



Published in final edited form as:

Magn Reson Med. 2020 August ; 84(2): 762–776. doi:10.1002/mrm.28172.

Fast submillimeter diffusion MRI using gSlider-SMS and SNR-enhancing joint reconstruction

Justin P. Haldar¹, Yunsong Liu¹, Congyu Liao², Qiuyun Fan², Kawin Setsompop²

¹Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA, USA

²Athinoula A. Martinos Center for Biomedical Imaging, Charlestown, MA, USA

Abstract

Purpose: We evaluate a new approach for achieving diffusion MRI data with high spatial resolution, large volume coverage, and fast acquisition speed.

Theory and Methods: A recent method called gSlider-SMS enables whole-brain submillimeter diffusion MRI with high signal-to-noise ratio (SNR) efficiency. However, despite the efficient acquisition, the resulting images can still suffer from low SNR due to the small size of the imaging voxels. This work proposes to mitigate the SNR problem by combining gSlider-SMS with a regularized SNR-enhancing reconstruction approach.

Results: Illustrative results show that, from gSlider-SMS data acquired over a span of only 25 minutes on a 3T scanner, the proposed method is able to produce 71 MRI images (64 diffusion encoding orientations with $b = 1500$ s/mm², and 7 images without diffusion weighting) of the entire in vivo human brain with nominal 0.66 mm spatial resolution. Using data acquired from 75 minutes of acquisition as a gold standard reference, we demonstrate that the proposed SNR-enhancement procedure leads to substantial improvements in estimated diffusion parameters compared to conventional gSlider reconstruction. Results also demonstrate that the proposed method has advantages relative to denoising methods based on low-rank matrix modeling. A theoretical analysis of the trade-off between spatial resolution and SNR suggests that the proposed approach has high efficiency.

Conclusions: The combination of gSlider-SMS with advanced regularized reconstruction enables high-resolution quantitative diffusion MRI from a relatively fast acquisition.

Keywords

constrained reconstruction; denoising; diffusion MRI; fast imaging; high resolution

Correspondence Justin P. Haldar, Hughes Aircraft Electrical Engineering Center (EEB) #442, M/C 2564, University of Southern California, University Park Campus, 3740 McClintock Avenue, Los Angeles, CA 90089, USA. jhaldar@usc.edu.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section.

1 | INTRODUCTION

It is challenging to acquire quantitative whole-brain in vivo human diffusion MRI data with submillimeter resolution. One of the main obstacles is the limited signal-to-noise ratio (SNR) associated with small voxel sizes. For illustration, a classical theoretical analysis¹ suggests that if we want the SNR of a 660 μm acquisition to match the SNR of a 2 mm acquisition with all other imaging parameters held equal, then the 660 μm data would need over 770 \times more data averaging, because SNR is proportional to the product between the voxel volume and the square-root of the number of averages. This degree of averaging is not practical for in vivo human applications.

Nevertheless, there has been considerable recent progress toward achieving higher resolution diffusion MRI.² On the data-acquisition side, some of the major advances have been driven by the use of simultaneous multi-slice (SMS)³⁻⁵ and volumetric⁶⁻¹² excitation methods, which yield substantially higher SNR efficiency relative to standard 2D spatial encoding methods. Related approaches acquire multiple sets of low-resolution high-SNR images with RF encoding, and computationally fuse these low-resolution images together into a “super-resolved” high-resolution image.¹³⁻²⁰ One recent approach, gSlider-SMS,¹⁸⁻²⁰ combines SMS and RF encoding ideas, and is able to achieve whole-brain in vivo human diffusion MRI with isotropic submillimeter resolution with relatively high SNR-efficiency. Unfortunately, the small size of the imaging voxels still means that gSlider-SMS often requires signal averaging.

Substantial progress has also been made in the use of signal processing methods to improve the SNR of noisy diffusion images (for example, see the many denoising techniques that are reviewed in Ref. 21, as well as more recent methods such as Refs. 22-26). One such approach, SNR-enhancing joint reconstruction (SER),^{21,27-29} is designed to spatially smooth the data to improve SNR while also preserving image edge features that are shared between different DWIs to avoid the loss of important high-resolution information. In contrast to many other denoising techniques that are difficult to characterize theoretically, the SER approach has a strong theoretical characterization that can provide the user with a precise understanding of the trade-offs associated with SER. It is known that SER is associated with a controllable trade-off between SNR improvement and spatial resolution,^{21,27-29} and the SER theory empowers the user to choose an appropriate balance between these two factors. Importantly, it has been shown that significant gains in SNR can be achieved using spatial smoothing with only modest corresponding losses in spatial resolution,^{21,27-32} and that the impact of these modest resolution losses can be largely mitigated because of the edge preservation properties of SER. In addition, previous analyses^{21,29} suggest that the regularization penalty behaves in a stable way in the presence of assumption violations, which can provide more confidence in its good performance when applied to data with nonstandard features. For example, previous work has shown that SER works well with images of both normal and injured tissue^{21,28} (as would be predicted by theory), while other popular nonlinear denoising approaches may work well with normal tissue but can yield problematic results when subtle injuries are present.²⁸

In this work, we propose and study the combination of gSlider-SMS with SER, and show that this combination enables state-of-the-art performance in achieving fast submillimeter diffusion MRI. Preliminary accounts of portions of this work were previously presented in Ref. [33].

2 | THEORY

2.1 | gSlider-SMS

For simplicity and brevity, this paper will describe gSlider-SMS¹⁸⁻²⁰ for the simplified case in which parallel imaging and SMS reconstruction have already been performed, and it remains to reconstruct the high-resolution spatial information along the slice dimension. Rather than exciting thin slices, conventional gSlider encoding excites slabs that are several times thicker than the desired high-resolution slice thickness, and uses multiple RF slab-encodings across consecutive TRs to resolve the constituent subslices that comprise the slab. For the purposes of our description and without loss of generality, we will assume that the thick slab is five times larger than the nominal slice thickness and that we seek to recover five constituent subslices. To achieve this, gSlider acquires a series of different images of the same slab, where the RF pulse is varied in each repetition to apply a distinct phase modulation pattern to the subslices, in a manner similar to Hadamard encoding and related methods.^{34,35}

If \mathbf{b} denotes the vector concatenating the collection of thick-slab images acquired with different RF encoding profiles and \mathbf{f} represents the high-resolution image, then the RF-encoding model (in the absence of motion-induced phase variations) can be written as $\mathbf{b} = \mathbf{A}\mathbf{f}$, where matrix \mathbf{A} captures the RF encoding procedure. Given this model, it is possible to recover the desired high-resolution slice information through standard linear Tikhonov-regularized reconstruction¹⁸:

$$\hat{\mathbf{f}} = (\mathbf{A}^H \mathbf{A} + \lambda \mathbf{I})^{-1} \mathbf{A}^H \mathbf{b}, \quad (1)$$

where λ is a regularization parameter and \mathbf{I} is the identity matrix. In practice, real data will possess motion-induced phase variations, and it is necessary to apply appropriate phase correction to \mathbf{b} prior to applying Equation 1. We refer to the phase-corrected version of Equation 1 as conventional gSlider reconstruction.¹⁸

The gSlider encoding matrix \mathbf{A} is designed to have two useful properties that allow it to outperform alternative RF encoding schemes like Fourier,³⁵ Hadamard,³⁵ and random^{36,37} RF encoding. First, the gSlider is designed to have high image SNR for all the RF encodings, unlike for example, Hadamard or Fourier encoding. This enables robust phase estimation and correction, which is important due to the phase inconsistencies of diffusion MRI.¹⁸ Second, the \mathbf{A} matrix is designed to be well conditioned, which ensures that the matrix inversion in Equation 1 does not substantially exacerbate noise, and actually leads to substantial SNR gains relative to 2D slice-by-slice acquisition.¹⁸ In addition, the use of RF encoding can offer sharper point-spread functions with less signal leakage/Gibbs ringing than k-space Fourier encoding when used for thin-slab encoding.

2.2 | SNR-enhancing joint reconstruction

While gSlider encoding is relatively SNR efficient, the use of small isotropic voxel sizes means that SNR is still a limiting factor. We propose to use SER^{21,27-29} to address this issue. SER uses SNR-efficient spatial smoothing to reduce noise perturbations, while leveraging the structural similarity between different diffusion weighted images (DWIs) to avoid blurring high-resolution image features. Since the data considered in this work is acquired with a partial Fourier readout, our SER approach also uses phase modeling to prevent the loss of in-plane resolution while also automatically accounting for the phase discrepancies associated with diffusion MRI.

