DTI-Guided Volumetric Spherical Harmonics Regression for Single-to-Multi-Shell dMRI Synthesis

Binghua Li[1][0000−0002−2595−4762], Christina Andica[1], Tong Liang[1], Ziqing Chang[1], Chao Li[1,2], Wataru Uchida[1], Kaito Takabayashi[1], Qibin Zhao[2,3], Toshihisa Tanaka[3,2], Zhe Sun[1,⋆], Shigeki Aoki[1]

[1] Juntendo University, Tokyo, Japan
b.li.qr@juntendo.ac.jp, z.sun.kc@juntendo.ac.jp
[2] RIKEN Center for Advanced Intelligence Project, Tokyo, Japan
[3] Tokyo University of Agriculture and Technology, Tokyo, Japan

Abstract. Multi-shell diffusion MRI (dMRI) unlocks more expressive microstructural modeling than single-shell scans, yet its longer acquisition time hinders deployment in large-scale cohorts and time-constrained clinical settings. Synthesizing an unobserved shell from a single-shell input is fundamentally ill-posed and further complicated by protocol mismatch, where source and target gradient direction sets may not align. We propose DTI-SHNet, a single-to-multi-shell synthesis framework that operates in the real symmetric spherical harmonics (SH) coefficient domain and performs spatially aware volumetric regression. Given a source shell, we estimate diffusion tensor imaging (DTI) and use direction-agnostic parametric maps along with a brain mask as conditioning priors to guide a 3D U-Net regressor for source-to-target SH coefficient regression. To couple coefficient accuracy with signal fidelity, we introduce a signal consistency regularization that reconstructs signals on randomly sampled canonical directions from predicted coefficients and enforces agreement in the signal domain. Experiments on UK Biobank and Cam-CAN data for b=1000 to b=2000 dMRI synthesis show that DTI-SHNet achieves competitive visual quality compared to advanced methods, while better preserving downstream diffusion measures. Our code is available at https://github.com/xiaovhua/dti-shnet.

Keywords: Diffusion MRI synthesis · multi-shell reconstruction · spherical harmonics · diffusion tensor imaging.

1 Introduction

Diffusion magnetic resonance imaging (dMRI) probes tissue microstructure in vivo and enables clinically relevant analyses such as white matter integrity assessment and tractography [2,20]. To reduce scan time, many studies rely on

⋆ Corresponding author

single-shell acquisitions that measure diffusion signals at a single b-value. Multi-shell protocols instead sample multiple b-values, expanding coverage and benefiting downstream models and tasks [9,22]. However, multi-shell imaging requires longer scans and is more motion sensitive, limiting feasibility in routine workflows and large-scale cohorts [1].

Synthesizing missing shells from single-shell dMRI is a practical alternative to lengthy multi-shell acquisitions, but remains challenging. A straightforward strategy is model-based extrapolation across b-values [2]. However, because single-shell data constrain the signal at only one b-value, and higher b-values involve stronger non-Gaussian attenuation and complex microstructural effects [8,13,7], the cross-b dependence is inherently underdetermined, making such extrapolation unreliable. Data-driven methods have therefore been explored, including diffusion-weighted volume prediction using closely matched gradient directions across shells [4]. Yet such image-domain synthesis is sensitive to gradient alignment and sampling density. SH-domain regression mitigates this issue by representing signals with spherical harmonics (SH) coefficients, providing a more interpretable formulation that is less sensitive to protocol mismatch [3,11,10]. Nevertheless, existing SH-based methods still suffer from two practical challenges: (i) Microstructural ambiguity, where learning the mapping between different b-values can be brittle and sensitive to the acquisition protocol without strong priors that do not depend on gradient directions; (ii) Signal-domain inconsistency, where coefficient-domain regression does not directly constrain reconstructed signals, potentially degrading the quantitative fidelity of downstream diffusion measures.

