



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

# Estimating Neural Orientation Distribution Fields on High Resolution Diffusion MRI Scans

Mohammed Munzer Dwedari<sup>\*1,2</sup>, William Consagra<sup>\*1</sup>, Philip Müller<sup>2</sup>, Özgün Turgut<sup>2</sup>, Daniel Rueckert<sup>2</sup>, and Yogesh Rathi<sup>1</sup>

<sup>1</sup> Psychiatry Neuroimaging Laboratory, Brigham and Women's Hospital, Harvard Medical School, Boston, USA {wconsagra,yogesh}@bwh.harvard.edu

<sup>2</sup> Technical University of Munich, Munich, Germany {munzer.dwedari,philip.j.mueller,oezguen.turgut,daniel.rueckert}@tum.de

**Abstract.** The Orientation Distribution Function (ODF) characterizes key brain microstructural properties and plays an important role in understanding brain structural connectivity. Recent works introduced Implicit Neural Representation (INR) based approaches to form a spatially aware continuous estimate of the ODF field and demonstrated promising results in key tasks of interest when compared to conventional discrete approaches. However, traditional INR methods face difficulties when scaling to large-scale images, such as modern ultra-high-resolution MRI scans, posing challenges in learning fine structures as well as inefficiencies in training and inference speed. In this work, we propose *HashEnc*, a grid-hash-encoding-based estimation of the ODF field and demonstrate its effectiveness in retaining structural and textural features. We show that HashEnc achieves a 10% enhancement in image quality while requiring 3x less computational resources than current methods. Our code can be found at <https://github.com/MunzerDw/NODF-HashEnc>.

**Keywords:** Orientation Distribution Function · Implicit Neural Representation · Diffusion MRI

## 1 Introduction

The Orientation Distribution Function (ODF) plays an important role for understanding brain structural connectivity and brain-based disorders [31]. It describes the angular probability distribution of water molecule diffusion in brain tissue [12], where water diffusion is strengthened along the direction of white matter fiber tracts. Thus, the ODF serves as an indirect characterization of the white matter fiber structure at a given voxel in the brain, which provides critical information for tractography and microstructure estimation [37].

Consagra *et al.* [5] introduced an Implicit Neural Representation (INR) framework (called NODF), using Sinusoidal Representation Networks (SIREN) [26], to model a spatially aware continuous ODF field, estimated from diffusion signal.

---

\* Authors contributed equally to this work

Their method enables resolution-agnostic estimation and uncertainty quantification of the ODFs. However, while large SIRENs can estimate the ODF field on individual slices and regions of interest, their ability to learn a continuous field for images on the scale of modern high-resolution whole brain scans [32] suffers from days-long training times [23], rendering them impractical for these applications. This issue arises because all the network weights need to be evaluated and updated during each pass. Moreover, the fine-tuning of important hyperparameters, including those for regularization and the sine frequency, becomes difficult or even computationally infeasible due to repeated network training.

To mitigate these issues, we investigate a solution, referred to as *HashEnc*, based on the grid-like local embeddings as proposed by Müller *et al.* [18] to replace SIREN in the NODF framework. The use of grid-like embeddings allows HashEnc to store local information of the training subject. As every region has designated embeddings, only the local neighborhood embeddings and the small MLP weights need to be updated during training. Therefore, the required MLP size to predict the final output signal for a coordinate becomes much smaller and thus efficient to train. We train HashEnc on a submillimeter resolution and low signal-to-noise-ratio diffusion MRI (dMRI) scan [32] and evaluate the results on highly detailed areas such as the cerebellum. While SIREN tends to over-smooth the estimated ODF, we demonstrate the capability of HashEnc to learn fine structural and textural details in significantly less training time. In summary, our contributions are:

- We propose *HashEnc*, a grid-hash-encoding-based INR that represents a “field” of ODFs in a spatially continuous manner across any ultra-high-resolution dMRI scan.
- We quantitatively and qualitatively compare HashEnc with SIREN, where HashEnc achieves a 10% enhancement in image quality while being up to 3x faster to train.
- We study the key characteristics of HashEnc through ablation studies.