SER is performed by solving an optimization problem that couples regularized denoising and partial Fourier reconstruction as in Ref. 21 (with appropriate modifications to incorporate the gSlider RF encoding model) with regularized phase estimation as in Refs. [38,39] (Our previous SER work²¹ used a two-step approach, in which the phase term for each image was estimated a priori using a low-resolution reconstruction of each image, and this phase estimate was subsequently used to estimate the image amplitude. This approach works well and is still used as an initialization in the present work. However, our new formulation allows for potentially fixing any errors that may exist in the initial phase estimate.). Specifically, we solve

$$\{\hat{\mathbf{f}}, \hat{\mathbf{p}}\} = \underset{\mathbf{f}, \mathbf{p}}{\operatorname{argmin}} \|\mathbf{b} - \mathbf{G}(e^{i\mathbf{p}} \odot \mathbf{A}\mathbf{f})\|_2^2 + \lambda_1 R(\mathbf{p}) + \lambda_2 J(\mathbf{f}). \quad (2)$$

The data vector $\mathbf{b} \in \mathbb{C}^{(5N_1N_2N_sN_d) \times 1}$ contains the complex-valued voxel values of the RF-encoded slab images (obtained after slice-GRAPPA reconstruction). When specifying the dimension of \mathbf{b} , we have assumed that there are 5 RF encodings, that each RF-encoded slab image contains $N_1 \times N_2$ voxels in-plane, that we have acquired slabs at N_s different slab positions to achieve full volume coverage, and that for diffusion encoding, we acquire data for N_d images with varying diffusion parameters. We use $N_3 = 5N_s$ to represent the total number of high-resolution thin subslices that we wish to recover. The optimization variable $\mathbf{f} \in \mathbb{R}^{(N_1N_2N_3N_d) \times 1}$ is the vector of the unknown high-resolution (ie, based on thin subslices) real-valued image voxel amplitudes for all of the diffusion weighted images, while $\mathbf{A} \in \mathbb{C}^{(5N_1N_2N_sN_d) \times (N_1N_2N_3N_d)}$ is the matrix modeling the RF-encoding procedure. The optimization variable $\mathbf{p} \in \mathbb{R}^{(5N_1N_2N_sN_d) \times 1}$ is used to model the unknown phase of each measured thick-slab acquisition (to enable phase correction to compensate for motion-induced phase discrepancies for gSlider-SMS reconstruction, as well as providing the phase constraints needed for resolution recovery from partial Fourier data acquisition), and $\mathbf{G} \in \mathbb{C}^{(5N_1N_2N_sN_d) \times (5N_1N_2N_sN_d)}$ is a matrix used to model the in-plane point-spread function of partial Fourier acquisition (ie, Fourier transformation of each image into k-space, setting the unmeasured portion of k-space equal to zero, followed by inverse Fourier transformation). The symbol \odot is used to denote the Hadamard product (ie, element-wise multiplication of two vectors) and $e^{i\mathbf{p}} \in \mathbb{C}^{(5N_1N_2N_sN_d) \times 1}$ is used to denote the vector of exponentiated phase values. A graphical depiction of this equation is shown in Figure 1.

The first term in Equation 2 is a standard least-squares data consistency penalty that encourages the reconstructed high-resolution image to be consistent with the low-resolution measured data. Partial Fourier constraints are imposed by requiring that \mathbf{f} is real-valued.²¹

The second term $R(\mathbf{p})$ in Equation 2 is a regularization penalty that encourages the estimated image phase to be smooth within each acquired image slab, but with no constraints on the phase behavior between different slabs or between different diffusion weighted images (as appropriate in the presence of random phase variations). To avoid difficulties associated with phase-wrapping and the nonuniqueness of phase, the regularization penalty is designed to regularize the exponentiated phase (a choice that is insensitive to 2π phase jumps^{38,39}), with

$$\begin{aligned} R(\mathbf{p}) &= \sum_{s=1}^{N_s} \sum_{q=1}^{N_d} \sum_{k=1}^5 \sum_{n=1}^{N_1 N_2} \sum_{m \in \Omega_n} |e^{i p_{sqkn}} - e^{i p_{sqkm}}|^2 \\ &= \|\mathbf{D}e^{i\mathbf{p}}\|_2^2. \end{aligned} \quad (3)$$

In this equation, the p_{sqkn} values denote the phase value from \mathbf{p} corresponding to the n th voxel in the s th slab with the k th RF encoding and the q th diffusion encoding, Ω_n corresponds to the set of 4 voxels that are immediate in-plane spatial neighbors to the n th voxel, and \mathbf{D} is the matrix representation of the finite differencing operation.

The third term $J(\mathbf{f})$ in Equation 2 is a regularization penalty that leverages the prior knowledge that \mathbf{f} is expected to be spatially smooth within each DWI volume, but also has edge structures that are common between different DWIs which ideally should be preserved by the reconstruction procedure.^{21,27-29} If we use f_n^q to denote the entry from \mathbf{f} corresponding to the n th voxel of the q th DWI, then we choose $J(\mathbf{f})$ following Refs. [21,22,27-29] as

$$J(\mathbf{f}) = \sum_{n=1}^{N_1 N_2 N_3} \sum_{m \in \Delta_n} \Psi \left(\sqrt{\sum_{q=1}^{N_d} |f_n^q - f_m^q|^2} \right), \quad (4)$$

where Δ_n corresponds to the set of 6 voxels that are immediate volumetric spatial neighbors of the n th voxel, and $\Psi(\cdot)$ is the convex Huber function:

$$\Psi(t) = \begin{cases} t^2, & t \leq \xi \\ 2\xi t - \xi^2, & t > \xi. \end{cases} \quad (5)$$

The Huber function is well known as a regularization penalty that is both convex and edge-preserving, and converges to a scaled version of the standard ℓ_2 -norm in the limit as ξ approaches zero. In the limit of small ξ , the penalty function of Equation 4 promotes joint sparsity of the image edges,²⁷ and becomes equivalent to a joint sparsity-promoting regularization penalty function that is popular in the multi-contrast compressed sensing MRI literature.^{40,41} We prefer to use larger values of ξ , since this enables quantitative theoretical characterizations of the resolution and noise characteristics of the reconstructed images

while still encouraging shared edge preservation.^{21,29} These theoretical characterizations offer a number of advantages as described previously.

In this work, we use global rescaling of the images in \mathbf{b} prior to image reconstruction, such that the median voxel intensity has the same magnitude across each of the DWIs. This normalization is useful to avoid the estimated shared-edge structure from being dominated by the images with the highest intensity (ie, the unweighted $b = 0$ images), and the normalization is easily removed once SER has been completed. This rescaling is equivalent to the use of additional hyperparameters that were present in the original formulation of SER,²¹ although leads to simpler equations.

As described in more detail in Ref. 21, we choose the reconstruction parameter ξ to be large enough that the estimated edge structures are relatively free of visible noise influences, and choose λ_2 based on the level of desired SNR improvement (eg, in the results we show later, we leverage the theoretical characterizations that are available for SER and choose λ_2 so that the SNR improvement in most regions of the brain is at least as good as $3\times$ averaging). The parameter λ_1 is chosen heuristically to obtain estimated phase maps that are not too noisy but which are also not visually oversmoothed.

For optimization of Equation 2, we choose an iterative alternation-based approach that guarantees monotonic decrease (and, therefore, convergence, although not necessarily to a global optimum) of the cost function value. In the first stage of the k th iteration, we estimate the value of the image magnitude $\hat{\mathbf{f}}$ given the previous estimate of the image phase according to

$$\hat{\mathbf{f}}^{(i)} = \underset{\mathbf{f}}{\operatorname{argmin}} \|\mathbf{b} - \mathbf{G}(e^{i\hat{\mathbf{p}}^{(i-1)}} \odot \mathbf{A}\mathbf{f})\|_2^2 + \lambda_2 J(\mathbf{f}), \quad (6)$$

subject to the constraint that \mathbf{f} be real-valued. The optimization problem in Equation 6 has identical form to that considered in Ref. 21, and we find a globally optimal solution using the simple iteratively-reweighted least-squares optimization algorithm that was previously proposed for this problem.²¹

In the second stage of the k th iteration, we estimate the value of the image phase $\hat{\mathbf{p}}$ given the estimate of the image amplitude according to

$$\hat{\mathbf{p}}^{(i)} = \underset{\mathbf{p}}{\operatorname{argmin}} \|\mathbf{b} - \mathbf{G}(e^{i\mathbf{p}} \odot \mathbf{A}\hat{\mathbf{f}}^{(i)})\|_2^2 + \lambda_1 \|\mathbf{D}e^{i\mathbf{p}}\|_2^2. \quad (7)$$

This is a nonlinear optimization problem that has exactly the same structure as the optimizations considered in previous work.^{38,39} Similar to the previous work^{38,39} we use the nonlinear conjugate gradient (NCG) algorithm with analytically-evaluated gradients to find a local minimum. In particular, we use a variation of NCG that uses the Polak-Ribière method constrained by the Fletcher-Reeves method.⁴² The gradient of Equation 7 can be computed analytically as

$$\nabla f(\mathbf{p}) = 2\text{imag} \left(e^{-i\mathbf{p}} \odot \mathbf{A}\hat{\mathbf{f}}^{(i)} \odot (\mathbf{G}^H \mathbf{G}(e^{i\mathbf{p}} \odot \mathbf{A}\hat{\mathbf{f}}^{(i)}) - \mathbf{G}^H \mathbf{b}) + \lambda_1 e^{-i\mathbf{p}} \odot (\mathbf{D}^H \mathbf{D}e^{i\mathbf{p}}) \right), \quad (8)$$

where $\text{imag}(\mathbf{x})$ denotes the imaginary part of a complex vector \mathbf{x} , and $_H$ denotes the standard conjugate transpose matrix operation.

