

Holistic Image Reconstruction for Diffusion MRI

Vladimir Golkov, Jorg M. Portegies, Antonij Golkov, Remco Duits,
and Daniel Cremers

Abstract Diffusion MRI provides unique information on the microarchitecture of biological tissues. One of the major challenges is finding a balance between image resolution, acquisition duration, noise level and image artifacts. Recent methods tackle this challenge by performing super-resolution reconstruction in image space or in diffusion space, regularization of the image data or of postprocessed data (such as the orientation distribution function, ODF) along different dimensions, and/or impose data-consistency in the original acquisition space. Each of these techniques has its own advantages; however, it is rare that even a few of them are combined. Here we present a holistic framework for diffusion MRI reconstruction that allows combining the advantages of all these techniques in a single reconstruction step. In proof-of-concept experiments, we demonstrate super-resolution on HARDI shells and in image space, regularization of the ODF and of the images in spatial and angular dimensions, and data consistency in the original acquisition space. Reconstruction quality is superior to standard reconstruction, demonstrating the feasibility of combining advanced techniques into one step.

V. Golkov (✉) • D. Cremers

Department of Informatics, Technische Universität München, Garching, Germany
e-mail: golkov@cs.tum.edu; cremers@tum.de

J.M. Portegies

Department of Mathematics and Computer Science, Eindhoven University of Technology,
Eindhoven, The Netherlands
e-mail: j.m.portegies@tue.nl

A. Golkov

Department of Mathematics, Augsburg University, Augsburg, Germany
e-mail: antonij.golkov@student.uni-augsburg.de

R. Duits

Department of Mathematics and Computer Science, and Department of Biomedical Engineering,
Eindhoven University of Technology, Eindhoven, The Netherlands
e-mail: r.duits@tue.nl

1 Introduction

Among the main problems in diffusion MRI are scan duration limits (thus a limited amount of data), image resolution limits, noise, and image artifacts. In recent years, a host of methods [1–9] have been developed to tackle these issues. These methods use (simplified) assumptions about the data, such as specific types of smoothness / transform-domain sparsity / low-rankedness, specific types of data similarity between different coordinates in the 3-D space of diffusion directions and weightings (q -space), accurate or simplified image acquisition models, in some cases combined with a tailored acquisition strategy.

Super-resolution in diffusion MRI allows increasing the resolution beyond the hardware limits. In the original super-resolution techniques for diffusion MRI [10, 11], there is no coupling of different q -space coordinates, i.e. each q -space coordinate is treated independently without taking advantage of common structure. It is performed from image space to image space, independently of the image reconstruction step. Recent methods [12–14] couple q -space coordinates and use the original data-acquisition space but regularize only in the reconstruction space—not in additional spaces.

The proposed method allows leveraging complementary information by coupling in q -space, while imposing data consistency in the original space and balancing regularization in several arbitrary representations simultaneously.

The rest of the paper is organized as follows. In Sect. 2.1, we describe the data formation model. In Sect. 2.2, we introduce holistic reconstruction (raw data consistency, several regularization spaces, super-resolution reconstruction in image and diffusion space) and give details on sampling in acquisition and reconstruction spaces, the regularizers, the optimization procedure and its implementation. We show results of holistic super-resolution reconstruction after artificial subsampling of Human Connectome Project data in Sect. 3 and conclude with a discussion in Sect. 4.

2 Methods

2.1 Image Acquisition Model

The image is modeled on a domain $\Omega \times \mathbb{R}^3$, where $\Omega \subset \mathbb{R}^3$ represents the domain in image space, and dimensions four to six of $\Omega \times \mathbb{R}^3$ represent the space consisting of three-dimensional diffusion directions and diffusion weightings (q -space) for which discrete samples are acquired. A complex-valued diffusion MRI image ρ is a mapping

$$\rho : \Omega \times \mathbb{R}^3 \rightarrow \mathbb{C} \text{ given by} \quad (1)$$

$$(y, q) \mapsto \rho(y, q) = r(y, q) \exp(i\varphi(r, q)), \quad (2)$$

where r is the image magnitude and φ is the image phase at spatial coordinate $y \in \Omega$ and q -space coordinate $q \in \mathbb{R}^3$. Magnitude r and phase φ are mappings

$$r: \Omega \times \mathbb{R}^3 \rightarrow \mathbb{R}, \quad (3)$$

$$\varphi: \Omega \times \mathbb{R}^3 \rightarrow S^1. \quad (4)$$

These images are not acquired directly. Acquisition is performed in k -space (more precisely: in the joint six-dimensional (k, q) -space), after Fourier transform $\mathcal{F}_{1,2}$ along the spatial dimensions 1 and 2 of Ω . When sampled at N data points, the resulting data $d \in \mathbb{C}^N$ forms from r and φ according to

