

# Neural Spherical Harmonics for Structurally Coherent Continuous Representation of Diffusion MRI Signal

**Citation for published version (APA):**

Hendriks, T., Vilanova, A., & Chamberland, M. (2024). Neural Spherical Harmonics for Structurally Coherent Continuous Representation of Diffusion MRI Signal. In M. Karaman, R. Mito, E. Powell, F. Rheault, & S. Winzeck (Eds.), *Computational Diffusion MRI: 14th International Workshop, CDMRI 2023, Held in Conjunction with MICCAI 2023, Vancouver, BC, Canada, October 8, 2023, Proceedings* (pp. 1-12). (Lecture Notes in Computer Science (LNCS); Vol. 14328). Springer. [https://doi.org/10.1007/978-3-031-47292-3\\_1](https://doi.org/10.1007/978-3-031-47292-3_1)

**Document license:**

TAVERNE

**DOI:**

[10.1007/978-3-031-47292-3\\_1](https://doi.org/10.1007/978-3-031-47292-3_1)

**Document status and date:**

Published: 07/02/2024

**Document Version:**

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

**Please check the document version of this publication:**

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

**General rights**

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

[www.tue.nl/taverne](http://www.tue.nl/taverne)

**Take down policy**




If you believe that this document breaches copyright please contact us at:

[openaccess@tue.nl](mailto:openaccess@tue.nl)

providing details and we will investigate your claim.



# Neural Spherical Harmonics for Structurally Coherent Continuous Representation of Diffusion MRI Signal

Tom Hendriks<sup>(✉)</sup>, Anna Vilanova, and Maxime Chamberland

Department of Computer Science and Mathematics, Eindhoven University  
of Technology, Groene Loper 5, 5612 AP Eindhoven, The Netherlands  
`t.hendriks@tue.nl`

**Abstract.** We present a novel way to model diffusion magnetic resonance imaging (dMRI) datasets, that benefits from the structural coherence of the human brain while only using data from a single subject. Current methods model the dMRI signal in individual voxels, disregarding the intervoxel coherence that is present. We use a neural network to parameterize a spherical harmonics series (NeSH) to represent the dMRI signal of a single subject from the Human Connectome Project dataset, continuous in both the angular and spatial domain. The reconstructed dMRI signal using this method shows a more structurally coherent representation of the data. Noise in gradient images is removed and the fiber orientation distribution functions show a smooth change in direction along a fiber tract. We showcase how the reconstruction can be used to calculate mean diffusivity, fractional anisotropy, and total apparent fiber density. These results can be achieved with a single model architecture, tuning only one hyperparameter. In this paper we also demonstrate how upsampling in both the angular and spatial domain yields reconstructions that are on par or better than existing methods.

**Keywords:** Diffusion MRI · Implicit Neural Representation · Spherical Harmonics

## 1 Introduction

The human brain is a highly structured organ. With the introduction of diffusion magnetic resonance imaging (dMRI) in vivo study of the structure of the brain became a possibility. The spatially coherent structures in the brain imply that spatial coherence should be present when modeling dMRI data. Diffusion tensor imaging (DTI) [3] fits a tensor for every voxel of the volume describing the diffusion in three primary directions. Constrained spherical deconvolution (CSD) [14] can describe the orientation and relative size of fiber bundles using fiber orientation distribution functions (fODFs). These are examples of methods that model the fiber orientation in every voxel independently, disregarding any

intervoxel coherence. Interpolating correctly between voxels using classical interpolation methods (e.g. cubic interpolation) is, therefore, difficult and susceptible to noise, and can discard anatomical details. Interpolation in the angular domain has proven to be a difficult task as well, as highlighted by recent challenges in the computational dMRI community [4, 9, 13]. Machine learning approaches for upsampling in both the angular and spatial domains are a promising avenue [1]. However, these methods often rely on a strong prior obtained by training on large amount of data. This is problematic when training data is scarce or if the model is applied to data inherently different from the data it was trained on (e.g. pathological data). Ideally, a continuous and structurally coherent model should be derived at the individual level (i.e.,  $n = 1$ ).

Neural radiance field (NeRF) [8] models have shown to be extremely effective at creating continuous 3-dimensional representations, known as implicit neural representations, of scenes given a limited number of 2-dimensional input images taken from limited angles. NeRF overfits a multi-layer perceptron to essentially capture a given scene in its parameters. Unseen angles can then be sampled from this network. This concept could translate well to dMRI, as we are trying to create a complete representation from an incompletely sampled angular domain. The difference with dMRI is that every angle in a dMRI-acquisition produces a complete 3-dimensional volume of data.

