Multi-band and In-plane Accelerated Diffusion MRI Enabled by Model-based Deep Learning in q-space And its Extension to Learning in the Spherical Harmonics Domain

Merry Mani[1,3], *Baolian Yang*[2], *Girish Bathla*[1], *Vincent Magnotta*[1,3,4], *Mathews Jacob*[5]

[1]*Department of Radiology, University of Iowa, Iowa City, Iowa*

[2]*GE Healthcare, Waukesha, WI*

[3]*Department of Biomedical Engineering, University of Iowa, Iowa City, Iowa*

[4]*Department of Psychiatry, University of Iowa, Iowa City, Iowa*

[5]*Department of Electrical and Computer Engineering, University of Iowa, Iowa City, Iowa*

October 13, 2021

Correspondence to :

Merry Mani

L309 Papajohn Biomedical Discovery Building,

169 Newton Road

Iowa City, Iowa, 52242

email: merry-mani@uiowa.edu

phone number: (319) 335-9569.

Word count : about 5000

figures+ tables count : 10

Abstract

Purpose: To propose a new method for the recovery of combined in-plane- and multi-band (MB)-accelerated diffusion MRI data.

Methods: Combining MB-acceleration with in-plane acceleration is crucial to improve the time-efficiency of high (angular and spatial) resolution diffusion scans. However, as the MB-factor and in-plane acceleration increase, the reconstruction becomes challenging due to the heavy aliasing. The new reconstruction utilizes an additional q-space prior to constrain the recovery, which is derived from the previously proposed qModeL framework. Specifically, the qModeL prior provides a pre-learned representation of the diffusion signal space to which the measured data belongs. We show that the pre-learned q-space prior along with a model-based iterative reconstruction that accommodate multi-band unaliasing, can efficiently reconstruct the in-plane and MB-accelerated data. The power of joint reconstruction is maximally utilized by using an incoherent under-sampling pattern in the k-q domain. We tested the proposed method on single and multi-shell data, with high/low angular resolution, high/low spatial resolution, healthy/abnormal tissues and 3T/7T field strengths. Further, the learning is extended to the spherical harmonic basis, to provide a rotational invariant learning framework.

Results:

The qModeL joint reconstruction is shown to simultaneously unalias and jointly recover DWIs with reasonable accuracy in all the cases studied. The reconstruction error from 18-fold accelerated multi-shell datasets was $< 3\%$. The micro-structural maps derived from the accelerated acquisitions also exhibit reasonable accuracy for both healthy and abnormal tissues. The deep learning (DL) enabled reconstructions are comparable to those derived using traditional methods.

Conclusion: qModeL enables the joint recovery of combined in-plane- and MB- accelerated dMRI utilizing deep learning.

Keywords: q-space Deep Learning, autoencoder q-space, machine learning diffusion MRI, NODDI deep learning, k-q acceleration, joint reconstruction, multi-band joint reconstruction, CAIPI sampling q-space, spherical harmonics, rotational invariant q-space learning

Introduction

The sensitivity of MRI to the microscopic diffusing motion of water molecules is exploited in diffusion weighted magnetic resonance imaging (dMRI) to study the tissue microstructure non-invasively. Classical tissue microstructure analysis methods have relied on simple models, such as the single tensor model that is widely employed in diffusion tensor imaging (DTI) (1). DTI-derived indices such as fractional anisotropy (FA) have been shown to be highly sensitive to tissue micro-structural changes arising from a number of factors such as neurodegeneration, demyelination, neurodevelopment and aging (2, 3).

Recently, advanced microstructure analysis using biophysically inspired multi-compartment models are gaining popularity because they possess greater specificity than DTI in relating the dMRI signal to the underlying cellular microstructure (4–8). However, advanced biophysical models are associated with a high acquisition burden compared to DTI. A greater number of diffusion encoded acquisitions, often spanning over multiple b-shells, are required for model fitting in these complex models. The increased number of acquisition lead to long scan time. Moreover, the longer TE associated with the diffusion preparations not only are detrimental in worsening the EPI artifacts, but also compromise the spatial resolution.

Simultaneous multi-slice imaging (SMS) (9) employing multi-band excitation pulses has been proposed as a viable solution to reduce the scan time (10) and is being widely employed in multi-shell dMRI studies (11). Reduction in TE can be achieved using partial Fourier methods or multi-shot methods, and can be combined with SMS imaging such that both TR and TE are reduced. A more desirable combination, however, is SMS imaging combined with in-plane acceleration methods that skip phase encoding (PE) lines, which can maximally accelerate the acquisition while minimizing the TE. However, this combination is challenging in practice (12). Both SMS and in-plane acceleration rely on the same coil hardware (and the coil sensitivity maps) to remove the aliasing artifacts arising from the respective accelerations. As the acceleration factor from both components increase, the reconstruction becomes ill-conditioned. Constrained reconstruction methods are required to overcome the ill-conditioning of the such reconstructions.

Previously, DAGER, a joint reconstruction method, was proposed for the reconstruction of dMRI data acquired with both SMS and in-plane acceleration (13). DAGER uses the similarity of information from closely sampled q-space data as an additional reconstruction prior. The similarity information is enforced using a spherical covariance function based on Gaussian process model (14). Further improvement in performance was obtained by utilizing a variable k-space sampling pattern among neighboring q-space points. Together with the variable k-q sampling pattern and joint reconstruction, very high acceleration factors close to 12-fold was achieved, by combining SMS acceleration of 4-fold and in-plane acceleration of 3-fold. A disadvantage of DAGER is that large number of q-space points (> 60) are needed to effectively impose

the similarity prior. NEATR is another DL-based method that was proposed for reconstructing dMRI data acquired with both MB and in-plane acceleration (15). NEATR relies on a pre-trained U-NET to provide the phase for solving the complex reconstruction and was only demonstrated in a limited experimental set-up, due to the highly demanding training process. Moreover, because of the limitations of the training process, there was no flexibility to modify the under-sampling pattern.

In this work, we demonstrate the utility of the previously proposed qModeL framework (16, 17) for the recovery of dMRI data acquired with both MB and in-plane acceleration. The qModeL framework synergistically combines the power of joint reconstruction and DL in the setting of an iterative model-based reconstruction. Unlike NEATR, qModeL training do not rely on in-vivo data for the learning of phase estimates. Instead, phase training is performed using experiments with simulated phase. Hence, the phase compensation embedded in qModeL is not tied to the specific acquisition settings of the training data and is generalizable to a broad range of acquisition parameters. Importantly, this relaxes the training data demands (more details in the next section). Additionally, because data-consistency to the measured k-space data is explicitly imposed during reconstruction, the qModeL framework can exploit random k-q sampling patterns, similar to DAGER. We achieve in-plane acceleration using under-sampled multi-shot technique, similar to DAGER and NEATR. However, unlike NEATR, we use randomized multi-shot technique so that we distribute the under-sampling artifacts incoherently. In contrast to DAGER, the incoherent under-sampling is not restricted to a neighborhood around a given q-space point. Instead, the k-q under-sampling is fully randomized for all the sampled points. The free design of the k-q under-sampling pattern and the ease of phase training provides greater flexibility for qModeL to be accommodated to a broad range of acquisition settings.

The contribution of this paper includes the following. (i) we demonstrate the efficacy of the qModeL framework to solve the problem of reconstructing MB-accelerated DWIs that are also under-sampled both in k-space and in q-space. We show very high accelerations achieved from the joint multi-shell reconstructions. Further, joint reconstruction from fewer q-space points (< 60) are shown, in conjunction with higher in-plane acceleration. (ii) We demonstrate the flexibility of the model-based DL formulation to accommodate a variety of acquisition settings, including different field strengths (including differing gradient hardware), and different multi-band and in-plane accelerated protocols (randomized k-q sampling patterns, various resolutions, single shell, multi-shell). Importantly, once the q-space has been pre-learned, the reconstructions from all these scenarios can be performed without having to re-train the neural network for each test case. (iii) We also show the translation of the qModeL DL framework to reconstruct healthy and abnormal data, without involving any additional training on abnormal data. (iv) Furthermore, we expand the scope

of qModeL learning beyond the domain of raw diffusion signal, to also train the DAEs on the spherical harmonic coefficients. The new formulation termed qModeL-SH enables the use of DAEs trained over a specific set of q-space samples to be re-used for another sampling pattern.

Theory

Brief Review of qModeL

qModeL is a joint reconstruction method built on a learning-based framework (17). It is designed to learn a latent representation of the q-space signals to be reconstructed, from appropriate training data. To learn the the latent representation or the low-dimensional manifold to which the q-space signal belongs, we employ a denoising autoencoder neural network (DAE). Q-space signals derived from a variety of tissue micro-structural composition, are provided as training data to the DAE. The learned q-space manifold prior is then embedded in an iterative model-based reconstruction, where the prior information about the q-space is used to constrain the recovery.

Of note, the training data for learning the q-space signal manifold can be extracted either from the q-dimension of in-vivo dMRI acquisitions (with matched q-space) from the different voxels, or can be generated numerically based on biophysical models (8, 18–24) as described in (17). While in-vivo data-driven training can have benefits of being model-free, biophysical models provide a convenient method to learn the q-space manifold at high SNR. This is especially useful when the acquisition setting becomes challenging to perform the experiments in-vivo. For example, high b-value acquisitions on MRI scanners with lower performance gradient systems require very long TEs thus requiring multiple averages to obtain the desired SNR for training purposes. This leads to long scan times or limited brain coverage. By simply matching the q-space sampling parameters in the biophysical models to the desired q-space, such training data can be easily generated. Moreover, biophysical models provide a convenient method to generate training data that span a wide range of parameters such as diffusivities and volume fractions and fiber orientations, whereas, the parameter span of data-driven training may be clustered around certain ranges (25). This convenient framework also allows for re-training the neural network to learn a new q-space manifold by simply modifying the q-space sampling parameters in the biophysical models. One disadvantage of previous data-driven DL methods is the generalizability of the learned information. In qModeL, instead of learning the image content, the goal is to learn the latent q-space representation. Hence, the specific imaging parameters of the acquisition used for the training data is not a concern for qModeL training as long as the q-space parameters are matched. Figure 1 shows an example of a case where the qModeL joint reconstruction is performed using data-driven and physics-driven training. Notably, the phase training can be performed on the in-vivo

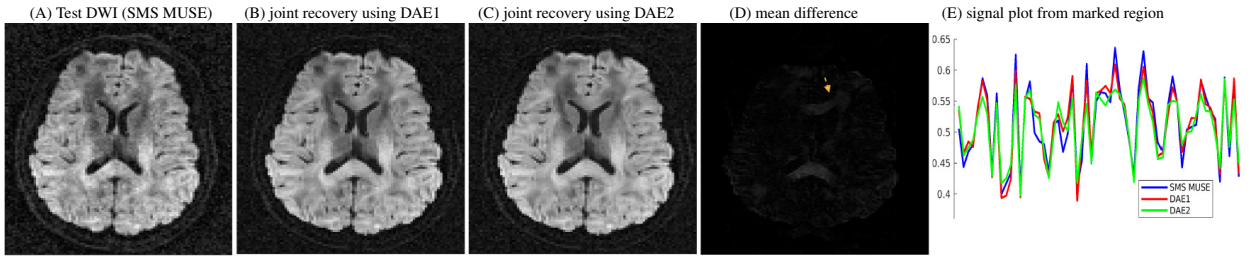

Figure 1: qModeL joint recovery utilizing DAEs that were trained using: (i) biophysical modeling (DAE1), and (ii) in vivo training data (DAE2). (A) shows a DWI that belongs to a test data acquired with 60 q-space samples at b-value of 1000 s/mm^2. (B) - (C) show the same DWI from the qModeL joint recovery based on DAE1 and DAE2 respectively. For the training of DAE2 in (C), training data was derived from a random set of two subjects from a pre-existing database with DWIs collected at b=1000 s/mm^2 and 60-directions, from which only magnitude data was available. The specific imaging parameters of the test data and the training database are very different. For example, the test data is acquired with a 4-shot acquisition, at 1.7mm resolution with a TE of 58.3ms, whereas the images from the pre-existing database was acquired with single-shot EPI at 2mm resolution, with a TE of 88ms. (D) shows the mean difference between the reconstruction from the two DAEs (magnified by 10 times). (E) shows the q-space signal plots from a region with relatively higher difference (marked in (D)) predicted by DAE1 and DAE2. The plot from both DAEs are noted to closely follow the q-space signal plot of the test data. This demonstrates that the learning involved in qModeL is pertinent only to the q-space information and not tied to the specific acquisition parameters of the training data.

magnitude training data as well as the numerically generated training samples.

