Accelerated Diffusion Spectrum Imaging with Compressed Sensing using Adaptive Dictionaries

Berkin Bilgic[1,*], **Kawin Setsompop**[2,3], **Julien Cohen-Adad**[2,3], **Anastasia Yendiki**[2], **Lawrence L. Wald**[2,3,4], and **Elfar Adalsteinsson**[1,4]

[1]Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

[2]A. A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Charlestown, Massachusetts, USA

[3]Harvard Medical School, Boston, Massachusetts, USA

[4]Harvard-MIT Division of Health Sciences and Technology, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

Abstract

Diffusion Spectrum Imaging (DSI) offers detailed information on complex distributions of intravoxel fiber orientations at the expense of extremely long imaging times (~1 hour). Recent work by Menzel et al. demonstrated successful recovery of diffusion probability density functions (pdfs) from sub-Nyquist sampled q-space by imposing sparsity constraints on the pdfs under wavelet and Total Variation (TV) transforms. As the performance of Compressed Sensing (CS) reconstruction depends strongly on the level of sparsity in the selected transform space, a dictionary specifically tailored for diffusion pdfs can yield higher fidelity results. To our knowledge, this work is the first application of adaptive dictionaries in DSI, whereby we reduce the scan time of whole brain DSI acquisition from 50 to 17 min while retaining high image quality. *In vivo* experiments were conducted with the 3T Connectome MRI. The root-mean-square error (RMSE) of the reconstructed 'missing' diffusion images were calculated by comparing them to a gold standard dataset (obtained from acquiring 10 averages of diffusion images in these missing directions). The RMSE from the proposed reconstruction method is up to 2 times lower than that of Menzel et al.'s method, and is actually comparable to that of the fully-sampled 50 minute scan. Comparison of tractography solutions in 18 major white-matter pathways also indicated good agreement between the fully-sampled and 3-fold accelerated reconstructions. Further, we demonstrate that a dictionary trained using pdfs from a single slice of a particular subject generalizes well to other slices from the same subject, as well as to slices from other subjects.

Introduction

Diffusion weighted MR imaging is a widely used method to study white matter structures of the brain. Diffusion Tensor Imaging (DTI) is an established diffusion weighted imaging method, which models the diffusion as a univariate Gaussian distribution (1). One limitation of this model arises in the presence of fiber crossings, and this can be addressed by using a more involved imaging method (2,3). Diffusion Spectrum Imaging (DSI) results in magnitude representation of the full q-space and yields a complete description of the

*Correspondence to: Berkin Bilgic, Massachusetts Institute of Technology, Room 36-776A, 77 Massachusetts Avenue, Cambridge, MA 02139, berkin@mit.edu, Fax: 617-324-3644, Phone: 617-866-8740.

diffusion probability density function (pdf) (4,5). While DSI is capable of resolving complex distributions of intravoxel fiber orientations, full q-space coverage comes at the expense of substantially long scan times (~1 hour).

Compressed Sensing (CS) comprises algorithms that recover data from undersampled acquisitions by imposing sparsity or compressibility assumptions on the reconstructed images (6). In the domain of DSI, acceleration with CS was successfully demonstrated by Menzel et al. (7) by imposing wavelet and Total Variation (TV) penalties in the pdf space. Up to an undersampling factor of 4 in q-space, it was reported that essential diffusion properties such as orientation distribution function (odf), diffusion coefficient, and kurtosis were preserved (7). A recent study focused on the problem of finding the best wavelet basis to represent the diffusion pdf by comparing various wavelet transforms (8).

The performance of CS recovery depends on the level of sparsity of the signal in the selected transform domain, as well as the incoherence of the aliasing artifacts in the transform domain and the amount of acceleration in the sampling space (6). While prespecified transformations such as wavelets and spatial gradients yield sparse signal representation, tailoring the sparsifying transform based on the characteristics of the particular signal type may offer even sparser results. K-SVD is an algorithm that designs a dictionary that achieves maximally sparse representation of the input training data (9). The benefit of using data-driven, adaptive dictionaries trained with K-SVD was also demonstrated in CS reconstruction of structural MR imaging (10,11).

In this work, we employ the K-SVD algorithm to design a sparsifying transform that yields a signal representation with increased level of sparsity. Coupling this adaptive dictionary with the FOcal Underdetermined System Solver (FOCUSS) algorithm (12), we obtain a parameter-free CS algorithm. With 3-fold undersampling of q-space, we demonstrate *in vivo* up to 2-fold reduced pdf reconstruction errors relative to our implementation of the CS algorithm that uses wavelets and variational penalties by Menzel et al. (7). At higher acceleration factors of 5 and 9, we still demonstrate up to 1.8-fold and 1.6-fold reduced errors relative to wavelet and TV reconstruction *at the lower acceleration factor of 3*. For additional validation, the RMSE of the reconstructed 'missing' diffusion images were calculated by comparing them to a gold standard dataset obtained with 10 averages. In this case, dictionary-based reconstructions were seen to be comparable to the fully-sampled 1 average data. For further validation, average Fractional Anisotropy (FA) and tract volume metrics obtained from 18 major white-matter pathways were compared between the fully-sampled and 3-fold accelerated dictionary reconstructions to yield good agreement. Additionally, we show that a dictionary trained on data from a particular subject generalizes well to reconstruction of another subject's data, still yielding up to 2-fold reduced reconstruction errors relative to using prespecified transforms. Hence, application of the proposed method might reduce a typical 50-minute DSI scan to 17 minutes (upon 3× acceleration) while retaining high image quality. In addition, we also investigate using a simple ℓ_1-norm penalty in the pdf space with the FOCUSS algorithm, and show that this approach gives comparable results to the more involved wavelet-and TV-based reconstruction by Menzel et al. (7), while being computationally more efficient.