The two optimization steps represented by Equations 6 and 7 are iterated until convergence, or until a maximum number of iterations has been met.

3 | METHODS

3.1 | Data acquisition and processing

Whole-brain gSlider-SMS diffusion MRI data were acquired on the 3T CONNECTOM system using a custom-built 64-channel array. Data were acquired corresponding to a nominal 660 μm isotropic resolution over a $220 \times 118 \times 151.8 \text{ mm}^3$ FOV (matrix size $332 \times 180 \times 230$), with 7 unweighted images (ie, $b = 0$) and 64 diffusion weighted images (DWIs) with $b = 1,500 \text{ s/mm}^2$ and uniformly distributed diffusion encoding orientations. Data were measured using an EPI readout with 2 simultaneously-acquired thick slabs (ie, each slab is 3.3 mm thick, which is 5 times the size of a 660 μm subslice) with blipped CAIPI³ and sagittal slab orientations (matrix size for each slab volume was $332 \times 180 \times 46$), 5 different RF encoding pulses, 6/8ths partial Fourier and $2 \times$ GRAPPA in-plane acceleration, ZOOPPA⁴³ outer volume suppression of the neck and phase-encoding along the superior-inferior axis, TE = 80 ms, and TR = 4.4 seconds per thick-slab volume (ie, 22 seconds per DWI). This entire measurement was acquired in ≈ 25 minutes, with 3 repetitions acquired to provide a gold standard reference (total ≈ 75 minutes). Prior to further processing steps, slice-GRAPPA⁴⁴ was used for SMS and parallel imaging reconstruction of the slab volumes, and in-plane registration was performed on these images using FSL.⁴⁵

gSlider reconstruction was performed for single-average data using the conventional gSlider approach (ie, phase estimation⁴⁶ followed by regularized reconstruction according to Equation 1) and the SER approach described above. A gold standard was obtained by applying conventional gSlider reconstruction (without SER) to each of the three repetitions, and the resulting phase-corrected real-valued images were then averaged together.

For comparison, we also applied denoising approaches based on low-rank matrix modeling to the images obtained from the conventional gSlider reconstruction. Low-rank modeling has become a relatively popular approach for modeling and denoising diffusion MRI data in recent years,^{22,24,29,47-49} and is based on the well-known denoising characteristics of principal component analysis (PCA). We considered three different types of low-rank modeling.

First, following Ref. [24], we applied sliding-window PCA denoising across the diffusion encoding dimension to overlapping spatial patches, and chose the rank value for each patch based on the Marchenko-Pastur distribution. This approach, which we denote as MPPCA, is based on locally low-rank modeling assumptions. Data processing for this approach was

performed using the MRtrix3 package (<http://www.mrtrix.org/>). We systematically evaluated a range of different patch sizes, up to a patch size of $19 \times 19 \times 19$ voxels (corresponding to a $12.5 \text{ mm} \times 12.5 \text{ mm} \times 12.5 \text{ mm}$ spatial region, matching the spatial coverage used in previous implementations for lower-resolution settings²⁴). Ultimately, we report results obtained using a patch size of $3 \times 3 \times 3$ voxels, which is substantially smaller than previously-used patch sizes (eg, which used $5 \times 5 \times 5$ voxels with 2.5 mm isotropic resolution or $7 \times 7 \times 7$ voxels with 1.5 mm isotropic resolution²⁴) but yielded the best denoising performance with respect to the quantitative error metrics described in the next section.

While various methods exist for estimating the optimal parameters to use in low-rank matrix denoising^{24,50,51} (including the previously-mentioned approach based on the Marchenko-Pastur distribution), we also were interested in understanding the best possible denoising performance that could be achieved using low-rank modeling constraints. As a result, the next two low-rank modeling methods, which we call local PCA (LPCA) and global PCA (GPCA), both choose rank parameter values to minimize the mean-squared error of the denoised images with respect to the gold standard images from 3 \times -averaged data. Note that an explicit objective of Refs. [50,51] is to choose parameters that minimize the mean-squared error, but these methods make potentially suboptimal choices because they do not have access to the true mean-squared error values. Since perfect adjustment of the rank parameter will not be available in practical experiments, the results obtained from LPCA and GPCA should be viewed as best-case scenarios for denoising with low-rank modeling constraints.

While LPCA and GPCA both use the same approach for rank selection, they differ from one another in the way they choose image patch sizes. Similar to MPPCA,²⁴ our LPCA implementation applies sliding-window PCA denoising to overlapping patches of size $3 \times 3 \times 3$ voxels, following a locally low-rank data model. Similar to Refs. [22,29], GPCA applies PCA denoising to the whole 3D image volume at once, following a globally low-rank data model.

3.2 | Performance metrics

Denoising performance was assessed using several different metrics. First, we computed the normalized root-mean-squared error (NRMSE) of the reconstructed DWIs with respect to the gold standard, defined as

$$NRMSE = \frac{\|\hat{\mathbf{f}} - \mathbf{f}_{\text{gold}}\|_2}{\|\mathbf{f}_{\text{gold}}\|_2} \quad (9)$$

where \mathbf{f}_{gold} denotes the gold standard result obtained from 3 \times averaging. While the gold standard data is not entirely noise-free, it does have substantially better quality than the data acquired without averaging, so smaller values of this NRMSE metric are expected to correlate well with improved denoising performance.

For many practical applications, the errors in the denoised images are probably less important than errors in the quantitative diffusion parameters that can be estimated from the

images. As a result, we computed several different diffusion parameters. Because diffusion data was collected with a single-shell q-space acquisition, we focused on quantitative parameters that can be estimated from single-shell data.

We started by using the software underlying the BrainSuite Diffusion Pipeline (<http://brainsuite.org/>) to estimate the diffusion tensor imaging (DTI) parameters of mean diffusivity (MD) and fractional anisotropy (FA). To obtain quantitative error measures, we computed NRMSE values within the brain for both the FA and MD maps.

We also used the software underlying the BrainSuite Diffusion Pipeline to estimate orientation distribution functions (ODFs) in two different ways: the Funk-Radon transform (FRT)⁵²⁻⁵⁴ and the Funk-Radon and Cosine Transform (FRACT).⁵⁵ The FRT and FRACT both estimate higher-order ODF representations that, unlike the simpler DTI model, are capable of capturing complicated fiber-crossing structures in white matter. Comparing these two methods, the FRT is expected to be less sensitive to noise than FRACT, while FRACT is expected to produce ODFs with higher angular resolution than the FRT. Both of these methods are usually designed to be used with high b-value diffusion data sampled with a large number of diffusion encoding directions. Since the data acquired in this work used a modest b-value with a modest number of encoding directions, the results of FRT and FRACT are not expected to be very impressive in this case, even for the gold standard data. Nevertheless, we believe that the errors observed in the FRT and FRACT ODF estimates before and after denoising should still be reflective of denoising performance. To obtain quantitative error measures, we computed NRMSE values within the white matter for the spherical harmonic coefficients used to represent the FRT and FRACT ODFs. Specifically, NRMSE was computed over white matter voxels where the gold standard FA was >0.3 .

4 | RESULTS

Gold standard, conventional gSlider, and SER images from several slices of a representative DWI are illustrated in Figure 2, with the differences with respect to the gold standard shown in Supporting Information Figure S1. As can be seen, the conventional gSlider result has a noisy appearance, particularly near the center of the brain. On the other hand, the SER has a much less noisy appearance, as should be expected.

More detail is available from inspecting the zoomed-in images presented in Figure 3, which also includes comparisons against MPPCA, LPCA, and GPCA. From this figure, it can be observed that all four denoising approaches have an apparent reduction in noise, although the extent of this varies from method to method. MPPCA appears to have the least amount of noise reduction among the four, while both GPCA and LPCA produce very crisp-looking images with minimal obvious noise content. We would argue that the SER result does not look as crisp or cosmetically pleasing as the LPCA and GPCA results, although still represents an improvement over MPPCA.

Looking at the NRMSE values for the DWIs, which are shown in Table 1, we observe that LPCA has the smallest NRMSE, followed by SER, GPCA, and MPPCA. The good performance of LPCA in this case is not surprising, since its parameters were tuned using

perfect knowledge of the gold standard to optimize this NRMSE value, and because LPCA has more flexibility than GPCA to adapt to the different properties of different image regions. On the other hand, it may be surprising that the SER result has smaller NRMSE than the GPCA result, given that the GPCA result was also designed to minimize this NRMSE value based on perfect prior knowledge that was not available to SER. This suggests that even in the ideal case, while global low-rank modeling can offer some degree of improvement, this model may be intrinsically more limited than some of these other approaches in this type of application. Interestingly, there is a major performance difference between LPCA and MPPCA, even though both of these results are based on the same underlying locally low-rank data model. This result illustrates the importance of choosing good rank parameters, although it should be noted that MPPCA uses a conservative automatic rank-selection rule that is not necessarily designed to achieve optimal NRMSE.