$$d = T(r, \varphi) + \varepsilon, \quad (5)$$

where ε is complex-valued i.i.d. Gaussian noise (thermal noise) and T is the encoding operator. The operator T composes r and φ pointwise into a complex-valued image via $C(r, \varphi) = r \odot \exp(i\varphi)$ where “ \odot ” is the pointwise product, followed by a Fourier transform into (k, q) -space and discrete sampling S :

$$T(r, \varphi) = S\mathcal{F}_{1,2}C(r, \varphi), \quad \text{with} \quad (6)$$

$$S: \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{C}^N \quad \text{given by} \quad (7)$$

$$(S\hat{\rho})_n = \int_{[-0.5, 0.5]^3} \hat{\rho}(k_n + v, q_n) \, dv, \quad (8)$$

where the $((k_n, q_n))_{n \in \{1, \dots, N\}}$ are the sampling points in (k, q) -space. Details can be found in [15, 16].

2.2 Holistic Reconstruction

Our goal is to reconstruct the image magnitude r and phase φ from the acquired data d . In order to improve image quality, such a reconstruction should include state-of-the-art image processing methods, such as denoising, super-resolution reconstruction and orientation distribution function¹ (ODF) enhancement. Rather than performing this in a classical manner, where each step is performed separately, we couple all transformations and regularizers into a single optimization problem. This allows performing the entire reconstruction in a single step, while having full control over the balance between all regularizers simultaneously. Furthermore, this avoids data-consistency formulations in intermediate spaces, where the noise

¹The ODF is a formalism that characterizes the strength of diffusion in different directions. It is defined formally below in Eq. (10).

distribution is difficult to model correctly (e.g. Rician signal distribution and other cases)—our least squares data term penalizes deviation from k -space measurements, where noise is Gaussian, while still reconstructing and regularizing in arbitrary spaces. Finally, a holistic formulation allows regularizing in additional spaces other than the acquisition and the reconstruction space. This allows for example using information from the ODF (otherwise calculated independently at a later step) to inform the super-resolution reconstruction in image space.

In our proof-of-concept holistic reconstruction experiments, we treat the entire six-dimensional data jointly (rather than treating each q -space coordinate independently during image space reconstruction, followed by treating each image coordinate y independently during q -space-based processing) and combine the following concepts into a single optimization problem:

- Data consistency in the original (k, q) -space,
- Reconstruction into (y, q) -space with super-resolution in both the spatial and diffusional dimensions,
- Spatial regularization of (y, q) -space data,
- Angular regularization of (y, q) -space data by treating each q -space shell independently as functions on the (uncoupled) space $\mathbb{R}^3 \times S^2$ of positions and orientations,
- Spatial and angular regularization of the ODFs which implicitly correspond to the reconstructed (y, q) -space data by treating them as functions on the (uncoupled) space $\mathbb{R}^3 \times S^2$ of positions and orientations.

The general form of holistic reconstruction into (y, q) -space is

$$\arg \min_{r, \varphi} \frac{1}{2} \|T(r, \varphi) - d\|^2 + R(r), \quad (9)$$

where $R(r)$ is a sum of regularization terms which may or may not transform the image magnitude r into another space, such as ODFs, prior to penalizing non-regularity.²

The “codomain” of our pipeline, i.e. the reconstruction space, can be extended into diffusion models, as in [17, 18]. These model-based methods can be complemented by our regularizers in additional spaces to yield a holistic framework.

Sampling Scheme in (k, q) -Space In order to verify the super-resolution reconstruction capability of our holistic reconstruction, we use data of uniquely high resolution from the Human Connectome Project [19–26], assuming it to be the ground truth underlying image data, and simulate a low-resolution k -space sampling of these ground truth images. In order to leverage complementarity of data in q -space, we employ a low-resolution (k, q) -space sampling scheme [13] in which high resolution components are left out alternatingly in vertical or horizontal

²The precise formula that we use for $R(r)$ will follow later in Eq. (12).