In this work, we propose to use a NeRF-like model to create a model of the dMRI data of a single subject that utilizes the structural coherence of the brain, while providing continuity in both the angular and spatial domain. We evaluate the resulting model in a number of downstream tasks, such as calculating microstructural metrics, and fODF estimation. We also demonstrate how the model can be used to upsample dMRI data in both the angular and spatial domain.

## 2 Methods and Experiments

### 2.1 Data

We sourced data from a single participant from the preprocessed Human Connectome Project dataset [18] consisting of 18  $b = 0$  s/mm<sup>2</sup> volumes, 90  $b = 1000$  s/mm<sup>2</sup> volumes, 90  $b = 3000$  s/mm<sup>2</sup> volumes, with 1.25 mm isotropic voxels.

### 2.2 Model

The neural spherical harmonics model (NeSH) is an adaptation from SH-NeRF [19]. NeSH outputs an approximation  $\hat{S}(x, y, z, \mathbf{b})$  for a diffusion signal  $S(x, y, z, \mathbf{b})$ . An input pair  $i \in I$  consists of a voxel midpoint coordinate  $(x, y, z)$  and a gradient direction vectors  $\mathbf{b}$ , where  $I$  is a set of all possible coordinate-direction pairs.  $I$  has size  $N = n_c \times n_d$  with  $n_c$  being the number of coordinates and  $n_d$  the number of directions. The input coordinates are scaled to lie in

$[-1, 1]$  and are positionally encoded using the generalization of the NeRF positional encoding [12] into input vector  $\mathbf{x}$ . Direction vector  $\mathbf{b}$  is converted into the corresponding polar angles  $\theta$  (azimuth,  $[0, 2\pi)$ ) and  $\phi$  (elevation,  $[0, \pi]$ ).

A simple multi-layer perceptron (MLP) maps  $\mathbf{x}$  into a coefficient vector  $\mathbf{k}$  that parameterizes a spherical harmonics (SH) series.

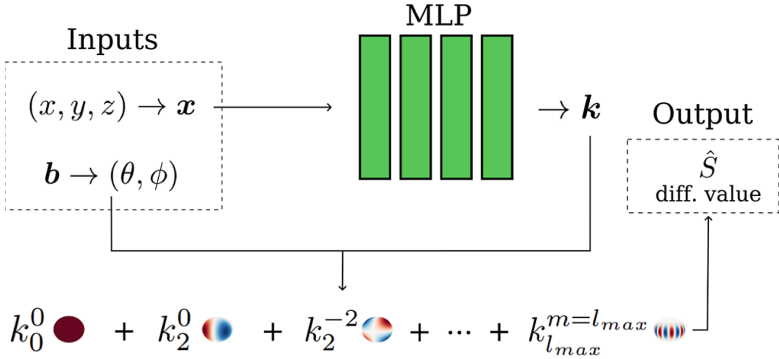
$$M_{\Psi} : \mathbf{x} \rightarrow \mathbf{k} \quad (1)$$

$$\mathbf{k} = (k_l^m)_{l:0 \leq l \leq l_{max}, m: -l \leq m \leq l} \quad (2)$$

where  $k_l^m$  is the coefficient for the SH component of degree  $l$  and order  $m$ ,  $l_{max}$  is the maximum degree of the SH-series, and  $m$  the order. For a given  $i$  we can now obtain an estimation of the dMRI signal:

$$\hat{S}_i = \hat{S}(x_i, y_i, z_i, \mathbf{b}_i) = \sum_{(k_l^m)^i \in \mathbf{k}_i} (k_l^m)^i Y_l^m(\theta, \phi) \quad (3)$$

where  $Y_l^m(\theta, \phi)$  is the SH component of degree  $l$  and order  $m$  for azimuth  $\theta$  and elevation  $\phi$  obtained from  $\mathbf{b}_i$ ,  $\mathbf{k}_i$  is the coefficient vector given by (1) for input  $i$ , and  $(k_l^m)^i$  is the coefficient of  $Y_l^m$ . The dMRI signal is reconstructed with a simplified real basis SH series, using only odd-numbered degrees. Different methods of simplifying the SH series exist [5, 14]; in this paper, the method described in MRtrix3 is used [17]. The full model is shown in Fig. 1.