Figures 4 and 5, respectively, show color-coded FA maps corresponding to the same images from Figures 2 and 3, and NRMSE values corresponding to the quantitative DTI parameters (MD and FA) are also shown in Table 1. Interestingly, SER and MPPCA are now the two best denoising methods with respect to these DTI parameters, and SER has a substantial advantage over MPPCA. Visually, Figure 5 shows that SER results in a color FA map that does not appear to be as noisy as the color FA maps from the other methods.

Surprisingly, even though LPCA had the smallest NRMSE with respect to the DWIs, it actually was one of the worst-performing methods with respect to these quantitative DTI parameters. In addition, the NRMSE values in these cases are not much improved over those obtained from conventional gSlider without any denoising. This likely occurs because quantitative model-fitting in diffusion MRI can be sensitive to small variations in the measured diffusion signal, while by its nature, low-rank modeling seeks to find a low-dimensional subspace representation that preserves signal energy as much as possible, which does not prioritize the preservation of smaller signal variations. This result also underscores the more general point that visually pleasant images with small error values will not necessarily yield good results when those images are used in quantitative applications, a concern that has also been raised in previous related work.^{21,22,29,30}

Another interesting observation is that the color FA map for GPCA from Figure 5 has a much noisier-looking appearance than might have been expected based on the relatively noise-free appearance of the corresponding DWIs from Figure 3. This might be explained by the fact that the DTI model is itself low-dimensional, such that the space of all possible signals that are compatible with the DTI model may already approximately reside in a low-dimensional subspace.²² Thus, the DTI fitting procedure may have an effect that is similar in some sense to low-rank modelling, and we might not expect low-rank denoising to yield a dramatic improvement in DTI parameters unless the explicit subspace constraints imposed by low-rank modeling are sufficiently distinct from the implicit subspace constraints associated with DTI modeling. In the case of global low-rank modeling (as used by GPCA), the subspace constraints are required to apply broadly to the data from all voxels, and so may not be very distinct from the implicit subspace constraints imposed by the DTI model.

Figures 6 and 7, respectively, show ODF estimation results for FRT and FRACT, with quantitative NRMSE values also reported in Table 1. The figures show that the spatial distribution of ODFs appears somewhat chaotic and disorganized when ODFs are estimated from the noisy conventional gSlider images, while the ODFs estimated from SER have more spatial coherence and match better with the characteristics of the ODFs estimated from the gold standard images. Quantitatively, SER still possesses the smallest NRMSE values for both FRT and FRACT ODFs. LPCA and MPPCA are virtually tied for second place with respect to FRT ODFs, and LPCA stands alone in second place for FRACT ODFs.

In addition to achieving good quantitative performance, SER also has the advantage that it is possible to theoretically characterize its noise and resolution characteristics.^{21,29} This allows us to have a deeper understanding of the limitations and trade-offs associated with SER reconstruction. While MPPCA, LPCA, and GPCA certainly should also be associated with various limitations and trade-offs, we are unaware of any convenient methods to compute characterizations of these trade-offs.

Figure 8 shows spatial maps of the expected reduction in noise variance obtained by using SER instead of conventional gSlider reconstruction, which were obtained using the exact same techniques described in previous SER publications.^{21,29} As noted previously, the regularization parameters for SER were designed to achieve at least a 3× reduction in noise variance across most brain regions. For reference, note that an expected 3× reduction in noise variance is commensurate with 3× averaging. As can be seen from this figure, the reduction in noise variance varies spatially, with larger improvements in SNR in smooth image regions, and smaller improvements in SNR near estimated image edge locations. This behavior is expected, and results from the fact that SER is designed to avoid smoothing across edge structures to preserve high-resolution image content.^{21,29} Notably, central regions of the brain are also observed to have less SNR-enhancement than outer portions of the brain, even though these regions are not necessarily associated with strong edge content. This occurs because the center of the brain has lower SNR than brain regions with closer proximity to the receiver coils, and SER has a more difficult time identifying the difference between actual image edge structure and noise in this region. Overall, the level of noise variance reduction ranged from approximately 3-5 across the entire image.

Importantly, even though SER employs edge-preserving regularization, the SNR improvement associated with SER is also associated with a degradation in spatial resolution, and this degradation can be characterized theoretically using point-spread functions/spatial response functions (SRFs).^{21,29} Using techniques described in previous publications,^{21,29,31} we calculated SRFs for one voxel where the noise variance reduction was approximately 3 and another voxel where the noise variance reduction was approximately 5, with results shown in Figure 9. For reference, we also show SRFs for conventional gSlider on the same axes. In both cases, the use of SER is associated with a small degradation in spatial resolution.

While there is no unique way of measuring spatial resolution,³² we have opted to use the full-volume at half-maximum (FVHM) of the SRF as a resolution measure. The FVHM is calculated by summing the volumes of all voxels for which the SRF is larger than half of its

maximum value, and can be viewed as a generalization of the conventional full-width at half maximum for one-dimensional point-spread functions to the three-dimensional setting. We observe that the FVHM of the SRF is fairly shift-invariant for conventional gSlider (as we had expected), with a value of roughly $0.380 \text{ mm}^3 = (724 \text{ }\mu\text{m})^3$ for all voxels. In contrast, the FVHMs for SER vary with the level of SNR improvement. For the voxel with $3\times$ noise variance reduction, the FVHM of the SRF for SER was roughly $0.426 \text{ mm}^3 = (752 \text{ }\mu\text{m})^3$, while the FVHM of the SRF was roughly $0.472 \text{ mm}^3 = (778 \text{ }\mu\text{m})^3$ for the voxel with $5\times$ noise variance reduction. Consistent with previous SER publications^{21,27-32} (and perhaps surprising from the perspective of conventional imaging expectations), the SNR improvement associated with SER is quite substantial compared to the relatively minor degradations observed in spatial resolution. In addition, it should also be remembered that SER is designed to avoid the blurring of information across image edges. In particular, the SRF for SER adapts itself to avoid signal from leaking across the edge structures of the image,^{29,30} such that even this minor loss in resolution may not be as deleterious as it might have been otherwise.

This figure also reveals that the spatial resolution characteristics for both SER and conventional gSlider are different along the RF-encoding dimension than they are along either of the two Fourier-encoding dimensions, as should be expected due to the use of substantially different spatial encoding mechanisms. In addition, we observe substantial asymmetry in the SRF along the RF-encoding dimension. While we have only shown results from two voxels, results from other voxels reveals that the shape of the SRF varies considerably depending on the position of the reconstructed voxel within the RF-encoded thick-slab that was used for acquisition. Reconstructed voxels from the center of a thick slab have SRFs that are more symmetric than voxels that are closer to the edges of the thick slabs. We also observe the tendency to smooth voxels more within the same slab rather than across different slabs. This is likely related to the structure of the inverse problem, where there is more ambiguity within the same slab than there is across slabs. The effects of this are also observed in the noise variance reduction maps from Figure 8, where some striping is observed along slab boundaries.

5 | DISCUSSION

Overall, while SER did not have the smallest NRMSE values with respect to the denoised DWIs, it did have the smallest NRMSE values in all other cases compared to globally and locally low-rank modeling methods. In addition, SER has useful theoretical characterizations that are generally not available for low-rank denoising approaches. While we cannot comprehensively compare SER against all existing diffusion MRI denoising methods, we should note that previous work has compared SER against non-local means denoising with similarly promising results.^{21,28,29}

While this paper presented a comparison of SER against low-rank modeling methods, it should be noted that SER and low-rank modeling are based on different principles and can potentially be combined synergistically. It has been previously demonstrated in a different context that combining SER with globally low-rank modeling can lead to further performance improvements.²² Based on our results, we expect that the combination of SER

with locally low-rank modeling could be even better. However, one challenge is that the use of any low-rank modeling can invalidate the theoretical characterizations of SER, while the use of SER can invalidate the noise modeling assumptions of automatic rank-selection rules. This can cause parameter selection for both approaches to become more difficult.

One of the key principles of SER is to estimate a single shared edge map that captures the collective features of the ensemble of DWIs, which is used to guide the edge-preserving reconstruction and denoising of all images simultaneously. Even though different images will almost certainly possess unique edge characteristics (eg, consider the edge-correlation matrix illustrated for different DWIs in Figure 10 of previous work²¹), the use of a shared edge map was intentionally chosen for SER because it ensures that the SRFs for a given voxel will be consistent across different DWIs.²¹ This is important because it implies that every DWI will be “blurred” in the same way (ie, the same voxel always represents the same spatial distribution of spins across different DWIs), which is important for the validity of subsequent quantitative analyses of the data. Notably, many other nonlinear denoising methods (including locally low-rank modeling approaches) do not possess the same kind of strong guarantees about the spatial consistency of different DWIs (and indeed, frequently do not possess theoretical spatial resolution characterizations of any kind). In addition, while SER should be expected to perform best when DWIs are properly coregistered and contain complementary edge information, it has been previously shown that the effects of inaccurate edge information or misregistration are expected to be relatively benign.²⁹ In particular, SER might lose some of its ability to correctly identify and preserve edge structures in such cases, but it will still possess an advantageous and well-characterized trade-off between spatial resolution and SNR.²⁹

It should also be noted that the gSlider approach can potentially suffer from image artifacts and blurring in the presence of large subject motion. Techniques have been previously developed to make gSlider robust to subject movements.¹⁹ The data shown in this paper did not have major motion issues, so we did not use such techniques in this work. However, it would be straightforward to combine motion-corrected gSlider methods with SER, with performance benefits anticipated for both gSlider and SER.

