
Coarse-to-Fine 3D MRI Reconstruction via 3D Neural Operators

Armeet Singh Jatyani* Jiayun Wang* Ryan Y. Lin
Valentin Duruisseaux Anima Anandkumar
California Institute of Technology
{armmeet,peterw,rylin,vduruiss,anima}@caltech.edu

Abstract

3D MRI reconstruction with deep learning is limited not only by GPU memory and voxel data resolution but also by the tendency of standard neural networks to overfit to training discretizations (resolutions), which makes them highly sensitive to variations in image resolution and sampling patterns. This forces downsampling or dimensional collapse and restricts generalization to higher-resolution volumes. We present a new neural operator framework for learning local features backed by 3D discrete-continuous convolutions (DISCO), which are inherently resolution-agnostic. Unlike conventional 3D convolutions or kernel-interpolated weights, the proposed 3D neural operator relies on filters in a continuous domain, while preserving local inductive biases. This design enables training on coarse, low-memory volumes with full backpropagation, and supports high-resolution zero-shot or few-shot inference without aliasing, while reducing memory cost. This coarse-to-fine regime allows memory-efficient 3D training and large-volume testing using inference only. We evaluate the proposed 3D local neural operator on the SKM-TEA dataset for accelerated 3D MRI reconstruction, demonstrating accurate reconstructions with strong runtime and memory efficiency. While we focus on 3D MRI, the proposed 3D DISCO-based operator is broadly applicable to other 3D imaging modalities and general 3D voxel-based data reconstruction tasks.

1 Introduction

Medical imaging plays a crucial role across numerous scientific fields [1–4], and building more robust models for modalities such as magnetic resonance imaging (MRI) [5] has therefore garnered tremendous attention [6–10]. Deep learning has proven extremely useful in improving medical imaging, particularly in MRI reconstruction from highly undersampled measurements [10–12]. However, these deep learning architectures are limited by GPU memory, and as a result most medical imaging studies in the literature focus on 2D reconstruction tasks [13, 14]. Deep learning architectures for 3D MRI reconstruction are increasingly desirable for both methodological and practical reasons. Compared to slice-based 2D approaches, 3D models exploit volumetric structure to improve depth consistency and reconstruction fidelity [13]. This shift is supported by the availability of large volumetric MRI datasets such as SKM-TEA [15] and by advances in hardware and GPU acceleration, which enable efficient whole-volume processing in a single forward pass.

3D imaging imposes substantially higher computational demands than its 2D counterpart, as algorithmic complexity significantly increases with the additional spatial dimension (the input resolution changes from $O(n^2)$ in 2D to $O(n^3)$ in 3D). Such an extension not only raises the number of trainable parameters but also amplifies memory usage and computational cost during both forward and backward passes. Consequently, naïve scaling of 2D architectures to 3D results in prohibitive resource

*Equal contribution.

