Article

Fast Fiber Orientation Estimation in Diffusion MRI from kq-Space Sampling and Anatomical Priors

Marica Pesce [1,*]**, Audrey Repetti** [1,2,*]**, Anna Auría** [3]**, Alessandro Daducci** [4]**, Jean-Philippe Thiran** [3,5]
and Yves Wiaux [1,*]

1. Institute of Sensors Signals and Systems (ISSS), Heriot-Watt University, Edinburgh EH14 4AS, UK
2. Department of Actuarial Mathematics & Statistics (AMS), Heriot-Watt University, Edinburgh EH14 4AS, UK
3. Signal Processing Lab (LTS5), École Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland; anna.auria@gmail.com (A.A.); jean-philippe.thiran@epfl.ch (J.-P.T.)
4. Department of Computer Science, University of Verona, 37134 Verona, Italy; alessandro.daducci@univr.it
5. Department of Radiology, Lausanne University Hospital (CHUV) and University of Lausanne, CH-1011 Lausanne, Switzerland
* Correspondence: pesce.marica@gmail.com (M.P.); a.repetti@hw.ac.uk (A.R.); y.wiaux@hw.ac.uk (Y.W.)

Abstract: High spatio-angular resolution diffusion MRI (dMRI) has been shown to provide accurate identification of complex neuronal fiber configurations, albeit, at the cost of long acquisition times. We propose a method to recover intra-voxel fiber configurations at high spatio-angular resolution relying on a 3D kq-space under-sampling scheme to enable accelerated acquisitions. The inverse problem for the reconstruction of the fiber orientation distribution (FOD) is regularized by a *structured sparsity* prior promoting simultaneously voxel-wise sparsity and spatial smoothness of fiber orientation. Prior knowledge of the spatial distribution of white matter, gray matter, and cerebrospinal fluid is also leveraged. A minimization problem is formulated and solved via a stochastic forward–backward algorithm. Simulations and real data analysis suggest that accurate FOD mapping can be achieved from severe kq-space under-sampling regimes potentially enabling high spatio-angular resolution dMRI in the clinical setting.

1. Introduction

Diffusion Magnetic Resonance Imaging (dMRI) is a unique non-invasive technique to infer the microscopic architecture of tissues in vivo. In recent years, dMRI has gained a lot of attention in neuroscience since it enables the mapping of the white matter fiber paths, revealing the existing connections between different brain areas [1–3]. In clinics, dMRI has shown to provide insights into many neurodegenerative diseases, such as Schizophrenia and Alzheimer's disease [4,5]. Diffusion MRI enables to map the restricted diffusion of the water molecules comprising the white matter tissue. The information captured by dMRI is then processed in order to infer the connectivity and the integrity of the white matter pathways.

A typical approach to trace the complex pathways of the white matter fiber bundles from dMRI signals relies on piecing together local fiber orientation information. Such local information is obtained by processing a multitude of MRI signals generated by applying various diffusion gradients during the acquisition. Each diffusion-weighted (DW) signal is sensitive to diffusion along a specific direction and at a specific intensity, identified by a so-called q-space point defined by the diffusion gradient applied [6].

Diffusion Tensor Imaging (DTI) [7] is the most widely-used technique in clinics for the intra-voxel fiber orientation estimation since it requires less than 7 diffusion volumes, i.e., 7 q-points. DTI relies on a Gaussian description of the water molecules diffusion phenomenon, which enables the presence of only one preferred diffusion direction. For

this reason, the presence of multiple fiber populations cannot be distinguished within the same voxel by using DTI, making this technique inappropriate for complex fiber architecture estimation.

High Angular Resolution Diffusion Imaging (HARDI) methods [8] have shown to provide more accurate fiber representations for complex fiber arrangements, by probing the observed brain with a large number of diffusion gradients. For instance, Diffusion Spectrum Imaging [9] considers 515 diffusion gradients allocated on a Cartesian grid, while q-ball approaches [10] use between 30 and 100 diffusion gradients distributed over a single shell. However, acquisitions characterized by large numbers of diffusion gradients result in long scan times, which are unfeasible in clinical applications. This limitation becomes even more obvious when the brain microstructure is investigated, encompassing not only fiber orientation, but also other microstructure parameters (e.g., diameter, etc.).