To this end, as shown in Fig. 1, we propose DTI-SHNet, a volumetric SH regression framework guided by diffusion tensor imaging (DTI) for single-to-multi-shell synthesis. Given single-shell measurements, we first fit a diffusion tensor model and derive scalar parametric maps, including fractional anisotropy (FA) and mean diffusivity (MD), which provide microstructural priors not tied to individual gradient directions [14]. Regularized real symmetric SH fitting is then used to obtain coefficient volumes, and a DTI-conditioned 3D regressor predicts target-shell SH coefficients from source-shell coefficients by jointly exploiting these priors and spatial context. To couple coefficient learning with signal fidelity, we further introduce a signal consistency regularizer that reconstructs signals along randomly sampled canonical directions and enforces agreement with the corresponding target signals in the signal domain. Our contributions are summarized as follows:

(1) We propose a DTI-guided volumetric SH regression framework for single-to-multi-shell dMRI synthesis.
(2) We develop DTI-SHNet, a spatially aware 3D SH coefficient regressor conditioned on direction-agnostic DTI-derived priors and optimized with signal consistency regularization.
(3) We evaluate DTI-SHNet on UK Biobank and Cam-CAN, demonstrating competitive synthesis quality and improved preservation of downstream diffusion measures over strong baselines.

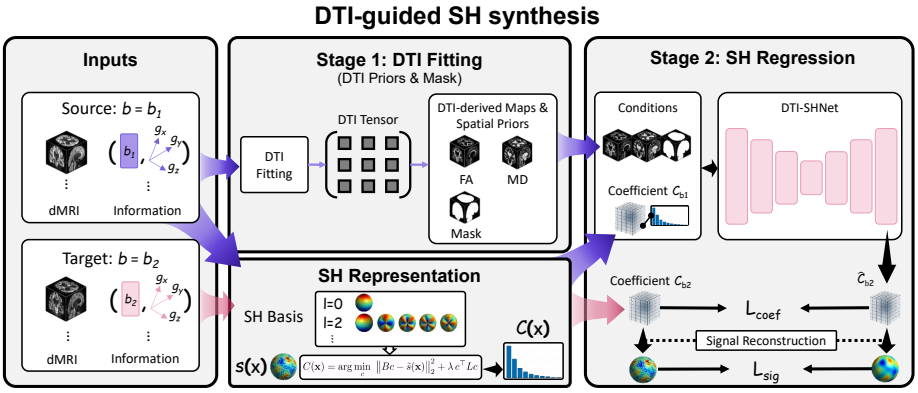

Fig. 1. Overview of our synthesis pipeline. Single-shell data at $b = b_1$ are represented as real symmetric SH coefficient volumes C_{b_1}. DTI-SHNet predicts target coefficient volumes $\widehat{C}_{b_2}$ guided by DTI priors. Target-shell signals are synthesized for desired directions by SH regression, with an additional signal consistency constraint $\mathcal{L}_{\text{sig}}$ applied on canonical directions sampled during training.

2 Methodology

2.1 Problem Formulation

Let $S(\mathbf{x}; b, \mathbf{g})$ denote the diffusion signal at voxel $\mathbf{x}$ acquired with b-value b and unit gradient direction $\mathbf{g}$. Given a reference baseline volume at $b = 0$ and N_1 DWIs at a source shell $b = b_1$ with directions $\{\mathbf{g}_1^{(i)}\}_{i=1}^{N_1}$, our goal is to synthesize the target-shell signal at $b = b_2$ for an arbitrary query direction $\mathbf{g}_2$. We denote the input set

$$\mathcal{X} = \{S(\mathbf{x}; 0)\} \cup \{S(\mathbf{x}; b_1, \mathbf{g}_1^{(i)})\}_{i=1}^{N_1}. \tag{1}$$

For each target query $(b_2, \mathbf{g}_2)$, we aim to produce $\widehat{S}(\mathbf{x}; b_2, \mathbf{g}_2)$ approximating $S(\mathbf{x}; b_2, \mathbf{g}_2)$.