Theory

CS Recovery with Prespecified Transforms

Letting $p \in \mathbb{C}^N$ represent the 3-dimensional diffusion pdf at a particular voxel as a column vector, and $q \in \mathbb{C}^M$ denote the corresponding undersampled q-space information, CS recovery with wavelet and TV penalties aim to solve the convex optimization problem at a single voxel,

$$min_{\boldsymbol{p}}\left\|\mathbf{F}_{\Omega}\boldsymbol{p}-\boldsymbol{q}\right\|_2^2+\alpha\cdot\left\|\boldsymbol{\Psi}\boldsymbol{p}\right\|_1+\beta\cdot\text{TV}(\boldsymbol{p}) \quad \text{Eq.1}$$

where $\mathbf{F}_\Omega$ is the undersampled Fourier transform operator, $\boldsymbol{\Psi}$ is a wavelet transform operator, TV($\cdot$) is the Total Variation penalty, and α and β are regularization parameters that need to be determined. CS recovery is applied on a voxel-by-voxel basis to reconstruct all brain voxels.

Training an Adaptive Transform with K-SVD

Given an ensemble $\mathbf{P} \in \mathbb{C}^{N \times L}$ formed by concatenating L example pdfs $\{\boldsymbol{p}_i\}_{i=1}^L$ collected from a training dataset as column vectors, the K-SVD algorithm (9) aims to find the best possible dictionary for the sparse representation of this dataset by solving,

$$min_{\mathbf{P},\mathbf{D}}\sum_{i=1}^{L}\left\|\boldsymbol{x}_i\right\|_0 \text{ subject to } \left\|\mathbf{P}-\mathbf{D}\mathbf{X}\right\|_F^2 \le \varepsilon \quad \text{Eq.2}$$

where $\mathbf{X}$ is the matrix that contains the transform coefficient vectors $\{\boldsymbol{x}_i\}_{i=1}^L$ as its columns, $\mathbf{D}$ is the adaptive dictionary, ε is a fixed constant that adjusts the data fidelity, and $\|\cdot\|_F$ is the Frobenius norm. The K-SVD works iteratively, first by fixing $\mathbf{D}$ and finding an optimally sparse $\mathbf{X}$ using orthogonal matching pursuit, then by updating each column of $\mathbf{D}$ and the transform coefficients corresponding to this column to increase data consistency.

CS Recovery with an Adaptive Transform using FOCUSS

The FOCUSS algorithm aims to find a sparse solution to the underdetermined linear system $\mathbf{F}_\Omega \mathbf{D} \boldsymbol{x} = \boldsymbol{q}$, where $\boldsymbol{x}$ is the vector of transform coefficients in the transform space defined by the dictionary $\mathbf{D}$ using the following iterations,

For iteration number $t = 1, \dots T$,

$$\mathbf{W}_{j,j}^t = diag\left(|x_j^t|^{1/2}\right) \quad \text{Eq.3}$$

$$\boldsymbol{s}^t = \text{argmin}_{\boldsymbol{s}}\left\|\boldsymbol{s}\right\|_2^2 \text{ such that } \mathbf{F}_\Omega\mathbf{D}\mathbf{W}^t\boldsymbol{s}=\boldsymbol{q} \quad \text{Eq.4}$$

$$\boldsymbol{x}^{t+1}=\mathbf{W}^t\boldsymbol{s}^t \quad \text{Eq.5}$$

Here, $\mathbf{W}^t$ is a diagonal weighting matrix whose j^{th} diagonal entry is denoted as $\mathbf{W}_{j,j}^t$, $\boldsymbol{x}^t$ is the estimate of transform coefficients at iteration t whose j^{th} entry is x_j^t. The final reconstruction in diffusion pdf space is obtained via the mapping $\boldsymbol{p} = \mathbf{D}\boldsymbol{x}^{T+1}$.

We note that it is possible to impose sparsity-inducing ℓ_1 penalty directly on the pdf coefficients by taking $\mathbf{D}$ to be the identity matrix $\mathbf{I}$.

Methods

Diffusion EPI acquisitions were obtained from three healthy volunteers (subjects A, B and C) using a novel 3T system (Magnetom Skyra Connectom, Siemens Healthcare, Erlangen, Germany) equipped with the AS302 "CONNECTOM" gradient with G_{max} = 300 mT/m (here we used G_{max} = 200 mT/m) and Slew = 200 T/m/s. A custom-built 64-channel RF

head array (13) was used for reception with imaging parameters of 2.3 mm isotropic voxel size, FOV = 220×220×130, matrix size = 96×96×57, b_{max} = 8000 s/mm^2, 514 directions full sphere q-space sampling (corners of q-space were zero-padded since they were not sampled) organized in a Cartesian grid with interspersed b=0 images every 20 TRs (for motion correction, 25 b=0 images in total), in-plane acceleration = 2× (using GRAPPA algorithm), TR/TE = 5.4 s/60 ms, total imaging time ~50 min. In addition, at 5 q-space points ([1,1,0], [0,2,−1], [0,0,3], [0,4,0], and [5,0,0]) residing on 5 different shells, 10 averages were collected for noise quantification. The corresponding b-values for these 5 points were 640, 1600, 2880, 5120, and 8000 s/mm^2. Eddy current related distortions were corrected using the reversed polarity method (14). Motion correction (using interspersed b=0) was performed using FLIRT (15) with sinc interpolation.

