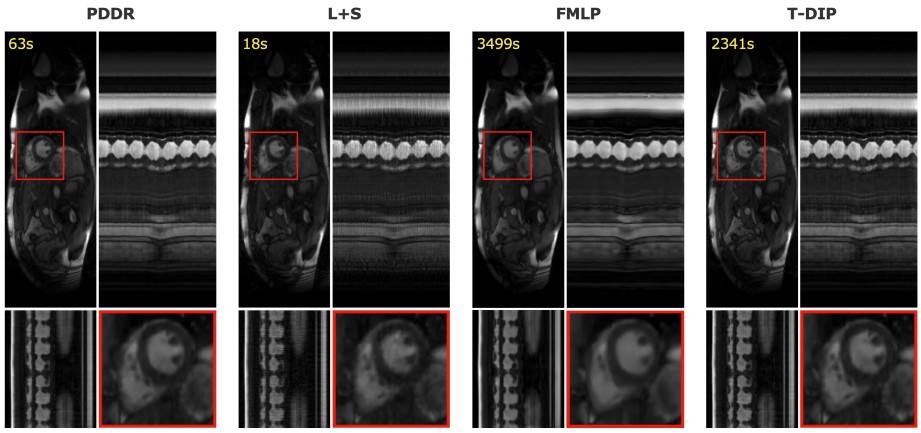

Fig. 2. Example reconstructions of prospective free-breathing data. Showing the full field-of-view, time profiles through the heart, and a close-up of the heart. The yellow label indicates the time for reconstruction of the respective video in seconds.

We provide the signal-to-error ratio (SER [dB] [13]) quantifying the estimation error with respect to holdout frequencies, and the motion smoothness metric temporal total variation (TTV). We report the mean over the prospective testset. The reference-free metrics amend the qualitative evaluation, e.g. for SDR, low signal estimation (SER:6.19) and high temporal variance (TTV:110.4) support the visual assessment, showing strong motion and flickering artifacts.

As shown in Figure 2, PDDR efficiently reconstructs cardiac videos. The figure shows reconstructions of a free-breathing cine MRI acquisition capturing 7 heartbeats in $N = 128$ video frames. Application of DPS exceeds available GPU memory, while piecewise regularization with PDDR enables efficient integration of the dynamic diffusion model for reconstruction without sacrificing video quality, achieving strong hold-out estimation (SER:16.05) and balanced temporal smoothness (TTV:31.4). While classical L+S is the fastest and shows similar signal estimation capabilities (SER:16.12), visual comparison shows strong noise in the reconstructions. The lower temporal variation (TTV:18.7) indicate lower motion dynamics in L+S, as low-rank enforcing can lead to reduced motion in the reconstructions. Due to its coordinate-based representation, reconstructions with FMLP are overly smooth (SER:14.45/TTV:7.65) and expensive to obtain. T-DIP provides good reconstructions (SER:16.03/TTV:33.7), but requires 39 minutes reconstruction time for the single video, compared to 63 seconds for PDDR with similar reconstruction quality.

Although not explicitly represented in the training data of PDDR, we can not observe visual degradation of reconstructions under respiratory motion. In terms of signal estimation capabilities, quantified by SER, we achieve similar or better performance compared to untrained references, indicating no performance decline in reconstruction of free-breathing data due to the train-test shift.

5 Conclusion

In this paper, we proposed a dynamic diffusion prior as piecewise regularizer in variational reconstruction of cardiac MRI videos. The method is shown to outperform classical, unsupervised, and spatial models in reconstruction of accelerated acquisitions. For the more challenging setting of prospective real-time reconstruction, our method yields slightly more accurate reconstructions while being computationally much more efficient than current unsupervised methods. Our method is designed to operate on long sequences in a piecewise manner, allowing for efficient handling of extended data while maintaining robust reconstruction quality. This design not only circumvents memory bottlenecks but also offers a principled approach for integrating long-duration, real-time acquisitions into the reconstruction pipeline. Fast and flexible reconstruction make PDDR a practical approach for cardiac cine, integrating problem-specific spatiotemporal diffusion priors into a highly underdetermined problem.

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

References

1. Blattmann, A., Rombach, R., Ling, H., Dockhorn, T., Kim, S.W., Fidler, S., Kreis, K.: Align your latents: High-resolution video synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22563–22575 (2023)
2. Chen, C., Liu, Y., Schniter, P., Tong, M., Zareba, K., Simonetti, O., Potter, L., Ahmad, R.: Ocmr (v1. 0)–open-access multi-coil k-space dataset for cardiovascular magnetic resonance imaging. arXiv preprint arXiv:2008.03410 (2020)
3. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. In: International Conference on Learning Representations (2023)
4. Chung, H., Ye, J.C.: Score-based diffusion models for accelerated mri. Medical Image Analysis **80**, 102479 (2022)
5. Contijoch, F., Rasche, V., Seiberlich, N., Peters, D.C.: The future of cmr: All-in-one vs. real-time cmr (part 2). Journal of Cardiovascular Magnetic Resonance **26**(1), 100998 (2024)
6. Daras, G., Nie, W., Kreis, K., Dimakis, A.G., Mardani, M., Kovachki, N.B., Vahdat, A.: Warped diffusion: Solving video inverse problems with image diffusion models. In: Advances in Neural Information Processing Systems. vol. 37, pp. 101116–101143 (2024)
7. 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 (2025)
8. Heckel, R., Jacob, M., Chaudhari, A., Perlman, O., Shimron, E.: Deep learning for accelerated and robust mri reconstruction. Magnetic Resonance Materials in Physics, Biology and Medicine **37**(3), 335–368 (2024)