Joint Reconstruction of Combined MB- and in-slice Accelerated data

In this work, we propose the utilization of the qModeL framework for the reconstruction of combined MB- and in-plane accelerated data. To efficiently combine in-plane acceleration (by skipping PE lines) with MB-acceleration, we make use of the incoherent k-q under-sampling method (13, 26, 27). Figure 2 illustrates the proposed sampling in the k-q domain. Specifically, for every TR, we randomize the PE lines that are used to sample the k-space of every q-space point. We achieve this by employing a modified blipped CAIPI technique with multi-shot readouts (28, 29) that use N_s shots in the non-accelerated case (figure 2(A)), and skipping random shots during each q-space point for the accelerated case (figure 2(B)). This result in incoherent under-sampling patterns along the q dimension, which can be exploited by joint reconstruction methods to afford very high acceleration factors.

Because the aliasing artifacts from the in-plane under-sampling is incoherent along the q-dimension, an additional reconstruction prior along the q-dimension is expected to significantly aid the unaliased recovery of the DWIs. We expect the pre-learned q-space signal prior to provide this additional constraint. Assuming MB bands to be simultaneously excited by the multi-band RF excitation, the qModeL joint reconstruction

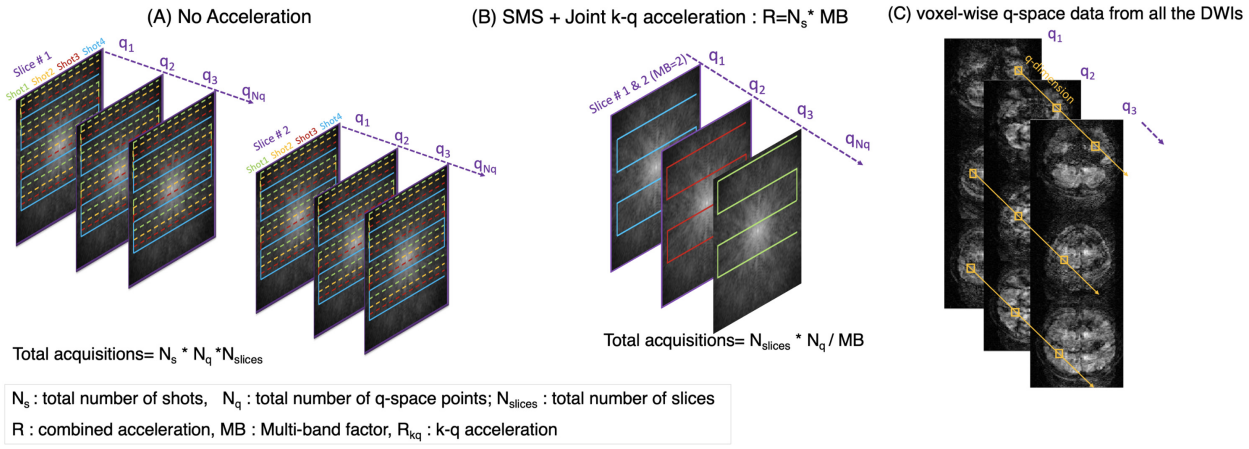

Figure 2: (A) shows the case with no acceleration involving the acquisition of two slices and (B) shows the case with combined MB- and in-plane acceleration. For the illustrated case in (B) with MB=2 and in-plane acceleration of 4-fold, a combined acceleration of R=8 is achieved ($\text{MB}\times\text{R}_{kq} = 2 \times 4$). To randomize the 4-fold in-plane under-sampling pattern for every q-space point in (B), we make use of multi-shot trajectories. A ($N_s =$) 4-shot EPI trajectory is designed to sample the k-space in the non-accelerated case in (A). For the accelerated case in (B), only 1 random shot (out of the 4 k-space shots) is used per q-space point (each shot is marked using a different color). (C) shows the extraction of the voxel-wise q-space signals to impose the q-space prior during the iterative joint reconstruction.

for the simultaneously excited set of DWIs from all q-space points can be written as:

$$\mathbf{S}^{*}_{\mathbf{mb}} = \underset{\tilde{\mathbf{S}_{mb}}}{\operatorname{argmin}} \left\| \mathcal{A}_{mb}(\mathbf{S}_{mb}) - \widehat{\mathbf{Y}} \right\|_2^2 + \lambda_1 \left\| \mathcal{P}_{\Theta}(\mathbf{S}_{mb}) \right\|_2^2 + \lambda_2 \left\| \mathbf{S}_{mb} \right\|_{TV} \qquad [1]$$

$$\text{where } \mathcal{P}_{\Theta}\big(\mathbf{S}_{mb}(r)\big) = \mathcal{D}_{\Theta}\big(\mathbf{S}_{mb}(r)\big) - \mathbf{S}_{mb}(r).$$

The notation, $\mathbf{S_{mb}}$ is used to denote the set of DWIs from the MB bands : $\{\mathbf{S_1}, \mathbf{S_2}..\mathbf{S}_{MB}\}$. $\mathbf{S_{mb}}$ is reconstructed jointly from all N_q q-space samples and is of size $N_x \times N_y \times MB \times N_q$. The data-consistency term makes use of a forward operator, $\mathcal{A}_{mb}(\mathbf{S}_{mb})$ that accounts for the slice folding due to the multi-band excitation from the MB bands and the specific k-space sampling pattern. Specifically, $\mathcal{A}_{mb}(\mathbf{S}_{mb})$ incorporates the following operations in a coil-by-coil fashion for every q-space point: $\mathcal{A}_{mb}(\mathbf{S}_{mb}) = \mathcal{M}_q\Big(\sum_b (P_b \circ \mathcal{F} \circ \mathcal{C}_b)(\mathbf{S}_b)\Big)$. Here, the symbol $\circ$ denotes the composition of the operators, $\mathcal{C}_b$ denotes the multiplication with coil sensitivities of the b^{th} band, $\mathcal{F}$ denotes the Fourier transform operation, P_b denotes the multiplication with linear phase factor that is used to achieve the desired shift in the images (i.e, controlled aliasing) for the b^{th} band, and $\sum_b$ denotes the summation of the phase modulated k-space data from the MB bands. $\mathcal{M}_q$ denotes the k-space sampling mask which accounts for the specific under-sampling pattern used for the q^{th} q-space point. $\widehat{\mathbf{Y}}$ is the measured multi-channel k-space data.

The second term in Eq [1] represents the projection error of the reconstructed q-space signal onto the learned

manifold. $\mathcal{D}_\Theta$ denotes the encoder-decoder mapping learned by the DAE. The projection error is computed in a voxel-wise manner by extracting the q-space signal $\mathbf{S}_{mb}(r)$ (as shown in figure 2(C)), which is projected to the pre-learned q-space as $\mathcal{D}_\Theta(\mathbf{S}_{mb}(r))$. Minimizing the ℓ_2-norm of the projection error, $\mathcal{P}_\Theta(\mathbf{S}_{mb}(r))$, constraints the reconstructed q-space signal at every voxel location r to be as close to the desired q-space measurement space. The third term imposes an additional total variation (TV) regularization on the DWIs.

qModeL-SH

As noted earlier, the qModeL pre-learning pertains to the q-space of the data to be reconstructed. One approach to pre-learn the q-space is to train the DAE by pre-matching the q-space points of the training signal to that of the data to be reconstructed. However, a limitation of this training approach is that the DAE, $\mathcal{D}_\Theta\big(\mathbf{S}_{mb}(r)\big)$, is unable to handle the reconstruction when the number of q-space points to be reconstructed differs from the number of q-space points at which the DAE was pre-trained. This is because the number of input neurons in $\mathcal{D}_\Theta$ is fixed to the number of reconstruction points in $\mathbf{S}_{mb}(r)$. Here, we propose an alternative training approach that offers more flexibility, termed qModeL-SH. In qModeL-SH, we make use of spherical harmonics (SH) transformations (30) and learn the q-space in the transformed domain. The SH basis allows the easy transformation from one set of q-space points to another and also enable a rotation-invariant learning. The qModeL-SH optimization problem can be written as

$$\mathbf{S}^{*}_{\mathbf{mb}} = \underset{\tilde{\mathbf{S}_{mb}}}{\operatorname{argmin}} \Big\| \mathcal{A}_{mb}(\mathbf{S}_{mb}) - \widehat{\mathbf{Y}} \Big\|_2^2 + \lambda_1 \left\| \mathcal{P}_\Theta(\mathbf{S}_{mb}) \right\|_2^2 + \lambda_2 \left\| \mathbf{S}_{mb} \right\|_{TV} \qquad [2]$$

$$\text{where } \mathcal{P}_\Theta\big(\mathbf{S}_{mb}(r)\big) = \mathcal{H}^{\dagger}_{q_r}\big\{\mathcal{D}^{q_t}_{\Theta}\big(\underbrace{\mathcal{H}_{q_r}\{\,\mathbf{S}_{mb}(r)\}}_{\mathbf{C}_{l,m}}\,\big)\big\} - \mathbf{S}_{mb}(r).$$

Here, the notation $\mathcal{H}_{q_r}$ stands for the spherical harmonics transformation that transforms the diffusion signal from the polar basis to the spherical harmonics basis, and $\mathcal{H}^{\dagger}_{q_r}$ denotes the transformation back to the polar basis. Unlike the previous DAE which was trained on the raw q-space signals, the new DAE, $\mathcal{D}^{q_t}_{\Theta}(\mathbf{C}_{l,m})$, is trained on the SH coefficients, $\mathbf{C}_{l,m}$, for a given order l and degree m. Here, the superscript q_t is used to emphasize the fact that the q-space points used during the training of the DAE can be different from the q-space points q_r used during reconstruction. During reconstruction, the operation $\mathcal{H}_{q_r}\{\,\mathbf{S}^{q_r}_{mb}(r)\}$, transforms the q-space points to be reconstructed to the SH coefficients. Then, the projection to the pre-learned space is performed making use of the SH coefficients. The resulting denoised SH coefficients are utilized to recover the predicted diffusion signal back as $\mathcal{H}^{\dagger}_{q_r}\big\{\mathcal{D}^{q_t}_{\Theta}(\mathbf{C}_{l,m})\big\}$. The projection error $\mathcal{P}_\Theta\big(\mathbf{S}_{mb}(r)\big)$ is computed as the residual of the predicted diffusion signal from the DAE and the CG-update, as before. Note that since the DAE, $\mathcal{D}^{q_t}_{\Theta}$, is trained on the SH coefficients, $\mathbf{C}_{l,m}$ of order l, the transformation $\mathcal{H}_{q_r}$ also employs the same order l for reconstruction.