Variable-density undersampling (using a power-law density function (6)) with $R = 3$ acceleration was applied in q-space on a 12×12×12 grid. Three different adaptive dictionaries were trained with data from slice 30 of subjects A, B and C. Reconstruction experiments were applied on test slices that are different than the training slices. In particular, two reconstruction experiments were performed. First, voxels in slice 40 of subject A were retrospectively undersampled in q-space, and reconstructed using 5 different methods: wavelet+TV method of Menzel et al. (7), ℓ_1-regularized FOCUSS, and Dictionary-FOCUSS with the three dictionaries trained on three different subjects. Second, voxels in slice 25 of subject B were retrospectively undersampled with the same $R = 3$ sampling pattern, and again reconstructed with wavelet+TV, ℓ_1-FOCUSS, and the three dictionaries trained on three different subjects. Slice 30 was selected for training and slices 25 and 40 were chosen for test based on their anatomical location, so that the test slices would reside on lower and upper parts of the brain, while the training slice was one of the middle slices. For Menzel et al.'s method, Haar wavelets in MATLAB's wavelet toolbox were used. The regularization parameters α and β in Eq.1 were chosen by parameter sweeping with values {10^{-4}, $3 \cdot 10^{-4}$, 10^{-3}, $3 \cdot 10^{-3}$} to minimize the reconstruction error of 100 randomly selected voxels in slice 40 of subject A. The optimal regularization parameters were found to be α = $3 \cdot 10^{-4}$ for wavelet and $\beta = 10^{-4}$ for the TV term. By taking the fully-sampled data as ground-truth, the fidelity of the five methods were compared using root-mean-square error (RMSE) normalized by the ℓ_2-norm of ground-truth as the error metric both in pdf domain and q-space.

Since the fully-sampled data are corrupted by noise, computing RMSEs relative to them will include contributions from both reconstruction errors and additive noise. To address this, the additional 10 average data acquired at the selected 5 q-space points were used. As a single average full-brain DSI scan takes ~50 min, it was not practical to collect 10 averages for all of the undersampled q-space points. As such, we rely on both error metrics, namely: the RMSE relative to one average fully-sampled dataset and the RMSE relative to gold standard data for 5 q-space points.

To compare the fully-sampled and 3-fold accelerated Dictionary-FOCUSS reconstructions in terms of tractography solutions, streamline deterministic DSI tractography on the two datasets was performed in trackvis (http://trackvis.org) and 18 white-matter pathways were labeled. The labeling was performed following the protocol described in (16), where two regions of interest (ROIs) are drawn for each pathway in parts of the anatomy that the pathway is known to traverse. To eliminate variability due to manual labeling in the two data sets and make our comparison as unbiased as possible, the ROIs used here were not drawn manually on the fully-sampled and 3-fold accelerated data. Instead, we obtained the ROIs from a different data set of 33 healthy subjects, where we had previously labeled the same pathways (17). We averaged the respective ROIs from the 33 subjects in MNI space (18) and mapped the average ROIs to the native space of the fully-sampled and $R = 3$

datasets using affine registration. In each data set we isolated the tractography streamlines going through the respective ROIs to identify the 18 pathways.

A Matlab toolbox that reproduces some of the *in vivo* results is available at: http://web.mit.edu/berkin/www/software.html

Results

Fig. 1 depicts the error of the different reconstruction methods in the pdf domain for each voxel in slice 40 of subject A. At $R = 3$ acceleration, reconstruction error of Menzel et al.'s method averaged over brain voxels in the slice was 15.8%, while the error was 15.0% for ℓ_1-regularized FOCUSS. Adaptive dictionary trained on subject A yielded 7.8% error. Similarly, reconstruction with dictionaries trained on pdfs of the other subjects B and C returned 7.8% and 8.2% RMSE, respectively. At $R = 5$, Dictionary-FOCUSS returned 8.9%, 8.9% and 9.3% error with training on subjects A, B and C, respectively. At $R = 9$ dictionary reconstruction with training on subjects A, B and C returned 10.0%, 10.0% and 10.4% RMSE.

In Fig. 2, reconstruction errors at $R = 3$ on slice 25 of subject B are presented. In this case, Menzel et al.'s method yielded 17.5% average RMSE, and ℓ_1-FOCUSS had 17.3% error. Dictionary trained on slice 40 of subject A returned 11.4% RMSE, while adaptive transforms trained on subjects B and C had 11.4% and 11.8% error, respectively. At a higher acceleration factor of $R = 5$, Dictionary-FOCUSS with training on subjects A, B and C returned 13.1%, 13.3% and 13.5% error. At $R = 9$ dictionary reconstruction with training on subjects A, B and C yielded 14.2%, 14.2% and 14.4% RMSE, respectively.

Fig. 3 presents RMSEs obtained on various slices of subject A using Dictionary- and ℓ_1-FOCUSS. Error bars that show the variation of the reconstruction errors are also included. RMSE maps on four selected slices are plotted for comparison. The same analysis is carried out on various slices of subject B, and the results are depicted in Fig. 4.

Reconstruction errors in q-space images of subject A obtained with Wavelet+TV, ℓ_1-FOCUSS and Dictionary-FOCUSS trained on the three subjects for the undersampled q-space directions are plotted in Fig. 5. For two particular diffusion directions [0,4,0] and [5,0,0], q-space reconstructions obtained with the three methods are also presented. In Fig. 5a, q-space images obtained with Wavelet+TV, ℓ_1-FOCUSS and Dictionary-FOCUSS (with training on subject B) are compared with the 10 average fully-sampled image at [5,0,0]. Fig. 5b presents the error images relative to the 10 average data for the three methods. Figs.5c and d depict the same analysis at direction [0,4,0].