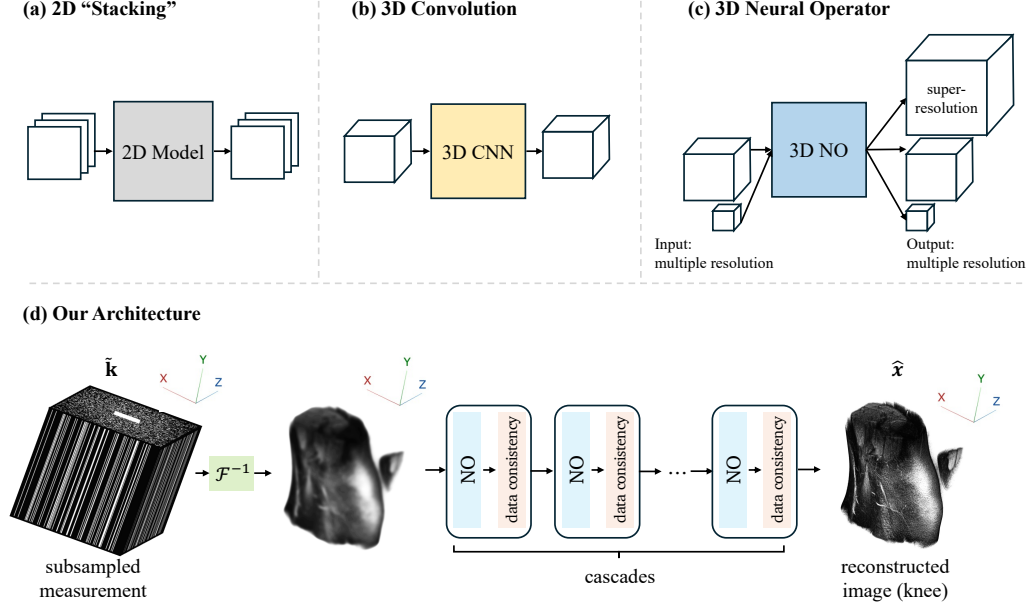


Figure 1: **(a)** Although 3D MRI data (e.g., SKM-TEA, knee MRI) are acquired using volumetric scanners, GPU memory constraints often necessitate training 2D “slice-based” models, which are later stacked for 3D volume inference. **(b)** 3D CNN-based approaches can only be trained on a single resolution due to their fixed kernel size, and the training resolution is thus constrained by GPU memory. **(c)** Our 3D neural operator (NO) model can be trained on any resolution and perform inference at any resolution. This capability allows for memory-efficient training on lower-resolution “coarse” samples while maintaining MRI reconstruction performance at “finer” high-resolution samples. **(d)** Our architecture is an end-to-end neural operator that learns local features to reconstruct high-frequency details missing in undersampled MRI measurements. Our neural operator layers learn continuous kernels as opposed to fixed-resolution conventional CNN kernels.

requirements, motivating the need for architectural modifications and efficiency-oriented methods. Beyond computational demands, ensuring generalization presents an additional challenge, particularly under the severe subsampling often required in 3D imaging. Standard neural networks tend to overfit to training discretizations, restricting their ability to generalize across different sampling patterns or resolutions. This is particularly problematic in 3D MRI where scan protocols vary widely. A popular workaround is to train on down-scaled samples, but traditional convolutions are highly sensitive to variations in resolution [14]. Kernel interpolation, involving upscaling kernels to match the original resolution, often introduces aliasing and hallucinated structures which are unacceptable in MRI since they can obscure critical anatomy, mislead clinical interpretation, and compromise patient safety.

Neural operators provide a principled alternative, as they are inherently resolution agnostic and can map functions to functions regardless of discretization [16, 17]. Neural operators can be trained on low-resolution *coarse* volumes with full backpropagation, and then applied at high resolution *finer* volumes during inference without additional memory burden from backpropagation. Recent studies have shown that neural operators for capturing local data features with localized integral improve computational imaging tasks like 2D MRI [11] and 3D photoacoustic tomography [18]. They outperform standard convolutions in adapting to diverse undersampling schemes and imaging resolutions. Such local neural operators are built on the basic block called discrete-continuous convolutions (DISCO) [19, 20], which mimics standard convolutions but in a resolution-agnostic manner. See Appendix C for more details about related neural operator literature.

Our approach. Motivated by the success of the local neural operator backed by DISCO in 2D MRI reconstruction [11], we propose its extension to 3D, creating a unified resolution-agnostic framework for efficient and scalable 3D MRI reconstruction. The proposed neural operator layers provide a resolution-agnostic alternative to conventional Conv3D layers. The resulting 3D neural operator can be trained on low-resolution samples while maintaining accurate reconstruction at higher resolutions, enabling super-resolution inference at scales beyond those seen during training.

On the SKM-TEA 3D knee MRI dataset [15], we compare our neural operator 3D MRI reconstruction models against the standard 3D CNN baselines in super-resolution tasks and find their performance is on par. The implemented 3D DISCO and corresponding neural operator architecture repository will be publicly available (after acceptance). Although we demonstrate on 3D MRI, the framework can thus be used for other 3D computational/medical imaging tasks and other scientific and engineering areas where 3D volumetric data is involved, such as seismic imaging, fluid dynamics simulations and 3D microscopy. Its flexibility enables seamless adaptation to different data modalities and physical constraints, making it a promising tool for a broad range of volumetric inverse problems.

2 Methods

While neural networks have had successes in medical imaging, they learn mappings between finite-dimensional vectors and are not resolution-agnostic. As a result, they usually perform well only at the training resolution, suffering a significant performance drop when measurements are taken at different resolutions [11]. Neural operators offer a principled generalization of neural networks to learn mappings between functions [16, 17], and have been successfully applied to a wide range of applications, including medical imaging [11, 21, 22]. By design, they can be applied to signals at any discretization or sampling pattern and their output can be queried on arbitrary discretizations. As a result, they offer a highly desirable unified approach to discretization-agnostic MRI reconstruction.

