

q-Space Deep Learning for Twelve-Fold Shorter and Model-Free Diffusion MRI Scans

Vladimir Golkov¹, Alexey Dosovitskiy², Philipp Sämann³, Jonathan I. Sperl⁴,
Tim Sprenger^{1,4}, Michael Czisch³, Marion I. Menzel⁴, Pedro A. Gómez^{1,4},
Axel Haase¹, Thomas Brox², and Daniel Cremers¹

¹ Technische Universität München, Garching, Germany

² University of Freiburg, Freiburg, Germany

³ Max Planck Institute of Psychiatry, Munich, Germany

⁴ GE Global Research, Garching, Germany

`golkov@cs.tum.edu`

Abstract. Diffusion MRI uses a multi-step data processing pipeline. With certain steps being prone to instabilities, the pipeline relies on considerable amounts of partly redundant input data, which requires long acquisition time. This leads to high scan costs and makes advanced diffusion models such as diffusion kurtosis imaging (DKI) and neurite orientation dispersion and density imaging (NODDI) inapplicable for children and adults who are uncooperative, uncomfortable or unwell. We demonstrate how deep learning, a group of algorithms in the field of artificial neural networks, can be applied to reduce diffusion MRI data processing to a single optimized step. This method allows obtaining scalar measures from advanced models at twelve-fold reduced scan time and detecting abnormalities without using diffusion models.

1 Introduction

Advanced diffusion MRI models such as DKI [1] and NODDI [2] are preferable over traditional diffusion MRI models because they provide more accurate characterization of tissue microstructure. However, they require long acquisition times. This can be problematic in clinical applications due to high scan costs or if the patient is uncooperative, uncomfortable or unwell.

In diffusion MRI, a number of diffusion-weighted images (DWIs) for different diffusion weightings and directions (constituting the so-called three-dimensional q-space) are acquired. The task in quantitative diffusion MRI is to find a mapping from a limited number of noisy signal samples to rotationally invariant scalar measures that quantify microstructural tissue properties. This inverse problem is solved in each image voxel. The classical approach consists of fitting [3] a diffusion model and calculating rotationally invariant measures from the fitted model parameters. Another approach to calculate scalar measures is approximation, particularly machine learning. Simulations of simplified tissue models with extensive sets of diffusion weightings [4, 5] indicate that standard model fitting procedures can be replaced by approximation methods. On the basis of

these observations, we apply deep learning [6–9] for accurate approximation and present a deep learning framework for different inputs (full and subsampled sets of regular DWIs, non-diffusion contrasts) and outputs (denoising, missing DWI reconstruction, scalar measure estimation, tissue segmentation). We term this framework *q-space deep learning* (q-DL). Scalar measure estimation from twelve-fold shortened acquisition is demonstrated on two advanced models: DKI and NODDI. By shortening the acquisition duration of advanced models by an order of magnitude, we strongly improve their potential for clinical use.

Recent applications of machine learning in diffusion MRI [10–12] use fitted model parameters as learning inputs, whereas we omit model fitting altogether by using the DWIs themselves, which allows the use of unprecedented subsampling factors that model fitting (Fig. 1a–d) cannot handle.

Besides, our framework allows tissue segmentation and lesion detection in diffusion MRI without using diffusion models.

2 Methods

Deep learning [6–9] is a set of algorithms for learning of input-output-mappings. It can outperform other machine learning methods and has recently been successfully applied in a variety of fields such as computer vision, natural language processing and drug discovery. It is based on artificial neural networks, where input data is propagated through several layers of hidden units (artificial neurons). Each layer is a data transformation step. The classical diffusion MRI pipeline also consists of several steps: in DKI, approximately 150 measurements are reduced to 22 model parameters, then to a few rotationally invariant measures, and finally (implicitly or explicitly) to the tissue property of interest. In every step, information is partly lost by reducing the degrees of freedom. However, the classical pipeline does not provide feedback to the earlier steps with regard to what part of the information should be retained or discarded and which transformations should be applied. Thus, the pipeline relies on handcrafting and fixing each step. Deep learning takes a more holistic approach: all layers are optimized jointly in terms of the final objective, namely minimizing the output error. This prevents the loss of information at intermediate steps.