Methods

To test the qModeL joint recovery for the MB- and in-plane-accelerated case, we used a Stejskal-Tanner diffusion encoding in a 4-shot EPI pulse sequence with multi-band capability. 90° sinc pulses and 180° VERSE (31) pulses respectively were frequency modulated and summed to generate the multi-band RF pulses. Experiments were performed on three different hardware to collect datasets at various settings.

Datasets

Table 1 shows the details of the datasets used in the study. All human experiments were performed following the Institutional Review Board requirements and obtaining informed written consent. Five datasets were used in the current study. Whereas dataset 1, 2 and 5 were from normal subjects, dataset 3 was from a volunteer diagnosed with multiple sclerosis and dataset 4 was a volunteer with a single large lesion.

Table 1: Details of Datasets used for Various Experiments

ID	Field Strength	FOV (mm)	Slice Th. (mm)	Matrix size	b-value(s/mm^2) # Directions	N_s/ MB	pF / TE (ms)	# coils	G_{max} (mT/m)/ SR_{max} (T/m/s)	Observation
Dataset 1	3T	220	2	128 x 88	1000	4	.68	32	33	normal
		.			20	3	57		150	
Dataset 2	3T	210	1.5	144 x 96	1000, 2000,3000	4	.71	48	70	normal
					60,60,60	3	60.7		150	
Dataset 3	3T	210	1.5	144 x 96	1000, 2000,3000	4	.71	48	70	Several small
					60,60,60	3	60.7		150	diffuse lesions
Dataset 4	7T	240	1.2	200 x 132	1000	4	.66	32	100	Single large
					60	2	48.4		200	lesion
Dataset 5	3T	210	1.5	144 x 96	1000	4	.71	48	70	normal
					60	3	60.7		150	

Joint k-q Under-sampling experiments performed on the different datasets

All datasets were acquired in a fully sampled manner. The term "fully sampled" refers to the case using the k-space data from all 4 shots. The datasets were retrospectively under-sampled by random shot under-sampling to test the recovery of the combined in-plane and MB-accelerated case. The notation R_{kq} denotes the k-q acceleration and R denotes the combined acceleration (MB $\times$ R_{kq}). Note that the "fully-sampled" (i.e, $R_{kq} = 1$) and under-sampled cases ($R_{kq} > 1$) have partial Fourier (pF) under-sampling, which is not counted in the acceleration. Because of the multi-shot nature of the acquisition, the T2* and the geometric distortions of the fully sampled data are matched to the under-sampled cases.

The following experiments were performed on the various datasets. The combined in-plane and MB-acceleration to achieve R= 12 (MB=3, R_{kq} =4) were experimented on datasets 1,2,3, and 5, whereas R= 8 (MB=2, R_{kq} =4) was tested on dataset 4. More aggressive joint k-q under-sampling experiments were

performed on the multi-shell datasets (datasets 2-3) with more q-space points. Here, in addition to the 4-fold in-plane under-sampling, about 20-33% of the total q-space samples were not sampled (17). Thus, combined acceleration of up to R=15-18 (MB=3, R_{kq} =5-6) can be achieved. The recovery of the entire 180 volumes, including the not-sampled volumes, were studied using this experiment. Datasets 5 and 2 was utilized to demonstrate the single-shell and the multi-shell qModeL-SH reconstructions respectively. Experiments were performed for the fully-sampled case of R=3 (MB=3, R_{kq} =1), and the accelerated case of R=12 (MB=3, R_{kq} =4), for both the datasets.

Reconstruction

The qModeL joint reconstruction is solved iteratively using the plug-and-play ADMM method (17, 32) and was implemented in Matlab. The DAE training was performed using biophysical modeling because of its many advantages compared to the data-driven training. The details of the biophysical models (8, 18–24, 33) are provided in the supporting information. Phase was simulated during training and was added to the magnitude q-space data generated by the biophysical models. Specifically, for every diffusion signal generated, a phase was added to the magnitude signal by multiplying with $exp(i\theta)$ where θ was randomly chosen to have a value between 0 to 2π. Furthermore, complex Gaussian noise at various levels were added to the data (17). The resulting complex data was used for training the qModeL DAE. For qModeL-SH, the DAE is trained on the spherical harmonic coefficients instead of on the raw diffusion signal. To achieve this, the SH transform is applied on the complex data generated from the biophysical models. SH order of $l = 6$ is used in all our experiments (30, 34). The complex SH coefficients are used as training inputs to the qModeL-SH DAE. xFor both qModeL and qModeL-SH DAE, the training for the multi-shell case was performed jointly for all shells. This is achieved by concatenating the q-space data from all the shells into a single vector. The DAE training was performed in TensorFlow. The training data were separated into real and imaginary parts for training. A 3-layer fully connected network was used for the encoder and for the decoder with 15 neurons in the bottleneck layer each for the real and the imaginary part for both qModeL and qModeL-SH training. The training procedure remained the same for 3T/7T datasets. The weights from the network, that was trained in TensorFlow, were saved and re-used in Matlab to implement $\mathcal{D}_\Theta(\mathbf{S})$ for the qModeL joint reconstruction.

While the training of the DAE on the raw diffusion signal is straightforward, the training on the SH coefficients are hard to achieve. This is due to the high dynamic range of the SH coefficients. The "DC component" (i.e., the SH coefficient corresponding to $l = 0$) is usually much stronger in amplitude compared to the rest of the coefficients, which makes the learning difficult. To overcome this problem, we developed

a modified DAE to learn only the non-DC terms. Hence, during qModeL-SH reconstruction, the denoising only pertains to the non-DC terms while the DC term remains the same. For multi-shell experiments, the SH transform of the raw diffusion signal was performed individually for each shell. The DC components are removed and the non-DC terms from all the shells are concatenated into a single vector to form the training data. During the joint recovery, the denoising only pertains to the non-DC terms for each shell, while the DC terms remains the same.

As detailed in the previous work (17), qModeL is a phase-compensated reconstruction. The forward model of qModeL in Eq 1-2 involves an initial phase estimation (performed from the data itself, accounted into $\mathcal{C}_b$ as detailed in (17)) similar to that used in previous methods such as MUSE (35) and its extension to the SMS setting (36). Hence, qModeL can be considered as a regularized SMS MUSE reconstruction, where regularization from the DAE and TV are additionally present. For comparison of the qModeL reconstruction, we utilize the SMS MUSE method (36) and the SMS MUSSELS method (29). These methods serve as a guidance to verify the correctness of diffusion contrasts generated by the DAEs. To quantify the performance of the under-sampled reconstructions, the root-mean squared error (RMSE) and structure similarity index (SSIM) were calculated with respect to the qModeL reconstruction with no under-sampling. The regularization parameters λ_1 and λ_2 were chosen by picking those that resulted in the lowest RMSE with respect to the fully-sampled case.

Results

We first demonstrate the performance of qModeL joint reconstruction to simultaneously unalias the MB aliasing and jointly recover the full set of q-space data. Figure 3 shows the qModeL joint reconstruction result on dataset 1 for the case with no under-sampling. For comparison, the SMS MUSE results from the individual reconstruction of the fully sampled data are shown. Two representative DWIs from the 20 diffusion directions are shown from all the three slices. The results confirm that the diffusion contrasts

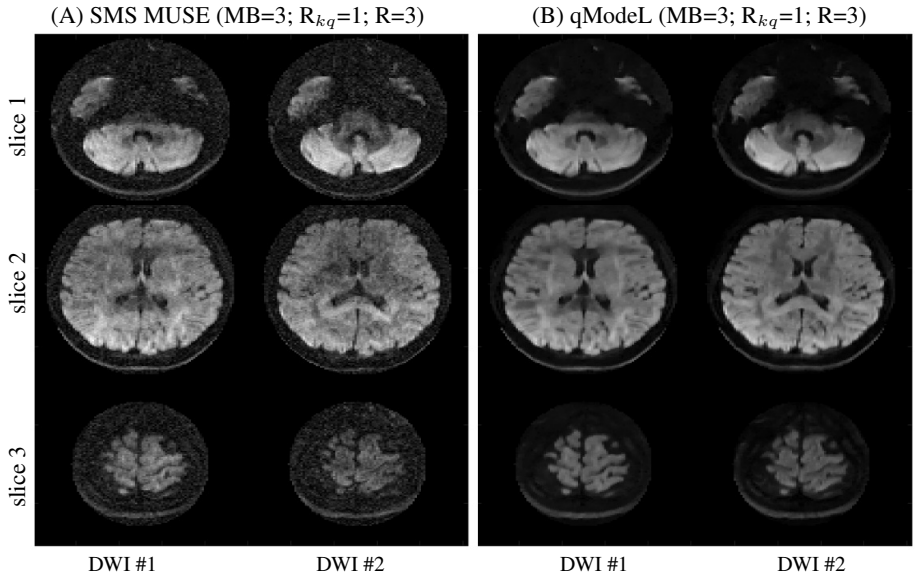

Figure 3: (A) SMS MUSE individual reconstruction, (B) qModeL joint reconstruction. Two representative DWIs from the full set of 20 DWIs are shown in each case. All images are displayed on the scale [0 1].

are correctly preserved by the DAEs. When compared to the SMS MUSE individual reconstruction of the DWIs (figure 3(A)), the improvement in SNR is evident for the qModeL joint reconstruction (figure 3(B)).

Figure 4 shows the results of the qModeL reconstruction for the under-sampled cases on the same dataset, corresponding to R = 6, 12, along with the residual error maps from the reconstructions. The RMSE and SSIM are 0.0125 and 0.994, respectively for R=6, and 0.0185 and 0.9771 for R=12. FA maps were computed by fitting a DTI model to the data recovered from the various accelerations. Figure S1 in the supporting information show the FA maps and the direction encoded color (DEC) maps. The average error in these maps compared to the case with no under-sampling are reported in the figure. The average angular error (computed for voxels with FA$>$0.2) is noted to be $< 12^{\circ}$ for R=12. Note that the same DAE was utilized in all of the above experiments involving the different randomized k-q under-sampling patterns and accelerations.