Given that accurate MRI reconstruction depends on capturing local anatomical structures, we propose a neural operator architecture that incorporates local inductive biases.

Local Features via Local Integration Operator. A common method of embedding local inductive bias into neural networks has been by using locally supported convolutional kernels, as in convolutional neural networks (CNNs). However, standard discrete convolutional kernels used in CNNs do not satisfy the resolution-agnostic properties of neural operators. Specifically, Liu et al. [20] show that CNN-style convolutional kernels converge to pointwise linear operators as the resolution is increased, instead of the desired local integration in the limit of infinite resolution. For a kernel κ and input function g defined over some compact subset $D \subset \mathbb{R}^d$, the *local convolution operator* in a standard convolution layer, which transforms input u to output v , is given by $(\kappa \star g)(v) = \int_D \kappa(u-v) \cdot g(u) du$.

Discrete-Continuous Convolution (DISCO). Given input points $(u_j)_{j=1}^m \subset D$ with corresponding quadrature weights q_j and output positions $v_i \in D$, we adopt the discrete-continuous convolution (DISCO) framework for operator learning [19, 20] and approximate the continuous convolution as

$$(\kappa \star g)(v_i) = \int_D \kappa(u - v_i) \cdot g(u) du \approx \sum_{j=1}^m \kappa(u_j - v_i) \cdot g(u_j) q_j. \quad (1)$$

We parameterize the kernel κ as a linear combination of pre-defined kernel basis functions κ^ℓ , i.e., $\kappa = \sum_{\ell=1}^L \theta^\ell \cdot \kappa^\ell$, where the coefficients θ^ℓ are learnable parameters. The convolutional kernel is thus parameterized by a finite number of parameters θ^ℓ , independently of the grid on which the kernel is evaluated. This makes the kernel resolution-agnostic because the resolution-agnostic basis is disentangled from the discrete learnable parameters. The basis κ^ℓ is defined in the function space, and will be discretized at the desired resolution. Since we are operating on an equidistant grid on a compact subset of \mathbb{R}^3 , we follow [20] and implement Eqn. (1) using standard convolutional kernels (thus enjoying the benefits of acceleration on GPUs using standard deep learning libraries). We make two crucial modifications however: **1)** the kernel itself is defined as a linear combination of basis functions κ^ℓ , and **2)** the size of the kernel scales with the input resolution so as to remain at fixed size with respect to the input domain. We use the local integration operator as the resolution-agnostic building block for the measurement space and image space operators.

3D Morlet-Wavelet Basis. We write a new parametrized 3D filter basis that can be queried at any resolution. It employs a 3D Morlet wavelet-like basis as the predefined kernel functions κ^ℓ , extending the 2D construction introduced in [20]. Internal kernel weights are obtained as a linear combination $\kappa = \sum_{\ell=1}^L \theta^\ell \cdot \kappa^\ell$ of these filter basis functions, with trainable parameters θ^ℓ [19, 20].

The Morlet wavelet-like basis is formulated by multiplying a Gaussian window with a complex exponential in 3D. More details on this construction are provided in Section A of the Appendix. For ease of use, our implementation follows the API of the `nn.Conv3D` layer in PyTorch. We visualize the first few Morlet wavelet basis filters as well as the construction of a DISCO 3D kernel in Figure 2.

Proposed 3D MRI NO Architecture. We denote the forward imaging operator as \mathcal{A} with process $\mathbf{k} = \mathcal{A}\mathbf{x} + \epsilon$, where ϵ is the error. The model, depicted in Fig. 1d), is an unrolled neural operator designed to solve the inverse problem. Classical compressed sensing methods reconstruct the image $\hat{\mathbf{x}}$ by solving an optimization problem $\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \frac{1}{2} \sum_i \|\mathcal{A}(\mathbf{x}) - \tilde{\mathbf{k}}\|_2^2 + \lambda \Psi(\mathbf{x})$. The unrolled architecture uses successive *cascades* to mimic different iterations in the optimization problem and to process subsampled MRI measurements with a soft data-consistency term that prevents the model from overriding the acquired measurements. Over multiple cascades, the model learns to *denoise* the sample image x . Specifically, we learn the prior in function space via discretization-agnostic neural operators in image space (NO). A cascade of unrolled layers is applied, each incorporating a data-consistency term and an image-space operator NO for prior learning:

$$\mathbf{x}^{t+1} \leftarrow \mathbf{x}^t - \eta^t \mathcal{A}^*(\mathcal{A}(\mathbf{x}^t) - \tilde{\mathbf{k}}) + \lambda^t \text{NO}^t(\mathbf{x}^t). \quad (2)$$

where η^t controls the weight of the physics prior, λ controls the weight of the data prior, and NO^t refers to the image-space neural operator at cascade t . Each cascade mimics the gradient descent step from \mathbf{x}^t to \mathbf{x}^{t+1} . The model is trained end-to-end and can thus also be considered as an entire model that reconstructs the final 3D MRI $\hat{\mathbf{x}}$ from subsampled MRI \mathbf{k} -space measurements.

3 Numerical Experiments

Details about the dataset and implementation are provided in Appendices D and E. We extended the E2E-VarNet [10] architecture to 3D as the baseline. E2E-VarNet is a state-of-the-art unrolled model which demonstrates superior performance over traditional optimization-based approaches [23, 24] for compressed sensing MRI reconstruction from undersampled measurements. We setup a zero-shot super-resolution task by downscaling full-resolution ($512 \times 512 \times 320$) samples to create a secondary half-resolution ($256 \times 256 \times 160$) dataset. First, both our model (NO3d) and the baseline model (E2E-VarNet3d) are trained on the half-resolution dataset, seeing no full-resolution examples. Then, after training, both models are evaluated on the original full-resolution dataset. Example (full-resolution) outputs of both our model and the baseline model are shown in Appendix F.

Table 1: Performance comparison of the baseline CNN-based E2E-VarNet3D and our DISCO-based NO3D models at half and full resolution. Models were trained on 1/2 resolution. Inference was performed at full resolution. Notably, our DISCO-based method is on par with the conventional CNN-based method. The metric is PSNR (dB).

Method	Half Resolution	Full Resolution
NO3d (DISCO, ours)	31.73	33.21
E2E-VarNet3D [10] (CNN, baseline)	31.73	33.47

4 Conclusion

We introduced a resolution-agnostic neural operator framework for 3D MRI reconstruction using DISCO (discrete-continuous convolutions) [19]. By representing convolutional kernels in a continuous function space while retaining local inductive biases, our approach overcomes the limitations of standard 3D convolutions, which overfit to specific discretizations and require substantial memory for high-resolution volumes. This coarse-to-fine paradigm enables training on low-resolution volumes with full backpropagation and zero-shot or few-shot inference at higher resolutions without aliasing or additional memory cost. Our experiments on the SKM-TEA 3D MRI knee reconstruction dataset demonstrate that 3D DISCO achieves accurate and reliable reconstructions while maintaining strong runtime and memory efficiency in super-resolution tasks. Beyond 3D MRI, this framework offers a robust solution for 3D imaging and volumetric data reconstruction, with potential applications across scientific and clinical domains where data resolution varies or computational resources are limited.

Future work could consider the following directions: 1) Implement and evaluate the performance of other filter bases for DISCO [19], such as piece-wise linear basis in [11]. 2) While this paper only reports zero-shot super-resolution reconstruction results, it is of interest to evaluate the proposed framework’s performance on different measurement sampling rates (acceleration rates) and measurement patterns. 3) Extend and evaluate the framework to other 3D MRI tasks (e.g. liver and brain reconstruction) as well as other 3D inverse problems like seismic imaging.

Acknowledgments and Disclosure of Funding

This work is supported in part by ONR (MURI grant N000142312654 and N000142012786). A.S.J. is supported in part by the Mellon Mays Undergraduate Fellowship and Undergraduate Research Fellowship (SURF) at Caltech. J.W. is supported in part by the Pritzker AI+Science initiative and Schmidt Sciences. A.A. is supported in part by Bren endowed chair and the AI2050 senior fellow program at Schmidt Sciences.