## 2 Related Work

**Orientation Distribution Function.** The estimation of ODFs from diffusion signals poses a challenging inverse problem, that has mostly been tackled voxel-wise [7,15] or by incorporating neighborhood information [2,3]. Furthermore, other lines of work introduce machine learning techniques to directly estimate ODFs from diffusion signal through supervised or unsupervised training [21,19,27]. Recently, [5] utilized INRs to continuously parameterize the ODF field and derive a conditional posterior distribution for uncertainty quantification.

**Implicit Neural Representations in Medical Imaging.** INRs are increasingly utilized in computer vision and medical imaging, enabling continuous modeling of discrete data with minimal memory usage [16,14,8,24,36,20,17,11]. Their flexibility and differentiability facilitate various tasks, such as image reconstruction [34,35,27,9], segmentation [38,10,1], and registration [33,28,4], effectively

addressing issues like scarce data and lengthy acquisition times. INRs are also used for inverse imaging tasks [22,5,29] or 3D volume reconstruction from sparse 2D images [6,13].

### 3 Method

#### 3.1 Background and Notation

**Orientation Distribution Function.** The Orientation Distribution Function (ODF) at a voxel  $\mathbf{v}$ ,  $g(\mathbf{v}, \cdot)$ , describes the angular distribution of water molecule diffusion and is connected to the diffusion signal, denoted as  $f(\mathbf{v}, \cdot) \mapsto \mathbb{R}^+$ , by the Funk-Radon transform (FRT). To compute the ODF, we truncate the spherical harmonic basis [31] at a finite rank  $K$ , modeling it as:

$$g(\mathbf{v}, \mathbf{p}) = \sum_{k=1}^K c_k(\mathbf{v}) \phi_k(\mathbf{p}), \quad (1)$$

where  $\phi_k(\mathbf{p})$  are the harmonic basis functions and  $c_k(\mathbf{v})$  are the harmonic coefficients. Measurements on  $M$  spherical locations  $\mathbf{p}_1, \dots, \mathbf{p}_M$ , called gradient directions, at each of a regular set of voxel locations  $\mathbf{v}_1, \dots, \mathbf{v}_N$ , translate to a noisy Gaussian model linking observed signals to the coefficients  $c_k(\mathbf{v})$ :

$$\mathbf{y}_i := (y_{i,1}, \dots, y_{i,M}) \sim \mathcal{N}(\Phi \mathbf{G} \mathbf{c}(\mathbf{v}_i), \sigma_e^2 \mathbf{I}_M), \quad (2)$$

where  $\mathbf{c}(\mathbf{v}) = (c_1(\mathbf{v}), \dots, c_K(\mathbf{v}))^\top$ ,  $\mathbf{G} \in \mathbb{R}^{K \times K}$  is the diagonal inverse matrix of the FRT,  $\Phi \in \mathbb{R}^{M \times K}$  is the evaluation of the  $K$  real-symmetric spherical harmonic basis functions along all  $M$  gradient directions,  $\mathbf{I}_M$  is the  $M$ -dimensional identity matrix, and  $\sigma_e^2$  is the measurement error variance.

**Neural Orientation Distribution Fields.** The NODF framework introduces an implicit model to capture the spatial correlation in the ODF field through a rank  $r$  spatial basis learned via an implicit neural representation  $\xi_\theta : \mathbb{R}^3 \mapsto \mathbb{R}^r$ . This basis is used to construct the harmonic coefficient fields via a multivariate linear basis expansion  $\mathbf{c}(\mathbf{v}) = \mathbf{W} \xi_\theta(\mathbf{v})$ , with  $\mathbf{W}$  being a matrix in  $\mathbb{R}^{K \times r}$ .

Following the framework, with a matrix normal prior on  $\mathbf{W}$  inducing a Gaussian process prior on  $g(\mathbf{v}, \cdot)$ , and given the normal likelihood (2), the posterior distribution can be derived as given in [5]:

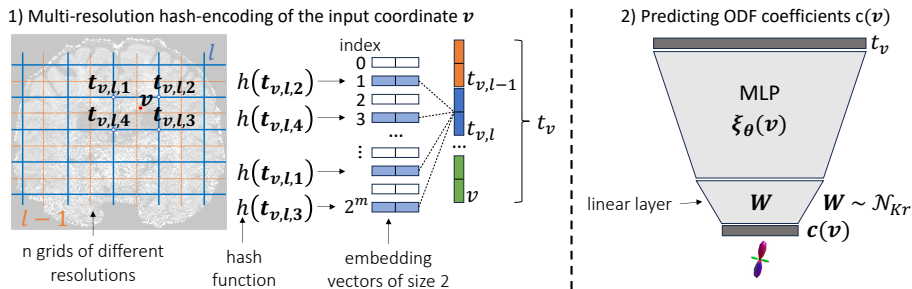
$$\text{vec}(\mathbf{W}) | \mathbf{V}, \mathbf{Y}, \theta, \gamma, \sigma_w^2, \sigma_e^2 \sim \mathcal{N}_{Kr} \left( \frac{1}{\sigma_e^2} \mathbf{A}_\theta^{-1} [\Xi_\theta^\top \otimes \Phi \mathbf{G}]^\top \text{vec}(\mathbf{Y}), \mathbf{A}_\theta^{-1} \right), \quad (3)$$

$$\mathbf{A}_\theta = \frac{1}{\sigma_e^2} \left( \frac{\sigma_e^2}{\sigma_w^2} \mathbf{I}_r \otimes \mathbf{R}_\gamma + \Xi_\theta \Xi_\theta^\top \otimes [\Phi \mathbf{G}]^\top \Phi \mathbf{G} \right) \quad (4)$$

where  $\text{vec}$  is the vectorization operator,  $\otimes$  denotes the Kronecker product, and  $\mathbf{R}_\gamma$  is the covariance matrix from a spherical Matern Gaussian process with parameters  $\gamma$ ,  $\mathbf{Y} = [\mathbf{y}_1^\top, \dots, \mathbf{y}_N^\top] \in \mathbb{R}^{M \times N}$ ,  $\Xi_\theta = [\xi_\theta^\top(\mathbf{v}_1), \dots, \xi_\theta^\top(\mathbf{v}_N)]^\top \in \mathbb{R}^{r \times N}$ ,

and  $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_N] \in \mathbb{R}^{N \times 3}$ . The unknown conditioning parameters of (3) are then estimated and plugged in for inference, i.e. point estimation and uncertainty quantification. Specifically, the network parameters  $\hat{\boldsymbol{\theta}}$  are estimated using stochastic gradient descent on a regularized variant of the negative log likelihood, where the regularization strength  $\lambda_c$  is selected using a Bayesian optimization scheme. We follow the same procedure as in [5] to estimate the remaining variance parameters  $\sigma_e^2, \sigma_w^2$ . When estimating a quantity of interest (QOI) from the ODFs for downstream tasks, e.g. fractional anisotropy or principal diffusion directions, we quantify the uncertainty in the QOI by sampling the ODF field (through the posterior (3)) and determining a confidence interval.

### 3.2 Grid-Hash-Encoding of Harmonic Coefficient Fields



**Fig. 1.** Overview of *HashEnc*. 1) Given point  $\mathbf{v}$ , for each resolution grid  $l$ , embedding vectors of the surrounding corner points are retrieved from the lookup table by hashing their grid coordinates  $t_{v,l,i}$ . Then, the corner embeddings are combined into one vector via linear interpolation. The final embedding vector  $t_v$  is obtained by concatenating the input coordinates  $\mathbf{v}$  and other grid vectors  $t_{v,l}$ . The grid is shown in 2D instead of 3D for the sake of clarity. 2)  $t_v$  is fed into a SIREN and processed by a linear layer  $W$  to output the spherical harmonic coefficients  $\mathbf{c}(\mathbf{v})$ .

Using a single SIREN MLP for spatial basis  $\boldsymbol{\xi}_\theta$  in large images leads to computational challenges. Due to the need for a large network rank  $r$ , the inversion of  $\mathbf{A}_\theta$  can become complex and unstable. In addition, gradient computation for  $\theta$  during backpropagation is slow, as it requires evaluating all parameters for every voxel. To address this, we propose adopting a grid-hash-encoding method (*HashEnc*) [18] with a much smaller MLP that trains more quickly, leveraging local embedding vectors for storing regional information.

