

LEARNING THE DOMAIN SPECIFIC INVERSE NUFFT FOR ACCELERATED SPIRAL MRI USING DIFFUSION MODELS

Trevor J. Chan¹, Chamith S. Rajapakse²

1. Department of Bioengineering, University of Pennsylvania, Philadelphia, United States
2. Department of Radiology, University of Pennsylvania, Philadelphia, United States

ABSTRACT

Deep learning methods for accelerated MRI achieve state-of-the-art results but largely ignore additional speedups possible with noncartesian sampling trajectories. To address this gap, we created a generative diffusion model-based reconstruction algorithm for multi-coil highly undersampled spiral MRI. This model uses conditioning during training as well as frequency-based guidance to ensure consistency between images and measurements. Evaluated on retrospective data, we show high quality (structural similarity >0.87) in reconstructed images with ultrafast scan times (0.02 seconds for a 2D image). We use this algorithm to identify a set of optimal variable-density spiral trajectories and show large improvements in image quality compared to conventional reconstruction using the non-uniform fast Fourier transform. By combining efficient spiral sampling trajectories, multicoil imaging, and deep learning reconstruction, these methods could enable the extremely high acceleration factors needed for real-time 3D imaging.

Index Terms— Accelerated MRI, Spiral MRI, Deep Learning, Image Reconstruction

1. INTRODUCTION

Despite the numerous advantages of MRI, inherent physical and hardware constraints cap acquisition speed and lead to long scan times. This creates myriad downstream obstacles, including low patient compliance, inefficient resource allocation, and image motion artifacts, among others. For this reason, methods to accelerate MRI acquisition have been and continue to be an active area of research. Acceleration can be achieved through a combination of using more efficient scanning sequences and reducing the number of measurements made in the frequency space of the image, k-space.

Considering the former, successful techniques for faster scanning include radial and spiral imaging methods, which exploit the unequal distribution of information across k-space by sampling lower frequencies more densely [1, 2].

Considering the latter, acquiring fewer measurements in k-space rewards a decrease in scan time proportional to the

ratio of undersampling, but comes at the cost of image quality. Sampling below the Nyquist limit introduces ambiguities during reconstruction which manifest as artifacts in the final image. Reconstruction algorithms must therefore leverage additional information, including multicoil data in the case of parallel imaging, and sparse priors, in the case of compressed sensing, in order to resolve these ambiguities [3, 4].

A third more recent approach to undersampled MRI reconstruction lies in deep learning methods, which essentially learn a set of image priors and use these to regularize solutions to the ill-posed reconstruction problem [5]. Within this category, diffusion models stand out for producing state-of-the-art results on image reconstruction tasks for faster scanning, motion correction, noise reduction, and others [6, 7, 8, 9]. Despite this, the vast majority of deep learning approaches, and to our knowledge, all diffusion-based approaches to image reconstruction, focus on Cartesian-sampled MRI, missing out on potential acceleration gains attained by using more efficient non-Cartesian sampling trajectories. **This work addresses this gap by introducing a diffusion model-based method for trajectory-agnostic undersampled MRI reconstruction of multicoil spiral MRI.**

Contributions:

- *Creation of a novel multi-conditioning strategy for solving the inverse non-uniform fast Fourier transform (nufft) using a learned conditional score function and weak frequency guidance during sampling*
- *Efficient hyperparameter search of the joint trajectory-reconstruction space and identification of optimal sampling trajectories*
- *Retrospective acquisition and reconstruction of a 2D, 256x256 pixel, 22x22 cm² image with a readout duration of 0.02 seconds.*

2. BACKGROUND

Canonical score-based models consider the mapping between a known distribution of independently and identically distributed samples of gaussian noise and an observed, but unknown distribution of data $p(x)$. These distributions bookend a Langevin diffusion process described by the following




Fig. 1. Example trajectories (A) and the corresponding readout gradients in k_x and k_y (B). All trajectories shown cover the frequency space of a 256x256 image and have a readout duration of 10.0 ms.

stochastic differential equation representing the trajectory of a sample from our data distribution into a sample from our noise distribution:

$$dx = f(x, t)dt + g(t)dw. \quad (1)$$

Here, functions $f(\cdot, t)$ and $g(\cdot)$ are the drift and diffusion coefficients of $x(t)$ respectively, and w is a standard Wiener process, or Brownian motion.