**Fig. 1.** A schematic representation of the Neural spherical harmonics (NeSH) model. Inputs coordinates are spatially encoded into  $(x)$ , directional vector  $\mathbf{b}$  is converted to  $\theta$  and  $\phi$ . Vector  $\mathbf{x}$  is passed through the multi-layer perceptron (MLP) to produce  $\mathbf{k}$ , which parameterizes the spherical harmonics series. This is sampled in direction  $\mathbf{b}$  to produce the final output  $\hat{S}$ .

We calculate the loss as an average over all inputs for the smooth L1 loss [6] between the value of  $\hat{S}_i$ , and the dMRI signal  $S_i$  defined as the dMRI signal measured at  $(x_i, y_i, z_i)$  in the direction of  $\mathbf{b}_i$ . Unregularized, NeSH could be susceptible to overfitting on noise, if the maximum degree of the SH series is larger than necessary to model the diffusion data in a given voxel. An L1 regularization

term is added as an incentive to minimize unnecessary coefficients. The resulting loss function is:

$$L = \frac{1}{N} \sum_{i \in I} \left( \text{smooth}_{L1}(S_i - \hat{S}_i) + \lambda \left( \sum_{k_l^m \in \mathbf{k}_i} |(k_l^m)^i| \right) \right) \quad (4)$$

where  $|(k_l^m)^i|$  is the absolute value of the coefficient. The loss is used to update the MLP parameters  $\Psi$ .

To reconstruct images from the trained model, a set  $C$  of  $(x, y, z)$  coordinates is generated at the desired spatial resolution, as well as a set  $B$  of directions in the desired angular resolution. A dMRI dataset is reconstructed by first positionally encoding, and mapping every coordinate  $\mathbf{c} \in C$  to  $\mathbf{k}_c$  using (1), and then sampling the SH-series parameterized with  $\mathbf{k}_c$  for every direction  $\mathbf{b} \in B$ . Effectively this applies (3) to every coordinate-direction pair, but only calculates  $\mathbf{k}_c$  once for every input coordinate.

The model has the following hyperparameters:  $l_{max}$  sets the maximum degree of the SH,  $l_{pos}$  sets the number of positional encodings,  $\sigma$  scales the positional encoding,  $n_{layers}$  sets the number of layers in the MLP,  $hidden\_dim$  sets the number of neurons in each layer,  $lr$  is the learning rate,  $\lambda$  scales the L1 regularization.

### 2.3 Implementation

The model is implemented in python version 3.9.16, with pytorch version 2.0.0. MRtrix3 version 3.0.4 is used to calculate DTI metrics and fODFs, and to visualize results. Scilpy<sup>1</sup> version 1.5 is used (with python version 3.10.10) to calculate fODF based metrics, and to create interpolated spherical functions. The BET of FSL version 6.0.6.4, is used for brain mask segmentation.

### 2.4 Experiments

**Reconstruction and Angular Upsampling of the dMRI Signal.** To assess if the proposed model can reproduce the original data, NeSH is fit on 30 gradient directions of the  $b = 1000$  s/mm<sup>2</sup> shell. A grid-search is performed over the hyperparameters. Visual inspection of the gradient images, as well as DTI metrics and fODF glyphs, determine which settings produce the best results. Then, to assess if these settings can be applied to a different set of gradient directions, the settings found in part one are used to fit the model on 90, 60, 45, 30, 15, 10 and 3 gradient directions for both the  $b = 1000$  s/mm<sup>2</sup> and  $b = 3000$  s/mm<sup>2</sup> dMRI acquisitions. As a comparison, spherical harmonics interpolation (SHI) [5] is fit on the same number of gradients. The root mean squared error (RMSE) is calculated for each of the models between the input gradient images and the reconstructed gradient images it produces:

$$\sqrt{\frac{1}{WHD|B|} \sum_{x=1}^W \sum_{y=1}^H \sum_{z=1}^D \sum_{\mathbf{b} \in B} (S(x, y, z, \mathbf{b}) - \hat{S}(x, y, z, \mathbf{b}))^2} \quad (5)$$

<sup>1</sup> <https://github.com/scilus/scilpy>.

where  $W$ ,  $H$ , and  $D$  are the width, height, and depth of the image,  $B$  is the set of gradient directions with size  $|B|$ ,  $S(x, y, z, \mathbf{b})$  is the measured signal, and  $\hat{S}(x, y, z, \mathbf{b})$  is the reconstructed signal at location  $x, y, z$  for gradient direction  $\mathbf{b}$ . Finally the capabilities of the model to upsample in the angular domain are assessed. The resulting models from the second part are sampled in all 90 gradient directions. The RMSE is calculated between the 90 original gradient images and the 90 reconstructed gradient images using (5). In all experiments the RMSE is only calculated within a brain mask.