Slow computation speed is one of the limitations of our current implementation, and for example, the results shown in this paper took multiple days to compute using a simple unoptimized Matlab implementation running on a Linux-based workstation with two quad-core Xeon 2.27 GHz processors and 48 GB of RAM. Most of this time (>90%) was associated with solving the phase-update step associated with Equation 7. While slow computation speed may be understandable given the very large size of this dataset (ie, the gSlider data vector **b** alone requires more than 15 GB of memory to store in its entirety in double precision!), it would of course be preferable if this computation were faster. Notably, although the phase-update step can lead to non-negligible phase changes as illustrated in Supporting Information Figure S4, we observed that, at least with this data, using the phase-update step resulted in changes to the estimated magnitude images that were largely negligible, both visually and quantitatively (eg, the change in NRMSE for the DTI parameters was $< 10^{-3}$). As such, while we believe that updating the phase using Equation 7 should be preferred, one practical approach to improving computational complexity would

be to skip the phase-update step and hold the phase fixed at its initial value. Alternatively, our group has recently been exploring faster algorithms for this type of optimization problem that have the potential to substantially reduce the amount of computation time,⁵⁶ although we believe that a thorough exploration of different algorithmic alternatives is beyond the scope of the present paper.

In the implementation of SER that was used for this paper, we relied on complex-valued images obtained from an initial parallel imaging and SMS reconstruction procedure, instead of performing SER from raw k-space data as advocated in previous SER papers.^{21,29} This choice was made because the raw k-space data for this acquisition was too big for easy manipulation and storage, and even if the k-space data had been stored, SER reconstruction from raw k-space data would have been practically unworkable without relying on specialized computational resources. As a result, the SER results were obtained from data possessing less information content than the original raw data, and for example, we did not have access to the complementary information provided by different channels or information about the spatially-varying noise characteristics associated with the images. We anticipate that denoising quality could potentially have been much better if at least a little more of the original information had been preserved. For example, if information about the spatially-varying noise variance had been preserved, it may have been possible to use this information to avoid the situation where the center of the brain experiences less denoising than exterior parts of the brain. While there exist various techniques for trying to estimate the characteristics of spatially-varying noise fields as an inverse problem,⁵⁷ there are also direct ways of precisely calculating and preserving this information at the time of image reconstruction,^{23,58} and we believe that exploring the use of such noise maps is a promising direction for future work.

One of our observations from this study was that basic SER reconstruction yields different resolution characteristics along the RF-encoded dimension than for the Fourier-encoded dimensions. It is not immediately clear whether this difference in resolution characteristics is problematic. However, if it is ultimately deemed to be problematic, it is interesting to note that techniques already exist for modifying regularized reconstruction methods to achieve more uniform resolution characteristics.^{59,60} The application of such ideas to SER may also be an interesting topic for future research.

We should also note that the results shown in this paper are based on the use of full RF-encoding for every DWI in acquisition together with spatial smoothness constraints for SNR-enhancement, but without any constraints on the structure of the data in q-space. In principal, several groups have previously shown that if a diffusion MRI acquisition requires multiple encodings for each DWI, it can be possible to eliminate certain encodings by exploiting the smoothness of the q-space signal.^{15,16,61} This approach is complementary to SER, and a preliminary exploration of combining SER with subsampled RF encoding has shown promising results.⁶² Further exploration of this kind of approach may enable even faster high-resolution diffusion MRI using gSlider-SMS.

Finally, we should note that the datasets shown in this paper were acquired on the 3T CONNECTOM system, which has unique capabilities that may not necessarily translate

directly to more common scanner hardware configurations. However, gSlider can still be implemented on standard 3T MRI hardware²⁰ and the combination of gSlider acquisition and SER-based image reconstruction can still offer substantial benefits in such settings, as we have demonstrated in recent preliminary work.⁶³ While we believe that a detailed exploration of this is beyond the scope of the present paper, we show illustrative results with 860 μm isotropic resolution obtained in 10 minutes from a standard (commercially-available) 3T Prisma scanner in Supporting Information Figure S5.

6 | CONCLUSION

This work proposed a new approach to fast diffusion MRI that uses a highly-efficient acquisition strategy together with an advanced and theoretically-characterizable denoising strategy to enable a relatively fast diffusion MRI experiment with state-of-the-art volume coverage, spatial resolution, and SNR. We believe that this kind of approach can prove useful across the full range of in vivo human diffusion MRI applications.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Funding information

National Institutes of Health, Grant/Award Number: R01-NS074980, R01-NS089212, R21-EB022951 and R01-MH116173; National Science Foundation, Grant/Award Number: CCF-1350563

REFERENCES

1. Macovski A. Noise in MRI. *Magn Reson Med*. 1996;36:494–497. [PubMed: 8875425]
2. Holdsworth SJ, O'Halloran R, Setsompop K. The quest for high spatial resolution diffusion-weighted imaging of the human brain in vivo. *NMR Biomed*. 2019;32:e4056. [PubMed: 30730591]
3. 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;67:1210–1224. [PubMed: 21858868]
4. Setsompop K, Kimmlingen R, Eberlein E, et al. Pushing the limits of in vivo diffusion MRI for the Human Connectome Project. *NeuroImage*. 2013;80:220–233. [PubMed: 23707579]
5. Sotiropoulos SN, Jbabdi S, Xu J, et al. Advances in diffusion MRI acquisition and processing in the Human Connectome Project. *NeuroImage*. 2013;80:125–143. [PubMed: 23702418]
6. Jeong EK, Kim SE, Parker DL. High-resolution diffusion-weighted 3D MRI, using diffusion-weighted driven equilibrium (DW-DE) and multishot segmented 3D-SSFP without navigator echoes. *Magn Reson Med*. 2003;50:821–829. [PubMed: 14523969]
7. Engström M, Skare S. Diffusion-weighted 3D multislabs echo planar imaging for high signal-to-noise ratio efficiency and isotropic image resolution. *Magn Reson Med*. 2013;70:1507–1514. [PubMed: 23359357]
8. Nguyen C, Fan Z, Sharif B, et al. In vivo three-dimensional high resolution cardiac diffusion weighted MRI: a motion compensated diffusion-prepared balanced steady-state free precession approach. *Magn Reson Med*. 2014;72:1257–1267. [PubMed: 24259113]
9. Frost R, Miller KL, Tijssen RHN, Porter DA, Jezzard P. 3D multislabs diffusion-weighted readout-segmented EPI with real-time cardiac reordered k-space acquisition. *Magn Reson Med*. 2014;72:1565–1579. [PubMed: 24347093]