In order to generate a novel sample from our data distribution, we can generate a random noise vector and attempt to solve the reverse-time SDE, but this is generally intractable. However, we can approximate this process by estimating the noise-conditioned score function, $\nabla_x \log p_t(x)$, which computes the likelihood of a sample x existing between the noise and image distributions. With this, the reverse-time SDE becomes:

$$dx = [f(x, t) - g(t)^2 \nabla_x \log p_t(x)] dt + g(t)d\bar{w} \quad (2)$$

The score function can be trained using a score matching with Langevin dynamics algorithm [10, 11].

3. METHODS

This research study was conducted retrospectively using human subject data made available in open access by [12]. Ethical approval was not required.

3.1. Data

We use the NYU FastMRI dataset [12] consisting of 6970 fully sampled 2D brain scans on hardware ranging from 4 to 24 coils. For training and testing, we consider axial T2 weighted turbo spin echo sequences characterized by the following sequence parameters: scan time=140 s, TR=6 s, TE=113 ms, slices=30, slice thickness=5 mm, field of view=22 cm, matrix size=320x320. Effective scan time for a 2D slice at 256² resolution is $140s/320 * 256/30 \approx 3.7s$. As this data is initially acquired using Cartesian sequences,




Fig. 2. Given measurements y_0 , reconstruction follows a modified diffusion sampling process. At each timestep, a noisy latent x_t is concatenated with a prior p_0 and passed to the denoising model to obtain $\tilde{x}_{t-1}$. To enforce consistency with y_0 , we compute a frequency gradient ∇y_{t-1} and solve for the image gradient using a modified iterative inverse nufft (section 3.3). A weighted sum of x_{t-1} and ∇x_{t-1} yields the corrected image x_{t-1} . This is repeated until $t = 0$.

we simulated spiral acquisition by retrospectively interpolating in k-space to attain complex-valued measurements along generated spiral trajectories.

3.2. Generating spiral trajectories

Following Kim et al. [13], we consider spiral trajectories of the form

$$k(\tau) = \int_0^\tau \frac{1}{\rho(\phi)} d\phi e^{j\omega\tau} \approx \lambda \tau^\alpha e^{j\omega\tau}. \quad (3)$$

Here, ρ denotes sampling density, τ is a function of time, ϕ is angular position, $\omega = 2\pi n$ is frequency, with n the number of turns in k-space, λ a scaling factor equal to $matrix\ size / (2 * FOV)$, and α is a bias term for oversampling the center of k-space relative to the edges. Solving this parametric equation under the constraints of capped gradient slew rate and capped gradient amplitude yields gradients ($g_x(t)$ and $g_y(t)$) as well as a spiral trajectory in the k_x, k_y -plane (figure 1). In doing so, we can tune sampling parameters to control for factors such as read out duration and dwell time, while varying the number of interleaves and ratio of low-to-high frequency oversampling.

3.3. Image reconstruction is inverse problem solving

MRI undersampled acquisition amounts to measuring an unknown signal x through some imperfect sampling function A : $y = Ax + \epsilon$. Here, y is the measured multicoil k-space data, and A is the non-uniform fourier transform. ϵ is measurement noise and exists in the same domain as the y ; in MRI, noise is gaussian-distributed across the real and imaginary components of y for each coil.

Reconstruction is an ill-posed inverse problem of recovering an image signal x from a set of incomplete k-space measurements y . As x and y exist in different domains, x is hidden behind a sampling operator A . Solving this problem necessitates prior knowledge. In our case, we learn an underlying conditional distribution of images and seek to reconstruct




Fig. 3. Representative reconstruction results for a single 2D 16 coil image. Retrospective k-space data was sampled with an optimized 23 interleave sequence with a total readout duration of 0.02 s. Rows 1 and 2 show the RSS-reconstructed images and log-scaled k-space magnitudes for the ground truth, inverse nufft, and proposed model reconstructions. Below are the individual coil magnitude and phase images for the fully sampled image, the inverse nufft reconstructions, and the model predictions.

samples from this distribution consistent with the measurements. Information is supplied in two forms: first, we learn a conditional score function $\nabla_x \log p_t(x_t | \tilde{A}^{-1}y_0)$, where y_0 is the measurement in frequency space and $\tilde{A}^{-1}$ is an approximate inverse of A , in our case the inverse nufft solved iteratively using conjugate gradients. We find that adding this supervision during training helps to constrain the model when faced with a large number of input image channels and the periodic ambiguity inherent when operating on complex numbers.