In an attempt to quantify the noise in q-space and separate it from CS reconstruction error, we take the 10 average data acquired at 5 q-space directions as ground truth and compute the RMSEs relative to them. Fig. 6 shows the error plots for the 1 average fully sampled data, Wavelet+TV, ℓ_1-FOCUSS, and Dictionary-FOCUSS reconstructions relative to the 10 average data for slices from subjects A and B.

Fig. 7a and b show tractography results of subject A for the labeled white-matter pathways in the fully-sampled and 3-fold accelerated Dictionary-FOCUSS reconstructions. Fig. 7c and d show plots of the average FA and volume of each pathway for the 18 white-matter pathways, as calculated from the two reconstructions.

Discussion

This work presented the first application of adaptive transforms to voxel-by-voxel CS reconstruction of undersampled q-space data. Relative to reconstruction with prespecified transforms, i.e. wavelet and TV, the proposed algorithm has up to 2 times reduced error in the pdf domain at the same acceleration factor ($R = 3$), while requiring no regularization parameter tuning. When the undersampling ratio was increased to $R = 5$ and even up to $R = 9$, the proposed method still demonstrated substantial improvement relative to using prespecified transforms at a lower acceleration factor of $R = 3$ (Figs.1 and 2). As demonstrated, a dictionary trained with pdfs from a single slice of a particular subject generalizes to other slices of the same subject, as well as to different subjects. However, further tests are needed to see if dictionaries can generalize across healthy and patient populations, or across age groups.

Since the acquired 1 average DSI data is corrupted by noise (especially in the outer shells), it is desired to obtain noise-free data for more reliable computation of CS reconstruction errors. Because even the 1 average full-shell acquisition takes ~50 min, it is practically not possible to collect multiple-average data at all q-space points. To address this, one representative q-space sample at each shell was collected with 10 averages to serve as "(approximately) noise-free" data. When the noise-free data were taken to be ground-truth, the dictionary reconstruction with 3-fold undersampling was comparable to the fully-sampled 1 average data for both subjects (Fig. 6).

RMSE in Fig. 2 was overall higher than in Fig. 1. A possible explanation is the inherently lower signal-noise-ratio (SNR) in the lower axial slice, particularly in the center area of the brain which is further away from the receive coils. In particular, the error is higher in the central region of the image where the SNR is expected to be lowest. Since the noisy 1 average datasets were taken to be the reference in RMSE computations in Figs. 1 and 2, we expect the errors to be influenced by noise in these lower SNR regions. As seen in Figs. 1 and 2, Wavelet+TV and ℓ_1-FOCUSS tend to yield larger error in the white matter, where the information is more critical for fiber tracking. Error maps from the dictionary reconstruction is more homogenous across white and gray matter, especially results on Fig. 2 resemble SNR maps where the middle of the brain is further away from the receive coils. As Fig. 6 demonstrates, dictionary reconstruction has a certain degree of denoising property, since it yields lower RMSE than the 1 average data relative to the 10 average data. This might be one possible explanation why the RMSE is relatively higher in the middle of the brain, which should be explored in future investigation.

As seen in Fig. 5, wavelet and TV penalized reconstruction and ℓ_1-FOCUSS yield especially poor quality results in estimating the high q-space samples. In particular, as depicted in Fig. 5a and b, these CS methods tend to underestimate the high q-space content. However, this is not a simple scaling problem, as they yield either flat (Wavelet+TV) or grainy (ℓ_1-FOCUSS) results. Because Wavelet+TV reconstruction imposes piece-wise smoothness assumption in compressed sensing reconstruction, it leads to loss of high frequency content. In the context of DSI, this corresponds to attenuated high q-space information (flat, underestimated outer shell). ℓ_1-penalized reconstruction encourages small number of non-zero coefficients, which is seen to be insufficient to model the diffusion pdfs. This also leads to underestimated high q-space content, but since there is no smoothness constraint in pdf domain, the reconstructed q-space is not flat. The RMSE plot in Fig. 5 also demonstrates that Wavelet+TV and ℓ_1-FOCUSS results are comparable to the adaptive reconstruction at lower q-space (Fig. 5c and d), and the difference becomes more pronounced as |q| increases.

Visual inspection of the tractography solutions from the fully sampled and 3-fold accelerated Dictionary-FOCUSS data sets (Fig. 7a and b) showed that the white-matter pathways reconstructed from the two acquisitions were very similar. When comparing average FA over each pathway, as calculated from the two reconstructions, there are two potential sources of differences: the tractography streamlines could be different, visiting different voxels in the brain for each data set, and/or the tensor, from which the FA value is calculated, could be different at same voxel for each data set. However, we found good agreement between the average FA values in the fully-sampled and 3-fold accelerated reconstructions (Fig. 7c and d). Some differences are to be expected in weaker pathways that only consist of very few streamlines and thus are more sensitive to noise and have lower test-retest reliability than the stronger pathways. This was the case particularly for the right inferior longitudinal fasciculus (R-ILF), which did not have any streamlines in the fully-sampled data set (Fig. 7d). Therefore it was not possible to extract an average FA value for the R-ILF from the fully sampled data. Apart from this pathway, the mean difference between the average FA values in the fully sampled and 3-fold accelerated data, as a percentage of the value in the fully sampled data, was 3%. For the volume estimates, the mean error was 16%. It is possible that even more stable FA and volume measurements could be obtained by manual labeling of the paths directly on each data set, instead of using the average ROIs. This is because the averaging of ROIs in MNI space is susceptible to misregistration errors, leading to average ROIs that are typically much larger than the individual ROIs than a rater would draw directly on the images. Thus the bundles that we obtained with the average ROIs are more likely to contain stray streamlines that would be eliminated in a careful individual manual labeling, leading to less noisy volume and FA estimates. However, we used the average ROIs here to avoid introducing variability due to manual labeling.