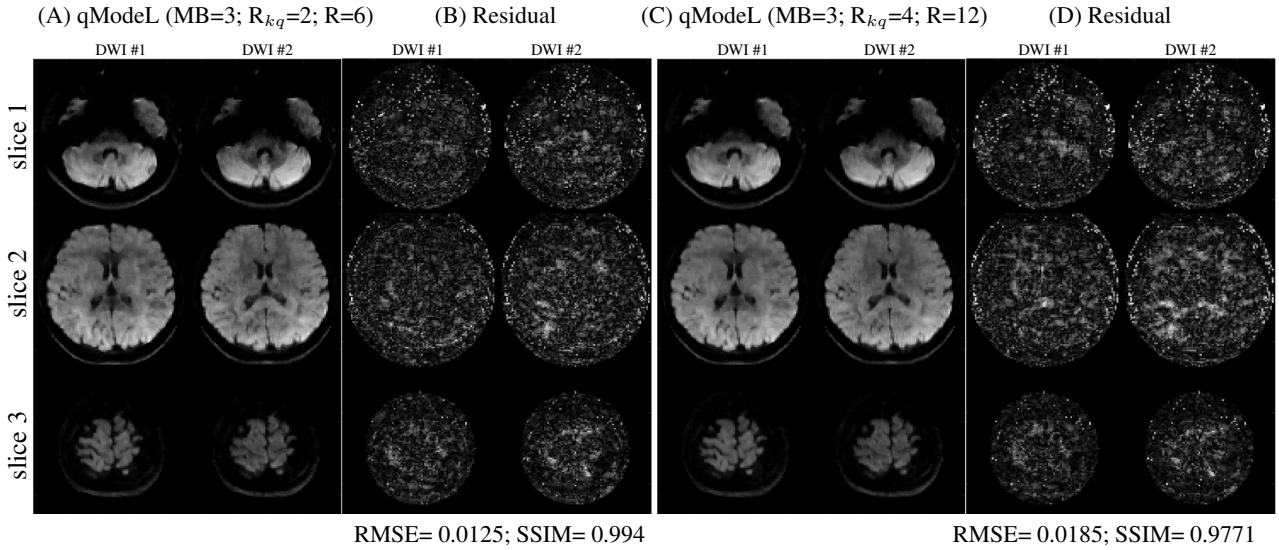

Figure 4: (A), (B) : qModeL joint recovery for R=6 and the residual. (C), (D) qModeL joint recovery for R=12 and the residual. Two representative DWIs from the full set of 20 DWIs are shown in each case. The normalized RMSE and SSIM for the above cases are reported with respect to the qModeL joint recovery at R=3. The residuals are scaled 10x for display. All images are displayed on the scale [0 1].

Multi-band Multi-shell Joint Reconstruction

For the experiments on dataset 2-3, a new DAE was trained to pre-learn the q-space signal space of the multi-shell data. Figure 5 show the experiments dataset 2 that was k-q under-sampled. DWIs from all q-space points that span the three shells are jointly reconstructed for all the simultaneously excited slices using qModeL joint recovery. Three DWIs from the above set, that belong to the b1000, b2000 and b3000 shells, are shown in figure 5 (the full set of simultaneously excited slices are given in figure S2). The SMS MUSE and SMS MUSSELS individual reconstruction of the DWIs are given in figures 5, S2. Notably, compared to these individual reconstructions, the SNR of the qModeL joint reconstruction are significantly improved

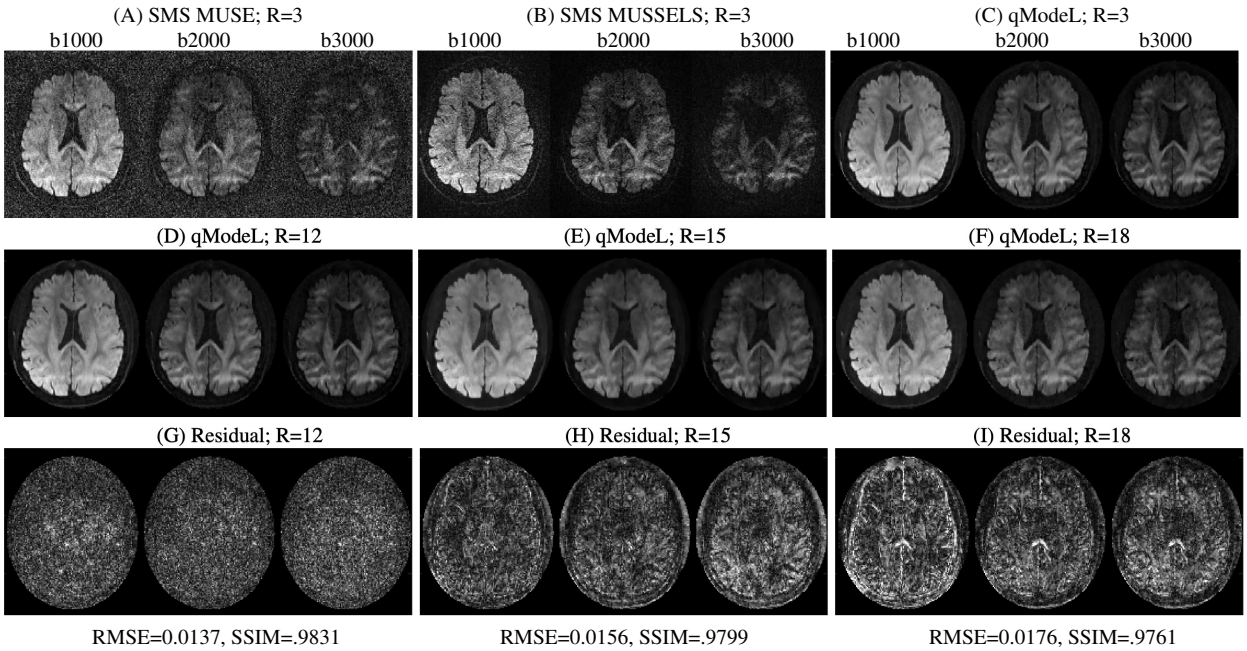

Figure 5: Multi-band multi-shell joint reconstruction (dataset 2): (A), (B) SMS MUSE and SMS MUSSELS individual reconstruction, and (C) qModeL joint reconstruction for R=3 (MB=3, Ns=4, R_{kq} =1) shown from 3 DWIs corresponding to b1000, b2000 and b3000. The images from one of the bands is given above. The images corresponding to all three bands in the multi-band acquisition are included in figures S2. (D), (G) shows the qModeL joint reconstruction at R=12 (MB=3, Ns=1, R_{kq} =4) and its residual error with respect to qModeL at R=3. (E), & (H) corresponds to R =15 (MB=3, R_{kq} =5), and (F), (I) corresponds to R=18 (MB=3, R_{kq} =6). All images are displayed on the scale [0 1]. The residuals is scaled 10x for display. The RMSE and the SSIM are reported for each case.

especially for the DWIs from the higher b-shells. Figure 5(D) shows the qModeL reconstruction results of the same DWIs from the MB-data with 4-fold in-plane acceleration. The reconstruction error for this case compared to the qModeL case at $R = 3$ is given in figure 5(G).

Further k-q acceleration is tested on the MB-accelerated setting by skipping random q-space points altogether. Figures 5(E), (H) correspond to the qModeL reconstruction and the residuals from the 15-fold accelerated case (MB=3, R_{kq}=5). Here 20% of the q-space points were not sampled (or equivalently, only 144 out of 180 q-space points were sampled), however all the 180 q-space points are reconstructed. The RMSE and SSIM for this case are 0.0156 and 0.9799 respectively. Similarly, figures 5(F), (I) corresponds to R=18 (MB=3, R_{kq}=6). Here, 60 random q-space points out of the 180 points (33%) are skipped. The full set of 180 DWIs are reconstructed using the joint recovery method in Eq [1]. The average error in all the 180 DWIs for the accelerated case at R=18 is 0.0176. The (10$\times$ magnified) reconstruction error for these cases are given in figure 5. Figure S2 in the supporting information show reconstruction from all the simultaneously excited slices.

Figure S3 in the supporting information show the breakdown of the RMSE and SSIM for the various gra-

dient directions and b-shells. Figure S3 (A)-(B) plots the overall RMSE and SSIM for each q-space point from the all tissue types, while figure S3 (C)-(D) plots the RMSE and SSIM for each q-space point from a white matter ROI (the ROI chosen for the plot is shown in figure S3(E)).

For the multi-shell DWIs reconstructed from the fully-sampled and under-sampled cases, microstructural maps were derived by making use of the NODDI model (6).The neurite density index (NDI) computed for R=3, 12 and 18 are shown in Figure S4 in the supporting information, while the orientation dispersion index (ODI) are shown in figure 6. The NDI and ODI maps for qModeL are noted to closely follow the SMS MUSE maps. Figure 6 also show the error in the ODI maps computed over the white matter (WM) region, while the unmasked ODI error maps as well as the NDI error maps are shown in Figure S4. The average error in these maps are also reported in the figure.

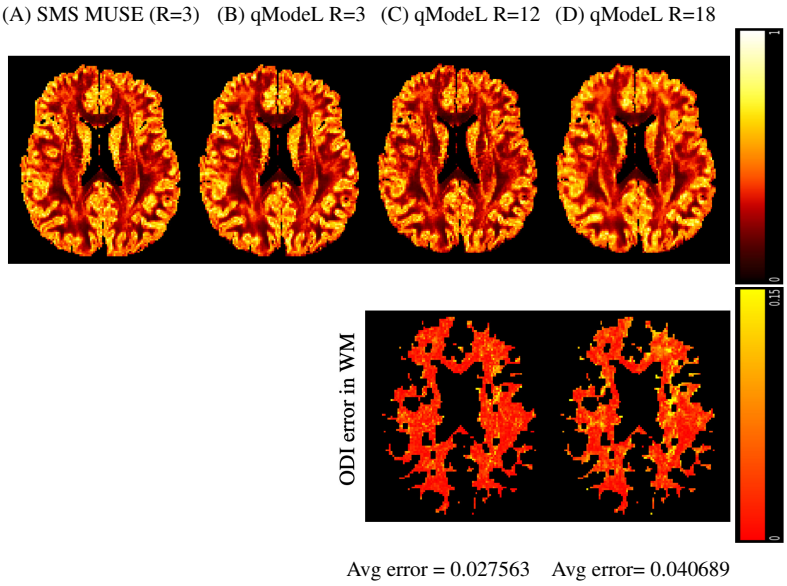

Figure 6: Orientation dispersion index (ODI) derived using the multi-shell data of dataset 2. (A) - (B) shows the ODI maps derived from DWIs reconstructed using SMS MUSE and qModeL respectively at R=3. (C) - (D) shows the ODI maps derived from DWIs reconstructed using qModeL respectively at R=12 and R=18. The error in the ODI maps windowed over the WM (defined as FA $>$ 0.2) are shown in the bottom row for the cases R=12 and R=18.