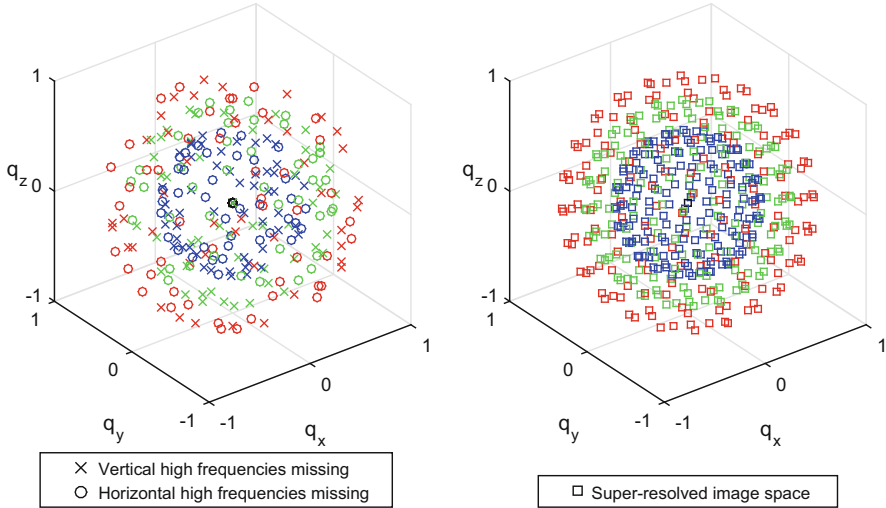


Fig. 1 Sampling scheme in q -space during acquisition (left) and reconstruction (right). The acquired data have alternating artificial subsampling in vertical/horizontal high frequencies in k -space. All high frequencies for all images are reconstructed. Colors encode the b -value: $\mathcal{B} = \{0, 1000, 2000, 3000\}$ s/mm²

image directions for different q -space coordinates. The q -space coordinates and the respective alternating vertical/horizontal k -space subsampling are shown in Fig. 1, left. Both acquisition and reconstruction (see next paragraph) use the set of b -values $\mathcal{B} = \{0, 1000, 2000, 3000\}$ s/mm².

Super-Resolution Sampling Scheme in Reconstruction Space While data are artificially subsampled in k -space for the experiments, the reconstruction space is discretized such that the original high image resolution is reconstructed. While 270 q -space coordinates are sampled (Fig. 1, left), 486 are reconstructed (Fig. 1, right). This scheme achieves a super-resolution reconstruction in image and diffusion space.

Regularization We will regularize several images of the type $U \in \mathbb{H}^2(\mathbb{R}^3 \times \mathbb{S}^2)$, namely the ODF and the spherical shells in q -space (where \mathbb{H}^k denotes the respective Sobolev space).

The ODF [27] for image r at image location $y \in \Omega$ and direction $n \in \mathbb{S}^2$ can be calculated as

$$\text{ODF}(r)(y, n) = \frac{1}{Z_\kappa} \int_0^\infty (\mathcal{F}_{4,5,6}r)(y, pn) p^\kappa dp \quad (10)$$

with the usual choice $\kappa = 2$, where Z_κ is a normalization constant and $\mathcal{F}_{4,5,6}$ is the Fourier transform along the diffusion dimensions four to six that calculates the diffusion propagator from q -space data in an idealized setting [28].

Let G_b be the linear operator that extracts a spherical q -space shell at a given b -value (diffusion weighting) from r :

$$(G_b(r))(y, n) = r(y, \sqrt{bn}). \quad (11)$$

In a proof-of-concept holistic reconstruction, the shells and the ODFs are regularized in the uncoupled space $\mathbb{R}^3 \times S^2$ of positions and orientations as follows:

$$\begin{aligned} R(r) = & \sum_{b \in \mathcal{B}} \int_{\mathbb{R}^3 \times S^2} \alpha_1 \|\nabla_y G_b(r)(y, n)\|^2 \\ & - \alpha_2 \langle G_b(r)(y, n), \Delta_{S^2} G_b(r)(y, n) \rangle + \alpha_3 |\Delta_{S^2} G_b(r)(y, n)|^2 dy d\sigma(n) \\ & + \int_{\mathbb{R}^3 \times S^2} \alpha_4 \|\nabla_y \text{ODF}(r)(y, n)\|^2 \\ & - \alpha_5 \langle \text{ODF}(r)(y, n), \Delta_{S^2} \text{ODF}(r)(y, n) \rangle + \alpha_6 |\Delta_{S^2} \text{ODF}(r)(y, n)|^2 dy d\sigma(n), \end{aligned} \quad (12)$$

where \mathcal{B} is the set of reconstructed b -values, the α_i are regularization parameters, σ is the usual surface measure on S^2 , Δ_{S^2} is the Laplace–Beltrami operator on the sphere and the negative inner products correspond to first-order regularization according to $\int -\langle U, \Delta U \rangle = \int \|\nabla U\|^2$ (i.e. Green’s identity with vanishing boundary conditions as we assume our functions U to vanish at the boundary).