In a previous study we evaluated the intra-rater and inter-rater reliability of the manual labeling procedure by performing manual labeling several times on the same data set. We found the average distance between pathways labeled by the same and different raters to be, respectively, in the order of 1 voxel and 2 voxels (17). In the present study we found that the distance between the pathways obtained from the fully-sampled and $R = 3$ data sets were comparable (median distance: 2.37mm, mean distance: 2.74mm with acquisition voxel size of 2.3mm isotropic). Further investigation with test-retest scans is warranted to determine how the differences between the fully sampled and 3-fold undersampled results compare to the test-retest reliability of each type of reconstruction.

In this study, fully-sampled q-space data were collected for comparison with the CS reconstruction methods. With the fully-sampled dataset, it was simple to apply the reverse polarity approach (14) to get good eddy current correction. We note that in a real random undersampling case where reverse pairs are not present, such eddy current correction method will not be applicable. However, various approaches exist in performing eddy current correction, such as linearly fitting the eddy-current distortions parameters (translation, scaling, shearing) using the available data, and then estimating the transformation for any given q-space data.

In our implementation, *per voxel* processing time of ℓ_1-FOCUSS was 0.6 seconds, while this was 12 seconds for Dictionary-FOCUSS and 27 seconds for Wavelet+TV method on a workstation with 12GB memory and 6 processors. Hence, full-brain reconstruction using the ℓ_1-FOCUSS algorithm would still take *days*. Because each voxel can be processed independently, parallel implementation is likely to be a significant source of performance gain. Dictionary training step (for subject A, using 3200 voxels inside the brain mask from a single slice) took 12 minutes. An additional research direction is to evaluate the change in

reconstruction quality when multiple slices are used for training. Increased processing times due to employing a larger dictionary may become a practical concern in this case.

The proposed CS acquisition/reconstruction can be combined with other techniques to further reduce the acquisition time. In particular, combining the proposed method with the Blipped-CAIPI Simultaneous MultiSlice (SMS) acquisition (19) could reduce a 50 minute DSI scan to 5.5 minutes upon 9-fold acceleration (3×3 CS-SMS).

Conclusion

By using a data-driven transform specifically tailored for sparse representation of diffusion pdfs, up to 2-fold reduction in reconstruction errors were obtained relative to using either prespecified wavelet and gradient transforms, or ℓ_1-norm penalty. Further, it was demonstrated that an adaptive dictionary trained on a particular subject generalizes well to other subjects, still yielding significant benefits in CS reconstruction performance. Coupled with the parameter-free FOCUSS algorithm, the proposed method can help accelerate DSI scans in the clinical domain.

Acknowledgments

References

1. Basser PJ, Mattiello J, LeBihan D. MR diffusion tensor spectroscopy and imaging. Biophys J. 1994; 66(1):259–267. [PubMed: 8130344]

2. Tuch DS. Q-ball imaging. Magn Reson Med. 2004; 52(6):1358–1372. [PubMed: 15562495]

3. Tournier JD, Calamante F, Gadian DG, Connelly A. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. Neuroimage. 2004; 23(3):1176–1185. [PubMed: 15528117]

4. Callaghan PT, Eccles CD, Xia Y. Nmr Microscopy of Dynamic Displacements - K-Space and Q-Space Imaging. J Phys E Sci Instrum. 1988; 21(8):820–822.

5. Wedeen VJ, Hagmann P, Tseng WY, Reese TG, Weisskoff RM. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. Magn Reson Med. 2005; 54(6): 1377–1386. [PubMed: 16247738]

6. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. Magn Reson Med. 2007; 58(6):1182–1195. [PubMed: 17969013]

7. Menzel MI, Tan ET, Khare K, Sperl JI, King KF, Tao X, Hardy CJ, Marinelli L. Accelerated diffusion spectrum imaging in the human brain using compressed sensing. Magn Reson Med. 2011; 66(5):1226–1233. [PubMed: 22012686]

8. Saint-Amant, E.; Descoteaux, M. Sparsity Characterisation of the Diffusion Propagator; 19th Annual ISMRM Scientific Meeting and Exhibition; 2011; Montreal, Canada.

9. Aharon M, Elad M, Bruckstein A. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. Ieee T Signal Proces. 2006; 54(11):4311–4322.

10. Ravishankar S, Bresler Y. MR image reconstruction from highly undersampled k-space data by dictionary learning. IEEE Trans Med Imaging. 2011; 30(5):1028–1041. [PubMed: 21047708]

11. Otazo, R.; Sodickson, DK. Adaptive Compressed Sensing MRI; 18th Annual ISMRM Scientific Meeting and Exhibition; 2010; Stockholm, Sweden.

12. Gorodnitsky IF, Rao BD. Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm. Ieee T Signal Proces. 1997; 45(3):600–616.

13. Keil, B.; Blau, JM.; Hoecht, P.; Biber, S.; Hamm, M.; Wald, LL. A 64-channel brain array coil for 3T imaging; 20th Annual ISMRM Scientific Meeting and Exhibition; 2012; Melbourne, Australia.