The input of the network is a 3D coordinate  $\mathbf{v} \in \mathbb{R}^3$  and the output is the real-symmetric spherical harmonic expansion coefficients of the ODF  $\mathbf{c}(\mathbf{v}) \in \mathbb{R}^K$ , where  $K = 45$  (Figure 1). Each resolution grid  $l$  takes the 3D coordinates as input and retrieves the grid coordinates of the 8 surrounding corners. These surrounding coordinates are hashed into index values and the corresponding

embedding vectors are retrieved from a dictionary of size  $2^m$  belonging to that grid. Via linear interpolation, at each resolution grid  $l$  the embedding vectors are interpolated into one vector  $t_{v,l}$ . The vectors of all  $n$  resolution grids are then concatenated together into one vector along with the input coordinates, forming  $t_v = (t_{v,1}, \dots, t_{v,n}, \mathbf{v})$ .  $t_v$  is then passed into an MLP head (SIREN of 2x64 hidden layers) to predict the spatial basis  $\xi_{\theta}(\mathbf{v})$ , which is subsequently multiplied with the linear layer  $\mathbf{W}$  to obtain the  $K$  ODF coefficients  $\mathbf{c}(\mathbf{v})$ . Using the inverse of the Funk-Radon transform  $\mathbf{G}$ , we obtain the coefficients of the signal  $f$  expansion over the real-symmetric spherical harmonic basis. From the  $\mathbf{p}$  points on the sphere, indicating the  $M$  gradient directions, we can obtain the diffusion signals  $f(\mathbf{v}, \mathbf{p}_1), \dots, f(\mathbf{v}, \mathbf{p}_M)$ .

To illustrate the computational advantage: for inference at a single voxel, HashEnc uses 13,000 parameters (31 input, 2x 64 hidden, 45 output size), <0.1% of the total number of parameters. In contrast, SIREN requires all its parameters for each voxel. Doubling HashEnc’s capacity from 2 to 4 embedding vector size adds <2,000 parameters (59 input size), while doubling SIREN’s doubles the parameters required per voxel, leading to significant computational challenges.

## 4 Experimental Setup

**Data.** We train on the publicly available high-resolution data ( $760 \mu\text{m}^3$ ) from [32]<sup>3</sup>. The data consists of multiple scan sessions with 420 gradient directions at  $b = 1,000\text{s}/\text{mm}^2$ . We train on one of the scan session data (with 70 gradient directions) which has very low SNR. As there is no ground truth data, we consider a 6 session average (across 420 directions) as a reasonable ground truth image and apply penalized Spherical Harmonics Least Square (SHLS) from [7] to derive ground truth ODFs. The dimension of the resulting image is  $190 \times 224 \times 178 \times M$ .

**Training and Evaluation.** We compare HashEnc and SIREN, the latter being an MLP with 10 layers of 1024 units each and sine activations. SIREN is trained with a learning rate of  $1\text{e-}6$  for 10,000 epochs. We use Bayesian optimization on a single slice to select  $\lambda_c$ . While potentially sub-optimal, this approach is computationally necessary, especially for SIREN, as full volume training takes days. HashEnc employs 14 resolution grid levels, starting from resolution size 6, with a  $2^{20}$ -sized lookup table per level. Both methods are trained on an RTX 4090 GPU with  $M = 70$ ,  $M = 40$ , and  $M = 20$  gradient directions using PyTorch.

We evaluate both methods on the Feature Similarity Index (FSIM)<sup>4</sup> [39], which mimics the human perception to images and focuses on features such as edges, corners, and textures. It consists of two components: Phase Congruency PC and Gradient Magnitude (GM). PC compares feature points regardless of brightness or contrast. GM captures the image edge information by measuring the gradient magnitude of the image. This constellation is important for

<sup>3</sup> License: <http://creativecommons.org/licenses/by/4.0/>

<sup>4</sup> Calculated via: <https://pypi.org/project/image-similarity-measures/>

distinguishing the tendency of SIREN to over-smooth and HashEnc to overfit noise. FSIM scores for gray scale General Fractional Anisotropy (GFA) and RGB Diffusion Tensor (DTI) images across all sagittal, axial and coronal slices are calculated against the 6-session average, with median values reported in Table 1 and sample images provided in the supplementary material. GFA calculates the degree of anisotropy of water diffusion at each voxel, where a higher value indicates stronger anisotropy. The Diffusion Tensor Image (DTI) indicates via RGB coloring the dominant fiber orientation direction at each voxel.

Furthermore, GFA images are visualized in Figure 2. DTI images and deconvolved ODFs (using Constrained Spherical Deconvolution [30]) on a small sagittal section in the Cerebellum are shown in Figure 3. To quantify the uncertainty of each method, the posterior is sampled 250 times. Then, the voxel-wise GFA is determined for each sampled ODF field. The uncertainty on the GFA is analysed via the standard deviation to mean ratio.