Defining appropriate inner products on the space $\mathbb{H}^2(\mathbb{R}^3 \times S^2) \ni U, V$ and on $\mathbb{H}^1(\mathbb{R}^3 \times S^2, \mathbb{R}^3) \ni \nabla_y U, \nabla_y V$ as

$$\langle U, V \rangle = \int_{\mathbb{R}^3 \times S^2} U(y, n) V(y, n) dy d\sigma(n), \quad (13)$$

$$\langle \nabla_y U, \nabla_y V \rangle = \sum_{i \in \{1, 2, 3\}} \int_{\mathbb{R}^3 \times S^2} (\nabla_y U(y, n))_i (\nabla_y V(y, n))_i dy d\sigma(n), \quad (14)$$

and using the induced norms, we can rewrite the problem (9, 12) as follows:

$$\begin{aligned} & \min_{r, \varphi} \frac{1}{2} \|T(r, \varphi) - d\|^2 \\ & + \sum_{b \in \mathcal{B}} \alpha_1 \|\nabla_y G_b(r)\|^2 - \alpha_2 \langle G_b(r), \Delta_{S^2} G_b(r) \rangle + \alpha_3 \|\Delta_{S^2} G_b(r)\|^2 \\ & + \alpha_4 \|\nabla_y \text{ODF}(r)\|^2 - \alpha_5 \langle \text{ODF}(r), \Delta_{S^2} \text{ODF}(r) \rangle + \alpha_6 \|\Delta_{S^2} \text{ODF}(r)\|^2. \end{aligned} \quad (15)$$

Reformulations To obtain a convenient min-max form with simpler expressions within the norms, we shall use the identity:

$$\|\hat{x}\|^2 = \sup_{\hat{y}} \langle \hat{x}, \hat{y} \rangle - \frac{1}{4} \|\hat{y}\|^2, \quad (16)$$

obtained by taking the convex biconjugate and completing the square. This reformulation introduces dual variables \hat{y} .

Optimization Procedure Our optimization problem (15) can be rewritten as a min-max problem of the form

$$\min_x \max_y G(x) + \langle K(x), y \rangle - F^*(y) \quad (17)$$

with convex G , F^* and a nonlinear K , which can be solved with the modified primal-dual hybrid gradient method for nonlinear K [15, 29, 30]:

$$x^{i+1} := (I + \tau \partial G)^{-1}(x^i - \tau [\nabla K(x^i)]^* y^i), \quad (18a)$$

$$x_\omega^{i+1} := x^{i+1} + \omega(x^{i+1} - x^i), \quad (18b)$$

$$y^{i+1} := (I + \sigma \partial F^*)^{-1}(y^i + \sigma K(x_\omega^{i+1})), \quad (18c)$$

where ∂f represents the subdifferential of a function f , defined as

$$\partial f(x_0) = \{v \mid f(x) - f(x_0) \geq \langle v, x - x_0 \rangle \forall x \in \text{dom} f\}, \quad (19)$$

and $(I + \lambda \partial f)^{-1}$ is the resolvent of the subdifferential, corresponding to the proximal operator [31]:

$$(I + \lambda \partial f)^{-1}x = \text{prox}_{\lambda f}(x) = \arg \min_z f(z) + \frac{1}{2\lambda} \|x - z\|^2. \quad (20)$$

The algorithm (18) has been applied [15] with the operator $T(r, \varphi)$ to non-diffusion MRI, and with another operator to diffusion MRI. The author announces combining $T(r, \varphi)$ with direct reconstruction of the diffusion tensor in a future study, while we present an application of $T(r, \varphi)$ to reconstruction in image \times diffusion space.

By rewriting all five norms in our problem (15) using the identity (16), we obtain the min-max form

$$\begin{aligned} \min_{r, \varphi} \max_{\lambda, (\zeta_b)_{b \in \mathcal{B}}, (\eta_b)_{b \in \mathcal{B}}, \xi, v} & \langle T(r, \varphi), \lambda \rangle - \langle d, \lambda \rangle - \frac{1}{2} \|\lambda\|^2 \\ & + \sum_{b \in \mathcal{B}} \alpha_1 \left(\langle \nabla_y G_b(r), \zeta_b \rangle - \frac{1}{4} \|\zeta_b\|^2 \right) \\ & - \alpha_2 \langle G_b(r), \Delta_{S^2} G_b(r) \rangle + \alpha_3 \left(\langle \Delta_{S^2} G_b(r), \eta_b \rangle - \frac{1}{4} \|\eta_b\|^2 \right) \\ & + \alpha_4 \left(\langle \nabla_y \text{ODF}(r), \xi \rangle - \frac{1}{4} \|\xi\|^2 \right) \\ & - \alpha_5 \langle \text{ODF}(r), \Delta_{S^2} \text{ODF}(r) \rangle + \alpha_6 \left(\langle \Delta_{S^2} \text{ODF}(r), v \rangle - \frac{1}{4} \|v\|^2 \right). \end{aligned} \quad (21)$$