To this purpose, more elaborate acquisition protocols are used, requiring more diffusion gradients, typically distributed over multiple shells in q-space [11,12]. Consequently, numerous HARDI approaches have been proposed in order to optimize the number and the spatial distribution of the q-space points to be considered for an accurate description on the white matter fiber arrangement [13,14]. More specifically, depending on the microstructure parameters of interest, different approaches have been developed to reduce the dMRI acquisition time while preserving the reconstruction quality.

We focus here on the problem of estimating white matter fiber orientations. A generalization to the study of additional microstructure parameters is beyond the scope of the present work. In this context, some of these methodologies are based on Spherical Deconvolution (SD) approaches, modeling the HARDI signal as a spherical convolution between the Fiber Orientation Distribution (FOD), representing the few active fiber orientations, and a kernel, representing the response signal of a single fiber [15–17]. Recently, SD has been utilized to recover the FOD from a reduced number of diffusion gradients (q-space under-sampling) [18–20].

The FOD recovery problem can be formulated as an inverse problem, which is ill-posed and highly sensitive to noise. In this context, an efficient approach consists in leveraging the convex optimization theory by defining the FODs as a solution to a convex regularized minimization problem, incorporating *a priori* information on the target FODs. In particular, the coefficients of interest are imposed to be non-negative and real valued, since FODs are functions on the unit sphere expressing the orientation and the volume fraction of the fiber populations contained in a single voxel [21].

When expressed as set of discrete orientations, FODs are characterized by coefficients summing up to one in each voxel. Finally, leveraging the fact that, at the imaging resolution available currently, each single voxel is assumed to be populated by only a few fiber bundles, sparsity priors can be incorporated in the model (see, e.g., [18–20,22–26]). In this context and inspired by the Compressed Sensing (CS) theory [27,28], many methods have obtained promising results recovering the FOD coefficients from a reduced number of diffusion gradients by leveraging sparsity-based priors.

Constraint Spherical Deconvolution (CSD) [21] represents the first attempt to regularize the FOD recovery problem. Based on the fact that the FOD coefficients correspond to the fiber volume fractions, they are assumed to be non-negative. CSD proposes to iteratively solve a sequence of weighted ℓ_2 norm problems whose weights, which depend on the FOD coefficients estimated at the previous iteration, are used to drive to zero negative and small coefficients. Inspired by the CS theory, refs. [18,23] suggested the use of the ℓ_1 norm for the recovery of the volume fractions coefficients. However, ref. [22] showed that the use of the ℓ_1 norm, meant as simple sum of the coefficients value, is inconsistent with the fact that the volume fractions coefficients sum up to 1 by definition.

Additionally, ref. [22] reformulated the FOD recovery in order to approach the ℓ_0 minimization. This is done by solving a sequence of weighted ℓ_1 norm problems whose weights at each iteration correspond to the inverse of the values of the solution of the previous problem [29]. At convergence, these weights lead the values of the coefficients to

be independent from the magnitude of the non-zero values, thus, mimicking the results of the ℓ_0 pseudo-norm.

Ref. [19] proposed the use of spatial regularization to solve the SD ill-posed problem. In this work a piece-wise smoothness of the FOD coefficients is promoted while encouraging the coefficients to provide high contrast. Ref. [19] observed that the contribution of each FOD coefficient is highly correlated to the coefficients associated with its neighbor directions. They proposed to regularize the FOD recovery by penalizing the presence of coefficients that exhibit large variations along similar orientations. In addition, they provided a way to discard the noisy contributions by driving to zero the coefficients that are not sufficiently distant from the mean value of the coefficients of interest.

In addition to the angular resolution, the spatial resolution of the acquired DW volumes is also important to achieve an accurate identification of the white matter paths. In principle, by considering small voxel sizes, the complexity of the inner fiber arrangements can be reduced. Thus, considering both high angular and spatial resolutions would be ideal to ensure an accurate fiber recovery [30–32].

Currently, single-shot (SS) DW echo-planar imaging (EPI) is the most popular technique used to perform the dMRI acquisition in clinics. SS-DW-EPI enables extremely fast scan times, which make DW signals nearly immune to patient motion and mostly sensitive to the water molecule movements. The SS-DW-EPI scheme is typically applied in 2D, where each slice is sequentially excited, diffusion-encoded and then collected using a unique trajectory. Despite being very fast, 2D SS-DW-EPI is very inefficient due to the procedure required for the generation of intense diffusion gradients responsible for sensitizing the MRI signal produced by the examined tissues with respect to the diffusion phenomenon occurring along the gradient direction. Indeed, more than 50% of the acquisition time is dedicated to the preparation of the diffusion gradients to be applied, which, in the 2D setting, is required for each single slice.