2.2 Signal Preprocessing

For each subject, we compute $S_0(\mathbf{x})$ as the median across all $b = 0$ volumes and normalize the diffusion signal as

$$\tilde{S}(\mathbf{x}; b, \mathbf{g}) = \frac{S(\mathbf{x}; b, \mathbf{g})}{S_0(\mathbf{x})}. \tag{2}$$

This normalization reduces inter-subject intensity variation and provides a stable signal representation for subsequent DTI fitting, SH coefficient estimation, and network regression.

2.3 Diffusion Tensor Imaging Priors

To inject microstructural cues, we estimate the diffusion tensor $\mathbf{D}(\mathbf{x})$ using the standard DTI signal model [2]:

$$\tilde{S}(\mathbf{x}; b, \mathbf{g}) = \exp\big(-b\,\mathbf{g}^\top \mathbf{D}(\mathbf{x})\mathbf{g}\big). \tag{3}$$

Taking $y = -\log(\tilde{S})$ yields a linear system in the tensor elements of $\mathbf{D}(\mathbf{x})$, which we solve by least squares using all diffusion-weighted measurements at $b = b_1$. From the fitted tensor $\hat{\mathbf{D}}(\mathbf{x})$, we compute standard DTI parametric maps, including fractional anisotropy $\mathrm{FA}(\mathbf{x})$ and mean diffusivity $\mathrm{MD}(\mathbf{x})$ [2,14,15]. We further derive a binary brain mask from $S_0(\mathbf{x})$ using percentile-based intensity thresholding

$$M(\mathbf{x}) = \mathbb{I}[S_0(\mathbf{x}) \geq \tau_M]\,, \tag{4}$$

where τ_M is empirically set to the 60th percentile of the positive S_0 intensities. These DTI-derived maps and the brain mask serve as auxiliary input channels for subsequent SH coefficient regression.

2.4 Spherical Harmonics Representation

A direct mapping is sensitive to protocol mismatch when source and target gradient direction sets do not align. We instead represent the angular diffusion signal at each shell using real symmetric spherical harmonics, enforcing antipodal symmetry by retaining only even degrees. For a shell with direction set $\{\mathbf{g}^{(i)}\}_{i=1}^{N}$, we estimate voxel-wise SH coefficients $C(\mathbf{x}) \in \mathbb{R}^K$ via regularized least squares [3]:

$$C(\mathbf{x}) = \arg\min_{c} \ \big\|Bc - \tilde{s}(\mathbf{x})\big\|_2^2 + \lambda\, c^\top L c, \tag{5}$$

where $\tilde{s}(\mathbf{x}) \in \mathbb{R}^N$ stacks the normalized signals $\tilde{S}(\mathbf{x}; b, \mathbf{g}^{(i)})$ over the N directions, $B \in \mathbb{R}^{N\times K}$ is the SH design matrix evaluated at these directions, L is the Laplace–Beltrami regularizer, and λ controls the regularization strength. With maximum degree $l_{\max}$ and antipodal symmetry, the coefficient dimension is $K = \sum_{l\in\{0,2,\ldots,l_{\max}\}}(2l+1)$. This representation remains stable when the direction sampling is not uniform. Let $C_{b_1}(\mathbf{x})$ and $C_{b_2}(\mathbf{x})$ denote the SH coefficient volumes fitted from the source and target shells, respectively.

2.5 DTI-SHNet for Volumetric Coefficient Regression

As shown in Fig. 2, DTI-SHNet uses a 3D U-Net [16] as the regressor f_θ to predict target-shell SH coefficients:

$$\widehat{C}_{b_2} = f_\theta\big(\mathrm{concat}(C_{b_1}, \mathrm{FA}, \mathrm{MD}, M)\big), \tag{6}$$

where the source-shell coefficients and conditioning maps are concatenated along the channel dimension, and f_θ outputs K target-shell coefficient channels. To exploit the spatial conditioning provided by the DTI-derived maps and brain mask while maintaining manageable memory usage, we perform volumetric regression using 3D patches. During inference, sliding-window prediction is used to obtain the full-volume target-shell coefficients.