**Spatial Upsampling.** The data modeled with NeSH can be sampled in any spatial resolution. This experiment assesses the quality of the data when upsampled in spatial domain. The HCP dataset is downsampled from 1.25 mm to 2.5 mm isotropic voxels. NeSH is fit on the downsampled dataset using 90 gradient directions, and then sampled at 1.25 mm isotropic resolution. The downsampled dataset is also upsampled to the original 1.25 mm isotropic resolution using cubic interpolation. For the resulting datasets a color encoded FA map is calculated and visualized to compare the results.

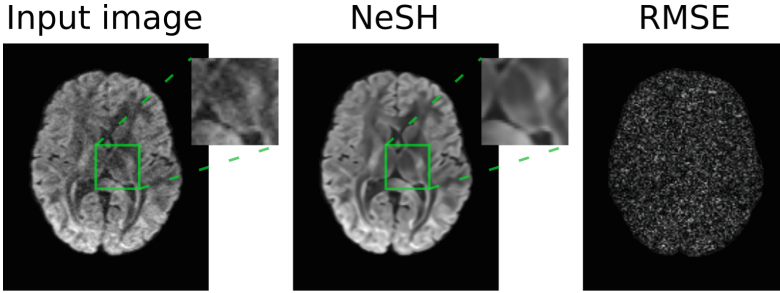
**DTI and fODF Metrics.** In this experiment we assess if the data modeled with NeSH can be used to produce three common dMRI microstructural metrics. Two DTI metrics: mean diffusivity (MD) and fractional anisotropy (FA), and one fODF metric: total apparent fiber density (AFD, [10]). The metrics are calculated for 90 gradients, 90 gradients reconstructed by NeSH fit on 90 gradients, 30 gradients, 90 gradients reconstructed by NeSH fit on 30 gradients, and 90 gradients reconstructed by SHI fit on 30 gradients. The  $b = 1000$  s/mm<sup>2</sup> shell was used for the DTI metrics, and the  $b = 3000$  s/mm<sup>2</sup> shell for AFD. The three measures are compared to the ones obtained from the full 90 gradients set by computing and visualizing a difference map.

**fODF Estimation.** This experiment is used to assess if fODFs can be generated from data modeled with NeSH. The same datasets as in the previous experiment are used. A response function is first extracted from the dMRI acquisitions using the single shell implementation of the algorithm by Tournier [15]. Secondly, the fODFs are calculated using single shell CSD [14]. For all five datasets, the  $b = 3000$  s/mm<sup>2</sup> shell is used. Results are visualized by showing fODF glyphs.

### 3 Results

**Reconstruction and Angular Upsampling of dMRI Signal.** The grid-search over the parameters resulted in the following hyperparameter settings:  $l_{max} = 8$  (for models trained  $\leq 10$  gradient directions  $l_{max} = 2$ ),  $l_{pos} = 12$  resulting in an input size of 75 (12 sine and cosine encodings for each dimension + raw coordinates),  $\sigma = 4$ ,  $n_{layers} = 4$ ,  $hidden\_dim = 2048$ ,  $lr = 10 \times 10^{-5}$ ,  $\lambda = 10 \times 10^{-6}$ . The Adam optimizer was used with default settings. The model

is trained for 5 epochs with a batch size of 1000. Figure 2 shows a slice of dMRI data for a single gradient direction, the output generated by NeSH, as well as the mean squared error (RMSE) between the two images. NeSH produces a smoother image, removing the noise from the input image. The noise appears to be randomly distributed, without anatomical residuals.