A number of methods have been proposed in order to acquire DW signals more efficiently and devote the most of the acquisition time to achieve higher resolution and enhance the image quality [33–35]. Parallel imaging and 3D diffusion MRI acquisitions are some of them [36]. In particular, moving from the 2D to the 3D approach enables reducing the number of times the diffusion preparation process is applied to cover the same brain volume.

Ideally, a single signal preparation would be enough to encode an entire 3D DW volume. However, this is not the case in practice due to the fast decay DW signals are subject to. Consequently, 3D DW EPI acquisitions typically perform the collection of multiple sub-volumes, which still requires the signal preparation to be applied fewer times than in the 2D setting. In this context, 3D acquisitions can be used in dMRI to reduce the acquisition time required to scan each DW volume or, alternatively, achieve higher signals quality within the same scan time.

In order to accelerate the dMRI acquisition process, previously proposed q-space methods have mainly focused on reducing the number of diffusion gradients from fully sampled DW images. However, the recovery of white matter structures can further benefit from the incomplete acquisition of each k-space volume.

Traditional approaches aiming at exploring the white matter microstructure from under-sampled kq-space data relies on two separate steps: first, DW images are recovered from incomplete k-space data, followed by the recovery of the fiber architecture from a limited number of q-space points.

Recently, methods recovering the diffusion profile directly from under-sampled kq-space data have started to gain significant attention [26,37–39]. The work of [37] developed the first approach to simultaneously recover the diffusion signal and the diffusion propagator from data sub-sampled in both the k and q 3D spaces. The regularization of the problem is performed both in the angular and in the spatial domain. More specifically, the proposed recovery problem promotes the sparsity of the DW images in the wavelet

domain, the sparsity of the diffusion signal in the angular domain and the smoothness of the DW images.

In the work of [37], the dictionary used to sparsify the signal in the angular domain was generated by a Dictionary Learning technique that enabled the creation of an adaptive dictionary to fully characterize the signal in each single voxel. The surfacelet transform has been shown to efficiently represent the directional information of the diffusion propagator by using only few coefficients. Inspired by these observations, ref. [40] proposed to recover the full diffusion propagator from measurements that are under-sampled in kq-space by leveraging two sparsity priors. The first prior promotes sparsity of the diffusion propagator in the surfacelet domain through the ℓ_1 norm. The second prior promotes sparsity of the diffusion propagator in the gradient domain by means of the TV penalty.

In the work of [41], they used rotation-invariant dictionaries in which only few atoms are active by adapting their orientation to the diffusion MRI data. The recovery framework proposed by [41] consists of alternating between the estimation of the coefficients identifying the dictionary atoms, the rotational transformation matrix, the phase contamination and the DW images. The sparsity of the vectors containing the coefficients individuating the dictionary atoms is leveraged to regularize the reconstruction. In addition, the corresponding DW images are recovered by promoting their sparsity in the wavelet domain.

As the method presented in this paper, ref. [26] focuses on the recovery of the FOD coefficients. In this work, the sparsity of the FOD coefficients is promoted through the ℓ_1 norm. Furthermore, a total variation (TV) penalty, acting on the DW images, indirectly promotes spatial fiber regularity within neighbor voxels.

In the present work, we develop a method to recover intra-voxel fiber configurations (FODs) at high spatio-angular resolution relying on a 3D kq-space under-sampling scheme. For each of a reduced number of q-space points, the time available for k-space sampling is used to acquire a sub-sampled k-space of a 3D sub-volume, rather than the full k-space of a 2D slice, thus, providing a potentially significant acceleration over the state-of-the-art.

Anatomical constraints are leveraged in order to recover the FOD coefficients. First, prior knowledge of the spatial distribution of the brain tissues, which can be inferred from the image acquired in the absence of diffusion, is explicitly incorporated in the FOD recovery problem. Secondly, the FOD estimate is defined as a solution to a regularized minimization problem, using the *structured sparsity* prior proposed in [24]. This regularization promotes simultaneously voxel-wise sparsity and spatial smoothness of fiber orientation.

Following [22,24], the resulting non-convex minimization problem is solved via a re-weighting scheme [29] involving a sequence of convex minimization problems with weighted sparsity prior. Specific to our contribution, a stochastic Forward–Backward (FB) algorithm is used to solve each of these convex problems, with convergence guarantees toward the minimum of the corresponding convex objective [42]. One of the advantages of the stochastic approach is that it offers the possibility to handle efficiently multi-coil acquisitions involving large volumes of data, minimizing both the memory requirement per iteration and the reconstruction time.