14. Bodammer N, Kaufmann J, Kanowski M, Tempelmann C. Eddy current correction in diffusion-weighted imaging using pairs of images acquired with opposite diffusion gradient polarity. Magnet Reson Med. 2004; 51(1):188–193.

15. Jenkinson M, Bannister P, Brady M, Smith S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. Neuroimage. 2002; 17(2):825–841. [PubMed: 12377157]

16. Wakana S, Caprihan A, Panzenboeck MM, Fallon JH, Perry M, Gollub RL, Hua K, Zhang J, Jiang H, Dubey P, Blitz A, van Zijl P, Mori S. Reproducibility of quantitative tractography methods applied to cerebral white matter. Neuroimage. 2007; 36(3):630–644. [PubMed: 17481925]

17. Yendiki A, Panneck P, Srinivasan P, Stevens A, Zollei L, Augustinack J, Wang R, Salat D, Ehrlich S, Behrens T, Jbabdi S, Gollub R, Fischl B. Automated probabilistic reconstruction of white-matter pathways in health and disease using an atlas of the underlying anatomy. Front Neuroinform. 2011; 5:23. [PubMed: 22016733]

18. Talairach, J.; Tournoux, P. Co-planar stereotaxic atlas of the human brain: 3-dimensional proportional system: an approach to cerebral imaging. Vol. viii. Stuttgart; New York: G. Thieme; New York: Thieme Medical Publishers; 1988. p. 122

19. Setsompop K, Gagoski BA, Polimeni JR, Witzel T, Wedeen VJ, Wald LL. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. Magn Reson Med. 2012; 67(5):1210–1224. [PubMed: 21858868]
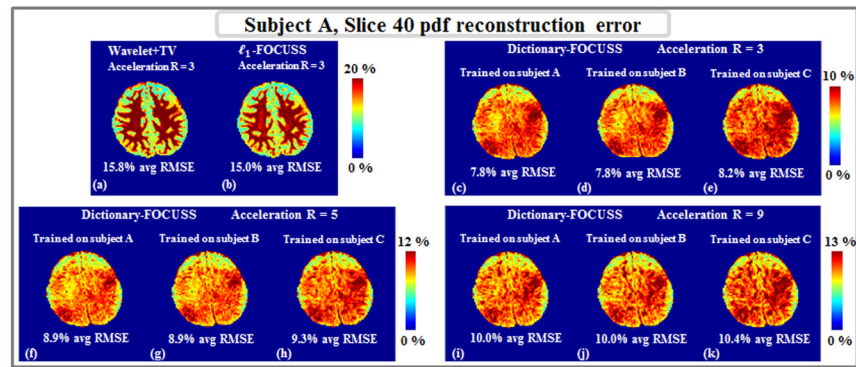
Fig. 1.