The primal variables are $x = (r, \varphi)$ and the dual ones are $y = (\lambda, (\zeta_b)_{b \in \mathcal{B}}, (\eta_b)_{b \in \mathcal{B}}, \xi, \nu)$, where for example η_b denotes the dual variable associated to $\|\Delta_{S^2} G_b(r)\|^2$. This can be regrouped into the standard form (17) as follows:

$$\begin{aligned}
G(x) &= \sum_{b \in \mathcal{B}} -\alpha_2 \langle G_b(r), \Delta_{S^2} G_b(r) \rangle - \alpha_5 \langle \text{ODF}(r), \Delta_{S^2} \text{ODF}(r) \rangle, \\
\langle K(x), y \rangle &= \langle T(r, \varphi), \lambda \rangle + \sum_{b \in \mathcal{B}} \alpha_1 \langle \nabla_y G_b(r), \zeta_b \rangle + \alpha_3 \langle \Delta_{S^2} G_b(r), \eta_b \rangle \\
&\quad + \alpha_4 \langle \nabla_y \text{ODF}(r), \xi \rangle + \alpha_6 \langle \Delta_{S^2} \text{ODF}(r), \nu \rangle, \\
\pm F^*(y) &= \pm \langle d, \lambda \rangle \pm \frac{1}{2} \|\lambda\|^2 \\
&\quad \pm \frac{1}{4} \left(\sum_{b \in \mathcal{B}} \alpha_1 \|\zeta_b\|^2 + \alpha_3 \|\eta_b\|^2 + \alpha_4 \|\xi\|^2 + \alpha_6 \|\nu\|^2 \right).
\end{aligned} \tag{22}$$

For the implementation of algorithm (18), we calculate the proximal operators [31]:

$$(I + \tau \partial G)^{-1} x = (I + \tau(Q + Q^*))^{-1} x, \tag{23}$$

$$Q = \sum_{b \in \mathcal{B}} G_b^* \Delta_{S^2} G_b + \text{ODF}^* \Delta_{S^2} \text{ODF}, \tag{24}$$

$$(I + \sigma \partial F^*)^{-1} y = \begin{pmatrix} (\lambda - \sigma d)/(\sigma + 1) \\ (\zeta_b/(1 + \alpha_1 \sigma/2))_{b \in \mathcal{B}} \\ (\eta_b/(1 + \alpha_3 \sigma/2))_{b \in \mathcal{B}} \\ \xi/(1 + \alpha_4 \sigma/2) \\ \nu/(1 + \alpha_6 \sigma/2) \end{pmatrix}. \tag{25}$$

Calculating $[\nabla K(x^i)]^*$ (18) for the nonlinear part $T(r, \varphi)$ (22) yields

$$[\nabla T(r, \varphi)]^* = (S\mathcal{F}_{1,2}[\nabla C(r, \varphi)])^* = [\nabla C(r, \varphi)]^* \mathcal{F}_{1,2}^* S^*, \tag{26}$$

$$[\nabla C(r, \varphi)]^* \hat{\lambda} = \begin{pmatrix} \Re(\hat{\lambda}) \cos(\varphi) + \Im(\hat{\lambda}) \sin(\varphi) \\ r(\Im(\hat{\lambda}) \cos(\varphi) - \Re(\hat{\lambda}) \sin(\varphi)) \end{pmatrix}. \tag{27}$$

Unbounded ODF Operator When writing out the Fourier transform $\mathcal{F}_{4,5,6}$ over $Q \in \mathbb{R}^3$, the ODF (10) contains the diverging term $\exp(-i\langle pn, Q \rangle) p^2$. Thus, the ODF operator is unbounded. Since an adjoint is required for the algorithm (18), the operator can be made bounded in the infinite-dimensional setting by including a Gaussian damping factor $\exp(-p^2/\zeta^2)$ as a mollifier. The operator bound of the discrete operator depends on the discretization, and in our discretization scheme no mollifier was needed in practice.