Results from both simulated and real data experiments showed that the proposed approach outperformed, in terms of its potential of accelerating the acquisition while preserving the FOD reconstruction quality, both the existing method in kq-space [26] and the traditional reconstruction approaches considering the TV prior for the reconstruction of the DW signals before the recovery of the FOD coefficients. Importantly, our experiments reveal that the optimal reduction of samples was achieved, not by simply reducing the number of q-space points while fully sampling the k-space, but rather by a combined kq-space under-sampling.

The remainder of the paper is organized as follows. The proposed approach is described in Section 2. In particular, the FOD reconstruction inverse problem is described in Section 2.2, and the proposed minimization problem and algorithm are given in Section 2.3. In Section 3, we provide the description of the experimental setup, and, in Section 4, we

present the results obtained from both synthetic and real data. Finally, we discuss the achieved reconstruction performances, and we conclude in Section 5.

2. Materials and Methods

2.1. Background

2.1.1. k and q Spaces Overview

The dMRI signal is captured by a multitude of MRI images, each of which is sensitive to water molecule diffusion occurring at a specific angle. As for the standard MRI, the dMRI signal is spatially encoded by the application of three spatial encoding gradients spanning the 3D space. The k-space corresponds with good approximation to the 2D or 3D Fourier transform of the signal. In addition to the application of the spatial encoding gradients, dMRI signals must be encoded for the diffusion process.

To this aim, additional gradients are applied during the acquisition, that are the diffusion gradient. Similarly to the k-space, encoded by the spatial gradients, the q-space is the 3D space defined by the diffusion gradients applied during the dMRI acquisition process. Each dMRI volume is associated with diffusion along a specific direction $\hat{\mathbf{q}}$ and, at a specific intensity $|\mathbf{q}|$, characterizing a point $\mathbf{q} = |\mathbf{q}|\hat{\mathbf{q}}$. The commonly used b-value parameter [10] associated with each DW volume is defined as $b = 4\pi^2|\mathbf{q}|^4 t$, with t as the diffusion time.

2.1.2. FOD Recovery via SD Framework

In this section, we describe the SD framework for FOD recovery, from q-space under-sampled measurements. We refer to [24] for further details of the global FOD recovery problem.

Let $X \in \mathbb{R}_+^{(n+2)\times N}$ be the unknown matrix containing the FOD coefficients of interest, where N is the number of imaged voxels and n is the number of dictionary atoms. Each column of X contains the $n + 2$ coefficients of the corresponding voxel. The objective is to find an estimate of the FOD field associated with all the voxels of the brain, from the degraded measurements $Y \in \mathbb{R}^{M\times N}$, given by $Y = \Phi X + \eta$, where $\Phi \in \mathbb{R}^{M\times(n+2)}$ is the observation matrix and $\eta \in \mathbb{R}^{M\times N}$ models the acquisition noise. Note that, for SNR > 2 the acquisition noise consists in a realization of an additive random i.i.d. Gaussian noise [43], where the Signal-to-Noise Ratio on the image acquired in absence of diffusion gradients (i.e., $\mathbf{s}_0$ image) is defined as SNR $= \frac{\langle \mathbf{s}_0 \rangle}{\sigma}$, where $\langle \cdot \rangle$ denotes the mean of its argument and σ is the standard deviation of the noise.

More precisely, Φ is a known dictionary whose columns (called dictionary atoms) correspond to the response signals on M q-points to n single fiber orientation [19]. In addition, two atoms representing isotropic diffusion (typically in gray matter or CSF) are considered in the dictionary. The corresponding atoms are invariant under rotation in q-space. The procedure adopted to generate Φ is provided in Section 3.3, based on the method developed in [22]. In particular, the first row of Φ is dedicated to atom values in the absence of diffusion gradient, equal to 1: $(\Phi_{1,d})_{1\leqslant d\leqslant(n+2)} = 1$.