Micro-structural maps of lesion using qModeL joint reconstruction of accelerated cases

We tested the capability of qModeL to also represent q-space signals from voxels with abnormality, using the DAEs trained on the biophysical models. For this purpose, dataset 3 was acquired on a subject diagnosed with multiple sclerosis. The T2 FLAIR data from the subject exhibit diffuse hyper-intensities in several regions that may indicate active neuroinflammation or other ongoing neurodegeneration. Figure S5 in the supporting information show the intra-axonal volume fraction (v_{ic}), the extra-axonal volume fraction (v_{ec}) and the NDI maps ($v_{ic}/(v_{ic} + v_{ec})$) derived from the MB-accelerated cases for the slice sets that were simultaneously excited. The maps derived from various acceleration factors are shown. Figure 7 shows the v_{ec} maps from two slices, along with the corresponding T2 FLAIR data, where both large and small lesions are indicated in blue (A more comprehensive set of images can be found in figure S6 in the supporting information). Figure S7 in the supporting information show the v_{ic} in comparison to the T2 FLAIR images

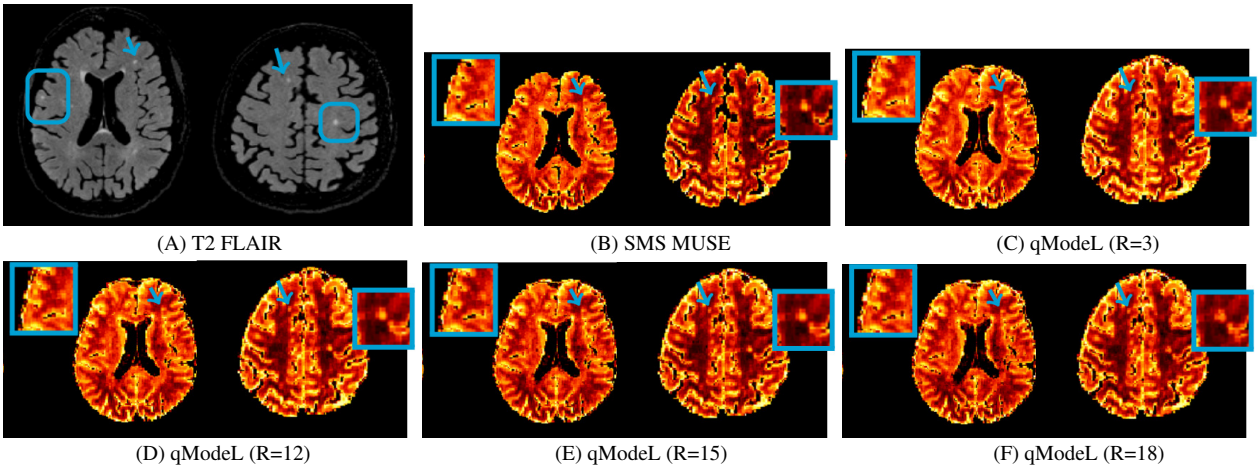

(A) T2 FLAIR (B) SMS MUSE (C) qModeL (R=3)

(D) qModeL (R=12) (E) qModeL (R=15) (F) qModeL (R=18)

Figure 7: The extra-cellular volume fraction (v_{ec}) computed using the NODDI model is shown for various acceleration factors in comparison to the T2 FLAIR data (shown from two slices). The hyper-intensities on the T2 FLAIR is seen as hyper-intensities in the v_{ec} maps also as indicated by the arrows. Some regions are zoomed for better visualization of the hyper-intensities. The qModeL reconstruction can faithfully reconstruct the lesion data at the various acceleration factors.

where the T2 hyper-intensities correspond to hypo-intensities. These images confirm the ability of qModeL to faithfully represent the lesion data in agreement to the SMS MUSE maps. A radiologist evaluated the various reconstructions in conjunction with the T2 FLAIR data. The report concurs to the above observation that the MS lesion features are preserved in the qModeL images at different accelerations. The radiologist also noted that at R=18, the subcortical lesions become less sharp, but is still conspicuous on the images. Figure S8 in the supporting information show the breakdown of the RMSE and SSIM for the various gradient directions and the b-shells for dataset 3.

Notably the same DAE was used for the reconstruction of datasets 2-3 for all the experiments, including the randomized under-sampling in q-space. No additional training was needed to reconstruct the data with the MS lesions.

Reconstruction of 7T data with lesion

While 7T diffusion data generally have different contrasts than 3T data due to the different relaxation times, the q-space relations remain the same. Hence DAEs pre-trained for 3T datasets can be utilized to reconstruct 7T datasets and vice-versa, with no modification needed. We demonstrate this flexibility using dataset 4 and dataset 5 (in the next section).

Dataset 4 is a lesion data that was collected on the 7T field strength. Figure S9 show the qModeL reconstruction of dataset 4 along with the SMS MUSE and the SMS MUSSELS reconstruction for the fully sampled case. The qModeL reconstructions from the under-sampled case at R=8 (MB=2, R_{kq}=4) are also shown along with its residual. The RMSE and SSIM are noted to be 0.0095 and 0.9982 respectively for

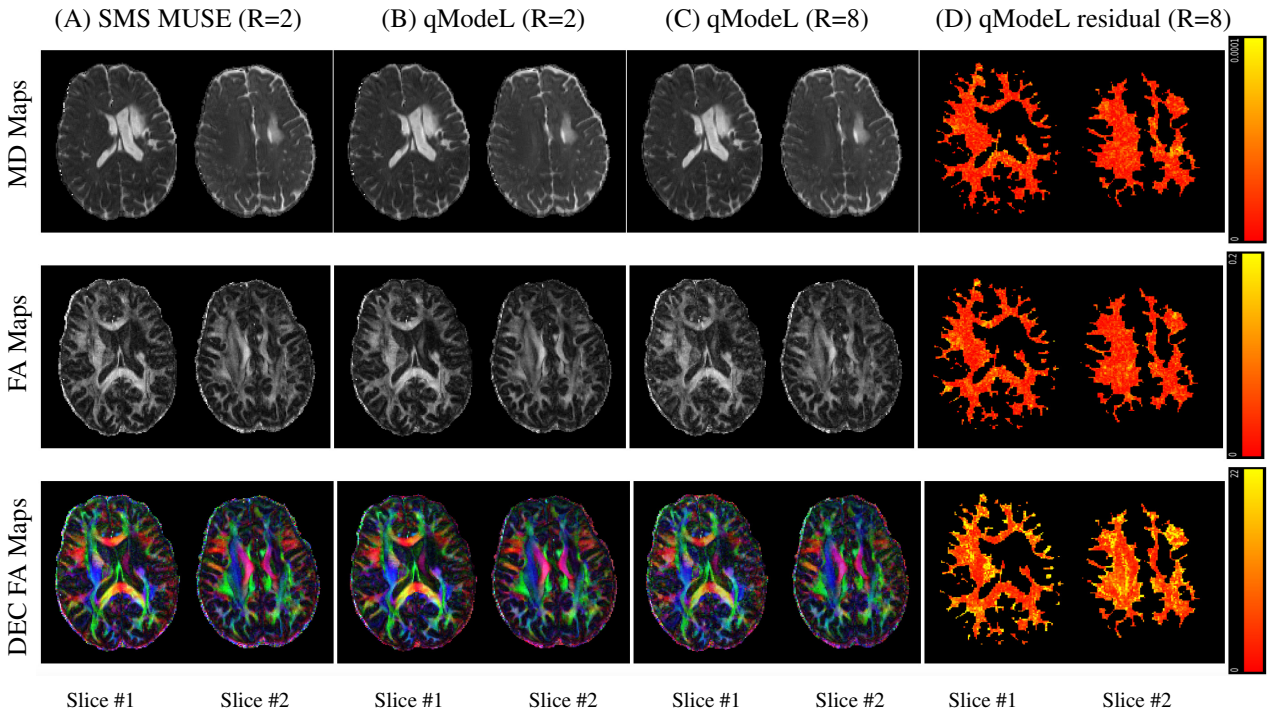

Figure 8: Dataset 4: (A) show the mean diffusivity (MD) maps, the FA maps and the DEC maps derived using the SMS MUSE reconstructed DWIs. (B)-(C) show the same maps derived using the qModeL reconstructions at R=2 and 8. The error in the microstructural parameter maps are given in (D), shown windowed over the white matter voxels (defined as FA>0.2).

R=8, computed with respect to R=2 from all DWIs.

The mean diffusivity (MD), FA and the DEC maps derived from the SMS MUSE and qModeL reconstruction of the above single-shell data are given in figure 8. The error in these microstructural maps for qModeL reconstruction at R=8 is also shown in the figure. The mean FA error for R=8 is 0.043, whereas the average angular error in the DEC FA map is 10.2°, computed for voxels where FA > 0.2. The qModeL reconstructed DWIs and the micro-structural maps are noted to closely follow the SMS MUSE DWIs and micro-structural maps at R=2, again confirming the utility of the physics-based qModeL training to represent q-space signals from abnormal tissues accurately. The radiologists' evaluation of the FA and MD maps also confirm the utility of the qModeL maps to be as good as the SMS MUSE maps.

Reconstruction using qModeL-SH

For qModeL-SH experiments, we first compared the qModeL-SH reconstruction to that using the qModeL DAE trained on the raw diffusion signal, on the fully sampled case. We utilized a single-shell dataset (dataset 5) for this purpose. The same DAE that was used to reconstruct the 7T 60-direction lesion data (dataset 4) was utilized for the qModeL reconstruction of dataset 5. Figure S10 (A) - (B) show the qModeL and qModeL-SH reconstruction for R=3 for this single-shell data. Here, we trained the qModeL-SH DAE

on the SH coefficients of the q-space samples that were used for training the qModeL DAE. The number of q-space samples utilized for training were matched to the q-space of the data to be reconstructed (i.e., q_t and q_r are same) for both qModeL and qModeL-SH DAEs, in this experiment. The difference image from the two reconstructions given in figure S10(C) show minimal difference, confirming that the qModeL-SH DAE is as good as the qModeL DAE for pre-learning the q-space.

Next, we show the utility of the qModeL-SH DAE to reconstruct q-space data when q_t is different from q_r. Two DAEs were trained for this purpose: (i) using SH coefficients derived from $q_t = 64$ q-space points, and (ii) using SH coefficient derived from $q_t = 90$ q-space points. These two DAE's were then utilized to reconstruct dataset 5 with $q_r = 60$ q-space points. Figure 9 show the qModeL-SH reconstruction employing these two DAE for R=3. Here also the residual show minimal error, confirming that the qModeL-SH using $q_t = 64, 90$ are equally good for reconstructing data with $q_r = 60$ points. Further, we also tested the qModeL-SH reconstruction for R=12. The comparison employing three DAEs trained using: $q_t = 60$, 64 and 90 are given in Figure S11 along with the residuals (computed with respect to the qModeL at R=3). The RMSE is noted to be less than .02 for all the cases.

Figures S12-S13 show the qModeL-SH experiments utilizing multi-shell data (dataset 2). Here also, the qModeL-SH DAE were trained in three settings: (i) $q_t = 180$ q-space points (60 points per shell), (ii) $q_t = 192$ q-space points (64 points per shell), and (iii) $q_t = 270$ q-space points (90 points per shell). These DAE's were then utilized to reconstruct dataset 2 with $q_r = 180$ q-space points (60 q-space points over three shells), at R=3 and R=12. The reconstruction quality was compared to that achieved using the qModeL multi-shell reconstruction for R=3 for dataset 2. Three representative DWIs from the multi-shell qModeL-SH reconstruction are included in figure S12 for R=3, and in figure S13 for R=12, using all three DAEs. The experiments show reasonable reconstructions achieved using qModeL-SH for the multi-shell case, with RMSE of 0.016 for R=12.