In an artificial neural network with L layers (particularly in a so-called multilayer perceptron), the data transformation in layer $i \in \{1, \dots, L\}$ depends on the weight matrix $\mathbf{W}^{(i)}$ and bias term $\mathbf{b}^{(i)}$ according to the rule $\mathbf{a}_j^{(i)} = s_i(\mathbf{W}^{(i)}\mathbf{a}_j^{(i-1)} + \mathbf{b}^{(i)})$, where $\mathbf{a}_j^{(i)}$ are the output vectors of layer i for data sample j , the $\mathbf{a}_j^{(0)}$ are the input *vectors* of the network, and s_i are nonlinearities (see below). During training, all weight matrices and bias terms are jointly adjusted such that the output vectors $\mathbf{a}_j^{(L)}$ for each training sample j (in our case: each image voxel j) well approximate the target output vectors $\mathbf{y}_j$. This adjustment is achieved by using the backpropagation algorithm (implementation [13]) to solve $\arg \min_{\mathbf{W}, \mathbf{b}} \sum_j \left\| \mathbf{a}_j^{(L)} - \mathbf{y}_j \right\|^2$, where the sum of errors is taken over all training samples j , and outputs $\mathbf{a}_j^{(i)}$ recursively depend on $\mathbf{W}, \mathbf{b}$. In all

experiments, training data originate from different human subjects than test data. The network thus does not “know” the true output vectors of the test data but rather estimates them based on the input-output-mapping learned from training data. Each voxel j is treated individually as a data sample. The algorithm does not know its location in the image. We introduce several input-output-mapping tasks. Different deep networks are trained for different tasks:

Denoising. For denoising, the signal $\mathbf{S}_j$ from all DWIs in voxel j is used as both the input $\mathbf{a}_j^{(0)}$ and target $\mathbf{y}_j$ of the neural network. The number of network inputs is the number n of used DWIs, i.e. each input vector $\mathbf{a}_j^{(0)} = \mathbf{S}_j$ has length n (for every j). The length of the output vector $\mathbf{a}_j^{(L)}$ is also n . A network trained to reconstruct its own inputs is known as an autoencoder [7]. Its approximate nature and dropout-based training [8] prevent overfitting and reduce noise.

Reconstruction of Missing DWIs. For q-DL-based Reconstruction of missing DWIs (q-DL-R), a neural network is trained to predict the signal $\mathbf{S}_j$ in all DWIs (voxel j) from a reduced subset $\mathbf{S}_{j,\alpha}$ where α is a pseudorandom subsampling multi-index (such that the q-space sampling is consistent across training and test data). The length of the input vector $\mathbf{a}_j^{(0)} = \mathbf{S}_{j,\alpha}$ is $|\alpha|$. Due to partial data redundancy in q-space, missing DWIs can be reconstructed from a reduced subset.

Estimation of Scalar Measures. A network is trained to predict microstructure-characterizing scalar measures $\mathbf{m}_j$ directly from the (reduced set of) DWIs $\mathbf{S}_{j,\alpha}$. In other words, inputs are $\mathbf{a}_j^{(0)} = \mathbf{S}_{j,\alpha}$ and targets are $\mathbf{y}_j = \mathbf{m}_j$. The length of the output vector is the number of considered scalar measures. Training targets $\mathbf{y}_j = \mathbf{m}_j$ are obtained from a fully sampled training dataset by model fitting.

