Unsupervised q-Space Interpolation Using Physics-Constrained Coordinate-Based Implicit Networks

Atakan Topcu[1,2], Abdallah Zaid Alkilani[1,2], Tolga Çukur[1,2,3], and Emine Ulku Saritas[1,2]
[1]Electrical and Electronics Department, Bilkent University, Ankara, Turkey, [2]National Magnetic Resonance Center (UMRAM), Bilkent University, Ankara, Turkey, [3]Neuroscience Graduate Program, Bilkent University, Ankara, Turkey

Synopsis

Motivation: Most diffusion MRI techniques require extensive sampling of q-space to effectively resolve fiber structures at a fine detail. The scan times become impractically long, especially for clinical settings.

Goal(s): Our goal is to arbitrarily interpolate the q-space data to enable downsampling of q-space, while maintaining high fidelity diffusion metrics.

Approach: We propose QUCCI, a subject-specific unsupervised implicit network model that utilizes both implicit and physics-driven explicit regularization to encode diffusion MRI signals with angular continuity.

Results: QUCCI achieves superior q-space interpolation, outperforming traditional and deep learning methods.

Impact: QUCCI provides high-fidelity diffusion MRI metrics via improving the angular interpolation of diffusion MRI signals under highly undersampled q-space cases, which may especially be beneficial in the clinical settings where excessively long scan times are impractical.

Introduction

Diffusion MRI (dMRI) has enabled the study of the brain's complex architecture in vivo. However, resolving fiber structures at a finer detail necessitates denser q-space sampling, which can be impractical due to excessively long scan times[1-3]. Deep learning methods for angular q-space interpolation show potential for reduced scan time by enabling q-space downsampling. However, these methods often demand large training data and pose challenges when dealing with non-conforming data, such as in pathological cases[4].

We propose a subject-specific, unsupervised Q-space Upsampling via physics-Constrained Coordinate-based Implicit network (QUCCI) that angularly interpolates q-space data. QUCCI combines implicit multilayer perceptron regularization with physics-driven spherical harmonic (SH) regularization and image regularization in the context of dMRI. In comparison to traditional and other deep learning methods, QUCCI demonstrates superior angular interpolation for highly undersampled q-space acquisitions.

Methods

Classical Spherical Harmonics Interpolation:
For each voxel, dMRI signal in q-space can be decomposed using SH basis functions via a least-squares fitting to the acquired q-space data[5]. For spherical harmonics interpolation (SHI), dMRI images at an unacquired q-space direction can be estimated via a linear combination of SHs with their corresponding SH basis functions.

Learning-based Approach:
A previous learning-based approach, NeSH, takes voxel-space coordinates as input, and utilizes coordinate-based networks to predict SH coefficients for each voxel[6]. For training, predicted SH coefficients are converted to estimated dMRI images, and estimated and acquired dMRI images are compared at sampled q-space directions. NeSH benefits from implicit regularization of coordinate-based networks to represent q-space in angular continuity, and utilizes l1-norm regularization on SH coefficients. While NeSH was shown to provide a coherent q-space representation across a range of undersampling rates, it performs suboptimally towards high undersampling rates.

Proposed Method:
We propose a subject-specific, unsupervised coordinate-based model that utilizes implicit and physics-driven regularization in voxel-space and q-space domains. QUCCI, outlined in Fig. 1, takes voxel-space coordinates as input, and sets them as centers of isotropic Gaussian distributions. Then, voxel coordinates sampled from these distributions are mapped to positional encodings[7]. This sampling scheme enables implicit voxel-space regularization by enforcing adjacent coordinates to have similar image intensities. To enforce data consistency, we adopt a scheduling system that sets the standard deviation of Gaussian distributions to zero near the end of training.

QUCCI predicts SH coefficients for each voxel. During training, the network minimizes mean-square-error (MSE) between the dMRI images estimated via SH coefficients and the acquired dMRI images at sampled q-space directions. This procedure enforces implicit q-space regularization, as the model learns to estimate SH coefficients with no explicit supervision. For explicit q-space regularization, Laplace-Beltrami (LB) regularization is applied on SH coefficients, as it was shown to be well-suited for single-shell q-space measurements distributed on a sphere[5]. A pre-trained plug-and-play denoiser is incorporated for explicit voxel-space regularization of the estimated dMRI images[8]. The overall loss is:

$$L_{QUCCI} = L_{MSE}(\mathcal{F}_{\theta,\phi}(k), y) + \mathcal{R}_{LB}(k) + \mathcal{R}_{PnP}(\mathcal{F}_{\theta,\phi}(k))$$

where k denotes SH coefficients, y denotes the acquired dMRI images, $\mathcal{F}_{\theta,\phi}$, is the physical model that estimates dMRI images from SH coefficients, L_{MSE} is MSE between the estimated and acquired dMRI images, $\mathcal{R}_{LB}$ is the LB regularizer over SH coefficients, and $\mathcal{R}_{PnP}$ is the plug-and-play denoiser in voxel-space.