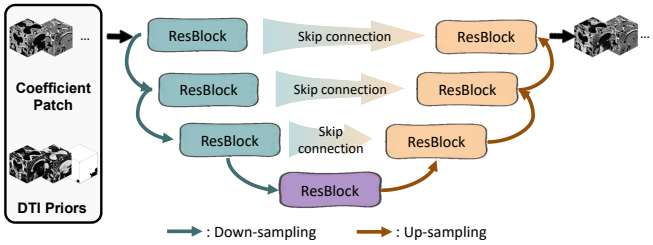

Fig. 2. Framework of DTI-SHNet for patch-based volumetric SH coefficient regression. A 3D U-Net predicts target-shell SH coefficient patches from source-shell coefficient patches conditioned on FA, MD, and the brain mask. Full-volume coefficients are obtained by sliding-window inference.

2.6 Signal Consistency Regularization

Coefficient regression may yield coefficients that are close numerically yet synthesize biased signals on directions of interest. To bridge coefficient learning and signal fidelity, we introduce a signal consistency regularization. In each iteration, we sample N_{can} directions uniformly on the sphere and build the corresponding SH basis matrix B_{can}. We then reconstruct signals on these canonical directions from the predicted and ground-truth coefficients, and penalize their discrepancy:

$$\mathcal{L}_{\text{sig}} = \frac{1}{N_{\text{can}}\,|\Omega|} \sum_{\mathbf{x}\in\Omega} \left\| B_{\text{can}} \widehat{C}_{b_2}(\mathbf{x}) - B_{\text{can}} C_{b_2}(\mathbf{x}) \right\|_1, \tag{7}$$

where Ω denotes the voxel set in the current training patch, $|\Omega|$ is the number of voxels, and $\|\cdot\|_1$ sums the absolute errors over the N_{can} canonical directions.

The overall training objective combines coefficient regression and signal consistency:

$$\mathcal{L} = \mathcal{L}_{\text{coef}} + \eta\, \mathcal{L}_{\text{sig}}, \tag{8}$$

where $\mathcal{L}_{\text{coef}}$ is a Huber loss on SH coefficients and η controls the contribution of signal consistency.

2.7 Synthesis for Arbitrary Target Directions

Given predicted coefficients $\widehat{C}_{b_2}$ and a target direction set $\{\mathbf{g}_2^{(i)}\}_{i=1}^{N_2}$, we synthesize the normalized target-shell signals via SH reconstruction:

$$\widehat{\tilde{S}}(\mathbf{x}; b_2, \mathbf{g}_2^{(i)}) = \left(B_2 \widehat{C}_{b_2}(\mathbf{x})\right)_i, \tag{9}$$

where B_2 is the SH basis matrix evaluated at the target direction set $\{\mathbf{g}_2^{(i)}\}_{i=1}^{N_2}$. Finally, we map the normalized signals back to the original intensity domain:

$$\widehat{S}(\mathbf{x}; b_2, \mathbf{g}_2^{(i)}) = S_0(\mathbf{x})\widehat{\tilde{S}}(\mathbf{x}; b_2, \mathbf{g}_2^{(i)}). \tag{10}$$

Together with the original shell at $b = b_1$, the synthesized shell at $b = b_2$ forms a multi-shell dMRI dataset.

3 Experiments and Results

3.1 Datasets and Experimental Settings

We evaluate our method on multi-shell dMRI from UK Biobank (UKB) [12] and the Cambridge Centre for Ageing and Neuroscience (Cam-CAN) [17,19]. UKB dMRI is acquired at 2 mm isotropic resolution with 5 b=0 volumes, 50 directions at b=1000, and 50 directions at b=2000, where the gradient direction sets are not shared across shells. Cam-CAN dMRI is acquired using a twice-refocused spin-echo sequence, including 3 b=0 volumes and 30 directions at each of b=1000 and b=2000, where the gradient directions are shared across shells. We use 822 subjects from UKB and 642 subjects from Cam-CAN. All dMRI data underwent gradient distortion correction, susceptibility distortion correction, and eddy-current and motion correction. Volumes are center-cropped and resized to $96 \times 96 \times 64$, and the coefficient regressor is trained on 3D patches of size $p = 32$.