Model-Free Segmentation. Tissue segmentation is achieved by training a neural network to discriminate between several tissue types. We propose modifying the approach [14] of multi-parametric MRI tissue characterization by artificial neural networks such that the DWIs are directly used as inputs rather than using scalar measures obtained from model fitting. State-of-the-art automatic segmentation [15, 16] (based on non-diffusion images with spatial priors) into healthy white matter (WM), grey matter (GM), cerebrospinal fluid (CSF) and multiple sclerosis lesions was used as ground truth for our proof-of-concept model-free segmentation (based on diffusion images without spatial priors). The q-DL framework allows incorporating additional contrasts other than DWIs as inputs to the learning algorithm. We used fluid-attenuated inversion recovery (FLAIR) signal as an additional input. The length of the output vector is the number of tissue classes (with each output representing a relative class membership “likeliness” using softmax, see below).

Additional Remarks. The trained network with optimized $\mathbf{W}^{(i)}$ and $\mathbf{b}^{(i)}$ is then applied to other datasets using the recursive formula for $\mathbf{a}_j^{(i)}$. In the case of




Fig. 1. Maps of radial kurtosis in the human brain for various methods and MRI scan acceleration factors. 148, 40, 25 and 12 randomly selected DWIs are used; required scan time for each scheme is shown in seconds per slice. (a–d) Standard processing (model fitting followed by radial kurtosis calculation). (e–h) q-DL-R, followed by model fitting and radial kurtosis calculation. (j–m) q-DL.

α -subsampling, a large number of DWIs is required only once (for the training dataset) in order to estimate target scalar measures $\mathbf{m}$ using model fitting. Subsequently, training is performed using few DWIs as inputs, and the trained network can be applied to any number of previously unseen test datasets which have only few DWIs (same subsampling scheme).

The deep learning toolbox [13] was used for experiments. In each of the described tasks, the neural network architecture used is a multilayer perceptron with three fully connected hidden layers, each consisting of 150 rectified linear units [6], i.e. $s_i(\mathbf{z}) = \max(\mathbf{0}, \mathbf{z})$. Linear units $s_L(\mathbf{z}) = \mathbf{z}$ are used in the output layer L for fitting tasks and softmax outputs $s_L(\mathbf{z}) = \exp(\mathbf{z}) / \|\exp(\mathbf{z})\|_1$ for classification tasks. Each input and output of the neural network is scaled to the interval $[0, 1]$ during training, and the same linear transformation is applied to test data. We found that results improve if the network is pre-trained [7] using training data or initialized with orthogonal random weights [9]. We use




Fig. 2. Error evaluation of different methods at different acceleration factors. Root-mean-square deviation from the radial kurtosis estimated by the standard pipeline (fully sampled model fitting) is used as error metric. Model fitting is outperformed by the proposed methods when less than 70 DWIs are used.

a dropout [8] fraction of 0.1, stochastic gradient descent with momentum 0.9, batch size 128, learning rate 0.01, warm-up learning rate 0.001 for first 10 epochs.

Datasets. The *in vivo* protocols were approved by our institutional review board and prior informed consent was obtained. Two data sets of healthy volunteers were acquired using a scheme optimized [17, 18] for DKI and suitable for NODDI: three shells ($b = 750, 1070, 3000 \text{ s/mm}^2$) with 25, 40, 75 directions, respectively, and eight $b = 0$ images. This is a non-radial multi-shell q-space acquisition scheme for which no missing data reconstruction algorithm exists to the best of our knowledge. Echo-planar imaging was performed using a 3T GE MR750 MR scanner (GE Healthcare, Waukesha, WI, USA) equipped with a 32-channel head coil ($T_E = 80.7 \text{ ms}$, $T_R = 2 \text{ s}$, $\text{FOV} = 24 \times 24 \times 4 \text{ cm}$, isotropic voxel size 2.5 mm, ASSET factor 2). All data underwent FSL topup distortion correction [19]. For the tissue segmentation experiments, six multiple sclerosis patients were scanned using a diffusion spectrum uniform sampling pattern [20] with 167 DWIs ($b_{\max} = 3000 \text{ s/mm}^2$, $T_E = 80.3 \text{ ms}$, $T_R = 5.4 \text{ s}$, $\text{FOV} = 24 \times 24 \times 12 \text{ cm}$, isotropic voxel size 2.5 mm, ASSET factor 2).