Implementation Details:
Single-shell dMRI data for five randomly chosen subjects from the preprocessed HCP dataset were utilized, comprising 18 b = 0 s/mm^2 volumes, and 90 b = 1000 s/mm^2 volumes[9]. QUCCI, SHI, and NeSH were implemented in Pytorch, and trained for 3000 epochs using Adam optimizer for a single slice with 145x145 matrix size. Competing methods were tested for 8-30 q-space directions, undersampled from the reference fully-sampled case of

90 q-space directions. Competing methods were used for estimating dMRI images at all 90 q-space directions. Dipy and MRtrix3 were used for calculating diffusion tensor imaging (DTI) metrics and fiber orientation distribution functions (fODFs), respectively.

Results

DTI maps obtained from interpolated q-space show stark differences among the competing methods, as shown in Figs. 2-3. Overall, QUCCI outperforms SHI and NeSH for highly undersampled cases, whereas NeSH's performance deteriorates. For example, with 8 acquired q-space directions, QUCCI boosts fractional anisotropy (FA) map fidelity by 7.17dB PSNR/33.05% SSIM over NeSH and 1.11dB PSNR/0.50% SSIM over SHI. To inspect the downstream capabilities, Fig. 4 displays representative fODFs of three different white matter regions, where fODFs from QUCCI display a high degree of match to the reference fODFs.

Conclusion

Even for highly undersampled acquisitions, QUCCI reconstructs dMRI-based metrics comparable to those derived from the reference fully-sampled q-space. Through subject-specific SH coefficient predictions and combined implicit and physics-driven explicit regularization, QUCCI enables reliable q-space interpolation with a significant reduction in dMRI scan times.

Acknowledgements

Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

1. Van Essen DC, Ugurbil K, Auerbach E, et al. The Human Connectome Project: a data acquisition perspective. Neuroimage. 2012;62(4):2222-2231. doi:10.1016/j.neuroimage.2012.02.018

2. Hutter J, Price AN, Cordero-Grande L, et al. Quiet echo planar imaging for functional and diffusion MRI. Magn Reson Med. 2018;79(3):1447-1459. doi:10.1002/mrm.26810

3. Howell BR, Styner MA, Gao W, et al. The UNC/UMN Baby Connectome Project (BCP): An overview of the study design and protocol development. Neuroimage. 2019;185:891-905. doi:10.1016/j.neuroimage.2018.03.049

4. Tax CM, Grussu F, Kaden E, et al. Cross-scanner and cross-protocol diffusion MRI data harmonisation: A benchmark database and evaluation of algorithms. Neuroimage. 2019;195:285-299. doi:10.1016/j.neuroimage.2019.01.077

5. Descoteaux M, Angelino E, Fitzgibbons S, Deriche R. Regularized, fast, and robust analytical Q-ball imaging. Magn Reson Med. 2007;58(3):497-510. doi:10.1002/mrm.21277

6. Hendriks T, Villanova A, Chamberland M. Neural Spherical Harmonics for structurally coherent continuous representation of diffusion MRI signal. arXiv preprint arXiv:2308.08210

7. Mildenhall B, Srinivasan PP, Tancik M, Barron JT, Ramamoorthi R, Ng R. Nerf: Representing scenes as neural radiance fields for view synthesis. Commun ACM. 2021 Dec 17;65(1):99-106

8. Zhang K, Li Y, Zuo W, Zhang L, Van Gool L, Timofte R. Plug-and-Play Image Restoration With Deep Denoiser Prior. IEEE Trans Pattern Anal Mach Intell. 2022 Oct 1;44(10):6360-6376. doi: 10.1109/TPAMI.2021.3088914

9. Van Essen DC, Smith SM, Barch DM, et al. The WU-Minn Human Connectome Project: an overview. Neuroimage. 2013;80:62-79. doi:10.1016/j.neuroimage.2013.05.041