Implementation Details The operators $\mathcal{F}_{1,2}$, S (6), ODF (10), G_b (11), ∇_y and Δ_{S^2} are linear. In the implementation, the spaces in which acquisition, regularization and reconstruction take place are discretized and thus the operators can be written as matrices. We obtain these matrices explicitly. Where not evident, an operator matrix is computed by applying the operator to all standard basis vectors of the discretized space, yielding the columns of the matrix. For pointwise operators, we compute and store repeating coefficients only once. When computing $[\nabla K(\cdot)]^*$ and $K(\cdot)$ in the algorithm (18), having the operator matrices explicitly has the advantages of rapid computation by matrix multiplication and easy computation of the adjoint operators. Besides, in the discretized setting, the ODF operator is not unbounded anymore and thus has an adjoint, as required by the algorithm. The norm $\|[\nabla K(\cdot)]^*\|$ of the operator $[\nabla K(\cdot)]^*$ explodes as the discretization becomes finer, but in our discretization settings there was no need to include a Gaussian mollifier in (10). The practical implementation of the ODF operator is given by generalized q -sampling imaging [32].

3 Results

Figure 2 shows the high-resolution “ground truth” image data from the Human Connectome Project (Fig. 2, left) alongside the results of two reconstruction methods applied to the same data that has been artificially subsampled according to the sampling scheme in (k, q) -space described in section “Sampling Scheme in (k, q) -Space” and illustrated in Fig. 1, left. This artificial subsampling procedure emulates a clinical setting where resolution is considerably lower than in the Human Connectome Project, and enables a comparison to this exceptionally high-resolution ground truth data. The two compared reconstruction methods are standard reconstruction ($\mathcal{F}_{1,2}$ -transformed subsampled data; Fig. 2, middle) and holistic




Fig. 2 High-resolution ground truth (*left*), standard reconstruction (*middle*), holistic super-resolution reconstruction (*right*)

image reconstruction (as described above, with super-resolution sampling as in Fig. 1, right; results in Fig. 2, right).

The employed parameters were $\alpha_1 = 0.3$, $\alpha_2 = 0.1$, $\alpha_3 = 0.1$, $\alpha_4 = 0.01$, $\alpha_5 = 0.3$, $\alpha_6 = 0.01$.

Holistic image reconstruction demonstrates considerably more detail than standard reconstruction. While standard reconstruction results have a visibly lower resolution, holistic reconstruction retrieves details that are present in the ground truth data due to its super-resolution scheme and regularization in image and diffusion space.

4 Discussion

The results of holistic reconstruction demonstrate considerably more detail than the standard reconstruction.

Among the numerous advanced diffusion MRI reconstruction methods existing in literature, many methods perform denoising, missing data reconstruction (q -space compressed sensing), enhancement, etc. as an intermediate post-processing step *after* image-space reconstruction. However, standard-reconstructed images can contain artifacts, intensity bias (e.g. Rician or more complicated), and irretrievably discard some parts of information present in the raw k -space data. Imposing data consistency in reconstructed image space can lead to these errors being propagated on into subsequent data processing steps, and/or introduce less tractable bias-correction terms. There is strong evidence that one-step pipelines are better than multi-step pipelines due to information loss in intermediate steps [33]. Particularly, imposing data consistency on the original raw data in k -space yields improved results compared to multi-step processing [13]. The holistic reconstruction framework presented herein allows imposing data consistency in the original data acquisition space, while also including regularization in several spaces (such as (y, q) -space and “ (y, ODF) -space”), and reconstructing into an arbitrary space, including super-resolution reconstruction sampling.

Super-resolution methods are beneficial for diffusion MRI due to their capability to exceed hardware limitations on resolution. In the presented holistic reconstruction framework, super-resolution is performed in image space *and* diffusion space simultaneously, cf. Fig. 1. At the same time, data consistency in the original space and regularizations in additional spaces are incorporated in a straightforward manner.

Many competing regularizers in different spaces exist in recent literature. Each of them incorporates certain assumptions and improves data quality at certain intermediate regularization strengths. Regularizations in different spaces can be combined into one procedure (including true data consistency and super-resolution) using holistic image reconstruction.

Reconstruction can be performed jointly with motion and distortion correction [5] in the future.

Finally, our choice of priors in (15) was based on isotropic Laplacians over the spatial and angular part, and as such defined on $\mathbb{R}^3 \times S^2$. Including anisotropies and alignment modeling in a crossing-preserving way via the *coupled* space $\mathbb{R}^3 \times S^2 = SE(3)/(\{0 \times SO(2)\})$, see [1] Theorem 2, and [34], is expected to give better results in future work.