References

- [1] V Seifert, M Zimmermann, C Trantakis, H-E Vitzthum, K Kühnel, A Raabe, F Bootz, J-P Schneider, F Schmidt, and J Dietrich. Open mri-guided neurosurgery. *Acta neurochirurgica*, 141:455–464, 1999.
- [2] DJ Husband, KA Grant, and CS Romaniuk. Mri in the diagnosis and treatment of suspected malignant spinal cord compression. *The British journal of radiology*, 74:15–23, 2001.
- [3] Denis Le Bihan. Looking into the functional architecture of the brain with diffusion mri. *Nature reviews neuroscience*, 4(6):469–480, 2003.
- [4] J Craig Richardson, Richard W Bowtell, Karsten Mäder, and Colin D Melia. Pharmaceutical applications of magnetic resonance imaging (mri). *Advanced drug delivery reviews*, 57(8):1191–1209, 2005.
- [5] Michael Lustig, David Donoho, and John M Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.*, 58(6):1182–1195, December 2007.
- [6] Mark A Griswold, Peter M Jakob, Robin M Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 47(6):1202–1210, 2002.
- [7] Michael Lustig, David L Donoho, Juan M Santos, and John M Pauly. Compressed sensing mri. *IEEE signal processing magazine*, 25(2):72–82, 2008.
- [8] Patricia M Johnson and Maria Drangova. Conditional generative adversarial network for 3d rigid-body motion correction in mri. *Magnetic resonance in medicine*, 82(3):901–910, 2019.
- [9] Hammernik K, Klatzer T, Kobler E, et al. Learning a variational network for reconstruction of accelerated MRI data, 2018.
- [10] Anuroop Sriram, Jure Zbontar, Tullie Murrell, Aaron Defazio, C Lawrence Zitnick, Nafissa Yakubova, Florian Knoll, and Patricia Johnson. End-to-end variational networks for accelerated mri reconstruction. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23*, pages 64–73. Springer, 2020.
- [11] Armeet Singh Jatyani, Jiayun Wang, Aditi Chandrashekar, Zihui Wu, Miguel Liu-Schiaffini, Bahareh Tolooshams, and Anima Anandkumar. A unified model for compressed sensing mri across undersampling patterns. In *Conference on Computer Vision and Pattern Recognition (CVPR) Proceedings*, 2025.
- [12] Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated mri. *Medical Image Analysis*, page 102479, 2022.
- [13] Ke Wang, Michael Kellman, Christopher M Sandino, Kevin Zhang, Shreyas S Vasanawala, Jonathan I Tamir, Stella X Yu, and Michael Lustig. Memory-efficient learning for high-dimensional mri reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 461–470. Springer, 2021.
- [14] Ken CL Wong, Hongzhi Wang, and Tanveer Syeda-Mahmood. Fnoseg3d: resolution-robust 3d image segmentation with fourier neural operator. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2023.

- [15] Arjun D Desai, Andrew M Schmidt, Elka B Rubin, Christopher M Sandino, Marianne S Black, Valentina Mazzoli, Kathryn J Stevens, Robert Boutin, Christopher Ré, Garry E Gold, et al. Skm-tea: A dataset for accelerated mri reconstruction with dense image labels for quantitative clinical evaluation. *arXiv preprint arXiv:2203.06823*, 2022.
- [16] Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, pages 1–9, 2024.
- [17] Julius Berner, Miguel Liu-Schiaffini, Jean Kossaifi, Valentin Duruisseaux, Boris Bonev, Kamyar Azizzadenesheli, and Anima Anandkumar. Principled approaches for extending neural architectures to function spaces for operator learning. *arXiv preprint arXiv:2506.10973*, 2025.
- [18] Jiayun Wang, Yousuf Aborahama, Arya Khokhar, Yang Zhang, Chuwei Wang, Karteekeya Sastry, Julius Berner, Yilin Luo, Boris Bonev, Zongyi Li, Kamyar Azizzadenesheli, Lihong V. Wang, and Anima Anandkumar. Accelerating 3d photoacoustic computed tomography with end-to-end physics-aware neural operators, 2025.
- [19] Jeremy Ocampo, Matthew A Price, and Jason D McEwen. Scalable and equivariant spherical cnns by discrete-continuous (disco) convolutions. *arXiv preprint arXiv:2209.13603*, 2022.
- [20] Miguel Liu-Schiaffini, Julius Berner, Boris Bonev, Thorsten Kurth, Kamyar Azizzadenesheli, and Anima Anandkumar. Neural operators with localized integral and differential kernels. In *Forty-first International Conference on Machine Learning*, 2024.
- [21] Jiayun Wang, Oleksii Ostras, Masashi Sode, Bahareh Tolooshams, Zongyi Li, Kamyar Azizzadenesheli, Gianmarco Pinton, and Anima Anandkumar. Ultrasound lung aeration map via physics-aware neural operators, 2025.
- [22] Bahareh Tolooshams, Lydia Lin, Thierry Callier, Jiayun Wang, Sanvi Pal, Aditi Chandrashekar, Claire Rabut, Zongyi Li, Chase Blagden, Sumner L. Norman, Kamyar Azizzadenesheli, Charles Liu, Mikhail G. Shapiro, Richard A. Andersen, and Anima Anandkumar. Vars-fusi: Variable sampling for fast and efficient functional ultrasound imaging using neural operators. *bioRxiv*, 2025.
- [23] Scott Shaobing Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.
- [24] D.L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- [25] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations, 2021.
- [26] Jean Kossaifi, Nikola Kovachki, Zongyi Li, David Pitt, Miguel Liu-Schiaffini, Valentin Duruisseaux, Robert Joseph George, Boris Bonev, Kamyar Azizzadenesheli, Julius Berner, et al. A library for learning neural operators. *arXiv preprint arXiv:2412.10354*, 2024.