RMSE at each voxel in slice 40 of subject A upon $R = 3$ acceleration and reconstruction with Menzel et al.'s method (a), ℓ_1-FOCUSS (b), Dictionary-FOCUSS trained on subjects A (c), B (d), and C (e). Dictionary-FOCUSS errors in (f), (g) and (h) are obtained at higher acceleration factor of $R = 5$ with training on subjects A, B and C, respectively. Results for the reconstructions at $R = 9$ are given in (i), (j) and (k).
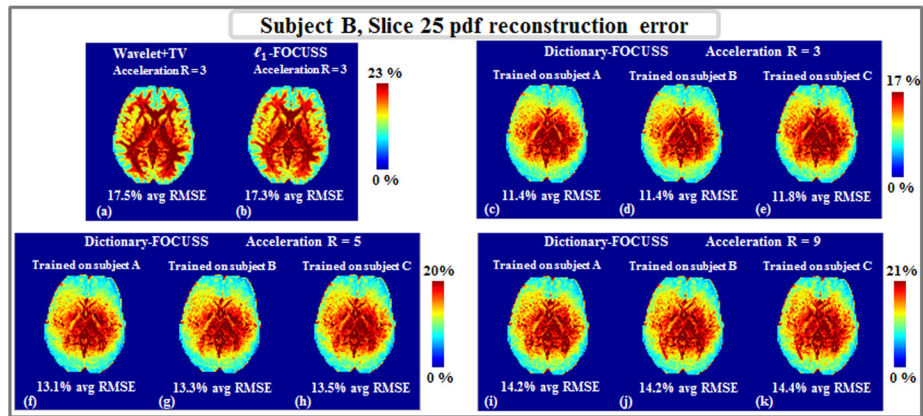
Fig. 2.
RMSE at each voxel in slice 25 of subject B upon $R = 3$ acceleration and reconstruction with Menzel et al.'s method (a), ℓ_1-FOCUSS (b), Dictionary-FOCUSS trained on subjects A (c), B (d), and C (e). Dictionary-FOCUSS errors in (f), (g) and (h) are obtained at higher acceleration factor of $R = 5$ with training on subjects A, B and C, respectively. Results for the reconstructions at $R = 9$ are given in (i), (j) and (k).
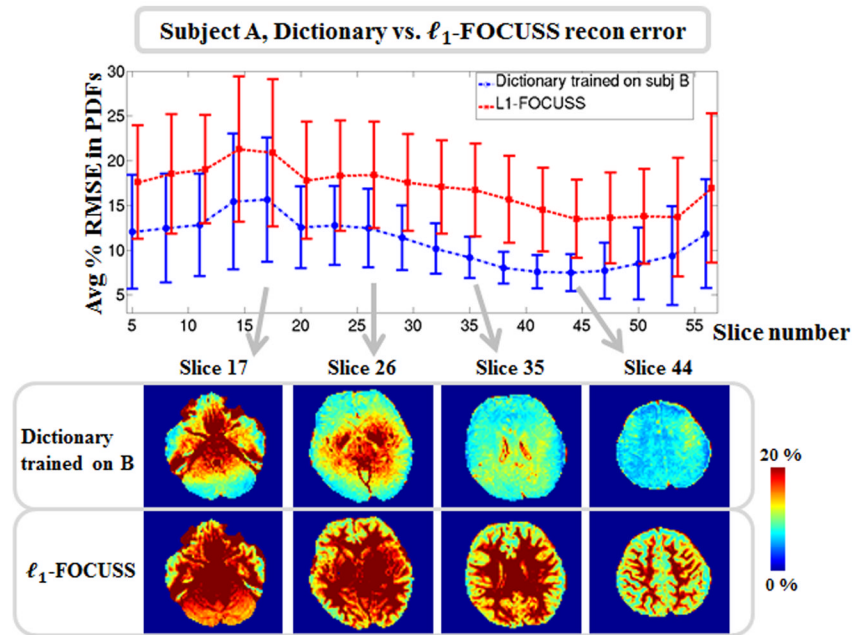
Fig. 3.
Mean and standard deviation of RMSEs computed on various slices of subject A using ℓ_1- and Dictionary-FOCUSS trained on subject B. Lower panel depicts RMSE maps for four selected slices.
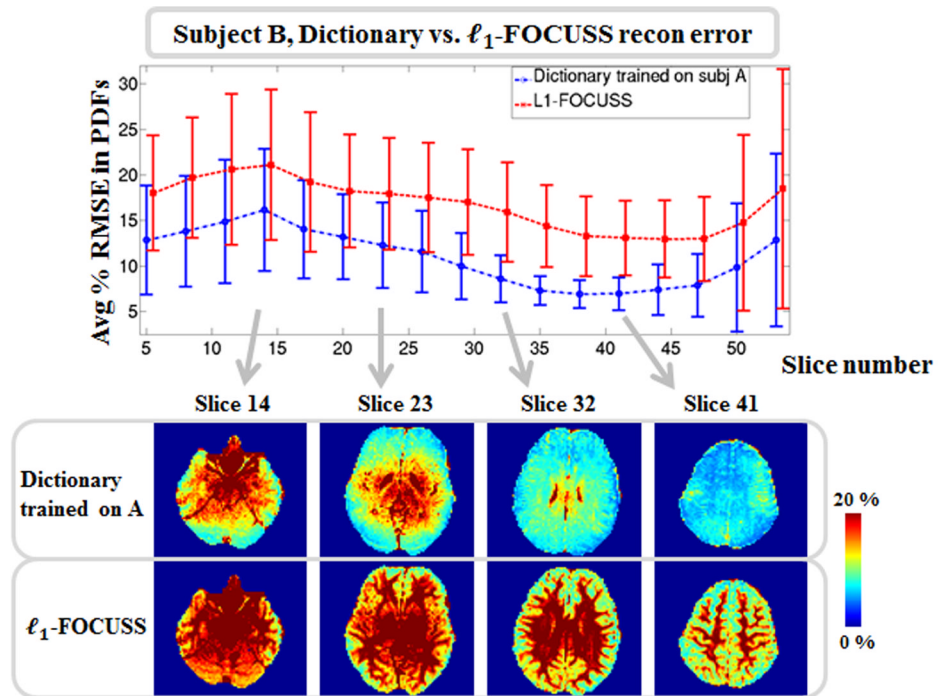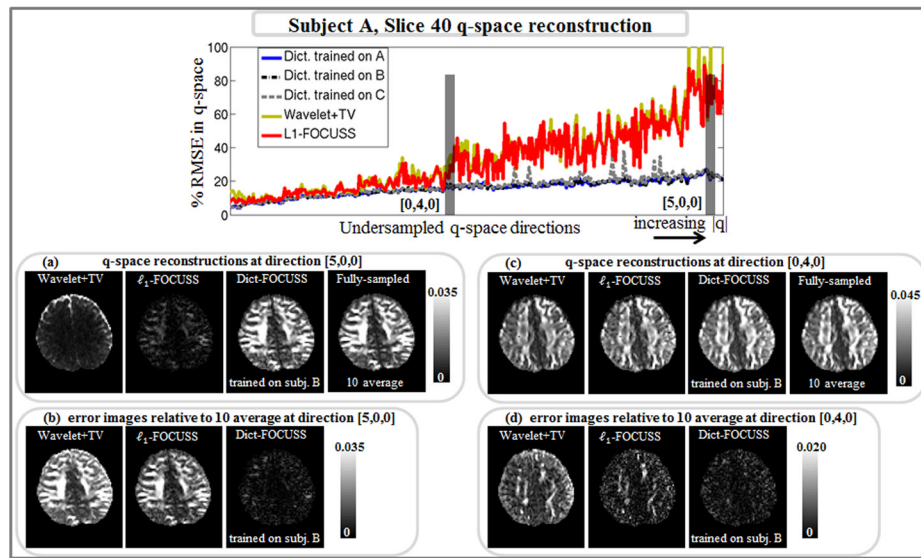
Fig. 4.
Mean and standard deviation of RMSEs computed on various slices of subject B using ℓ_1- and Dictionary-FOCUSS trained on subject A. Lower panel depicts RMSE maps for four selected slices.