Discussion

In this work, we showed the adaptability of the qModeL learnable framework to simultaneously unalias and jointly recover diffusion scans that are multi-band accelerated and under-sampled. The under-sampling of the PE lines and the multi-band aliasing, together result in an increased number of aliased signals that needs to be separated during reconstruction. The pre-learned q-space manifold prior, made available from the qModeL framework, aids with the removal of the aliased signals. This is achieved by minimizing the projection error of the aliased signal to the pre-learned manifold. To further ease the complexity of the multi-band accelerated and under-sampled reconstruction, we made use of incoherent under-sampling patterns in

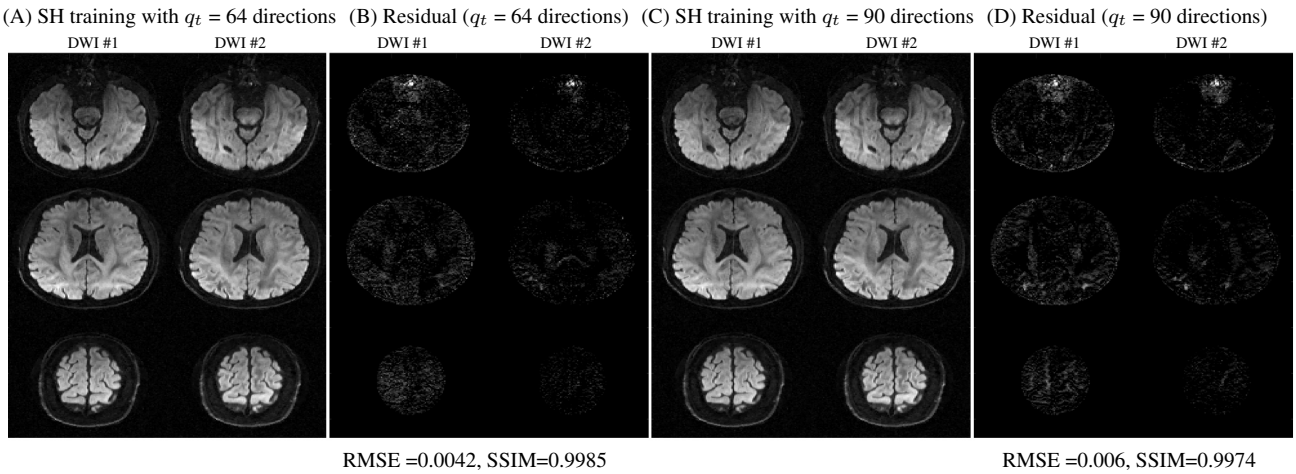

Figure 9: Dataset 5: Two representative DWIs reconstructed using qModeL-SH at R=3. (A) show reconstruction where the SH coefficient training was performed using 64 q-space points, and (B) show the residual with respect to qModeL reconstruction (amplified 20X). (C) show reconstruction where the SH coefficient training was performed using 90 q-space points, and (D) show the residual with respect to qModeL reconstruction (amplified 20X). All images are displayed on the scale [0 1].

the k-q space. These patterns help to shift the aliases with respect to each other along the q-dimension. In comparison to the individual reconstruction of the DWIs, the proposed joint reconstruction optimally make use of the shared information available in the combined k-q space and help to compensate for the SNR loss associated with the under-sampling.

Of note, once the DAE is trained to learn the q-space manifold to which the data needs to be reconstructed, it does not need to be re-trained for a different level of acceleration, under-sampling pattern or resolution or field strength. The explicit accounting of the acquisition physics in the forward model will handle these changes. The qModeL-SH prior introduced above further enhances the flexibility of the trained framework by enabling the pre-trained DAE to be used for a wide range of q-space points. This relatively high flexibility is partly made possible by the biophysical models that are used for training the DAE. The DAEs can efficiently learn the q-space signal relations modeled by the biophysical models. Moreover, re-training of the DAEs to learn a new q-space measurement space is a trivial process when utilizing biophysical models, as we have shown with the several DAEs we have used in this work. In contrast, training using in-vivo data will heavily depend on the quality of the training data available. Hence, in-vivo training data would require carefully curated data, to ensure the availability of various noise levels, resolution, accommodation of heterogeneous tissue types etc.

On the contrary, training using biophysical models also have caveats. The relationships learned by the DAEs are only as good as those predicted by the training models. It is to be noted that biophysical models are only approximations of the tissue types and are not exact in terms of modeling the behavior of the diffusion sig-

nal in every tissue type. Even very elaborate models such as the 3-compartment model may not model gray matter or mixed tissue types accurately for all acquisition conditions. The inaccuracy in the models can bias the reconstruction and can become more apparent at high acceleration factors. At low accelerations, this may not pose a big challenge. This is because the qModeL reconstruction is not solely dictated by the accuracy of the biophysical models. The data consistency term ensures that the recovered diffusion signal is close enough to the "measured true" diffusion signal. Moreover, the pre-learned signal model is not used as a constraint, rather it is only used as a prior. Hence, as long as the forward prediction given by the training model is good enough, the DAE can learn the diffusion measurement space well enough to help constrain the reconstruction to an acceptable sub-space out of all possible solutions. However, at high acceleration factors, when the available data points become fewer, the priors play a bigger role in determining the final solution. Hence, it is important to have accurate models to have accurate reconstruction especially at high acceleration factors. As more advanced models emerge, the learning can be adapted to such models. A balanced approach of augmenting the training using in-vivo data may also provide improved learning of the diffusion signal space to further improve the reconstruction. Alternately, learning of a more complete basis or dictionary that could represent all tissue types during reconstruction may be explored in future works. We note that the qModeL framework is general enough to work in a range of scenarios: (i) a model-free approach, where in-vivo training data may be used to train the deep network, instead of biophysical models, (ii) any signal model can be utilized for training.

On a related note, degeneracy in parameter estimation of the biophysical models is a widely known issue. Parameter degeneracy implies that multiple solutions are possible for the tissue microstructure parameter estimates, while solving the inverse problem of fitting the model to the diffusion signal. However, we note here that the issue of parameter degeneracy does not affect the proposed scheme, because qModeL reconstruction do not consider solving the above inverse problem at any point during the reconstruction. Only the accuracy of forward mapping of the tissue microstructure parameters to the diffusion signal is relevant for qModeL framework. Once the DWIs are reconstructed and are ready for model fitting using biophysical models, parameter degeneracy becomes relevant. We do not address this issue in this work. A detailed study of the parameter degeneracy problem and the landscape of the inverse fitting, given in (33), demonstrates the impact of SNR on the fitting accuracy. Of note, the proposed joint reconstruction with incoherent k-q under-sampling pattern provides improved SNR as demonstrated from several test cases in the present study. Further studies are needed to evaluate the impact of the under-sampling patterns and joint reconstructions on the accuracy of model fitting.

One advantage of the joint reconstruction method combined with incoherent k-q sampling patterns, is the gain in SNR achieved from the maximal utilization of shared k-q information available (figure S14). This also means that with more q-space points, a higher acceleration can be achieved, and vice versa. When only fewer q-space points are available, the number of randomized patterns as well as the available number of k-q points are reduced. Hence, blurry or noisy reconstructions are possible at high accelerations in such cases. This trend was noted in other joint reconstruction methods such as DAGER (13) also, where the need for $>>$ 60 q-space samples are noted to recover data with R=12 (where the in-plane under-sampling was even lower). We suspect that the blurring noted in the microstructural maps of dataset 1 (given in figure S1(D)) is due to the heavy under-sampling. Note that this data was reconstructed from only 20 q-space points. We believe that sharper reconstructions could have been possible if more q-space points were available, or if the sampling patterns were less restrictive (as is currently the case with a fixed partial Fourier under-sampling pattern). Alternatively, sharper reconstructions may have been achieved by using more sophisticated reconstructions that can account for the partial Fourier sampling also (15).

In the current work, we evaluated the performance of the qModeL reconstruction accuracy as a function of gradient directions, b-values and also focused on white matter regions only. From the plots given in S3 and S8, no gradient directions were noted to perform preferentially better or worse. The SSIM seems to be higher for low b-shells and progressively decreases for higher b-shells (figures S3(B), S8(b)). The RMSE is also slightly higher for higher b-shells, but only show a small change across the different shells. Within the white matter ROI, the average RMSE is noted to be slightly higher than the average over all tissue types. However, the variation of RMSE with respect to the b-value is not present within the white matter ROI. (Of note, the white matter ROI chosen to plot the RMSE could possibly have more than one fiber population.) These patterns are consistent across the two datasets studied. The choice of the regularization parameters λ_1 and λ_2 can partly explain this trend. The optimal regularization parameters were chosen by picking those that resulted in the lowest overall RMSE with respect to the fully sampled case. Nevertheless, the overall RMSE as well as the RMSE within the WM ROI, is noted to be < 0.03 for every shell (indicated by the dotted line shown in the plots).

In this work, we introduced the qModeL-SH prior to improve the flexibility of the q-space pre-learning. In qModeL-SH, the pre-learning of q-space is achieved in the SH domain. Hence, arbitrary modifications of the test q-space points can still be handled by the DAEs trained on the SH coefficients, as long as the degree and order of the SH transform are adequate to represent both the testing and training q-space points well. For example, we showed reconstruction where the DAE trained with 90 q-space samples was used

to reconstruct a dataset with 60 q-space samples. When the q-space relations pre-learned during training is not applicable to a test data, re-training of the DAEs are needed. Such scenarios include reconstructing datasets with different b-values than originally trained or the number of sampling points are drastically different. However, this is not a major limitation of the method since the re-training of the DAEs are fast (20 mins to less than 2 hours) and can be pre-trained and stored for a range of b-values / q-space samples. The present work also demonstrates the reconstruction of multi-shell data using the SH bases. Other bases that utilize different choices for radial basis function for representing multi-shell data are available (37–39). The utility of such bases could be explored in future to improve the accuracy of the multi-shell q-space manifold learning.

We also tested the ability of the biophysically-driven training method to accurately represent the q-space signal in brain tissues with abnormalities. We noted that the reconstructed images, and the micro-structural indices derived from the images generated via the qModeL framework, closely follow those reconstructed without using constraints on the q-space signals. Further, we tested the generalization capability of the learning-based prior on abnormal data (and healthy data) collected at different field strengths.

In the current work, we showed combined acceleration factors upto R=18 for multi-shell data, by choosing the multi-band acceleration to a factor of 3. It is possible to increase the multi-band factor to 4 using methods such as multiPINS RF excitation (13, 40). This can increase the achievable acceleration even further without additional SNR loss. Future works will explore this option.

We have studied the contribution of the q-space manifold regularizer and the TV regularizer separately on the qModeL reconstruction in (17). Experiments with and without the TV regularization are reported in (17). Specifically for TV regularizer, concerns about blurring can be alleviated by appropriately tuning λ_2.

Conclusion: We presented a new reconstruction method for the recovery of multi-band accelerated dMRI data that is also in-plane accelerated. The method introduces a complementary reconstruction prior in the q-domain to constrain the reconstruction. The joint reconstruction helps to compensate for the SNR loss from the in-plane under-sampling. The method was tested on single- and multi-shell acquisitions as well as on healthy and abnormal tissues. The reconstructed images and the micro-structural indices derived from the proposed method is shown to provide good accuracy to be used for routine micro-structural studies.