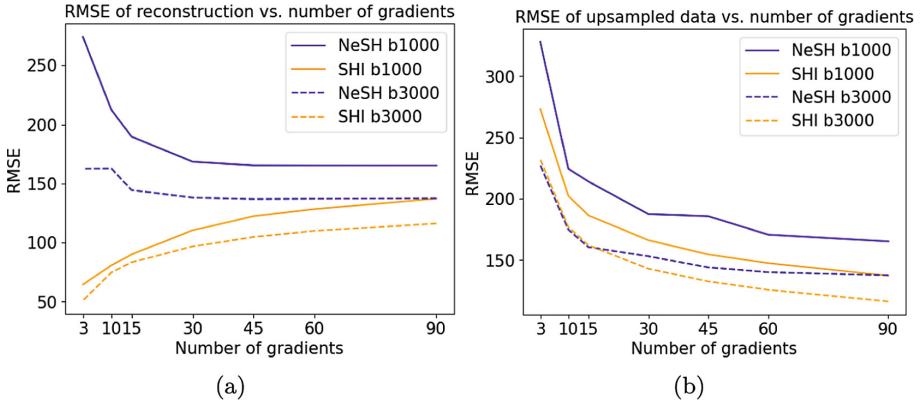
**Fig. 2.** A single axial slice of dMRI data, from a single gradient direction shown as baseline (column one) and as reconstructed by NeSH (column two). The root mean squared error (RMSE) between the two images is shown in column three. A two time magnification of the area in the green boxes shows the denoising effect more clearly. (Color figure online)

The comparisons of reconstruction error between NeSH and SHI are shown in Fig. 3a. NeSH has a higher RMSE when reconstructing a lower number of gradients, which lowers with an increasing number of gradients. SHI has a lower RMSE when reconstructing a lower number of gradients, which increases with an increasing number of gradients. SHI has a consistently lower RMSE compared to NeSH on both  $b = 1000$  s/mm<sup>2</sup> and  $b = 3000$  s/mm<sup>2</sup>.

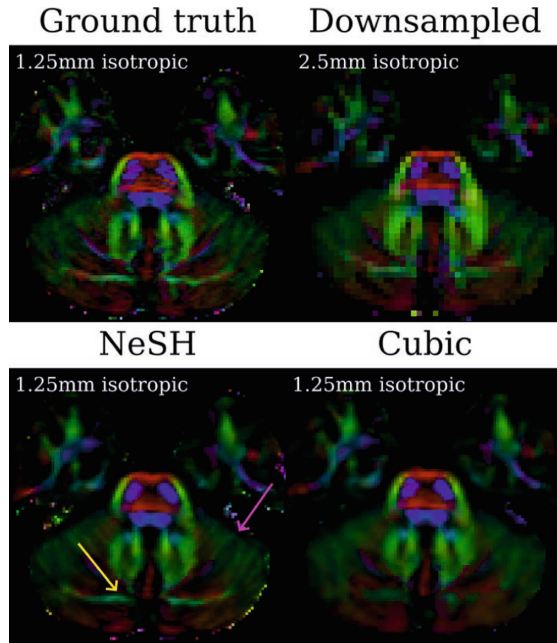
The comparisons of upsampling error between NeSH and SHI are shown in Fig. 3b. For both NeSH and SHI the RMSE of the upsampled data lowers when the model is fit on more gradient directions. SHI has a lower RMSE for all gradient subsets for  $b = 1000$  s/mm<sup>2</sup>, and for all subsets with more than 15 gradients for  $b = 3000$  s/mm<sup>2</sup>.

**Spatial Upsampling.** Figure 4 shows the color encoded FA maps for this experiment. The dataset reconstructed by NeSH fit on 90 gradients is able to reconstruct details in the cerebellar cortex and cerebellar white matter that are lost in cubic interpolation. Finer-grained details of the 1.25 mm isotropic voxel data are lost in both upsampling methods.

**DTI Metrics and fODF Metrics.** Figure 5 shows the results of this experiment. The DTI metrics (MD, FA) show low error when downsampling, and in NeSH and SHI reconstructions on both 90 and 30 gradients, when compared to