## 5 Results


### 5.1 Comparison with Current Methods

**Table 1.** Median value of Feature Similarity Index (FSIM) to the 6 session average of all sagittal, axial and coronal slices. FSIM-GFA is calculated on gray scale GFA images, and FSIM-DTI is calculated on RGB DTI images. 1 means perfect similarity. Visuals are provided in the supplementary materials. HashEnc performs better in all settings except for  $M = 20$  on FSIM-DTI.

Gradient Directions ( $M$ )	Model	FSIM-GFA	FSIM-DTI
70	SIREN	0.55	0.61
	HashEnc	<b>0.66</b>	<b>0.68</b>
40	SIREN	0.62	0.63
	HashEnc	<b>0.66</b>	<b>0.66</b>
20	SIREN	0.62	<b>0.64</b>
	HashEnc	<b>0.64</b>	0.62


Compared to SIREN, HashEnc shows a higher structural similarity to the 6 session average volume in terms of GFA gray scale and DTI RGB images. This is especially apparent for  $M = 70$ , where SIREN shows worse performance to  $M = 40$  and  $M = 20$  (Table 1). One reason for the lower similarity scores of SIREN is the over-smoothing effect that SIREN shows, which creates blurry spots in the image (see Figures 2 and 3). We note that this over-smoothing effect in SIREN might be mitigated by globally tuning its hyperparameters, i.e., selecting  $\lambda_c$  by training on the whole image rather than just a slice. However, automatically doing this through Bayesian Optimization is highly computationally problematic

due to the excessively long training times. This further underscores the advantage of HashEnc’s faster training times. On the other hand, SIREN shows a robust performance as  $M$  gets lower and outperforms HashEnc on FSIM-DTI for  $M = 20$ . Further visuals are provided in the supplementary material.



**Fig. 2.** GFA reconstruction images from a sagittal cerebellum slice at different training times with  $M = 70$  gradient directions. The scale on the top right indicates the degree of anisotropy of water diffusion. Included are GFA of the training image (1 session) and the 6 session average for reference. HashEnc fits a much more detailed ODF field after significantly less training time compared to SIREN.

The dominant advantage of HashEnc are the training and inference times, allowing for much faster estimation of the ODF field and evaluation for downstream tasks. HashEnc can already produce fine-grained estimations after 100 epochs of training, which is not the case for SIREN (Figure 2). After 1000 training epochs, HashEnc shows similar but more detailed results whereas SIREN is yet to fit fine-grained regions. The time efficiency comes from using multi-resolution grid embeddings that store local volume information at various resolution levels, allowing for a smaller MLP head and the use of a larger learning rate. As for inference speed, HashEnc requires about 4 seconds to infer the ODF coefficients on the whole brain while SIREN requires 173 seconds (tested on CPU). When SIREN is trained for 10,000 epochs, visually, both methods produce results with comparable performance but different characteristics (Figure 3). SIREN produces overly-smooth estimates of the ODF field, resulting in slightly blurry but less noisy look, as demonstrated in Figure 3. This can be observed specifically in areas with high contrast, such as between white and gray matter. The gradual transition of SIREN, which is visible on the borders of the blue fiber tracts in the ODF images for  $M = 70$  and  $M = 40$ , is also reflected in the low FSIM scores in Table 1. HashEnc on the other hand learns individual details better in fine-grained regions as can be seen in the width of the blue fiber tracts in the DTI images of  $M = 70$  and  $M = 40$ . However, it tends to overfit to noise easier, such as in the cases of  $M = 40$  and  $M = 20$ . For these (lower) gradient directions, SIREN is more robust to noise. As for uncertainty quan-



**Fig. 3.** Qualitative reconstruction examples on a small sagittal Cerebellum section, showcasing DTI images, deconvolved ODFs, and GFA uncertainty. The scale on the bottom left indicates variability in the GFA of the ODF samples. SIREN and HashEnc are trained for 10,000 and 3,000 epochs, respectively, and compared for  $M = [70, 40, 20]$  gradient directions. SIREN tends to over-smooth the ODF field, particularly at  $M = 70$ , but is more robust with fewer gradient directions. HashEnc matches the structural and textural details of the 6-session average better and exhibits less uncertainty.

tification, HashEnc shows a consistent and lower uncertainty, whereas SIREN exhibits larger uncertainty especially in the border regions between white and gray matter at  $M = 40$  and  $M = 20$  gradients. Additionally, HashEnc computes posterior (3) means and variances faster due to its smaller  $\mathbf{W}$  matrix size.