Appendix

A Discrete-Continuous Convolution Implementation

Our implementation of the 3D neural operator is built upon the discrete-continuous-convolution framework [11]. We implement 3D filter basis functions in an anonymous code repository, which will be public after acceptance. We also implement 3D local neural operator layers that construct continuous filters to perform discrete-continuous convolutions in the same repository.

We implement a 3D Morlet-wavelet-like filter basis. Each basis function is constructed by multiplying three separate 1D Fourier bases along the x , y , and z axes to create 3D wave patterns as a Fourier basis H . Higher frequency basis functions multiply higher frequency harmonics. The result is multiplied by a Hann window W . Specifically,

$$\psi_{p,m,n}(\vec{x}) = W(r) \cdot H_{p,m,n}(\vec{x})$$

where

$$W(r) = \begin{cases} \cos^2\left(\frac{\pi r}{2R_{\text{cutoff}}}\right) & \text{if } r \leq R_{\text{cutoff}} \\ 0 & \text{if } r > R_{\text{cutoff}} \end{cases}$$

and

$$H_{p,m,n}(x, y, z) = h_n(x) \cdot h_m(y) \cdot h_p(z) \quad \text{with} \quad h_n(x) = \begin{cases} \cos\left(\frac{k_n \pi x}{w}\right) & \text{if } n \text{ is even} \\ \sin\left(\frac{k_n \pi x}{w}\right) & \text{if } n \text{ is odd} \end{cases}$$

B Visualizations of the 3D Morlet Waveless Basis Filters and DISCO Kernel

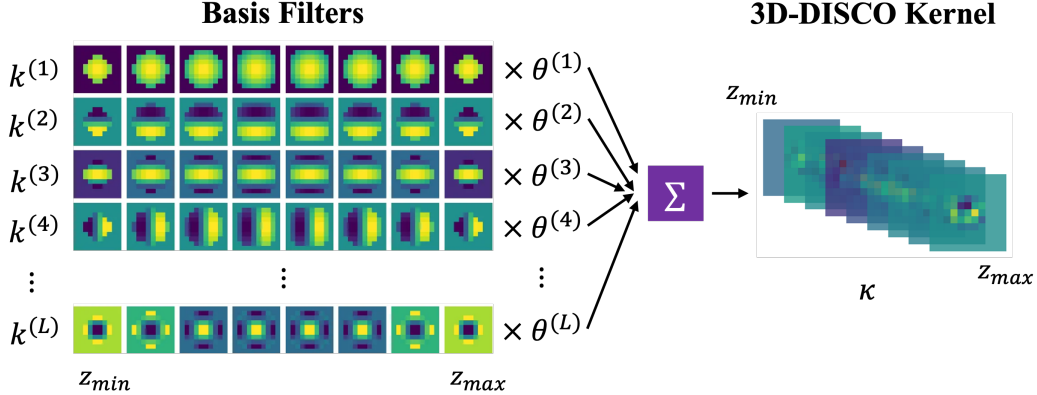


Figure 2: **Left:** first few 3D Morlet wavelet basis filter functions. Each row represents a basis function, with columns showing cross-section slices of the 3D basis function. **Right:** computed 3D DISCO kernel, as a linear combination of the basis filters. Note that this visualization is made with an arbitrarily chosen resolution. Since the filter-basis functions are continuous, we can discretize them at any desired resolution and compile the same kernel to any resolution at training or inference time.