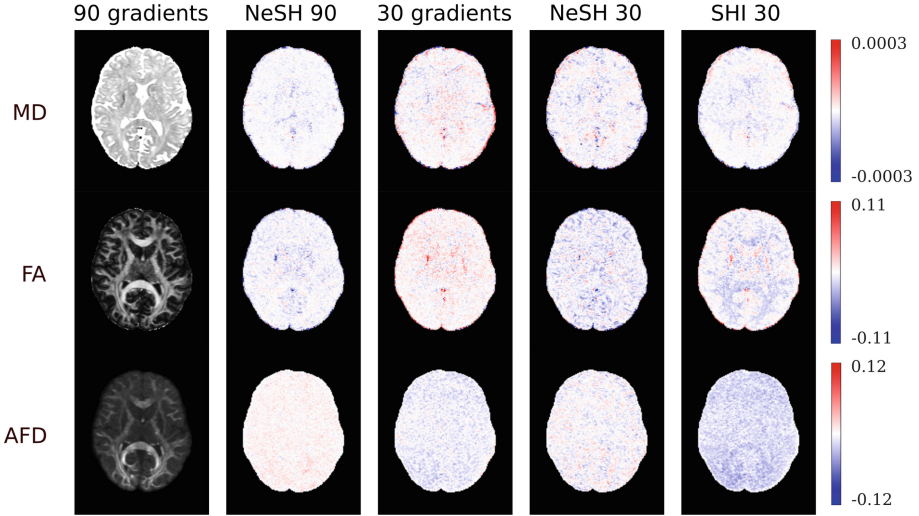


**Fig. 3.** Root mean squared error (RMSE) of the reconstructed dMRI volumes by NeSH and spherical harmonics interpolation (SHI) compared to the gradient images used to fit the model (a) and compared to the full set of 90 gradient images (b), for both  $b = 1000$  s/mm<sup>2</sup> and  $b = 3000$  s/mm<sup>2</sup>.



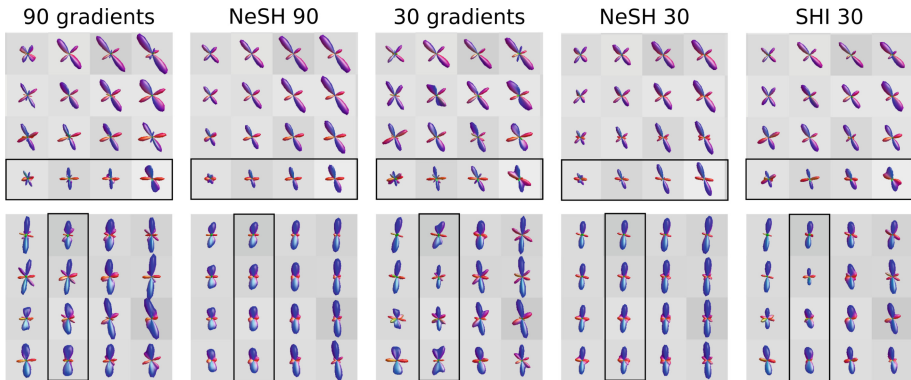
**Fig. 4.** A slice of color encoded FA maps at the cerebellar level, generated for different datasets. Clockwise starting top left: the 1.25 mm isotropic original image (ground truth), the downsampled 2.5 mm isotropic image, 1.25 mm isotropic image upsampled using cubic interpolation, 1.25 mm isotropic image upsampled using NeSH. The arrows show two areas where NeSH is able to reconstruct details of the cerebellum which are lost in cubic interpolation.

the metrics calculated on 90 gradients. Downsampling and SHI both over- and underestimate the metric, with most errors located in the white matter areas. NeSH more frequently underestimates the metrics and the errors are located more in the grey matter areas. The AFD is overestimated by NeSH reconstructions, and underestimated when downsampling or reconstructing using SHI. In AFD the distribution of error is similar for all methods, but downsampling and SHI show an underestimation of the AFD, while NeSH shows errors in both directions.



**Fig. 5.** Visualization of mean diffusivity (MD), fractional anisotropy (FA) and total apparent fiber density (AFD). In the first column the metrics are shown as a map for a single axial slice of the volume when calculated on the full set of 90 gradient images. The remaining columns show the difference map for the other datasets when compared to the 90 gradient images. Blue signifies negative difference, red signifies positive difference. (Color figure online)

**fODF Estimation.** The visualization of a group of fODFs in the centrum semi-ovale and the descending part of the CST can be seen in Fig. 6. The glyphs created from the NeSH reconstructions show a smooth, structurally-coherent change, while maintaining the important information, i.e. the crossing of fibers in the centrum semi-ovale. In presence of a big fiber tract such as the CST, the NeSH reconstruction shows a decrease in amplitude in other directions. In other methods, the fODFs exhibit more noise and less alignment between voxels, while the peaks appear sharper.



**Fig. 6.** Magnified coronal view showing the fiber orientation distributions glyphs for the different datasets on a background of a T1-weighted image. First row: centrum semi-ovale, highlighted region shows the increased intervoxel consistency in NeSH modelled data. Second row: descending part of the corticospinal tract, highlighted region shows increased intervoxel consistency in NeSH modelled data, as well as a decrease in the size of the crossing fibers.