## 5.2 Ablation Studies

**How do grid resolution levels and lookup table size affect the characteristics of the ODF field?** Different resolution levels  $n$  (12–14) and lookup table sizes  $2^m$  ( $m = 19$  and  $m = 20$ ) are analysed. A longer and thinner estimation of the fiber tracts can be observed for  $m = 20$  (see Figure 2 in supplementary material). On the other hand, having a higher number of resolution levels shows finer but noisier details, whereas for  $n = 12$  resolution levels the image looks smoother with some information lost (e.g. the tip of the thin blue fiber tracts). Quantitatively,  $n = 14$  resolution levels and  $2^{20}$  lookup table size shows the highest feature similarity score to the 6 session average.

**How does the MLP head affect HashEnc?** In this experiment, we try three types of MLP heads, including SIREN [26], WIRE [25], and ReLU [18]. Our experiments show that there is no significant difference both visually (DTI images) and quantitatively (FSIM score) (see Figure 3 in supplementary material).



## 6 Discussion and Conclusion

In this work, we propose *HashEnc*, a solution based on grid-like local embeddings and replace SIREN in the NODF framework of [5] to estimate the ODF field on high-resolution diffusion MRI scans. While SIREN suffers from over-smoothing high contrast regions, HashEnc learns better fine-grained structural features with significantly less training time, making it feasible for downstream tasks as reflected in the Feature Similarity Index (FSIM). We want to acknowledge that HashEnc is limited in its ability to adapt to different noise levels in the image. Our training image contains varying levels of noise across different regions, which HashEnc does not consider, as the number of multi-resolution grids is fixed for all regions and  $\sigma_c^2$  is assumed to be spatially constant. We encourage further research to address this limitation in future studies.

**Acknowledgments.** The authors would like to acknowledge the following grants which supported this work: R01MH119222, R01MH125860, R01MH132610, R01NS125307 and R01NS125781

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

## References

1. Barrowclough, O.J., et al.: Binary segmentation of medical images using implicit spline representations and deep learning. *Computer Aided Geometric Design* **85**, 101972 (2021)
2. Becker, S.M.A., et al.: Position-orientation adaptive smoothing of diffusion weighted magnetic resonance data (poas). *Medical image analysis* **16**(6), 1142–1155 (2012)
3. Becker, S.M.A., et al.: Adaptive smoothing of multi-shell diffusion weighted magnetic resonance data by mspoas. *Neuroimage* **95**, 90–105 (2014)
4. Byra, M., et al.: Exploring the performance of implicit neural representations for brain image registration. *Scientific Reports* **13**(1), 17334 (2023)
5. Consagra, W., Ning, L., Rathi, Y.: Neural orientation distribution fields for estimation and uncertainty quantification in diffusion mri. *Medical Image Analysis* p. 103105 (2024)
6. Corona-Figueroa, A., et al.: Mednerf: Medical neural radiance fields for reconstructing 3d-aware ct-projections from a single x-ray. In: 2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC). IEEE (2022)
7. Descoteaux, M., et al.: Regularized, fast, and robust analytical q-ball imaging. *Magnetic Resonance in Medicine* **58**(3), 497–510 (2007)
8. Esmailzadeh, S., et al.: Meshfreeflownet: A physics-constrained deep continuous space-time super-resolution framework. In: SC20: International Conference for High Performance Computing, Networking, Storage and Analysis. IEEE (2020)
9. Ewert, C., Kügler, D., Stirnberg, R., Koch, A., Yendiki, A., Reuter, M.: Geometric deep learning for diffusion mri signal reconstruction with continuous samplings (discus). *Imaging Neuroscience* **2**, 1–18 (2024)