Fig. 3. Comparison of (a) model fitting and (b-e) q-DL for neurite orientation dispersion index based on NODDI. q-DL allows strong scan time reduction with moderate contrast loss.

3 Results

In Fig. 1, we show the radial kurtosis [21] measure based on DKI because of its particular susceptibility to noise and because no analytical solutions are known. In the case of denoising and q-DL-R, model parameters are estimated from the reconstructed DWIs by standard model fitting [3] for comparison.

DKI results of denoising for fully sampled data (Fig. 1e) exhibit slightly less noise than the standard pipeline (Fig. 1a) without sacrificing spatial resolution. Moreover, compared with the standard pipeline (Fig. 1a–d), results of q-DL-R (Fig. 1e–h) and of q-DL (Fig. 1j–m) exhibit feasibility of scan time reduction by a factor of twelve. Fig. 2 compares the methods in terms of root-mean-square deviation from model-fitting results of fully sampled data. Despite their approximate nature leading to disparity with model fitting, the proposed methods outperform model fitting when less than 70 DWIs are used. NODDI results (Fig. 3) demonstrate strong scan time reduction with moderate contrast loss when using q-DL. Segmentation results are shown in Fig. 4. The area under the curve (AUC) of the receiver operating characteristic (ROC) for lesions ranged between 0.869 and 0.934 for six different patients. AUC for WM, GM and CSF was consistently above 0.894 for all patients. Thus, q-space data can be used for segmentation directly without a diffusion model, i.e. without the intermediate information loss detailed above. Results were slightly worse when not using FLAIR (between 0.859 and 0.934 for lesions; above 0.892 for WM, GM, CSF). Tailoring the protocol to optimal results in specific applications is subject of future research.




Fig. 4. Direct model-free segmentation. Slices from datasets with the best (upper row, 0.934) and worst lesion AUC (lower row, 0.869) are shown.

4 Discussion

The presented short-scan and model-free protocols open interesting research directions. While state-of-the-art methods [12, 22] require 30 DWIs for NODDI and 64 for radial kurtosis, we require only 8 and 12, respectively. Our results indicate that a considerable amount of information is contained in a limited number of DWIs, and that this information can be better retrieved by deep learning than by model fitting. The network architecture is quite simple compared to other works in deep learning, and yet it works surprisingly well for diffusion MRI. In all presented applications, neural network training takes about one minute on a desktop computer. The network needs to be trained only once and can be applied to any number of datasets, taking 0.03 seconds per dataset, as opposed to several minutes per dataset required by most model fitting methods. Analytical solutions of scalar measure estimation [23] provide reduction of scan time and processing time comparable to q-DL, but are limited to specific scalar measures and acquisition schemes, as opposed to q-DL. A combination with multi-slice imaging is straightforward, yielding additional scan time reduction. Benefits of using spatial neighborhoods [12] for q-DL and (Rician-)noise-robust training [7] can be explored in the future. Our methods require full sampling of the training data only; subsequently, the network can be applied to new subsampled data.

Acknowledgments. We are grateful to S. Pölsterl and B. Menze (TU München) for discussions. V. Golkov is supported by the Deutsche Telekom Foundation.

References

1. Jensen, J.H., Helpert, 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)
2. 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)
3. Veraart, J., Sijbers, J., Sunaert, S., Leemans, A., Jeurissen, B.: Weighted linear least squares estimation of diffusion MRI parameters: strengths, limitations, and pitfalls. *NeuroImage* 81, 335–346 (2013)
4. Nilsson, M., Alerstam, E., Wirestam, R., Ståhlberg, F., Brockstedt, S., Lätt, J.: Evaluating the accuracy and precision of a two-compartment Kärger model using Monte Carlo simulations. *Journal of Magnetic Resonance* 206, 59–67 (2010)
5. Nedjati-Gilani, G., Hall, M.G., Wheeler-Kingshott, C.A.M., Alexander, D.C.: Learning microstructure parameters from diffusion-weighted MRI using random forests. *Joint Annual Meeting ISMRM-ESMRMB*, 2626 (2014)
6. Jarrett, K., Kavukcuoglu, K., Ranzato, M.A., LeCun, Y.: What is the best multi-stage architecture for object recognition? In: *IEEE 12th International Conference on Computer Vision*, pp. 2146–2153 (2009)
7. Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., Manzagol, P.A.: Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research* 11, 3371–3408 (2010)

8. Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.R.: Improving neural networks by preventing co-adaptation of feature detectors. arXiv:1207.0580 (2012)
9. Saxe, A.M., McClelland, J.L., Ganguli, S.: Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In: International Conference on Learning Representations, Banff, Canada (2014)
10. Schultz, T.: Learning a reliable estimate of the number of fiber directions in diffusion MRI. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012, Part III. LNCS, vol. 7512, pp. 493–500. Springer, Heidelberg (2012)
11. Nedjati-Gilani, G.L., Schneider, T., Hall, M.G., Wheeler-Kingshott, C.A.M., Alexander, D.C.: Machine learning based compartment models with permeability for white matter microstructure imaging. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) MICCAI 2014, Part III. LNCS, vol. 8675, pp. 257–264. Springer, Heidelberg (2014)
12. Alexander, D.C., Zikic, D., Zhang, J., Zhang, H., Criminisi, A.: Image quality transfer via random forest regression: applications in diffusion MRI. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) MICCAI 2014, Part III. LNCS, vol. 8675, pp. 225–232. Springer, Heidelberg (2014)
13. Palm, R.B.: Prediction as a candidate for learning deep hierarchical models of data. Master’s thesis, Technical University of Denmark (2012)
14. Bagher-Ebadian, H., Jafari-Khouzani, K., Mitsias, P.D., Lu, M., Soltanian-Zadeh, H., Chopp, M., Ewing, J.R.: Predicting final extent of ischemic infarction using artificial neural network analysis of multi-parametric MRI in patients with stroke. PLoS ONE 6(8) (2011)
15. Ashburner, J., Friston, K.J.: Unified segmentation. NeuroImage 26, 839–851 (2005)
16. Van Leemput, K., Maes, F., Vandermeulen, D., Colchester, A., Suetens, P.: Automated segmentation of multiple sclerosis lesions by model outlier detection. IEEE Transactions on Medical Imaging 20, 677–688 (2001)
17. Poot, D.H.J., den Dekker, A.J., Achten, E., Verhoye, M., Sijbers, J.: Optimal experimental design for diffusion kurtosis imaging. IEEE Transactions on Medical Imaging 29(3), 819–829 (2010)
18. Veraart, J., Van Hecke, W., Sijbers, J.: Constrained maximum likelihood estimation of the diffusion kurtosis tensor using a Rician noise model. Magnetic Resonance in Medicine 66(3), 678–686 (2011)
19. Andersson, J.L.R., Skare, S., Ashburner, J.: How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. NeuroImage 20(2), 870–888 (2003)
20. Menzel, M.I., Tan, E.T., Khare, K., Sperl, J.I., King, K.F., Tao, X., Hardy, C.J., Marinelli, L.: Accelerated diffusion spectrum imaging in the human brain using compressed sensing. Magnetic Resonance in Medicine 66(5), 1226–1233 (2011)
21. Hui, E.S., Cheung, M.M., Qi, L., Wu, E.X.: Towards better MR characterization of neural tissues using directional diffusion kurtosis analysis. NeuroImage 42(1), 122–134 (2008)
22. Paquette, M., Merlet, S., Gilbert, G., Deriche, R., Descoteaux, M.: Comparison of sampling strategies and sparsifying transforms to improve compressed sensing diffusion spectrum imaging. Magnetic Resonance in Medicine 73, 401–416 (2015)
23. Hansen, B., Lund, T.E., Sangill, R., Jespersen, S.N.: Experimentally and computationally fast method for estimation of a mean kurtosis. Magnetic Resonance in Medicine 69(6), 1754–1760 (2013)