Second, we use frequency space gradients to weakly guide the sampling process. At each time step during sampling, we compute the forward nufft of an uncorrected $\tilde{x}_{t-1}$ and take a difference between that and the measurements y_0 . To minimize this difference, the approximate gradient in image space is calculated by solving a modified approximate inverse nufft $\tilde{A}_t^{-1}$, which corrects for low frequency biases and applies time step-dependent noising determined by the noise schedule $\sigma(t)$, which is necessary for sampling with langevin diffusion.

$$\tilde{A}_t^{-1}x_t = \frac{\tilde{A}^{-1}x_t}{c_1 e^{-c_2 r^2}} + \mathcal{N}(0, \sigma(t)^2) \quad (4)$$

Following [11], we choose a linear noise schedule and observe that the underlying ordinary differential equation describing transit from latent to image is locally linear, so summation of $\tilde{x}_{t-1}$ and ∇x_{t-1} to obtain a frequency-corrected image x_{t-1} is akin to gradient descent (figure 2).




Fig. 4. We performed a grid hyperparameter search over a 2D trajectory space. We fixed readout duration at 0.02 seconds and varied the number of interleaves from 1 to 125 and alpha from 1 to 4. Based on structural similarity of the model-reconstructed images, we found multiple trajectories that yield improved image quality. In comparison, the naive Archimedean spiral, corresponding to 1 interleave and $\alpha = 1$, performs very poorly.

In practice, due to the non-invertibility of the nufft, imperfections in the approximate inverse nufft bleed into the final image reconstruction, introducing artifacts and reducing quality. To avoid this, we anneal the guidance signal following an empirically chosen linear schedule $\gamma(t) = \beta(1 - t)$, ensuring strong guidance at the outset of sampling and minimal artifacts at the end of sampling. A consequence of this choice is that we do not strongly enforce that $Ax_0 = y_0$ at time 0.

4. RESULTS

Model reconstruction performance was evaluated on a held-out test dataset. Test trajectories have a fixed readout duration of 0.02 seconds, in which time the measurements needed to reconstruct a 256x256 pixel, 22x22 cm² 2D image are acquired. Reconstructed image quality was scored using structural similarity (SSIM) (figure 3).

To investigate the effect choosing different scanning trajectories has on the quality of reconstructed images, we also perform a grid hyperparameter search of spiral trajectories with a fixed readout duration of 0.02 seconds (figure 4) and varying α and interleaves. Surprisingly, the common ‘naive’ trajectory, a single interleave Archimedean spi-




Fig. 5. (A) Effect of sampling trajectory optimization, model reconstruction without frequency guidance, and model reconstruction with frequency guidance. For the non-optimized trajectory, we used a single interleave Archimedean spiral with a readout duration of 0.02 s. The optimized trajectory uses a 23 interleave, $\alpha = 1.23$ sequence with an identical readout duration. (B) Snapshots of the image latent x_t and the gradient signal ∇x_t taken during a diffusion sampling process.

ral, corresponding to $\alpha = 1$, performs very poorly when sampled below the Nyquist limit. Trajectories which perform better tended to lie along two logarithmic curves roughly characterized by $\alpha = 1.33 \log(0.39 \text{ interleaves})$ and $\alpha = 0.87 \log(0.54 \text{ interleaves})$.

Finally, we conduct an ablation study to disentangle the effects of optimal sampling trajectory without model reconstruction, model reconstruction without frequency guidance, and model reconstruction with frequency guidance (figure 5). We find that all three contribute to noticeable increases in image quality, both visual, and quantitative based on SSIM. The combination of choosing an optimal trajectory, performing model reconstruction with conditioning, and using annealed frequency guidance results in large improvements in image quality, up to and exceeding a 0.15 boost in SSIM.

5. DISCUSSION

While initial results are promising, the main limitation of this project is the reliance on retrospective, Cartesian-sampled data. Implementing the sequences outlined in this work will likely require customizing spiral sequences so as to match the contrast and signal of the original Cartesian sequences, which will constrain the space of realizable trajectories. Until a dataset of raw non-Cartesian MRI data becomes available, this will continue to be an obstacle. For a similar reason, it is difficult to make head-to-head comparisons between the original sequence and the proposed sequences without prospective validation. For this reason, the immediate task is

to acquire prospective data with custom sequences and use it to validate image reconstruction. Currently, a concrete direct comparison is between the proposed sequences and their Nyquist-sampled counterparts, which run roughly 3x longer.