Figures

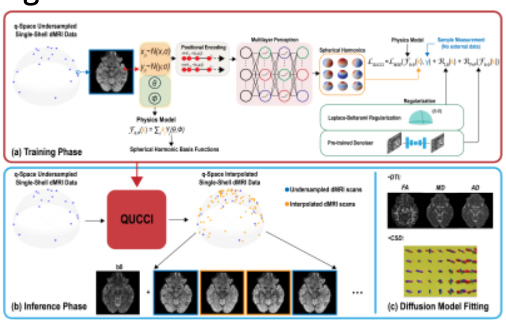

Figure 1. (a) The proposed QUCCI model for subject-specific, unsupervised implicit q-space interpolation. **(b)** During inference, QUCCI is used to interpolate the undersampled q-space. **(c)** Estimated dMRI images can then be used for a variety of downstream tasks such as computing DTI metrics and constrained spherical deconvolution (CSD) metrics.

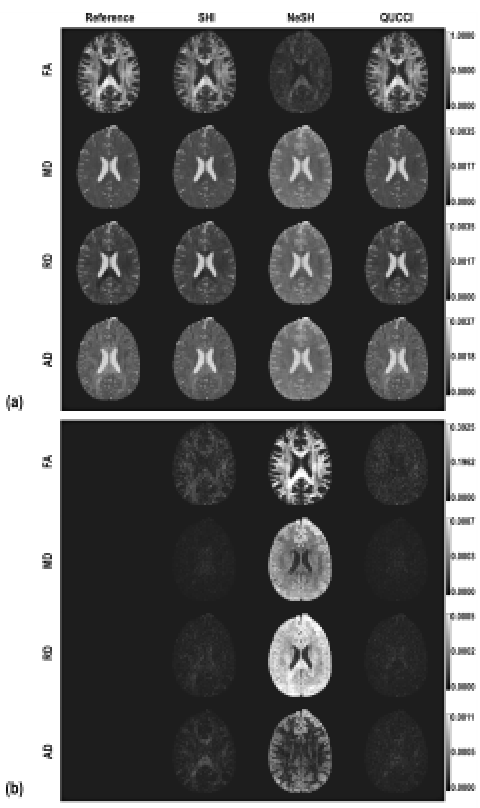

Figure 2. For the highly undersampled case of 15 q-space directions, **(a)** representative DTI metrics obtained from dMRI images estimated for all 90 q-space directions using QUCCI and the competing methods. **(b)** The corresponding error maps with respect to the reference metrics from fully-sampled 90 q-space directions. QUCCI demonstrates visibly improved performance for all DTI metrics. (FA: fractional anisotropy, MD: mean diffusivity, RD: radial diffusivity, AD: axial diffusivity).

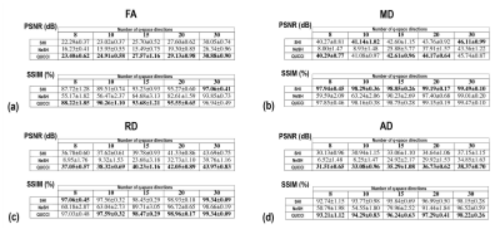

Figure 3. Quantitative performance evaluations using PSNR and SSIM, reported as mean±standard deviation across 5 subjects for DTI metrics: **(a)** FA, **(b)** MD, **(c)** RD, and **(d)** AD. QUCCI outperforms SHI and NeSH for highly undersampled cases. (FA: fractional anisotropy, MD: mean diffusivity, RD: radial diffusivity, AD: axial diffusivity).

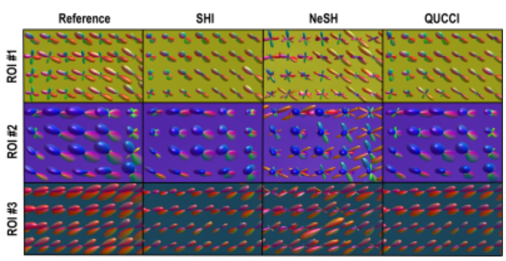

Figure 4. Representative fiber orientation distribution function (fODF) glyphs for the highly undersampled case of 15 q-space directions, displayed for three different 7x4 ROIs corresponding to white matter regions of a single subject. fODFs from QUCCI display a high degree of match to the reference fODFs.