We use {b=0, b=1000} as input and synthesize the target shell at b=2000 for the desired gradient directions. For SH representation, we use Laplace–Beltrami regularization with $\lambda = 0.006$ and set the maximum SH degree to $l_{\max} = 8$ following prior work [3,11]. We set the signal consistency weight to $\eta = 0.1$. We adopt a fixed subject-wise split with 90% of subjects for training and 10% for evaluation, with no subject overlap, and use the same splits for all methods. We train the model for 30 epochs using AdamW with a learning rate of 2×10^{-4} and batch size 4. The number of epochs is fixed based on preliminary convergence experiments without using the evaluation set, and the final checkpoint is used for all reported results. All experiments are implemented in PyTorch and run on an NVIDIA A100-PCIE-40GB GPU.

We compare against extrapolation [2], SH-DNN [11], a direction-matched 3D U-Net (DirUNet) [4], and a denoising diffusion probabilistic model (DDPM) [5,18] conditioned on b and $\mathbf{g}$. We report the structural similarity index measure (SSIM) [21], peak signal-to-noise ratio (PSNR) [6], and normalized mean squared error (NMSE). Moreover, to assess downstream diffusion measures, we compute the mean absolute errors of FA, MD, and NODDI-derived neurite density index (NDI), isotropic volume fraction (ISO), and orientation dispersion index (ODI) estimated from the synthesized multi-shell scans and their ground-truth counterparts. The FA and MD maps used for evaluation are re-estimated from the synthesized scans, rather than taken from the single-shell DTI priors used for conditioning in DTI-SHNet.

3.2 Results Analysis

Tab. 1 reports quantitative results on UKB and Cam-CAN for synthesizing the b=2000 shell from b=1000. Overall, deterministic learning-based methods outperform DTI extrapolation, indicating that the Gaussian assumption becomes insufficient at higher b-values. DDPM does not consistently outperform deterministic regressors in this setting, suggesting that accurate conditional genera-

Table 1. Quantitative reconstruction results on UKB and Cam-CAN for synthesizing b=2000 from b=1000.

Method	UKB			Cam-CAN		
	PSNR↑	SSIM↑	NMSE↓	PSNR↑	SSIM↑	NMSE↓
Extrapolation [2]	21.9 ± 0.6	0.825 ± 0.014	0.180 ± 0.028	22.0 ± 0.5	0.805 ± 0.014	0.134 ± 0.023
DDPM [5]	21.5 ± 0.8	0.824 ± 0.022	0.229 ± 0.030	24.0 ± 0.8	0.883 ± 0.023	0.109 ± 0.021
SH-DNN [11]	$\mathbf{25.6 \pm 0.7}$	$\mathbf{0.930 \pm 0.007}$	$\mathbf{0.079 \pm 0.016}$	26.7 ± 0.5	0.951 ± 0.005	0.045 ± 0.007
DirUNet [4]	24.1 ± 0.6	0.895 ± 0.010	0.109 ± 0.022	29.0 ± 0.6	0.968 ± 0.004	0.026 ± 0.004
DTI-SHNet	25.3 ± 0.7	0.922 ± 0.008	0.083 ± 0.018	$\mathbf{29.3 \pm 0.6}$	$\mathbf{0.970 \pm 0.004}$	$\mathbf{0.025 \pm 0.004}$

Table 2. Downstream diffusion-measure errors on UKB and Cam-CAN for synthesizing b=2000 from b=1000. ΔMD is scaled by 10^{-4}.