The rows of the data matrix $Y \in \mathbb{R}^{M\times N}$ span the unfolded diffusion volumes, acquired with M gradients and normalized by the intensities of the volumes acquired in the absence of diffusion, denoted by $\mathbf{s}_0$. The $\mathbf{s}_0$ volume is assumed to be obtained from a separate acquisition and is used here as prior information. The first row of Y is devoted to the normalized $\mathbf{s}_0$ volume, i.e., all its coefficients are equal to 1. Thus, for each voxel $v \in \{1, \ldots, N\}$, $Y_{1,v} = \sum_{i=1}^{n+2} \Phi_{1,i}X_{i,v} = \sum_{i=1}^{n+2} X_{i,v} = 1$. In summary, introducing unitary lines in both Y and Φ intrinsically forces the sum of the FOD coefficients to be equal to 1, injecting directly this prior information into the inverse problem.

2.2. Inverse Problem Formulation

2.2.1. Proposed Measurement Model

In this section, we describe the proposed under-sampled kq-space measurement model. We consider acquisitions obtained from C coil receivers, with M diffusion gradients, sampling K k-space points. The measurement matrix $\hat{Y} \in \mathbb{C}^{MC \times K}$ is expressed as follows:

$$\hat{Y} = [\hat{Y}_{q,c}]_{1 \leqslant q \leqslant M, 1 \leqslant c \leqslant C}, \quad \text{with} \quad \hat{Y}_{q,c} = \mathcal{A}_{q,c}(X) + \eta_{q,c}, \tag{1}$$

where $\hat{Y}_{q,c} \in \mathbb{C}^{1 \times K}$ is the under-sampled k-space of the diffusion volume acquired with gradient indexed by q and by the receiver coil indexed by c, and $\eta_{q,c} \in \mathbb{C}^{1 \times K}$ is a realization of an independent and identically distributed (i.i.d) Gaussian noise. Indeed, while the noise contamination on the magnitude signals is characterized by a Rician distribution, the original noise in k-space is Gaussian [44]. The matrix $X \in \mathbb{R}^{(n+2) \times N}$ represents the unknown FOD field of interest, and the linear operator $\mathcal{A}_{q,c}$ is given by:

$$\mathcal{A}_{q,c}(X) = \Phi_q X S_0 U^{(c)} H^{(q,c)} F M^{(q)}. \tag{2}$$

The various terms defining $\mathcal{A}_{q,c}$ are described as follows. $\Phi_q \in \mathbb{R}^{1 \times (n+2)}$ is the qth row of the dictionary Φ that spans the response of a single fiber oriented along n different directions, to which two isotropic compartments are added, to represent the gray matter and the CSF. In order to implicitly force the FOD coefficients of each voxel to sum up to one, the Fourier coefficients of the $\mathbf{s}_0$ volume are intentionally introduced in the first row of $\hat{Y}$. Here again, the $\mathbf{s}_0$ volume is assumed to be obtained from a separate acquisition and is used as prior information.

Measurements cannot be normalized as the diffusion volumes themselves are not accessible but only an incomplete k-space counterpart. Thus, a matrix $S_0 \in \mathbb{R}^{N \times N}$, whose diagonal elements correspond to the $\mathbf{s}_0$ volume coefficients, is explicitly introduced in the full measurement operator. The acquisition of the diffusion signal from multiple channels is taken into account through the diagonal matrix $U^{(c)} \in \mathbb{C}^{N \times N}$, which contains the sensitivity map of the corresponding receiver coil c. Moreover, motion and magnetic field inhomogeneities generate phase distortions that are taken into account in the diagonal matrix $H^{(q,c)} \in \mathbb{C}^{N \times N}$.

We refer to Sections 3.5 and 3.4 for the description of the procedures used to estimate the coil sensitivity and the phase contamination, respectively. Finally, $F \in \mathbb{C}^{N \times N}$ represents the 3D Fourier matrix and $M^{(q)} \in \mathbb{R}^{N \times K_q}$ is a realization of a binary mask that under-samples each slice of the acquired volume. It is important to notice that different realizations of $M^{(q)}$ may be considered for each applied diffusion gradient. In addition, $M^{(1)}$ is chosen to be the identity matrix of $\mathbb{R}^N$, to fully acquire the Fourier coefficients of the $\mathbf{s}_0$ signal (stored in $Y_{1,c}$) for normalization, segmentation and calibration purposes.