9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Advances in Neural Information Processing Systems. vol. 33, pp. 6840–6851 (2020)
10. Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A.G., Tamir, J.: Robust compressed sensing mri with deep generative priors. In: Advances in Neural Information Processing Systems. vol. 34, pp. 14938–14954 (2021)
11. Knoll, F., Hammernik, K., Zhang, C., Moeller, S., Pock, T., Sodickson, D.K., Akcakaya, M.: Deep-learning methods for parallel magnetic resonance imaging reconstruction: A survey of the current approaches, trends, and issues. IEEE Signal Processing Magazine **37**(1), 128–140 (2020)
12. Krainovic, A., Ruschke, S., Heckel, R.: Resolution-robust 3d mri reconstruction with 2d diffusion priors: Diverse-resolution training outperforms interpolation. arXiv preprint arXiv:2412.18584 (2024)
13. Kunz, J.F., Ruschke, S., Heckel, R.: Implicit neural networks with fourier-feature inputs for free-breathing cardiac mri reconstruction. IEEE Transactions on Computational Imaging **10**, 1280–1289 (2024)
14. Mardani, M., Song, J., Kautz, J., Vahdat, A.: A variational perspective on solving inverse problems with diffusion models. In: International Conference on Learning Representations. vol. 2024, pp. 28027–28053 (2024)
15. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: International Conference on Machine Learning. pp. 8162–8171 (2021)
16. Otazo, R., Candes, E., Sodickson, D.K.: Low-rank plus sparse matrix decomposition for accelerated dynamic mri with separation of background and dynamic components. Magnetic Resonance in Medicine **73**(3), 1125–1136 (2015)
17. Ozturkler, B., Mardani, M., Vahdat, A., Kautz, J., Pauly, J.M.: Regularization by denoising diffusion process for mri reconstruction. In: NeurIPS 2023 Workshop on Deep Learning and Inverse Problems (2023)
18. Poddar, S., Jacob, M.: Dynamic mri using smoothness regularization on manifolds (storm). IEEE Transactions on Medical Imaging **35**(4), 1106–1115 (2015)
19. Rajiah, P.S., François, C.J., Leiner, T.: Cardiac mri: state of the art. Radiology **307**(3), e223008 (2023)
20. Schlemper, J., Caballero, J., Hajnal, J.V., Price, A.N., Rueckert, D.: A deep cascade of convolutional neural networks for dynamic mr image reconstruction. IEEE Transactions on Medical Imaging **37**(2), 491–503 (2017)
21. Song, J., Vahdat, A., Mardani, M., Kautz, J.: Pseudoinverse-guided diffusion models for inverse problems. In: International Conference on Learning Representations (2023)
22. 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 (2014)
23. Uecker, M., Ong, F., Tamir, J.I., Bahri, D., Virtue, P., Cheng, J.Y., Zhang, T., Lustig, M.: Berkeley advanced reconstruction toolbox. In: Proc. Intl. Soc. Mag. Reson. Med. vol. 23, p. 9 (2015)
24. Vornehm, M., Chen, C., Sultan, M.A., Arshad, S.M., Han, Y., Knoll, F., Ahmad, R.: Multi-dynamic deep image prior for cardiac mri. Magnetic Resonance in Medicine **94**(6), 2668–2679 (2025)
25. Vornehm, M., Wetzl, J., Giese, D., Fürnrohr, F., Pang, J., Chow, K., Gebker, R., Ahmad, R., Knoll, F.: Cinevn: Variational network reconstruction for rapid functional cardiac cine mri. Magnetic Resonance in Medicine **93**(1), 138–150 (2025)
26. Wang, C., Lyu, J., Wang, S., Qin, C., Guo, K., Zhang, X., Yu, X., Li, Y., Wang, F., Jin, J., et al.: Cmrxrecon: A publicly available k-space dataset and benchmark to advance deep learning for cardiac mri. Scientific Data **11**(1), 687 (2024)

27. Wang, Z., Huang, J., Huang, M., Wang, C., Yang, G., Qu, X.: Robust cardiac cine mri reconstruction with spatiotemporal diffusion model. IEEE Transactions on Computational Imaging **11**, 1258–1270 (2025)
28. Wang, Z., Wang, F., Qin, C., Lyu, J., Ouyang, C., Wang, S., Li, Y., Yu, M., Zhang, H., Guo, K., et al.: Cmrxrecon2024: A multimodality, multiview k-space dataset boosting universal machine learning for accelerated cardiac mri. Radiology: Artificial Intelligence **7**(2), e240443 (2025)
29. Yoo, J., Jin, K.H., Gupta, H., Yerly, J., Stuber, M., Unser, M.: Time-dependent deep image prior for dynamic mri. IEEE Transactions on Medical Imaging **40**(12), 3337–3348 (2021)