Dataset	Method	ΔFA ↓	ΔMD ↓	ΔNDI ↓	ΔODI ↓	ΔISO ↓
UKB	Extrapolation [2]	0.025 ± 0.002	1.922 ± 0.211	0.180 ± 0.010	0.088 ± 0.005	0.081 ± 0.008
	DDPM [5]	0.256 ± 0.062	5.604 ± 1.550	0.116 ± 0.036	0.116 ± 0.020	0.057 ± 0.016
	SH-DNN [11]	0.028 ± 0.002	0.352 ± 0.050	0.060 ± 0.006	0.038 ± 0.003	0.033 ± 0.003
	DirUNet [4]	0.041 ± 0.003	0.387 ± 0.081	0.065 ± 0.013	0.050 ± 0.004	0.037 ± 0.009
	DTI-SHNet	$\mathbf{0.022 \pm 0.001}$	$\mathbf{0.281 \pm 0.051}$	$\mathbf{0.038 \pm 0.005}$	$\mathbf{0.028 \pm 0.002}$	$\mathbf{0.021 \pm 0.003}$
Cam-CAN	Extrapolation [2]	0.034 ± 0.003	2.940 ± 0.252	0.197 ± 0.014	0.090 ± 0.007	0.100 ± 0.009
	DDPM [5]	0.196 ± 0.052	4.574 ± 1.370	0.096 ± 0.021	0.078 ± 0.012	0.054 ± 0.014
	SH-DNN [11]	0.047 ± 0.003	0.676 ± 0.072	0.114 ± 0.011	0.064 ± 0.006	0.057 ± 0.007
	DirUNet [4]	0.032 ± 0.003	0.530 ± 0.081	0.085 ± 0.013	0.041 ± 0.004	0.046 ± 0.007
	DTI-SHNet	$\mathbf{0.032 \pm 0.002}$	$\mathbf{0.426 \pm 0.062}$	$\mathbf{0.068 \pm 0.009}$	$\mathbf{0.038 \pm 0.004}$	$\mathbf{0.038 \pm 0.005}$

tion at high b remains challenging. SH-DNN benefits from learning in the spherical harmonics coefficient domain and achieves the best reconstruction metrics on UKB. DirUNet leverages volumetric regression, but it relies on direction-wise pairing and becomes more sensitive when the gradient direction sets are not shared across shells, which is the case in UKB. Our DTI-SHNet achieves competitive reconstruction fidelity on UKB and the best overall reconstruction performance on Cam-CAN, reflecting the benefit of combining SH-domain regression with DTI-derived spatial priors.

Tab. 2 evaluates downstream diffusion measures by reporting errors of FA, MD, and NODDI-derived NDI, ODI, and ISO estimated from the synthesized multi-shell scans. DTI-SHNet yields the lowest or tied-lowest errors across these downstream measures, indicating improved preservation of diffusion-derived information.

Fig. 3 presents representative visual comparisons for b=2000 synthesis on UKB, including two selected diffusion-weighted images, DTI-derived FA and MD maps, and NODDI-derived NDI, ODI, and ISO maps. Consistent with the quantitative results, DTI-SHNet produces more spatially coherent diffusion contrast and downstream parametric maps, with lower normalized errors than competing baselines.