## 4 Discussion

We have introduced a novel method to model only a single acquisition of dMRI data using a neural representation of spherical harmonics, called NeSH. We show that dMRI data reconstructed by NeSH appears to be denoised compared to the original data (Fig. 2). We hypothesise that the model is able to capture the continuous structures of the brain, but not the erratic nature of the noise. The RMSE lies consistently higher for both the reconstruction and upsampling compared to SHI (Fig. 3), which could partly be explained by the removal of noise. For the reconstruction of the input gradients, NeSH performs worse with a decreasing number of input gradients. Possibly this indicates that to find a good representation NeSH needs a minimum number of gradients, which appears to be around 15. Both NeSH and SHI show an increase in RMSE when upsampling to 90 gradients from a decreasing number of gradients, as is to be expected.

We also show that NeSH can also be used to upsample in the spatial domain. Figure 4 shows that NeSH is able to reconstruct details that are not clearly visible in the 2.5 mm isotropic data, but are present in the 1.25 mm ground truth. This strengthens the hypothesis that using multiple gradient directions, NeSH can model a continuous representation of the dMRI data. While the achievable level of detail is lower than achieved by Alexander et al. [2], it does not rely on a prior learned from a large dataset.

Furthermore, we show that the reconstructed dMRI volumes can be used to calculate MD, FA, and AFD. Compared to the metrics calculated using 90 gradient direction, NeSH differs mostly in the gray matter areas, while down-sampled volumes and SHI reconstructed volumes have differences in the white matter areas. This supports our hypothesis that NeSH benefits from the struc-

tural continuity of the fiber bundles to model the data. As the white matter areas are usually the areas of interest, this could be seen as a benefit of NeSH. The increased brightness in the posterior commissure and surrounding tissue can be explained by the lack of bias field correction in pre-processing.

Finally, we show that fODFs generated from NeSH reconstructions have a smooth change in fiber directions between voxels. This is also supportive of the structural continuity hypothesis. The 90 gradient, 30 gradient, and SHI reconstructed FODs show a more erratic pattern, in which the FODs are less aligned overall. The decrease in size of the crossing fibers in the descending part of the CST shows that NeSH prioritizes the major bundle in this area. Some information on possible smaller bundles is now lost, however, which is something that should be looked into in future versions.

**Limitations.** The lack of a gold-standard in dMRI complicates the interpretation of the experiments. The denoising effect shown in Fig. 2 is an example. We cannot be certain if the representation modelled by NeSH is a more realistic one than the more noisy representation or just a smoother one. In the last experiment NeSH consistently shows lower values in grey matter areas. The fODFs in the grey matter area correspondingly show a less peaked, smaller amplitude. Compared to existing techniques this can be interpreted as an error, but anatomically it makes sense as there are no large fiber bundles in the grey matter. Additionally, with no ground truth data, it is difficult to assess how good the representation outside of the voxels actually is. Synthetic datasets with known ground truths can provide a better idea. Furthermore, both the architecture and positional encodings used in this paper are simple. Many developments in the field NeRF have taken place since [12, 20]. Architectural and methodological changes to NeSH could lead to further improvement. Finally, we choose to model the dMRI signal directly through an SH-series, in order to evaluate the data quality with a variety of downstream tasks. This is not a necessity. Anything that can be transformed into dMRI signal can be modelled by NeSH (e.g. peak directions or fODFs, which can be convolved into a diffusion weighted signal).

**Future Work.** Future work will further investigate the advantages of modelling dMRI data in a continuous space, as well as further evaluate the findings of the experiments. First, the quality and usability of the denoising properties of NeSH should be compared to other existing denoising methods. Second, using clinical datasets of lower angular and spatial resolution can provide insight into the ‘real-world’ clinical applicability of NeSH. This is especially interesting in MRI acquisitions of pathology (e.g. a glial-cell tumor) in the brain, as models relying on a prior learned on outside data might fail here. The harmonization of dMRI datasets across scanners and protocols [9, 13], is another area of research where NeSH can be applied. Third, fiber tracking is a common use-case for dMRI. We have performed fibertracking using iFOD2 [16] with tract masks and begin- and endpoint inclusion for different numbers of gradient directions. This showed no major differences between the different methods for all inspected tracts. An

explanation for this is the high spatial resolution of the HCP data, which allows tracts to be generated even with a downsampled angular resolution. Further research on datasets with lower spatial resolution will have to show the value of using NeSH reconstructions for fiber tracking. Fourth, a recent paper by Mancini et al. [7] has shown how compression of dMRI data using sinusoidal representation networks (SIRENs) [11] does not lead to reduced quality in downstream tasks. Using a SIREN architecture could also prove useful for the SH-based approach we have described. Finally, the generalization of the model to other subjects, protocols, and scanners has to be evaluated. We have performed a preliminary experiment which showed comparable results for signal reconstruction and fODFs.