2.2.2. Tissue Segmentation Constraints

The spatial distribution maps of the different brain tissues can be obtained from the segmentation of a structural image of the imaged brain. Such images are generally acquired by default in clinical practice; alternatively, the spatial distribution maps can be estimated from the $\mathbf{s}_0$ signal, which is always fully acquired for normalization and calibration purposes. An illustration of a tissue distribution is given in the first row of Figure 1. In our approach, we propose to explicitly take advantage of the prior knowledge of tissue distributions over the space in order to further regularize the ill-posed FOD estimation problem. We assume that the number of voxels containing white matter, gray matter and CSF tissues, are known and indicated by N_1, N_2 and N_3, respectively.

In order to fully exploit the tissue distribution information, we propose a novel formulation for the FOD global problem in the kq-space. To this aim, we define three different variables $S_1 \in \mathbb{R}_+^{n \times N_1}$, $S_2 \in \mathbb{R}_+^{1 \times N_2}$ and $S_3 \in \mathbb{R}_+^{1 \times N_3}$, related to the white matter tissue, the gray matter tissue and the CSF, respectively, as shown in Figure 1.

More specifically, S_1 contains the effective FOD coefficients associated with the white matter fiber, S_2 models the isotropic behavior characterizing the gray matter tissue, and S_3 takes into account the isotropic behavior of the CSF. The object X is fully characterized by $S = (S_1, S_2, S_3)$ through the linear mapping $X = \mathcal{Z}(S)$, where the operator $\mathcal{Z}$ concatenates the matrices resulting from the expansions of S_1, S_2 and S_3 with zero-valued columns in the places of the voxels that are known to not contain the corresponding tissue.

A schematic representation of $X = \mathcal{Z}(S)$ is reported in Figure 1. In accordance with the white matter distribution maps (blue map on the top of Figure 1), the columns of S_1 (blue columns on the bottom) span the FOD coefficients of the voxels that are characterized by nonzero values in the map. In an analogous way, the coefficients of S_2 and S_3 (green and yellow columns on the bottom of Figure 1, respectively) represent the isotropic compartments associated with the voxels corresponding to the nonzero values in the gray matter and CSF maps (green and yellow map on the top of Figure 1, respectively). Using the proposed notation, problem (1) can be rewritten as:

$$\hat{Y}_{q,c} = \mathcal{A}_{q,c}(\mathcal{Z}(S)) + \eta_{q,c}. \tag{3}$$

It is worth noticing that, in the particular case when the global FOD estimation formulation is applied only to white matter voxels, we have $X = S = S_1$, and $\mathcal{Z}$ is the identity operator.
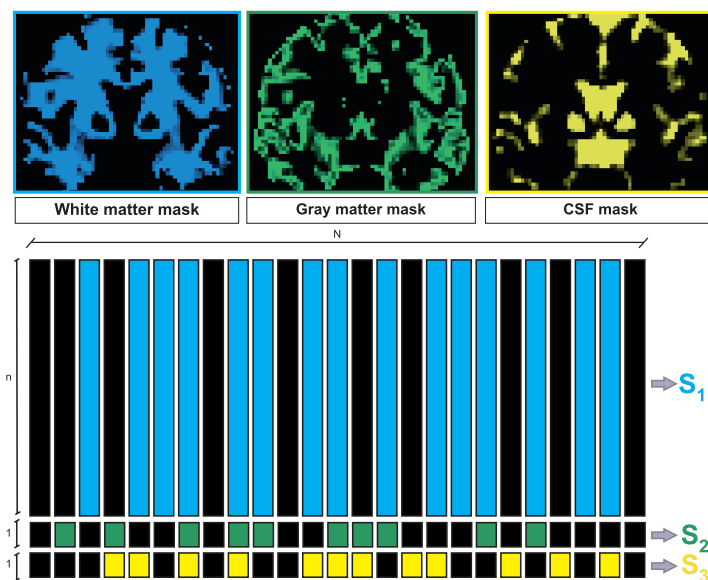

Figure 1. Schematic representation of the matrices S_1, S_2 and S_3 involved in the FOD reconstruction. Each column of S_1 spans the FOD coefficients of the voxels containing white matter tissue (blue columns). Coefficients modeling gray matter and CSF are stored in S_2 (green columns) and in S_3 (yellow columns), respectively. The FOD reconstruction is performed by estimating S_1, S_2 and S_3, whose active columns are provided by the tissue segmentation maps (top).