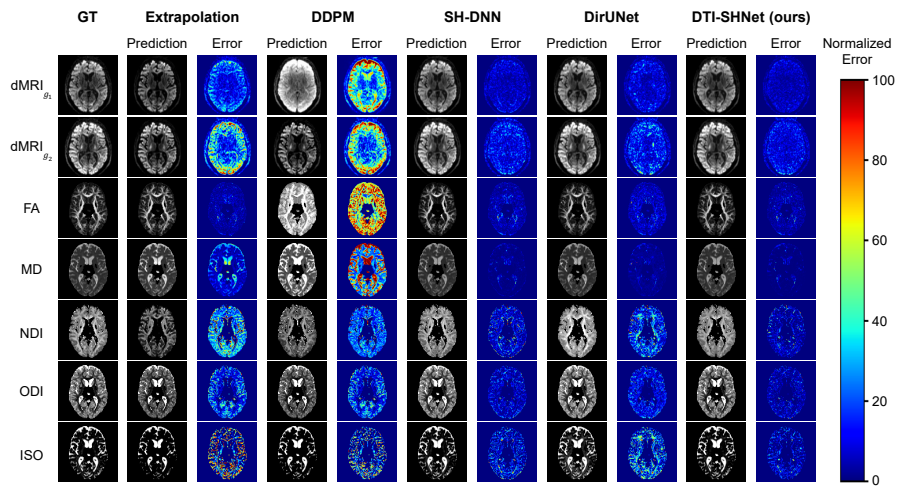

Fig. 3. Qualitative comparison for b=2000 synthesis from b=1000 on UKB. Representative slices include two diffusion-weighted images, DTI-derived FA and MD maps, and NODDI-derived NDI, ODI, and ISO maps. DTI-SHNet better preserves anatomical structures and downstream diffusion measures with reduced normalized errors compared with competing baselines.

3.3 Ablation Study

We ablate two components in DTI-SHNet: the DTI-derived priors and the signal consistency loss $\mathcal{L}_{\text{sig}}$. As shown in Tab. 3, DTI-derived priors mainly improve reconstruction fidelity, while $\mathcal{L}_{\text{sig}}$ provides an additional constraint in the signal domain and strengthens the preservation of diffusion-derived measures. Combining both components achieves the best overall balance between synthesis quality and downstream reliability, supporting their complementary roles in accurate multi-shell dMRI synthesis and downstream reliability.

Table 3. Ablation study on DTI-derived priors and the signal consistency loss $\mathcal{L}_{\text{sig}}$ for synthesizing b=2000 from b=1000. ΔMD is scaled by 10^{-4}.

Dataset	Priors	$\mathcal{L}_{\text{sig}}$	PSNR↑	SSIM↑	NMSE↓	ΔFA ↓	ΔMD ↓	ΔNDI ↓	ΔODI ↓	ΔISO ↓
UKB	×	×	25.33	0.9218	0.0829	0.0235	0.337	0.0497	0.0303	0.0263
	✓	×	**25.42**	**0.9235**	**0.0814**	0.0224	0.322	0.0402	0.0293	0.0220
	×	✓	25.29	0.9219	0.0838	0.0233	0.282	0.0458	0.0292	0.0256
	✓	✓	25.33	0.9216	0.0830	**0.0220**	**0.281**	**0.0381**	**0.0282**	**0.0212**
Cam-CAN	×	×	27.90	0.9547	0.0341	0.0604	0.903	0.0923	0.0434	0.0447
	✓	×	28.64	0.9627	0.0288	0.0418	0.656	0.0776	0.0456	0.0408
	×	✓	28.61	0.9639	0.0290	0.0343	0.467	0.0808	**0.0368**	0.0432
	✓	✓	**29.32**	**0.9700**	**0.0246**	**0.0319**	**0.426**	**0.0681**	0.0375	**0.0378**

4 Conclusion

We introduced DTI-SHNet, a DTI-guided volumetric SH regression framework for synthesizing a high-b shell from single-shell dMRI. By predicting real symmetric SH coefficients, DTI-SHNet supports synthesis for arbitrary target direction sets within the trained target shell and is less sensitive to gradient direction mismatch. DTI-derived priors and a signal consistency regularization stabilize this ill-posed mapping and improve signal fidelity. Experiments on UKB and Cam-CAN show competitive reconstruction quality and improved preservation of downstream diffusion measures. A current limitation is that the model is trained for a specific source–target b-value pair and therefore does not directly generalize to unseen b-values without retraining. Future work will explore b-value-conditioned or continuous q-space synthesis to improve generalization across acquisition protocols.

Acknowledgments. This research was conducted using the UK Biobank Resource under Application Number 135200. It was also supported by the Japan Society for the Promotion of Science (JSPS) KAKENHI under Grant Numbers JP23K28109, JP24K03005, JP25KJ1215 and JP26K14888.