Acknowledgements V.G. is supported by the Deutsche Telekom Foundation. The research leading to the results of this article has received funding from the European Research Council under the ECs 7th Framework Programme (FP7/2007-2014)/ERC grant agr. no. 335555. Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

1. Duits, R., Franken, E.: Left-invariant diffusions on the space of positions and orientations and their application to crossing-preserving smoothing of HARDI images. *Int. J. Comput. Vis.* **92**(3), 231–264 (2010)
2. Lustig, M., Donoho, D., Pauly, J.M.: Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.* **58**(6), 1182–1195 (2007)
3. Khare, K., Hardy, C.J., King, K.F., Turski, P.A., Marinelli, L.: Accelerated MR imaging using compressive sensing with no free parameters. *Magn. Reson. Med.* **68**(5), 1450–1457 (2012)
4. Paquette, M., Merlet, S., Gilbert, G., Deriche, R., Descoteaux, M.: Comparison of sampling strategies and sparsifying transforms to improve compressed sensing diffusion spectrum imaging. *Magn. Reson. Med.* **73**, 401–416 (2015)
5. Tao, S., Trzasko, J.D., Shu, Y., Huston, J., Bernstein, M.A.: Integrated image reconstruction and gradient nonlinearity correction. *Magn. Reson. Med.* **74**(4), 1019–1031 (2015)
6. Feng, L., Grimm, R., Block, K.T., Chandarana, H., Kim, S., Xu, J., Axel, L., Sodickson, D.K., Otazo, R.: Golden-angle radial sparse parallel MRI: combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric MRI. *Magn. Reson. Med.* **72**, 707–717 (2014)
7. Cauley, S.F., Xi, Y., Bilgic, B., Xia, J., Adalsteinsson, E., Balakrishnan, V., Wald, L.L., Setsompop, K.: Fast reconstruction for multichannel compressed sensing using a hierarchically semiseparable solver. *Magn. Reson. Med.* **73**, 1034–1040 (2015)
8. Mani, M., Jacob, M., Guidon, A., Magnotta, V., Zhong, J.: Acceleration of high angular and spatial resolution diffusion imaging using compressed sensing with multichannel spiral data. *Magn. Reson. Med.* **73**, 126–138 (2015)
9. Rathi, Y., Michailovich, O., Laun, F., Setsompop, K., Grant, P.E., Westin, C.F.: Multi-shell diffusion signal recovery from sparse measurements. *Med. Image Anal.* **18**(7), 1143–1156 (2014)
10. Scherrer, B., Gholipour, A., Warfield, S.K.: Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisitions. *Med. Image Anal.* **16**(7), 1465–1476 (2012)
11. Poot, D.H.J., Jeurissen, B., Bastiaensen, Y., Veraart, J., Van Hecke, W., Parizel, P.M., Sijbers, J.: Super-resolution for multislice diffusion tensor imaging. *Magn. Reson. Med.* **69**(1), 103–113 (2013)