10. Chang HC, Sundman M, Petit L, et al. Human brain diffusion tensor imaging at submillimeter isotropic resolution on a 3 Tesla clinical MRI scanner. *NeuroImage*. 2015;118:667–675. [PubMed: 26072250]
11. Van AT, Aksoy M, Holdsworth SJ, Kopeinigg D, Vos SB, Bammer R. Slab profile encoding (PEN) for minimizing slab boundary artifact in three-dimensional diffusion-weighted multislabs acquisition. *Magn Reson Med*. 2015;73:605–613. [PubMed: 24691843]
12. Holtrop JL, Sutton BP. High spatial resolution diffusion weighted imaging on clinical 3 T MRI scanners using multislabs spiral acquisitions. *J Med Imaging*. 2016;3:023501.
13. Greenspan H. Super-resolution in medical imaging. *Comput J*. 2009;52:43–63.
14. Scherrer B, Gholipour A, Warfield SK. Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisitions. *Med Image Anal*. 2012;16:1465–1476. [PubMed: 22770597]
15. Van Steenkiste G, Jeurissen B, Veraart J, et al. Super-resolution reconstruction of diffusion parameters from diffusion-weighted images with different slice orientations. *Magn Reson Med*. 2016;75:181–195. [PubMed: 25613283]
16. Ning L, Setsompop K, Michailovich O, et al. A joint compressed-sensing and super-resolution approach for very high-resolution diffusion imaging. *NeuroImage*. 2016;125:386–400. [PubMed: 26505296]
17. Hennel F, Tian R, Engel M, Pruessmann KP. In-plane "superresolution" MRI with phaseless sub-pixel encoding. *Magn Reson Med*. 2018;80:2384–2392. [PubMed: 29656440]
18. Setsompop K, Fan Q, Stockmann J, et al. High-resolution in vivo diffusion imaging of the human brain with generalized slice dithered enhanced resolution: simultaneous multislice (gSlider-SMS). *Magn Reson Med*. 2018;79:141–151. [PubMed: 28261904]
19. Wang F, Bilgic B, Dong Z, et al. Motion-robust sub-millimeter isotropic diffusion imaging through motion corrected generalized slice dithered enhanced resolution (MC-gSlider) acquisition. *Magn Reson Med*. 2018;80:1891–1906. [PubMed: 29607548]
20. Liao C, Stockmann J, Tian Q, et al. High-fidelity, high-isotropic resolution diffusion imaging through gSlider acquisition with $<10^{-3}$ & T₁ corrections and integrated δB_0 /Rx shim array. Preprint 2018;ArXiv:1811.05473.
21. Haldar JP, Wedeen VJ, Nezamzadeh M, et al. Improved diffusion imaging through SNR-enhancing joint reconstruction. *Magn Reson Med*. 2013;69:277–289. [PubMed: 22392528]
22. Lam F, Babacan SD, Haldar JP, Weiner MW, Schuff N, Liang ZP. Denoising diffusion-weighted magnitude MR images using rank and edge constraints. *Magn Reson Med*. 2014;71:1272–1284. [PubMed: 23568755]
23. Varadarajan D, Haldar JP. A majorize-minimize framework for Rician and non-central chi MR images. *IEEE Trans Med Imaging*. 2015;34:2191–2202. [PubMed: 25935028]
24. Veraart J, Novikov DS, Christiaens D, Ades-Aron B, Sijbers J, Fieremans E. Denoising of diffusion MRI using random matrix theory. *NeuroImage*. 2016;142:394–406. [PubMed: 27523449]
25. Sperl JI, Sprenger T, Tan ET, Menzel MI, Hardy CJ, Marinelli L. Model-based denoising in diffusion-weighted imaging using generalized spherical deconvolution. *Magn Reson Med*. 2017;78:2428–2438. [PubMed: 28244188]
26. Kafali SG, Cukur T, Saritas EU. Phase-correcting non-local means filtering for diffusion-weighted imaging of the spinal cord. *Magn Reson Med*. 2018;80:1020–1035. [PubMed: 29427379]
27. Haldar JP, Liang ZP. Joint reconstruction of noisy high-resolution MR image sequences. In: *IEEE International Symposium on Biomedical Imaging, Paris, 2008*:752–755.
28. Kim JH, Song SK, Haldar JP. Signal-to-noise ratio-enhancing joint reconstruction for improved diffusion imaging of mouse spinal cord white matter injury. *Magn Reson Med*. 2016;75:852–858. [PubMed: 25824472]
29. Haldar JP. *Constrained Imaging: Denoising and Sparse Sampling* [PhD thesis]. Urbana, IL, USA: University of Illinois at Urbana-Champaign; 2011.
30. Haldar JP, Hernando D, Song SK, Liang ZP. Anatomically constrained reconstruction from noisy data. *Magn Reson Med*. 2008;59:810–818. [PubMed: 18383297]

31. Haldar JP, Sakaie K, Liang ZP. Resolution and noise properties of linear phase-constrained partial Fourier reconstruction. In: Proceedings of the International Society for Magnetic Resonance in Medicine, Honolulu, 2009:2862.
32. Haldar JP, Liang ZP. On MR experiment design with quadratic regularization. In: Proceedings of the IEEE International Symposium on Biomedical Imaging, Chicago, 2011:1676–1679.
33. Haldar JP, Fan Q, Setsompop K. Whole-brain quantitative diffusion MRI at 660 μm resolution in 25 minutes using gSlider-SMS and SNR-enhancing joint reconstruction. In: Proceedings of the International Society for Magnetic Resonance in Medicine, Singapore, 2016:102.
34. Maudsley AA. Multiple-line-scanning spin density imaging. *J Magn Reson.* 1980;41:112–126.
35. Oh CH, Park HW, Cho ZH. Line-integral projection reconstruction (LPR) with slice encoding techniques: Multislice regional imaging in NMR tomography. *IEEE Trans Med Imaging.* 1984;MI-3:170–178.
36. Haldar JP, Hernando D, Liang ZP. Compressed-sensing MRI with random encoding. *IEEE Trans Med Imaging.* 2011;30:893–903. [PubMed: 20937579]
37. Kim TH, Haldar JP. SMS-LORAKS: Calibrationless simultaneous multislice MRI using low-rank matrix modeling. In: Proceedings of the IEEE International Symposium on Biomedical Imaging, New York City, 2015:323–326.
38. Haldar JP, Wang Z, Popescu G, Liang ZP. Deconvolved spatial light interference microscopy for live cell imaging. *IEEE Trans Biomed Eng.* 2011;58:2489–2497. [PubMed: 21622067]
39. Zhao F, Noll DC, Nielsen JF, Fessler JA. Separate magnitude and phase regularization via compressed sensing. *IEEE Trans Med Imaging.* 2012;31:1713–1723. [PubMed: 22552571]
40. Trzasko J, Manduca A. Group sparse reconstruction of vector-valued images. In: Proceedings of the International Society for Magnetic Resonance in Medicine, Montreal, 2011:2839.
41. Huang J, Chen C, Axel L. Fast multi-contrast MRI reconstruction. *Magn Reson Imaging.* 2014;32:1344–1352. [PubMed: 25193110]
42. Gilbert JC, Nocedal J. Global convergence properties of conjugate gradient methods for optimization. *SIAM J Optim.* 1992;2:21–42.
43. Heidemann RM, Anwander A, Feiweier T, Knösche TR, Turner R. k-space and q-space: combining ultra-high spatial and angular resolution in diffusion imaging using ZOOPPA at 7 T. *NeuroImage.* 2012;60:967–978. [PubMed: 22245337]
44. Cauley SF, Polimeni JR, Bhat H, Wald LL, Setsompop K. Interslice leakage artifact reduction technique for simultaneous multislice acquisitions. *Magn Reson Med.* 2014;72:93–102. [PubMed: 23963964]
45. Jenkinson M, Beckmann CF, Behrens TEJ, Woolrich MW, Smith SM. FSL. *NeuroImage.* 2012;62:782–790. [PubMed: 21979382]
46. Eichner C, Cauley SF, Cohen-Adad J, et al. Real diffusion-weighted MRI enabling true signal averaging and increased diffusion contrast. *NeuroImage.* 2015;122:373–384. [PubMed: 26241680]
47. Manjón JV, Coupé P, Concha L, Buades A, Collins DL, Robles M. Diffusion weighted image denoising using overcomplete local PCA. *PLOS ONE.* 2013;8:e73021. [PubMed: 24019889]
48. Gao H, Li L, Zhang K, Zhou W, Hu X. PCLR: phase-constrained low-rank model for compressive diffusion-weighted MRI. *Magn Reson Med.* 2014;72:1330–1341. [PubMed: 24327553]
49. Ma S, Nguyen C, Christodoulou A, et al. Accelerated cardiac diffusion tensor imaging using joint low-rank and sparsity constraints. *IEEE Trans Biomed Eng.* 2018;65:2219–2230. [PubMed: 29989936]
50. Candès EJ, Sing-Long CA, Trzasko JD. Unbiased risk estimates for singular value thresholding and spectral estimators. *IEEE Trans Signal Process.* 2013;61:4643–4657.
51. Gavish M, Donoho DL. Optimal shrinkage of singular values. *IEEE Trans Inf Theory.* 2017;63:2137–2152.
52. Tuch DS. Q-ball imaging. *Magn Reson Med.* 2004;52:1358–1372. [PubMed: 15562495]
53. Hess CP, Mukherjee P, Han ET, Xu D, Vigneron DB. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magn Reson Med.* 2006;56:104–117. [PubMed: 16755539]

54. Descoteaux M, Angelino E, Fitzgibbons S, Deriche R. Regularized, fast, and robust analytical Q-ball imaging. *Magn Reson Med*. 2007;58:497–510. [PubMed: 17763358]
55. Halдар JP, Leahy RM. Linear transforms for Fourier data on the sphere: application to high angular resolution diffusion MRI of the brain. *NeuroImage*. 2013;71:233–247. [PubMed: 23353603]
56. Liu Y, Halдар JP. NAPALM: An algorithm for MRI reconstruction with separate magnitude and phase regularization. In: *Proceedings of the International Society for Magnetic Resonance in Medicine, Montreal, 2019*:4764.
57. Aja-Fernandez S, Vegas-Sanchez-Ferrero G. *Statistical Analysis of Noise in MRI: Modeling, Filtering and Estimation*. Springer; 2016.
58. Robson PM, Grant AK, Madhuranthakam AJ, Lattanzi R, Sodickson DK, McKenzie CA. Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. *Magn Reson Med*. 2008;60:895–907. [PubMed: 18816810]
59. Ahn S, Leahy RM. Analysis of resolution and noise properties of nonquadratically regularized image reconstruction methods for PET. *IEEE Trans Med Imaging*. 2008;27:413–424. [PubMed: 18334436]
60. Stayman JW, Fessler JA. Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction. *IEEE Trans Med Imaging*. 2000;19:601–615. [PubMed: 11026463]
61. Bhushan C, Joshi AA, Leahy RM, Halдар JP. Improved B_0 -distortion correction in diffusion MRI using interlaced q-space sampling and constrained reconstruction. *Magn Reson Med*. 2014;72:1218–1232. [PubMed: 24464424]
62. Halдар JP, Setsompop K. Fast high-resolution diffusion MRI using gSlider-SMS, interlaced subsampling, and SNR-enhancing joint reconstruction. In: *Proceedings of the International Society for Magnetic Resonance in Medicine, Honolulu, 2017*:174.
63. Liu Y, Liao C, Setsompop K, Halдар JP. Whole-brain DTI at 860 μm isotropic resolution in 10 minutes on a commercial 3T scanner. In: *Proceedings of the International Society for Magnetic Resonance in Medicine, Montreal, 2019*:3352.