Acknowledgements

Financial support for this study was provided by grant 1R01EB031169, R01AG067078, 1R01EB022019-01A and 1 I01 RX002987. This work was conducted on MRI instruments funded by 1S10OD025025-01 and 1S10RR028821-01. This project was supported in part by a 2020 NARSAD Young Investigator Grant from the Brain & Behavior Research Foundation.

Legends:

Fig 1: qModeL joint recovery utilizing DAEs that were trained using: (i) biophysical modeling (DAE1), and (ii) in vivo training data (DAE2). (A) shows a DWI that belongs to a test data acquired with 60 q-space samples at b-value of 1000 s/mm^2. (B) - (C) show the same DWI from the qModeL joint recovery based on DAE1 and DAE2 respectively. For the training of DAE2 in (C), training data was derived from a random set of two subjects from a pre-existing database with DWIs collected at b=1000 s/mm^2 and 60-directions, from which only magnitude data was available. The specific imaging parameters of the test data and the training database are very different. For example, the test data is acquired with a 4-shot acquisition, at 1.7mm resolution with a TE of 58.3ms, whereas the images from the pre-existing database was acquired with single-shot EPI at 2mm resolution, with a TE of 88ms. (D) shows the mean difference between the reconstruction from the two DAEs (magnified by 10 times). (E) shows the q-space signal plots from a region with relatively higher difference (marked in (D)) predicted by DAE1 and DAE2. The plot from both DAEs are noted to closely follow the q-space signal plot of the test data. This demonstrates that the learning involved in qModeL is pertinent only to the q-space information and not tied to the specific acquisition parameters of the training data.

Fig 2: (A) shows the case with no acceleration involving the acquisition of two slices and (B) shows the case with combined MB- and in-plane acceleration. For the illustrated case in (B) with MB=2 and in-plane acceleration of 4-fold, a combined acceleration of R=8 is achieved (MB$\times$R$_{kq} = 2 \times 4$). To randomize the 4-fold in-plane under-sampling pattern for every q-space point in (B), we make use of multi-shot trajectories. A (N_s =) 4-shot EPI trajectory is designed to sample the k-space in the non-accelerated case in (A). For the accelerated case in (B), only 1 random shot (out of the 4 k-space shots) is used per q-space point (each shot is marked using a different color). (C) shows the extraction of the voxel-wise q-space signals to impose the q-space prior during the iterative joint reconstruction.

Fig 3: A) SMS MUSE individual reconstruction, (B) qModeL joint reconstruction. Two representative DWIs from the full set of 20 DWIs are shown in each case. All images are displayed on the scale [0 1].

Fig 4: (A), (B) : qModeL joint recovery for R=6 and the residual. (C), (D) qModeL joint recovery for R=12 and the residual. Two representative DWIs from the full set of 20 DWIs are shown in each case. The normalized RMSE and SSIM for the above cases are reported with respect to the qModeL joint recovery at R=3. The residuals are scaled 10x for display. All images are displayed on the scale [0 1].

Fig 5: Multi-band multi-shell joint reconstruction (dataset 2): (A), (B) SMS MUSE and SMS MUSSELS individual reconstruction, and (C) qModeL joint reconstruction for R=3 (MB=3, Ns=4, R_{kq} =1) shown from 3 DWIs corresponding to b1000, b2000 and b3000. The images from one of the bands is given above. The images corresponding to all three bands in the multi-band acquisition are included in figures S2. (D), (G) shows the qModeL joint reconstruction at R=12 (MB=3, Ns=1, R_{kq} =4) and its residual error with respect to qModeL at R=3. (E), & (H) corresponds to R =15 (MB=3, R_{kq} =5), and (F), (I) corresponds to R=18 (MB=3, R_{kq} =6). All images are displayed on the scale [0 1]. The residuals is scaled 10x for display. The RMSE and the SSIM are reported for each case.

Fig 6: Orientation dispersion index (ODI) derived using the multi-shell data of dataset 2. (A) - (B) shows the ODI maps derived from DWIs reconstructed using SMS MUSE and qModeL respectively at R=3. (C) - (D) shows the ODI maps derived from DWIs reconstructed using qModeL respectively at R=12 and R=18. The error in the ODI maps windowed over the WM (defined as FA > 0.2) are shown in the bottom row for the cases R=12 and R=18.

Fig 7: The extra-cellular volume fraction (v_{ec}) computed using the NODDI model is shown for various acceleration factors in comparison to the T2 FLAIR data (shown from two slices). The hyper-intensities on the T2 FLAIR is seen as hyper-intensities in the v_{ec} maps also as indicated by the arrows. Some regions are zoomed for better visualization of the hyper-intensities. The qModeL reconstruction can faithfully reconstruct the lesion data at the various acceleration factors.

Fig 8: Dataset 4: (A) show the mean diffusivity (MD) maps, the FA maps and the DEC maps derived using the SMS MUSE reconstructed DWIs. (B)-(C) show the same maps derived using the qModeL reconstructions at R=2 and 8. The error in the microstructural parameter maps are given in (D), shown windowed over the white matter voxels (defined as FA> 0.2).

Fig 9: Dataset 5: Two representative DWIs reconstructed using qModeL-SH at R=3. (A) show reconstruction where the SH coefficient training was performed using 64 q-space points, and (B) show the residual with respect to qModeL reconstruction (amplified 20X). (C) show reconstruction where the SH coefficient training was performed using 90 q-space points, and (D) show the residual with respect to qModeL reconstruction (amplified 20X). All images are displayed on the scale [0 1].

Table 1: Details of Datasets used for Various Experiments.

Fig S1: (A)-(D) FA maps and direction encoded color (DEC) maps from the SMS MUSE for R=3 and qModeL joint recovery for R=3, 6 and 12. For this dataset, only 20 diffusion directions are available. Due to the noisy SMS MUSE reconstruction for this case, the SMS MUSE FA and DEC maps (A) are also noisy. The qModeL maps at R=3 (B) show better SNR due to the joint recovery. The accelerated results at R=12 (D) show some blurring. The average error in the respective maps compared to R=3 is given in the bottom of the figures. The errors are computed for voxels where FA > 0.2.

Fig S2: Multi-band multi-shell joint reconstruction (dataset 2): (A), (B), (C) SMS MUSE reconstruction, SMS MUSSELS reconstruction and qModeL joint reconstruction of R=3 (MB=3, Ns=4, R_{kq} =1) shown from 3 DWIs corresponding to b1000, b2000 and b3000. (D), (G) shows the qModeL joint reconstruction at R=12 (MB=3, Ns=1, R_{kq} =4) and its residual error with respect to qModeL at R=3. (E), & (H) corresponds to R =15 (MB=3, R_{kq} =5), and (F), (I) corresponds to R=18 (MB=3, R_{kq} =6). All images are displayed on the scale [0 1]. The residuals is scaled by a factor of 4 for display. The RMSE and the SSIM are reported for each case.

Fig S3: (A) Plot of RMSE as a function of the gradient directions and b-value is shown from dataset 2 for accelerations of R=12, 15 and 18. (B) shows the plot of SSIM as a function of the gradient directions and b-value. The RMSE and SSIM plotted in (A)-(B) are the average over all tissue types from the slice shown in (E). (C)-(D) show plots of RMSE and SSIM from a white matter ROI that is shown in (E). The x-axis show the gradient directions with points 1 to 60 corresponding to the 60 directions of b1000 shell, points 61 to 120 corresponding to the 60 directions of b2000 shell, and points 121 to 180 corresponding to the 60 directions of b3000 shell. The data from the different shells are shown in different colors. The order of gradient directions remain constant for all the shells. The mean from each plot is indicated using the dashed lines.

Fig S4: Micro-structural maps derived using the multi-shell data of dataset 2 from one slice is shown for the neurite density index (NDI) and the orientation dispersion index (ODI). (A) -(B) shows the micro-structural maps derived from DWIs reconstructed using SMS MUSE and qModeL respectively at R=3. (C)- (D) shows the micro-structural maps derived from DWIs reconstructed using qModeL respectively at R=12 and R=18. The error in the micro-structural maps at R=12 and R=18 are shown.

Fig S5: The micro-structural maps intra-cellular volume fraction, extra-cellular volume fraction and NDI are shown for the simultaneously excited slices from dataset 3 with multiple sclerosis. The NDI values in the voxels that are at the boundary between tissue and ventricles are less consistent in the SMS MUSE reconstruction. The high noise in the SMS MUSE reconstruction, especially at the high b-shells, might be causing this.

Fig S6: The extra-cellular volume fraction (v_{ec}) computed using the NODDI model is shown for various acceleration factors in comparison to the T2 FLAIR data. The hyper-intensities on the T2 FLAIR is seen as hyper-intensities in the v_{ec} maps also as indicated by the arrows. Some regions are zoomed for better visualization of the hyper-intensities. The qModeL reconstruction can faithfully reconstruct the lesion data at the various acceleration factors.

Fig S7: The intra-cellular volume fraction maps are shown in comparison to the T2 FLAIR images from dataset 3 with multiple sclerosis. The hyper-intensities in the T2 FLAIR corresponds to hypo-intensities in the intra-cellular volume fraction maps derived from qModeL, and closely follow the pattern observed in the SMS MUSE based reconstruction. At high acceleration also, the hypo-intensities are preserved.

Fig S8: Dataset 3: (A) Plot of RMSE as a function of the gradient directions and b-value is shown for accelerations of R=12, 15 and 18. (B) shows the plot of SSIM as a function of the gradient directions and b-value. (A)-(B) are plotted from the whole slice shown in (E). (C)-(D) show plots of RMSE and SSIM from a white matter ROI. The white matter ROI that was chosen for the above plot is shown in (E). The data from the different b-values are shown in different colors. The order of gradient directions remain constant for all the shells. The mean from each plot is indicated using the dashed lines.

Fig S9: (A) Two representative DWIs reconstructed using the SMS MUSE (A) and the SMS MUSSELS (B) method on the data with lesion (dataset 4) are given in the top row. (C) show the corresponding DWIs from the qModeL joint reconstruction at R=2. (D) - (E) show the qModeL reconstruction of the same DWIs at R=8 and the residual computed from the white matter region (defined as FA > 0.2). The RMSE and the SSIM computed from the white matter voxels are reported in the figure. The residual was multiplied 5 times to magnify the errors. All images are displayed on the same scale [0 1].

Fig S10: Two representative DWIs reconstructed at R=3 from dataset 5 : (A) using qModeL, with DAE trained on raw diffusion signal, (B) using qModeL-SH with DAE trained on the SH coefficients, and (C) their difference. Both training were done using the same set of q-space points. In (A), the same DAE that was utilized to reconstruct dataset 4 was employed.

Fig S11: Two representative DWIs reconstructed using qModeL-SH at R=12. (A) show reconstruction where the SH coefficient training was performed using the same 60 q-space points as used for reconstruction, and (D) show the residual with respect to qModeL reconstruction (amplified 5 times). (B) show qModeL-SH reconstruction where the SH coefficient training was performed using 64 q-space points, and (E) show the residual with respect to qModeL reconstruction (amplified 5 times). (C) show reconstruction where the SH coefficient training was performed using 90 q-space points, and (F) show the residual with respect to qModeL reconstruction (amplified 5 times). The RMSE and SSIM computed with respect to qModeL at R=3 is noted in the figure.