12. Tobisch, A., Neher, P.F., Rowe, M.C., Maier-Hein, K.H., Zhang, H.: Model-based super-resolution of diffusion MRI. In: Schultz, T., Nedjati-Gilani, G., Venkataraman, A., O'Donnell, L., Panagiotaki, E. (eds.) *Computational Diffusion MRI and Brain Connectivity, MICCAI Workshops 2013. Mathematics and Visualization*, pp. 25–34. Springer International Publishing Switzerland (2014)
13. Golkov, V., Sperl, J.I., Menzel, M.I., Sprenger, T., Tan, E.T., Marinelli, L., Hardy, C.J., Haase, A., Cremers, D.: Joint super-resolution using only one anisotropic low-resolution image per q-space coordinate. In: O'Donnell, L., Nedjati-Gilani, G., Rathi, Y., Reisert, M., Schneider, T. (eds.) *Computational Diffusion MRI, MICCAI Workshop 2014*, pp. 181–191. Springer International Publishing Switzerland (2015)
14. Van Steenkiste, G., Jeurissen, B., Veraart, J., den Dekker, A.J., Parizel, P.M., Poot, D.H.J., Sijbers, J.: Super-resolution reconstruction of diffusion parameters from diffusion-weighted images with different slice orientations. *Magn. Reson. Med.* **75**(1), 181–195 (2016)
15. Valkonen, T.: A primal-dual hybrid gradient method for non-linear operators with applications to MRI. *Inverse Prob.* **30**(5), 055012 (2014)
16. Brown, R.W., Cheng, Y.C.N., Haacke, E.M., Thompson, M.R., Venkatesan, R.: *Magnetic Resonance Imaging: Physical Principles and Sequence Design*, 2nd edn. Wiley, Hoboken (published simultaneously in Canada) (2014)
17. Welsh, C.L., Dibella, E.V.R., Adluru, G., Hsu, E.W.: Model-based reconstruction of under-sampled diffusion tensor k-space data. *Magn. Reson. Med.* **70**(2), 429–440 (2013)
18. Valkonen, T., Bredies, K., Knoll, F.: TGV for diffusion tensors: a comparison of fidelity functions. *J. Inverse Ill-Posed Prob.* **21**(3), 355–377 (2013)
19. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K.: The WU-Minn Human Connectome Project: an overview. *NeuroImage* **80**, 62–79 (2013)
20. Feinberg, D.A., Moeller, S., Smith, S.M., Auerbach, E., Ramanna, S., Glasser, M.F., Miller, K.L., Ugurbil, K., Yacoub, E.: Multiplexed echo planar imaging for sub-second whole brain fMRI and fast diffusion imaging. *PLoS ONE* **5**(12), e15710 (2010)
21. Setsompop, K., Cohen-Adad, J., Gagoski, B.A., Raj, T., Yendiki, A., Keil, B., Wedeen, V.J., Wald, L.L.: Improving diffusion MRI using simultaneous multi-slice echo planar imaging. *NeuroImage* **63**(1), 569–580 (2012)
22. Xu, J., Li, K., Smith, R.A., Waterton, J.C., Zhao, P., Chen, H., Does, M.D., Manning, H.C., Gore, J.C.: Characterizing tumor response to chemotherapy at various length scales using temporal diffusion spectroscopy. *PLoS ONE* **7**(7), e41714 (2012)
23. Sotiropoulos, S.N., Jbabdi, S., Xu, J., Andersson, J.L., Moeller, S., Auerbach, E.J., Glasser, M.F., Hernandez, M., Sapiro, G., Jenkinson, M., Feinberg, D.A., Yacoub, E., Lenglet, C., Van Essen, D.C., Ugurbil, K., Behrens, T.E.J.: Advances in diffusion MRI acquisition and processing in the Human Connectome Project. *NeuroImage* **80**, 125–143 (2013)
24. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M.: The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* **80**, 105–124 (2013)
25. Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M.: FSL. *NeuroImage* **62**(2), 782–790 (2012)
26. Fischl, B.: FreeSurfer. *NeuroImage* **62**(2), 774–781 (2012)
27. Lin, C.P., Wedeen, V.J., Chen, J.H., Yao, C., Tseng, W.Y.I.: Validation of diffusion spectrum magnetic resonance imaging with manganese-enhanced rat optic tracts and ex vivo phantoms. *NeuroImage* **19**, 482–495 (2003)
28. Stejskal, E.O.: Use of spin echoes in a pulsed magnetic-field gradient to study anisotropic, restricted diffusion and flow. *J. Chem. Phys.* **43**(10), 3597–3603 (1965)
29. Pock, T., Cremers, D., Bischof, H., Chambolle, A.: An algorithm for minimizing the Mumford-Shah functional. In: 2009 IEEE 12th International Conference on Computer Vision (ICCV). Number 813396, IEEE, pp. 1133–1140 (2009)

30. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imaging Vis.* **40**(1), 120–145 (2011)
31. Parikh, N., Boyd, S.: Proximal Algorithms. *Found. Trends Optim.* **1**, 123–231 (2014)
32. Yeh, F.-C., Wedeen, V.J., Tseng, W.-Y.I.: Generalized q -sampling imaging. *IEEE Trans Med Imaging* **29**(9), 1626–1635 (2010)
33. Golkov, V., Dosovitskiy, A., Sämann, P., Sperl, J.I., Sprenger, T., Czisch, M., Menzel, M.I., Gómez, P.A., Haase, A., Brox, T., Cremers, D.: q -Space deep learning for twelve-fold shorter and model-free diffusion MRI scans. In: MICCAI (2015)
34. Portegies, J.M., Fick, R.H.J., Sanguinetti, G.R., Meesters, S.P.L., Girard, G., Duits, R.: Improving fiber alignment in HARDI by combining contextual PDE flow with constrained spherical deconvolution. *PLoS ONE*. See <http://bmia.bmt.tue.nl/people/RDuits/mainJorg.pdf> (2015, submitted). Available on arXiv 2015