10. Gu, J., Tian, F., Oh, I.S.: Retinal vessel segmentation based on self-distillation and implicit neural representation. *Applied Intelligence* **53**(12), 15027–15044 (2023)
11. Hendriks, T., Vilanova, A., Chamberland, M.: Neural spherical harmonics for structurally coherent continuous representation of diffusion mri signal. In: *International Workshop on Computational Diffusion MRI*. pp. 1–12. Springer (2023)
12. Karimi, D., et al.: Learning to estimate the fiber orientation distribution function from diffusion-weighted mri. *NeuroImage* **239**, 118316 (2021)
13. Maas, K.W., et al.: Nerf for 3d reconstruction from x-ray angiography: Possibilities and limitations. In: *VCBM 2023: Eurographics Workshop on Visual Computing for Biology and Medicine*. Eurographics Association (2023)
14. Mescheder, L., et al.: Occupancy networks: Learning 3d reconstruction in function space. In: *Proceedings of the IEEE/CVF conference on CVPR* (2019)
15. Michailovich, O., Rathi, Y.: On approximation of orientation distributions by means of spherical ridgelets. *IEEE Transactions on Image Processing* **19**(2), 461–477 (2009)
16. Mildenhall, B., et al.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
17. Molaei, A., et al.: Implicit neural representation in medical imaging: A comparative survey. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023)
18. Müller, T., et al.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)* **41**(4), 1–15 (2022)
19. Nath, V., et al.: Deep learning reveals untapped information for local white-matter fiber reconstruction in diffusion-weighted mri. *Magnetic resonance imaging* **62**, 220–227 (2019)
20. Park, J.J., et al.: Deepsdf: Learning continuous signed distance functions for shape representation. In: *Proceedings of the IEEE/CVF conference on CVPR* (2019)
21. Patel, K., Groeschel, S., Schultz, T.: Better fiber odfs from suboptimal data with autoencoder based regularization. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer International Publishing (2018)
22. Reed, A.W., et al.: Dynamic ct reconstruction from limited views with implicit neural representations and parametric motion fields. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021)
23. Reiser, C., et al.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021)
24. Saito, S., et al.: Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In: *Proceedings of the IEEE/CVF international conference on computer vision* (2019)
25. Saragadam, V., et al.: Wire: Wavelet implicit neural representations. In: *Proceedings of the IEEE/CVF Conference on CVPR* (2023)
26. Sitzmann, V., et al.: Implicit neural representations with periodic activation functions. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 7462–7473 (2020)
27. Spears, T., et al.: Learning spatially-continuous fiber orientation functions (2023)
28. Sun, S., et al.: Mirnf: Medical image registration via neural fields. *arXiv preprint arXiv:2206.03111* (2022)
29. Sun, Y., et al.: Coil: Coordinate-based internal learning for imaging inverse problems. *arXiv preprint arXiv:2102.05181* (2021)

30. Tournier, J.D., et al.: Resolving crossing fibres using constrained spherical deconvolution: validation using diffusion-weighted imaging phantom data. *Neuroimage* **42**(2), 617–625 (2008)
31. Tuch, D.S.: Q-ball imaging. *Magnetic Resonance in Medicine* **52**(6), 1358–1372 (2004)
32. Wang, F., et al.: In vivo human whole-brain connectom diffusion mri dataset at 760  $\mu\text{m}$  isotropic resolution. *Scientific data* **8**(1), 122 (2021)
33. Wolterink, J.M., Zwienenberg, J.C., Brune, C.: Implicit neural representations for deformable image registration. In: *International Conference on Medical Imaging with Deep Learning*. PMLR (2022)
34. Wu, Q., et al.: An arbitrary scale super-resolution approach for 3d mr images via implicit neural representation. *IEEE Journal of Biomedical and Health Informatics* **27**(2), 1004–1015 (2022)
35. Xu, J., et al.: Nesvor: Implicit neural representation for slice-to-volume reconstruction in mri. *IEEE Transactions on Medical Imaging* (2023)
36. Yu, A., et al.: pixelnerf: Neural radiance fields from one or few images. In: *Proceedings of the IEEE/CVF Conference on CVPR* (2021)
37. Zhang, F., et al.: Quantitative mapping of the brain’s structural connectivity using diffusion mri tractography: A review. *Neuroimage* **249**, 118870 (2022)
38. Zhang, H., et al.: Nerd: Neural representation of distribution for medical image segmentation. *arXiv preprint arXiv:2103.04020* (2021)
39. Zhang, L., et al.: Fsim: A feature similarity index for image quality assessment. *IEEE Transactions on Image Processing* **20**(8), 2378–2386 (2011)