Fig S12: Three representative DWIs from the multi-shell data that was reconstructed using qModeL-SH reconstruction at R=3. (A) show reconstruction where the SH coefficient training was performed using the same 180 q-space points as used for reconstruction, and (D) show the difference image with respect to qModeL reconstruction at R=3 (shown in figure S2(C)) (amplified 4 times). The difference is quantified as the RMSE and is noted in the figure. (B)-(C) show qModeL-SH reconstruction where the SH coefficient training was performed using 192 q-space points and 270 q-space points respectively, and (E)-(F) show the respective difference image with respect to qModeL reconstruction (amplified 4 times). The quantified difference for all the choices of q_t is close to 0.0095 confirming the suitability of any of the q_t to reconstruct the above data. In all the cases, the difference is noted to be higher for the higher shells.

Fig S13: Three representative DWIs from the multi-shell data that was reconstructed using qModeL-SH reconstruction at R=12. (A) show reconstruction where the SH coefficient training was performed using the same 180 q-space points as used for reconstruction, and (D) show the *residual* with respect to qModeL reconstruction (amplified 4 times). (B), (C) show qModeL-SH reconstruction where the SH coefficient training was performed using 192 q-space points, and 270 q-space points respectively. (D), (E) show the corresponding *residual* with respect to qModeL reconstruction (amplified 4 times). The RMSE and SSIM computed with respect to qModeL at R=3 (shown in figure S2(C)) is noted for each case. Note that in this case, the residual displays a combined effect of the difference noted in figure S12 between the two recon-

structions, and the effect of under-sampling.

Fig S14: Comparison of reconstructions at R=12 from dataset 1. SMS MUSE individual reconstruction (left), SMS MUSSELS individual reconstruction accounting for partial Fourier (middle) and qModeL joint reconstruction (right) are shown. The improvement in reconstruction arising from the joint recovery of the 20 q-space points are obvious from the qModeL reconstruction results.

References

1 J Basser P, Mattiello J, and LeBihan D. Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance Series B*, 103(3):247–254, 1994.

2 Beaulieu C. The basis of anisotropic water diffusion in the nervous system - a technical review. *NMR in Biomedicine*, 15(7-8):435–455, 2002.

3 Assaf Y and Pasternak O. Diffusion Tensor Imaging (DTI)-based White Matter Mapping in Brain Research: A Review. *Journal of Molecular Neuroscience*, 34(1):51–61, 2008.

4 Assaf Y and Basser PJ. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. *NeuroImage*, 27(1):48–58, 2005.

5 Panagiotaki E, Schneider T, Siow B, Hall MG, Lythgoe MF, and Alexander DC. Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison. *NeuroImage*, 59(3):2241–2254, 2012.

6 Zhang H, Schneider T, Wheeler-Kingshott CA, and Alexander DC. NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage*, 61(4):1000–1016, 2012.

7 Fieremans E, Jensen JH, and Helpern JA. White matter characterization with diffusional kurtosis imaging. *NeuroImage*, 58(1):177–188, 2011.

8 Jelescu IO, Veraart J, Adisetiyo V, Milla SS, Novikov DS, and Fieremans E. One diffusion acquisition and different white matter models: How does microstructure change in human early development based on WMTI and NODDI? *NeuroImage*, 107:242–256, 2015.

9 Barth M, Breuer F, Koopmans PJ, Norris DG, and Poser BA. Simultaneous multislice (SMS) imaging techniques. *Magnetic resonance in medicine*, 75(1):63–81, 2016.

10 Setsompop K, Gagoski BA, Polimeni JR, Witzel T, Wedeen VJ, and Wald LL. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magnetic Resonance in Medicine*, 67(5):1210–1224, 2012.

11 Sotiropoulos SN, Jbabdi S, Xu J, Andersson JL, Moeller S, Auerbach EJ, Glasser MF, Hernandez M, Sapiro G, Jenkinson M, Feinberg DA, Yacoub E, Lenglet C, Van Essen DC, Ugurbil K, and Behrens TE. Advances in diffusion MRI acquisition and processing in the Human Connectome Project. *NeuroImage*, 80:125–143, 2013.

12 Uğurbil K, Xu J, Auerbach EJ, Moeller S, Vu AT, Duarte-Carvajalino JM, Lenglet C, Wu X, Schmitter S, Van de Moortele PF, Strupp J, Sapiro G, De Martino F, Wang D, Harel N, Garwood M, Chen L, Feinberg DA, Smith SM, Miller KL, Sotiropoulos SN, Jbabdi S, Andersson JLR, Behrens TEJ, Glasser MF, Van Essen DC, Yacoub E, and WU-Minn HCP Consortium. Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. *NeuroImage*, 80:80–104, 2013.

13 Wu W, Koopmans PJ, Andersson JL, and Miller KL. Diffusion Acceleration with Gaussian process Estimated Reconstruction (DAGER). *Magnetic Resonance in Medicine*, 82(1):107–125, 2019.

14 Rasmussen CE and Williams CKI. *Gaussian Processes for Machine Learning*. ISBN 9780262182539.

15 Bilgic B, Chatnuntawech I, Manhard MK, Tian Q, Liao C, Iyer SS, Cauley SF, Huang SY, Polimeni JR, Wald LL, and Setsompop K. Highly accelerated multishot echo planar imaging through synergistic machine learning and joint reconstruction. *Magnetic Resonance in Medicine*, 82(4):mrm.27813, 2019.

16 Mani MP, Aggarwal HK, Ghosh S, and Jacob M. Model-Based Deep Learning for Reconstruction of Joint k-q Under-sampled High Resolution Diffusion MRI. In *Proceedings - International Symposium on Biomedical Imaging*, volume 2020-April, pages 913–916. IEEE Computer Society, 2020. ISBN 9781538693308.

17 Mani M, Magnotta VA, and Jacob M. qModeL: A plug-and-play model-based reconstruction for highly accelerated multi-shot diffusion MRI using learned priors. *Magnetic Resonance in Medicine*, 86(2):835–851, 2021.

18 Novikov DS, Kiselev VG, and Jespersen SN. On modeling. *Magnetic Resonance in Medicine*, 79(6):3172–3193, 2018.

19 Novikov DS, Fieremans E, Jespersen SN, and Kiselev VG. Quantifying brain microstructure with diffusion MRI: Theory and parameter estimation. *NMR in Biomedicine*, 32(4):e3998, 2019.

20 Panagiotaki E, Schneider T, Siow B, Hall MG, Lythgoe MF, and Alexander DC. Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison. *NeuroImage*, 59(3):2241–2254, 2012.

21 Reisert M, Kellner E, Dhital B, Hennig J, and Kiselev VG. Disentangling micro from mesostructure by diffusion MRI: A Bayesian approach. *NeuroImage*, 147:964–975, 2017.

22 Behrens TEJ, Berg HJ, Jbabdi S, Rushworth MFS, and Woolrich MW. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage*, 34(1):144–55, 2007.

23 Landman BA, Bogovic JA, Wan H, El Zahraa ElShahaby F, Bazin PL, and Prince JL. Resolution of crossing fibers with constrained compressed sensing using diffusion tensor MRI. *NeuroImage*, 59(3):2175–86, 2012.

24 Mani M, Jacob M, Guidon A, Magnotta V, and Zhong J. Acceleration of high angular and spatial resolution diffusion imaging using compressed sensing with multichannel spiral data. *Magnetic Resonance in Medicine*, 73 (1):126–138, 2015.

25 Gyori NG, Palombo M, Clark CA, Zhang H, and Alexander DC. Training Data Distribution Significantly Impacts the Estimation of Tissue Microstructure with Machine Learning. *bioRxiv*, page 2021.04.13.439659, 2021.

26 Mani M, Jacob M, Guidon A, Liu C, Song A, Magnotta V, and Zhong J. Acceleration of high angular and spatial resolution diffusion imaging using compressed sensing. In *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 326–329. IEEE, 2012. ISBN 978-1-4577-1858-8.

27 Mani M, Jacob M, Guidon A, Magnotta V, and Zhong J. Acceleration of high angular and spatial resolution diffusion imaging using compressed sensing with multichannel spiral data. *Magnetic Resonance in Medicine*, 73 (1):126–138, 2015.

28 Dai E, Ma X, Zhang Z, Yuan C, and Guo H. Simultaneous multislice accelerated interleaved EPI DWI using generalized blipped-CAIPI acquisition and 3D K-space reconstruction. *Magnetic Resonance in Medicine*, 77(4): 1593–1605, 2017.

29 Mani M, Jacob M, McKinnon G, Yang B, Rutt B, Kerr A, and Magnotta V. SMS MUSSELS: A navigator-free reconstruction for simultaneous multi-slice-accelerated multi-shot diffusion weighted imaging. *Magnetic Resonance in Medicine*, 83(1):154–169, 2020.

30 Christiaens D and Tournier JD. Modeling Fiber Orientations Using Diffusion MRI. In *Advances in Magnetic Resonance Technology and Applications*, volume 1, pages 509–532. Academic Press, 2020.

31 Conolly S, Nishimura D, Macovski A, and Glover G. Variable-rate selective excitation. *Journal of Magnetic Resonance (1969)*, 78(3):440–458, 1988.

32 Chan SH, Wang X, and Elgendy OA. Plug-and-Play ADMM for Image Restoration: Fixed-Point Convergence and Applications. *IEEE Transactions on Computational Imaging*, 3(1):84–98, 2016.

33 Jelescu IO, Veraart J, Fieremans E, and Novikov DS. Degeneracy in model parameter estimation for multi-compartmental diffusion in neuronal tissue. *NMR in Biomedicine*, 2016.

34 Descoteaux M, Angelino E, Fitzgibbons S, and Deriche R. Regularized, fast, and robust analytical Q-ball imaging. *Magnetic Resonance in Medicine*, 58(3):497–510, 2007.

35 Chen NK, Guidon A, Chang HC, and Song AW. A robust multi-shot scan strategy for high-resolution diffusion weighted MRI enabled by multiplexed sensitivity-encoding (MUSE). *NeuroImage*, 72:41–47, 2013.

36 Herbst M, Deng W, Ernst T, and Stenger VA. Segmented simultaneous multi-slice diffusion weighted imaging with generalized trajectories. *Magnetic Resonance in Medicine*, 78(4):1476–1481, 2017.

37 Ozarslan E, Koay CG, and Basser PJ. Simple harmonic oscillator based estimation and reconstruction for one-dimensional q-space {MR}. *Proc. 16th Annual Meeting ISMRM, Toronto, Ontario, Canada*, 16:35, 2008.

38 Özarslan E, Koay CG, and Basser PJ. Simple Harmonic Oscillator Based Reconstruction and Estimation for One-Dimensional q-Space Magnetic Resonance (1D-SHORE). In *Excursions in Harmonic Analysis, Volume 2*, volume 2, pages 373–399. Birkhäuser Boston, Boston, 2013.

39 Christiaens D, Cordero-Grande L, Hutter J, Price AN, Deprez M, Hajnal JV, and Tournier JD. Learning Compact q-Space Representations for Multi-Shell Diffusion-Weighted MRI. *IEEE Transactions on Medical Imaging*, 38 (3):834–843, 2019.

40 Eichner C, Wald LL, and Setsompop K. A low power radiofrequency pulse for simultaneous multislice excitation and refocusing. *Magnetic Resonance in Medicine*, 72(4):949–958, 2014.