C Related Works

A common neural operator architecture is the Fourier neural operator (FNO) [25], which applies a Fourier transform to the input, truncates high-frequency modes, and multiplies the result pointwise with a learned weight tensor. By the convolution theorem, this operation is equivalent to a global convolution. FNOseg3D [14] presents a specialized implementation of the FNO for resolution-robust 3D MRI segmentation. However, MRI segmentation and reconstruction from undersampled measurements are fundamentally different tasks. Reconstruction requires capturing intricate anatomical structures, favoring architectures that can learn local inductive biases. Because FNO truncates high frequencies, it may lose information crucial for reconstruction. In contrast, DISCO does not truncate frequencies and can learn the local inductive biases necessary for accurate reconstruction. Jatyani et al. [11] demonstrate that DISCO-based neural operators are robust in super-resolution and extended field-of-view tasks in 2D MRI reconstruction from undersampled measurements. Building on this framework, we extend DISCO to 3D Cartesian coordinates.

D Dataset and Setup

The SKM-TEA dataset [15] is a multi-coil knee volume MRI dataset, sized at 1.3 TB. Acquisitions are captured with the qDESS technique and each sample has two echoes. We choose the first echo in all of our experiments. Raw k -space data is captured at a $512 \times 512 \times 160$ resolution. In image-space, this corresponds to MR volumes of the same $512 \times 512 \times 160$ resolution. As described in the dataset source paper, voxel spacing/resolution at this configuration is $0.3125 \times 0.3125 \times 0.8$. That is, the third dimension was acquired at roughly half the resolution as the first two dimensions.

To achieve near-isotropic resolution, we preprocess the dataset by interpolating the $512 \times 512 \times 160$ dataset over the last dimension to $512 \times 512 \times 320$. We denote this the full-resolution dataset. By downscaling these samples by a factor of 2, we get the $1/2$ resolution $256 \times 256 \times 160$ dataset.

Both the baseline CNN and our neural operator (NO) model are trained on the half-resolution $256 \times 256 \times 160$ samples. Then, we perform zero-shot inference on out-of-domain full super-resolution $512 \times 512 \times 320$ samples. Both models see no prior examples of full-resolution samples before evaluation, making this a true zero-shot setup.

The implementation of our neural operator architectures leverages the `neuraloperator` library [26] (github.com/neuraloperator/neuraloperator) and `torch-harmonics` library (available at github.com/NVIDIA/torch-harmonics), both which we extended with new 3D layers and filter basis functions.

E Neural Operator for 3D MRI Reconstruction from Undersampled Measurements

Our neural operator architecture is trained across $4 \times \text{GH200}$ nodes for quicker training. Each node handles a batch size of 2, for a total batch size of 8. Each neural operator "NO" layer in Fig. 1 is a U-shaped DISCO Neural Operator, as first introduced in [11] with 4 encoder layers and 4 decoder layers. A DISCO kernel radius cutoff of 0.02 is used, with kernel shape parameter of $[6, 6, 6]$.

The baseline 3D CNN architecture closely follows the state-of-the-art for 2D MRI reconstruction, i.e. the E2E-VarNet [10], where 2D convolution layers are replaced with fully 3D convolution layers, and U-Nets are initialized with 4 encoders and 4 decoders.

Both models are trained with the Adam optimizer for up to 50 epochs. Finally, both models are trained with 0.001 learning rate. Sweeps were ran to choose these optimal hyperparameters.

F Visualizations

We visualize 3D MRI reconstruction in Fig. 3, with projections on different coordinate planes.




Figure 3: A visualization of three different viewpoints (slices along different planes) for a particular sample of the MRI dataset. The leftmost column depicts the target image, the middle column depicts the predictions of the trained neural operator, and the rightmost column depicts the output from the CNN baseline. All models were trained on 1/2 resolution, with a 4x measurement sampling rate. Inference was performed at full resolution.