2.3. Minimization Problem and Algorithm

2.3.1. Constrained Weighted-ℓ_1 Minimization Problem

Recently, convex and non-convex optimization methods have gained a great deal of attention to solve ill-posed or ill-conditioned inverse problems. In this context, the estimate of the sought object is defined to be a minimizer of a sum of two functions: the data-fidelity term associated with the signal model, and the regularization term incorporating *a priori* information of the target solution. On the one hand, when having an additive i.i.d. Gaussian noise, the data-fidelity of choice is the least squares criterion. On the other hand, sparsity-aware models, especially leveraging a CS framework [45], are particularly suitable for solving ill-posed or ill-conditioned inverse problems.

The most suitable way to promote sparsity consists in using the ℓ_0 pseudo-norm, which counts for the number of non-zero coefficients of the signal of interest [46]. However, the ℓ_0 pseudo-norm, being neither smooth nor convex, can be difficult to handle efficiently and it is often replaced by its convex relaxation, namely the ℓ_1 norm [45]. In the context of FOD reconstruction, the signal X is sparse in the q-space as a consequence of the low number of fiber populations (i.e., non-zero FOD coefficients) that are expected to be contained within a voxel.

However, when SD problems are considered, it was demonstrated in [22] that the use of the ℓ_1 norm is inconsistent with the physical constraint that the volume fractions of each voxel sum up to one. To overcome this difficulty, [22] proposed to use a reweighting ℓ_1 approach [29] to approximate the ℓ_0 pseudo-norm in the context of the voxel-wise FOD estimation. Such an approach was subsequently generalized in [24] to solve the inverse problem described in Section 2.1.2. This consists of solving sequentially several constrained weighted ℓ_1 problems. The weights are chosen to simultaneously promote voxel-wise sparsity and spatial smoothness of fiber orientation. Formally, the problem is written as follows:

$$\min_{X} \|\Phi X - Y\|_2^2 \text{ s.t. } X \in \mathcal{B}_{1,W}^+(\kappa), \tag{4}$$

where $\mathcal{B}_{1,W}^+(\kappa)$ denotes the intersection of the real positive orthant $\mathbb{R}_+^{(n+2)\times N}$ with the weighted ℓ_1 ball of radius $\kappa \geqslant 0$ centered in 0 with weighting matrix $W = (W_{d,v})_{d,v} \in [0,+\infty[^{(n+2)\times N}$. Precisely, this is defined as

$$\mathcal{B}_{1,W}^+(\kappa) = \left\{ X \in \mathbb{R}_+^{(n+2)\times N} \mid \|X\|_{1,W} \leqslant \kappa \right\}, \tag{5}$$

$\|\cdot\|_{1,W}$ is the weighted ℓ_1 norm given by

$$\|X\|_{1,W} = \sum_{v=1}^{N} \sum_{d=1}^{n+2} W_{d,v}|X_{d,v}|. \tag{6}$$

In this work, we generalize this weighting scheme to solve the kq-space inverse problem (3). In particular, we aim to solve

$$\min_{S=(S_1,S_2,S_3)} \left\| \mathcal{A}(\mathcal{Z}(S)) - \hat{Y} \right\|_2^2 \text{ s.t. } \begin{cases} S_1 \in \mathcal{B}_{1,W}^+(\kappa), \\ S_2 \in \mathbb{R}_+^{1\times N_2}, \\ S_3 \in \mathbb{R}_+^{1\times N_3}, \end{cases} \tag{7}$$

where $\mathcal{A}\colon \mathbb{R}^{(n+2)\times N} \to \mathbb{C}^{MC\times k}$ consists of the concatenation of the operators $\mathcal{A}_{q,c}$, for all (q,c).

2.3.2. Algorithm for Constrained Weighted ℓ_1 Minimization

As proposed by the reweighting framework, problem (7) is solved several times, considering different weighting matrices W, in order to mimic the ℓ_0 pseudo-norm. Note that the reweighting framework is underpinned by recent theoretical results [47–50] showing that sequentially minimizing convex problems with weighted-ℓ_1 priors corresponds to solving for a critical point of a non-convex problem with a log-sum prior, itself a close approximation to the target ℓ_0 prior. Considering one weighting cycle, a simple algorithm to solve problem (7) is the FB algorithm [51].

At each iteration $j \in \mathbb{N}$, it updates the variables $S^{(j+1)} = \left(S_1^{(j+1)}, S_2^{(j+1)}, S_3^{(j+1)}\right)$ by performing a gradient step on the data-fidelity term, and three projection steps to handle the three constraints described in (7). However, computing the gradient step can be computationally unaffordable when processing high dimensional data, which is particularly the case with multi-coil kq-space raw data.