## 5 Conclusion

Modeling dMRI data using NeSH produces results in downstream tasks with similar or possibly better results than established methods. It also shows promising results in the field of angular and spatial upsampling. NeSH can make use of the structural coherence in the brain, and does not rely on a prior learned on other datasets. The experiments in this paper provide an interesting avenue for modeling dMRI data, which should be further explored in future research.

## References

1. Aja-Fernandez, S., et al.: Validation of deep learning techniques for quality augmentation in diffusion MRI for clinical studies (2023)
2. Alexander, D.C., et al.: Image quality transfer and applications in diffusion MRI. *Neuroimage* **152**, 283–298 (2017)
3. Basser, P.J., Mattiello, J., LeBihan, D.: Estimation of the effective self-diffusion tensor from the NMR spin echo. *J. Magn. Reson. Ser. B* **103**(3), 247–254 (1994)
4. Bonet-Carne, E., Grussu, F., Ning, L., Sepehrband, F., Tax, C.M.: Computational Diffusion MRI: International MICCAI Workshop, Granada, Spain, September 2018. Springer, Cham (2019). <https://doi.org/10.1007/978-3-030-05831-9>
5. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, fast, and robust analytical q-ball imaging. *Magn. Reson. Med.: Off. J. Int. Soc. Magn. Reson. Med.* **58**(3), 497–510 (2007)
6. Girshick, R.: Fast R-CNN. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1440–1448 (2015)
7. Mancini, M., Jones, D.K., Palombo, M.: Lossy compression of multidimensional medical images using sinusoidal activation networks: an evaluation study. In: Cetin-Karayumak, S., et al. (eds.) CDMRI 2022. LNCS, vol. 13722, pp. 26–37. Springer, Cham (2022). [https://doi.org/10.1007/978-3-031-21206-2\\_3](https://doi.org/10.1007/978-3-031-21206-2_3)
8. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: representing scenes as neural radiance fields for view synthesis. *Commun. ACM* **65**(1), 99–106 (2021)
9. Ning, L., et al.: Cross-scanner and cross-protocol multi-shell diffusion MRI data harmonization: algorithms and results. *Neuroimage* **221**, 117128 (2020)

10. Raffelt, D., et al.: Apparent fibre density: a novel measure for the analysis of diffusion-weighted magnetic resonance images. *Neuroimage* **59**(4), 3976–3994 (2012)
11. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions (2020)
12. Tancik, M., et al.: Fourier features let networks learn high frequency functions in low dimensional domains. *Adv. Neural. Inf. Process. Syst.* **33**, 7537–7547 (2020)
13. Tax, C.M., et al.: Cross-scanner and cross-protocol diffusion MRI data harmonisation: a benchmark database and evaluation of algorithms. *Neuroimage* **195**, 285–299 (2019)
14. Tournier, J.D., Calamante, F., Connelly, A.: Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. *Neuroimage* **35**(4), 1459–1472 (2007)
15. Tournier, J.D., Calamante, F., Connelly, A.: Determination of the appropriate b value and number of gradient directions for high-angular-resolution diffusion-weighted imaging. *NMR Biomed.* **26**(12), 1775–1786 (2013)
16. Tournier, J.D., Calamante, F., Connelly, A., et al.: Improved probabilistic streamlines tractography by 2nd order integration over fibre orientation distributions. In: *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 1670. *Isrmr* (2010)
17. Tournier, J.D., et al.: MRtrix3: a fast, flexible and open software framework for medical image processing and visualisation. *NeuroimageD* **202**, 116137 (2019)
18. Van Essen, D.C., et al.: The WU-Minn human connectome project: an overview. *Neuroimage* **80**, 62–79 (2013)
19. Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: Plenotrees for real-time rendering of neural radiance fields. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5752–5761 (2021)
20. Zhu, F., Guo, S., Song, L., Xu, K., Hu, J., et al.: Deep review and analysis of recent NeRFs. *APSIPA Trans. Signal Inf. Process.* **12**(1) (2023)