Apart from the short-term task of matching sequence parameters between spiral and Cartesian sequences, the choice of spiral sequence leaves a considerable amount of flexibility even within the space of optimized interleave and density pairs identified above. Variation in number of interleaves, and by extension the duration of a single interleave, allows for tailoring of sequence contrast, signal, and speed to task-specific requirements. An area in which these sequences could provide additional benefit, even outside of sheer acceleration, would be in imaging tissues with a very short $T2^*$, as acceleration within an interleave allows for proportionally more data acquisition to occur before signal has decayed.

6. CONCLUSION

Here we introduce a new method and show preliminary results for reconstructing spiral MRI using a diffusion model. Combining multicoil imaging, spiral scanning, and under-sampling enables dramatically faster imaging speeds. Applications of this work are widespread; in addition to the numerous typical benefits associated with faster scanning, including better patient compliance and fewer motion artifacts, these methods have the potential to reach the extremely high acceleration factors necessary to achieve high resolution real-time 3D imaging.

7. ACKNOWLEDGMENTS

No funding was received for conducting this study. The authors have no relevant financial or non-financial interests to disclose.

8. REFERENCES

- [1] Mark J Blum, Michael Braun, and Dov Rosenfeld, “Fast magnetic resonance imaging using spiral trajectories,” in *Medical Imaging*. SPIE, 1987, vol. 767, pp. 40–46.
- [2] Stefanie Winkelmann, Tobias Schaeffter, Thomas Koehler, Holger Eggers, and Olaf Doessel, “An optimal radial profile order based on the golden ratio for time-resolved mri,” *IEEE transactions on medical imaging*, vol. 26, no. 1, pp. 68–76, 2006.
- [3] Daniel K Sodickson and Warren J Manning, “Simultaneous acquisition of spatial harmonics (smash): fast imaging with radiofrequency coil arrays,” *Magnetic resonance in medicine*, vol. 38, no. 4, pp. 591–603, 1997.
- [4] Michael Lustig, David Donoho, and John M Pauly, “Sparse mri: The application of compressed sensing for rapid mr imaging,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [5] Julio A Oscanoa, Matthew J Middione, Cagan Alkan, Mahmut Yurt, Michael Loecher, Shreyas S Vasanaawala, and Daniel B Ennis, “Deep learning-based reconstruction for cardiac mri: A review,” *Bioengineering*, vol. 10, no. 3, pp. 334, 2023.
- [6] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon, “Solving inverse problems in medical imaging with score-based generative models,” *arXiv preprint arXiv:2111.08005*, 2021.
- [7] Patricia M Johnson, Michael P Recht, and Florian Knoll, “Improving the speed of mri with artificial intelligence,” in *Seminars in musculoskeletal radiology*. Thieme Medical Publishers, 2020, vol. 24, pp. 012–020.
- [8] Zhuo-Xu Cui, Chentao Cao, Jing Cheng, Sen Jia, Hairong Zheng, Dong Liang, and Yanjie Zhu, “Spirit-diffusion: Self-consistency driven diffusion model for accelerated mri,” *arXiv preprint arXiv:2304.05060*, 2023.
- [9] Asad Aali, Marius Arvinte, Sidharth Kumar, and Jonathan I Tamir, “Solving inverse problems with score-based generative priors learned from noisy data,” *arXiv preprint arXiv:2305.01166*, 2023.
- [10] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole, “Score-based generative modeling through stochastic differential equations,” *arXiv preprint arXiv:2011.13456*, 2020.
- [11] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine, “Elucidating the design space of diffusion-based generative models,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 26565–26577, 2022.
- [12] Florian Knoll, Jure Zbontar, Anuroop Sriram, Matthew J Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J Geras, Joe Katsnelson, Hersh Chandarana, et al., “fastmri: A publicly available raw k-space and dicom dataset of knee images for accelerated mr image reconstruction using machine learning,” *Radiology: Artificial Intelligence*, vol. 2, no. 1, pp. e190007, 2020.
- [13] Dong-hyun Kim, Elfar Adalsteinsson, and Daniel M Spielman, “Simple analytic variable density spiral design,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 50, no. 1, pp. 214–219, 2003.