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

References

1. Alexander, D.C., Dyrby, T.B., Nilsson, M., Zhang, H.: Imaging brain microstructure with diffusion MRI: Practicality and applications. NMR in Biomedicine **32**(4), e3841 (2019)
2. Basser, P.J., Mattiello, J., LeBihan, D.: MR diffusion tensor spectroscopy and imaging. Biophysical Journal **66**(1), 259–267 (1994)
3. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Apparent diffusion coefficients from high angular resolution diffusion imaging: Estimation and applications. Magnetic Resonance in Medicine **56**(2), 395–410 (2006)
4. Dugan, R., Carmichael, O.: Multi-shell dMRI estimation from single-shell data via deep learning. In: International Workshop on Machine Learning in Clinical Neuroimaging. pp. 14–22. Springer (2023)
5. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems (NeurIPS) **33**, 6840–6851 (2020)
6. Hore, A., Ziou, D.: Image quality metrics: PSNR vs. SSIM. In: 2010 20th International Conference on Pattern Recognition. pp. 2366–2369. IEEE (2010)
7. Jelescu, I.O., Veraart, J., Fieremans, E., Novikov, D.S.: Degeneracy in model parameter estimation for multi-compartmental diffusion in neuronal tissue. NMR in Biomedicine **29**(1), 33–47 (2016)
8. Jensen, J.H., Helpern, J.A., Ramani, A., Lu, H., Kaczynski, K.: Diffusional kurtosis imaging: The quantification of non-Gaussian water diffusion by means of magnetic resonance imaging. Magnetic Resonance in Medicine **53**(6), 1432–1440 (2005)

9. Jeurissen, B., Tournier, J.D., Dhollander, T., Connelly, A., Sijbers, J.: Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data. NeuroImage **103**, 411–426 (2014)
10. Jha, R.R., Jaswal, G., Bhavsar, A., Nigam, A.: Single-shell to multi-shell dMRI transformation using spatial and volumetric multilevel hierarchical reconstruction framework. Magnetic Resonance Imaging **87**, 133–156 (2022)
11. Koppers, S., Haarburger, C., Merhof, D.: Diffusion MRI signal augmentation: From single-shell to multi-shell with deep learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 61–70. Springer (2016)
12. Miller, K.L., Alfaro-Almagro, F., Bangerter, N.K., Thomas, D.L., Yacoub, E., Xu, J., Bartsch, A.J., Jbabdi, S., Sotiropoulos, S.N., Andersson, J.L., et al.: Multimodal population brain imaging in the UK Biobank prospective epidemiological study. Nature Neuroscience **19**(11), 1523–1536 (2016)
13. Novikov, D.S., Fieremans, E., Jespersen, S.N., Kiselev, V.G.: Quantifying brain microstructure with diffusion MRI: Theory and parameter estimation. NMR in Biomedicine **32**(4), e3998 (2019)
14. Pierpaoli, C., Basser, P.J.: Toward a quantitative assessment of diffusion anisotropy. Magnetic Resonance in Medicine **36**(6), 893–906 (1996)
15. Pierpaoli, C., Jezzard, P., Basser, P.J., Barnett, A., Di Chiro, G.: Diffusion tensor MR imaging of the human brain. Radiology **201**(3), 637–648 (1996)
16. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 234–241. Springer (2015)
17. Shafto, M.A., Tyler, L.K., Dixon, M., Taylor, J.R., Rowe, J.B., Cusack, R., Calder, A.J., Marslen-Wilson, W.D., Duncan, J., Dalgleish, T., et al.: The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) study protocol: A cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing. BMC Neurology **14**(1), 204 (2014)
18. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: Proceedings of the International Conference on Machine Learning (ICML). pp. 2256–2265. PMLR (2015)
19. Taylor, J.R., Williams, N., Cusack, R., Auer, T., Shafto, M.A., Dixon, M., Tyler, L.K., Henson, R.N., et al.: The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: Structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. NeuroImage **144**, 262–269 (2017)
20. 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)
21. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. IEEE Transactions on Image Processing **13**(4), 600–612 (2004)
22. Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C.: NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. NeuroImage **61**(4), 1000–1016 (2012)