Fig. 5.
Top panel shows RMSEs in 'missing' q-space directions that are estimated with Wavelet +TV, ℓ_1-FOCUSS and Dictionary-FOCUSS with training on subjects A, B and C at R=3. q-space images at directions [5,0,0] (a) and [0,4,0] (c) are depicted for comparison of the reconstruction methods. In panels (b) and (d), reconstruction errors of Wavelet+TV, ℓ_1-FOCUSS and dictionary reconstructions relative to the 10 average fully-sampled image at directions [5,0,0] and [0,4,0] are given.
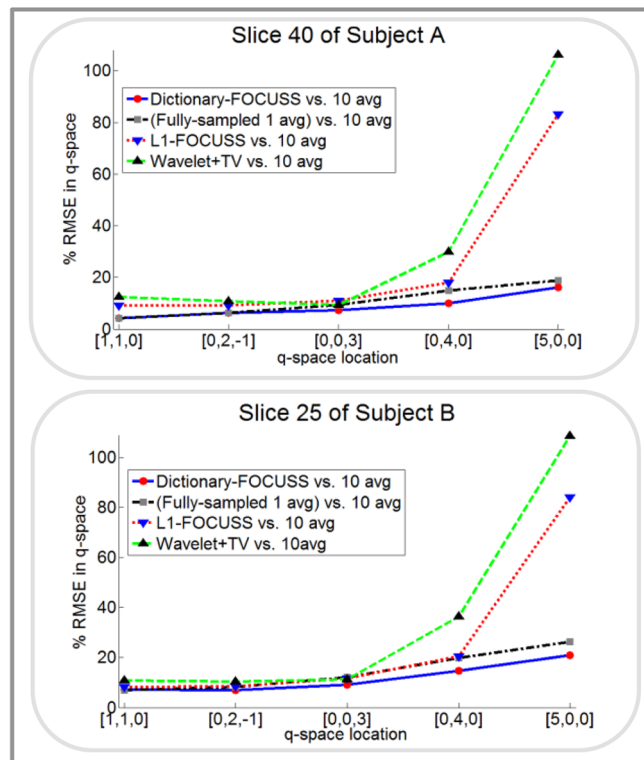
Fig. 6.
Panel on top depicts RMSEs of Wavelet+TV, ℓ_1-FOCUSS and Dictionary-FOCUSS at $R = 3$ and fully-sampled 1 average data computed in 5 q-space locations relative to the 10 average data for subject A. Panel on the bottom shows the same comparison for the slice belonging to subject B.
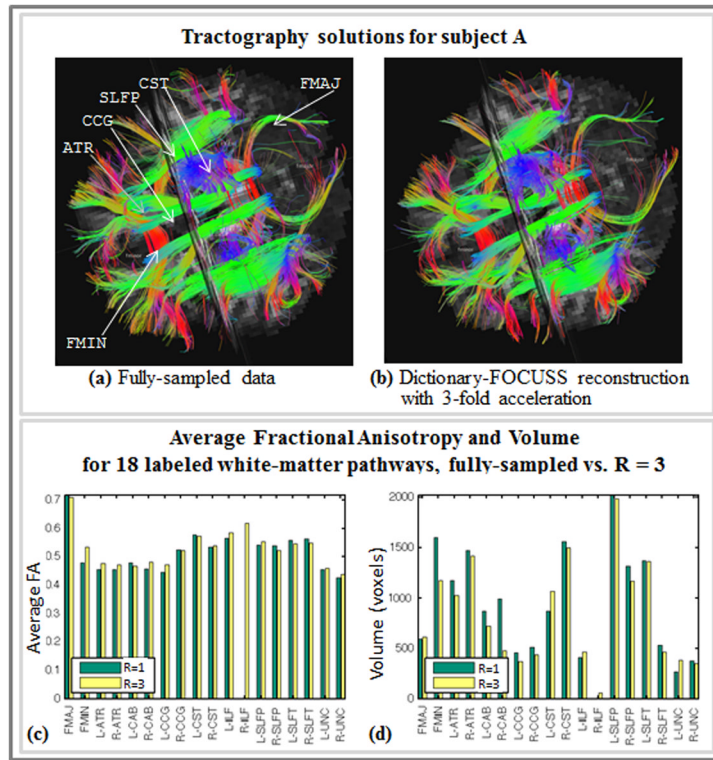
Fig. 7.
Axial view of white-matter pathways labeled from streamline DSI tractography in fully-sampled data (a) and Dictionary-FOCUSS reconstruction at $R = 3$ (b). The following are visible in this view: corpus callosum - forceps minor (FMIN), corpus callosum - forceps major (FMAJ), anterior thalamic radiations (ATR), cingulum - cingulate gyrus bundle (CCG), superior longitudinal fasciculus - parietal bundle (SLFP), and the superior endings of the corticospinal tract (CST). Average FA (c) and volume in number of voxels (d) for each of the 18 labeled pathways, as obtained from the fully-sampled (R=1, green) and Dictionary-FOCUSS reconstructed with 3-fold undersampling (R=3, yellow) datasets belonging to subject A. Intra-hemispheric pathways are indicated by "L-" (left) or "R-" (right). The pathways are: corpus callosum - forceps major (FMAJ), corpus callosum - forceps minor (FMIN), anterior thalamic radiation (ATR), cingulum - angular (infracallosal) bundle (CAB), cingulum - cingulate gyrus (supracallosal) bundle (CCG), corticospinal tract (CST), inferior longitudinal fasciculus (ILF), superior longitudinal fasciculus - parietal bundle (SLFP), superior longitudinal fasciculus - temporal bundle (SLFT), uncinate fasciculus (UNC).