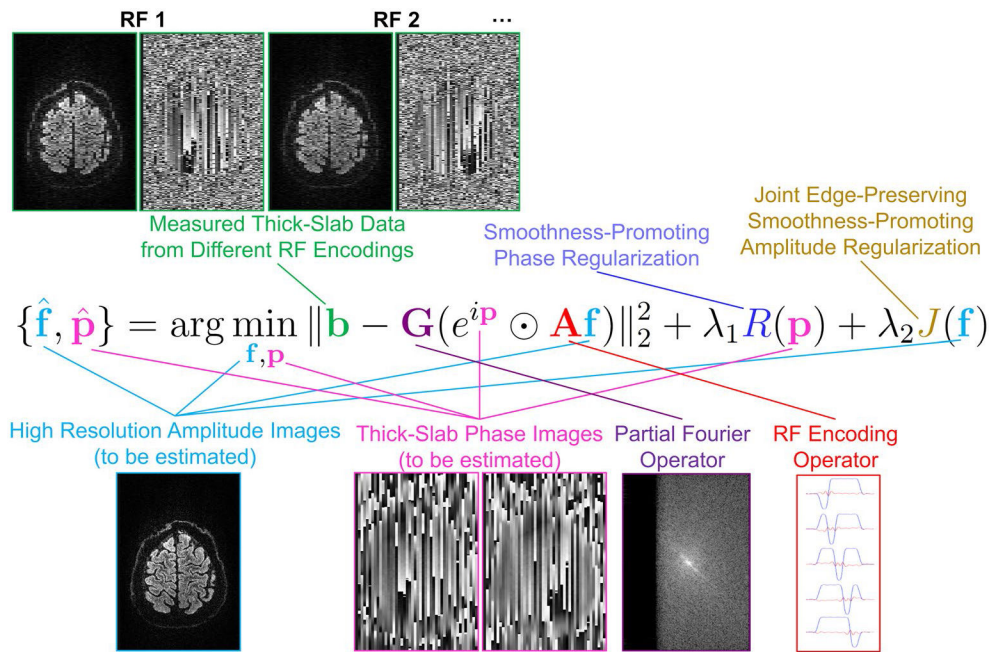


FIGURE 1. A graphical overview of the proposed image reconstruction formulation, depicting all of the relevant inputs, outputs, and operators. In this figure, RF encoding is performed using sagittal slabs, and the phase smoothness prior is applied in 2D within each of the sagittal slabs

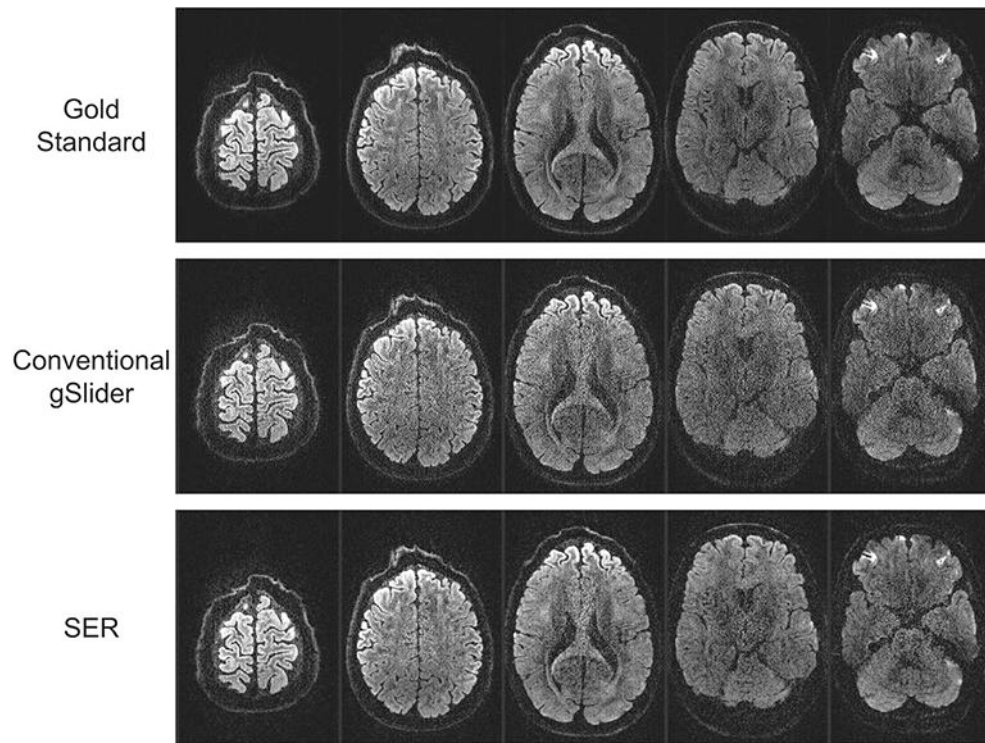


FIGURE 2. Gold standard, conventional gSlider, and SER images from five slices of a representative DWI. For reference, the differences with respect to the gold standard are also shown in Supporting Information Figure S1




FIGURE 3.

Comparison of a slice from a representative DWI obtained by different reconstruction and denoising approaches. Zoomed-in images (from the region corresponding to the yellow box shown in the gold standard image) are also shown for easier visualization of fine image details

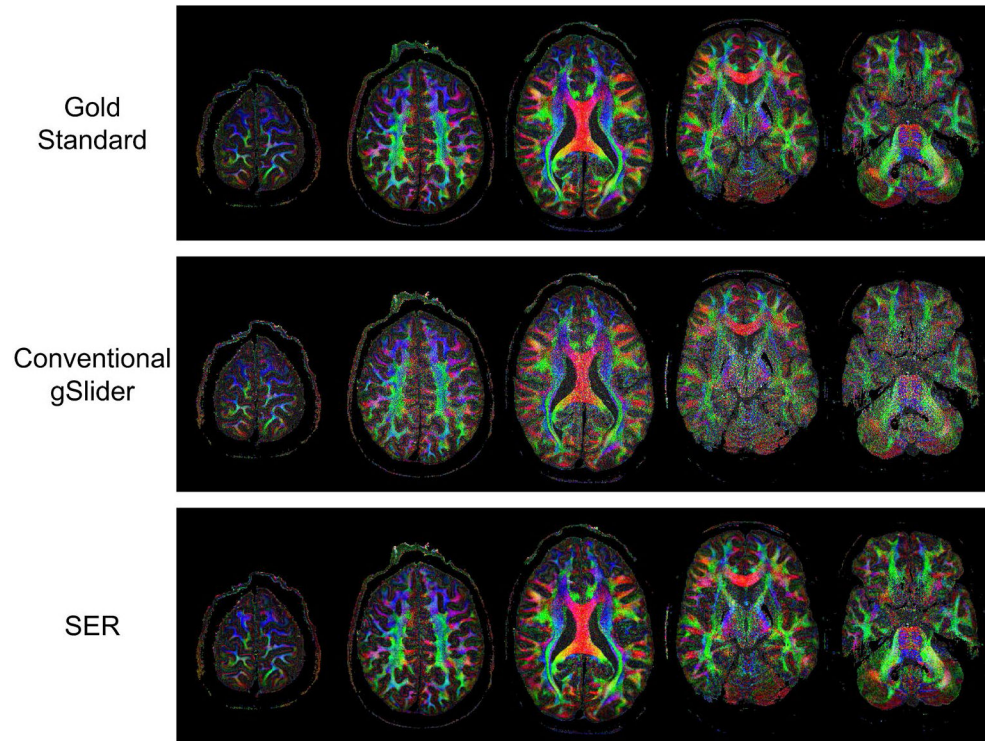


FIGURE 4. Color FA maps computed based on gold standard, conventional gSlider, and SER images from five slices of the brain volume

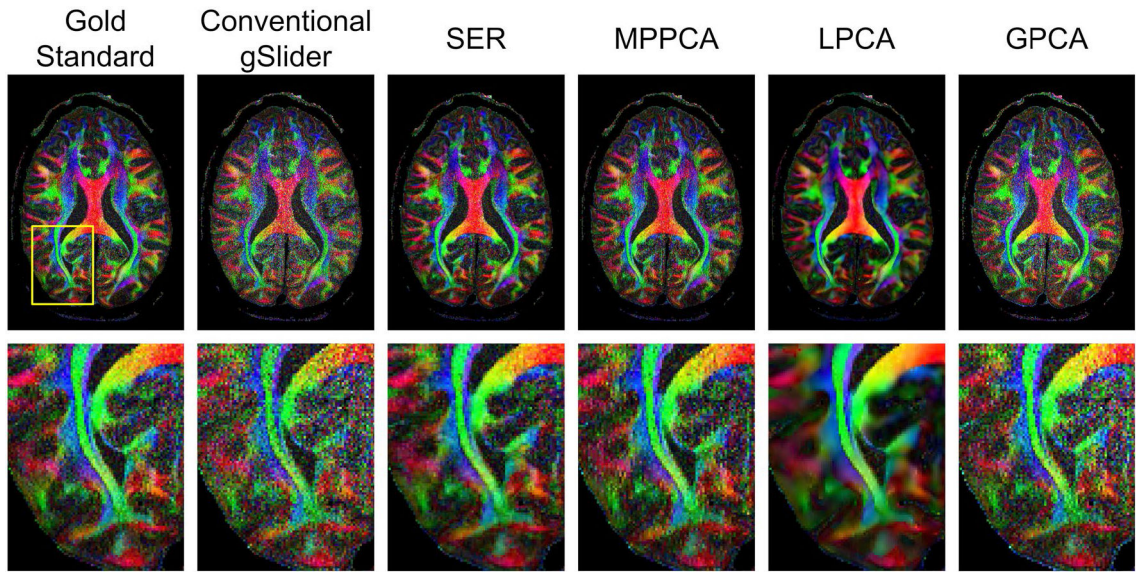


FIGURE 5. Comparison of color FA maps obtained by different reconstruction and denoising approaches. Zoomed-in images (from the region corresponding to the yellow box shown in the gold standard map) are also shown for easier visualization of fine details

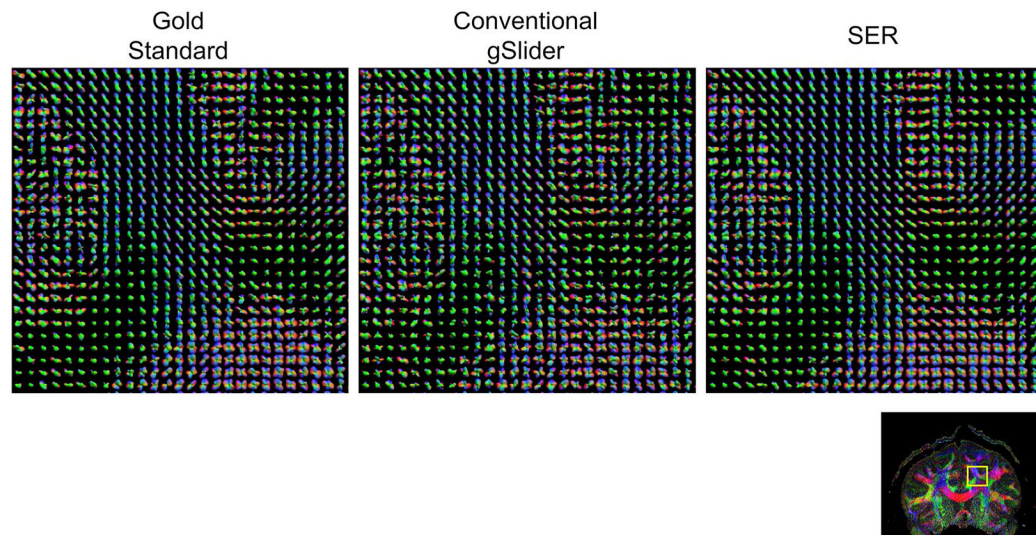


FIGURE 6.

ODFs estimated using the FRT from gold standard, conventional gSlider, and SER images. ODFs are shown from the brain region indicated by the yellow box in the reference color FA image. We have chosen to show a relatively large ROI to make it easier to visualize the higher level white-matter connectivity patterns that are present in the data. In addition, a higher resolution version of this figure is provided in Supporting Information Figure S2 for readers interested in the details of individual ODFs

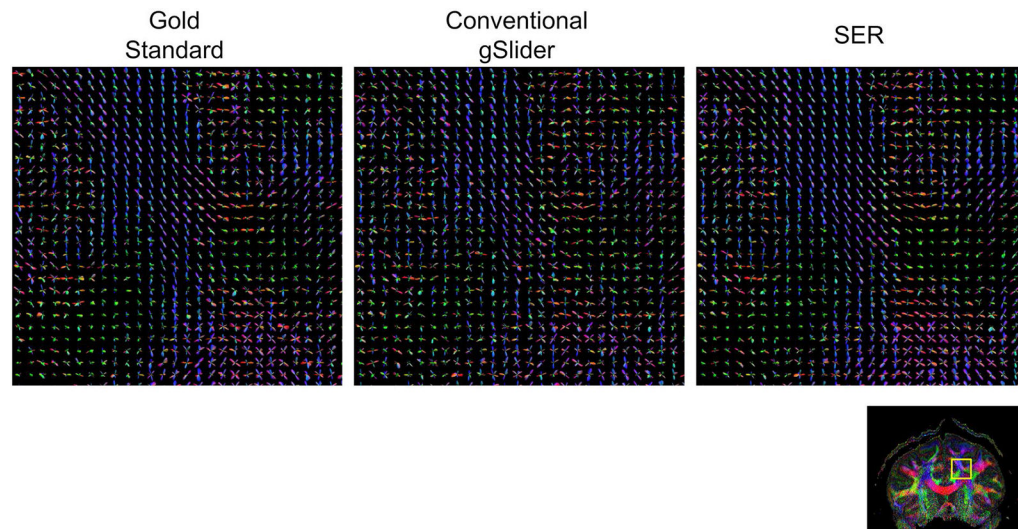


FIGURE 7.

ODFs estimated using the FRACT from gold standard, conventional gSlider, and SER images. ODFs are shown from the brain region indicated by the yellow box in the reference color FA image. We have chosen to show a relatively large ROI to make it easier to visualize the higher level white-matter connectivity patterns that are present in the data. In addition, a higher resolution version of this figure is provided in Supporting Information Figure S3 for readers interested in the details of individual ODFs




FIGURE 8. Spatial maps of the expected reduction in noise variance obtained by using SER instead of conventional gSlider reconstruction. Maps are shown corresponding to the same image slices from Figures 2 and 4

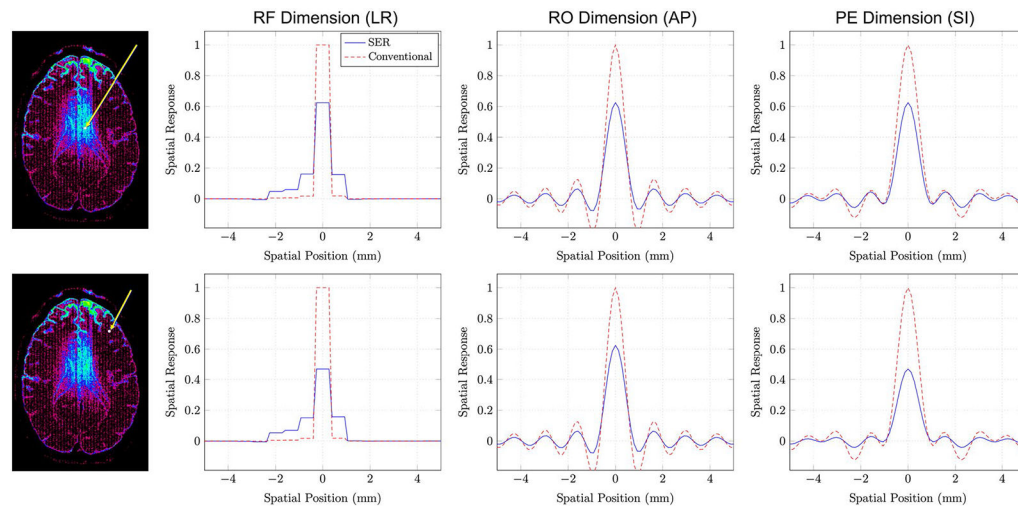


FIGURE 9.

SRFs for SER and conventional gSlider obtained from (top row) a voxel where the SNR improvement associated with SER was approximately 3 and (bottom row) a voxel where the SNR improvement associated with SER was approximately 5. The voxel positions are indicated as shown in the images on the left. The SRFs in this case are three-dimensional functions, which are hard to display. For easier visualization, we have shown one-dimensional plots passing through the peak of the SRF along different orientations. Specifically, we show SRF plots along the RF-encoding dimension (left-right anatomically), the readout encoding dimension (anterior-posterior anatomically), and the phase encoding dimension (superior-inferior anatomically)

Quantitative error measures obtained with different reconstruction and denoising methods

TABLE 1

	Conventional	<i>g</i> Slider	SER	MPPCA	LPCA	GPCA
NRMSE of DWIs	0.28	0.22	0.22	0.25	0.21	0.23
NRMSE of MD	0.17	0.06	0.06	0.15	0.17	0.15
NRMSE of FA	0.36	0.27	0.27	0.30	0.35	0.40
NRMSE of FRT	0.14	0.10	0.10	0.11	0.11	0.12
NRMSE of FRACT	0.87	0.62	0.62	0.77	0.65	0.78

For each metric, the best NRMSE value is indicated in bold.