To overcome this issue, we rely on a more evolved algorithm, namely the stochastic FB algorithm [42], where only an approximation of the gradient of interest is computed at

each iteration. Specifically, we propose to select randomly at each iteration a subset of the available coils, and approximate the gradient using only the data from the selected coils. The proposed method is described in Algorithm 1.

Steps 6–10 describe the computation of the approximated gradient. Below is an explanation of these steps. Let $\hat{Y}_c$ be the data measured by coil c and $\mathcal{A}_c$ be its associated measurement operator. At iteration $j \in \mathbb{N}$, the gradient of the data fidelity term $\|\mathcal{A}(\mathcal{Z}(S^{(j)})) - \hat{Y}\|_2^2$, denoted by $\overline{G}^{(j+1)}$, can be written as the sum $\sum_{c=1}^{C} G_c^{(j+1)}$, where $G_c^{(j+1)}$ corresponds to the gradient of $\|\mathcal{A}_c(\mathcal{Z}(S^{(j)})) - \hat{Y}_c\|_2^2$. Using the stochastic approach described in [42], the gradient $\overline{G}^{(j)}$ can be approximated by only updating a subset $\mathcal{S}_c^{(j)} \subset \{1, \dots, C\}$ of the gradients $G_c^{(j)}$. More specifically, we define the approximation $\widetilde{G}^{(j)}$ of $\overline{G}^{(j)}$ as follows: $\widetilde{G}^{(j)} = \sum_{c \in \mathcal{S}_c^{(j)}} G_c^{(j)} + \sum_{c \notin \mathcal{S}_c^{(j)}} G_c^{(j)}$, where $G_c^{(j)} = \mathcal{Z}^{\dagger}\left(\mathcal{A}_c^{\dagger}(\mathcal{A}_c(\mathcal{Z}(S^{(j)})) - \hat{Y}_c)\right)$ if $c \in \mathcal{S}_c^{(j)}$, and $G_c^{(j)} = G_c^{(j-1)}$; otherwise, note that $(\cdot)^{\dagger}$ denotes the adjoint operator of its argument. Step 10 corresponds to a reformulation of this approximation.

Algorithm 1 Stochastic FB to solve (7)

1: **Input:** Let $S^{(0)} = (S_1^{(0)}, S_2^{(0)}, S_3^{(0)})$, with $S_1^{(0)} \in \mathbb{R}_+^{n \times N_1}$, $S_2^{(0)} \in \mathbb{R}_+^{1 \times N_2}$, and $S_3^{(0)} \in \mathbb{R}_+^{1 \times N_3}$.

2: Let $W \in \mathbb{R}_+^{n \times N_1}$, and $(\kappa, \gamma) \in [0, \infty[^2$.

3: Let $G^{(0)} = \left(\mathcal{Z}^{\dagger}\left(\mathcal{A}_c^{\dagger}(\mathcal{A}_c(\mathcal{Z}(S^{(j)})) - \hat{Y}_c)\right)\right)_{1 \leq c \leq C}$ and $\widetilde{G}^{(0)} = \sum_{c=1}^{C} G_c^{(0)}$

4: **Iterations:**

5: **For** $j = 0, 1, 2, \dots$

6:     select randomly $\mathcal{S}_c^{(j)} \subset \{1, .., C\}$

7:     **For** $c = 1, \dots, C$

8:         $G_c^{(j+1)} = \begin{cases} \mathcal{Z}^{\dagger}\left(\mathcal{A}_c^{\dagger}(\mathcal{A}_c(\mathcal{Z}(S^{(j)})) - \hat{Y}_c)\right), & \text{if } c \subset \mathcal{S}_c^{(j)} \\ G_c^{(j)} & \text{otherwise} \end{cases}$

9:     **end for**

10:     $\widetilde{G}^{(j+1)} = \widetilde{G}^{(j)} - \sum_{c \in \mathcal{S}_c^{(j)}} (G_c^{(j)} - G_c^{(j+1)})$

11:     $\widetilde{S}^{(j)} = S^{(j)} - \gamma \widetilde{G}^{(j+1)}$

12:     $S_1^{(j)} = \mathcal{P}_{\mathcal{B}_{1,W}^+(\kappa)}\left(\widetilde{S}_1^{(j)}\right)$

13:     $S_2^{(j)} = \mathcal{P}_{\mathbb{R}_+^{1 \times N_2}}\left(\widetilde{S}_2^{(j)}\right)$

14:     $S_3^{(j)} = \mathcal{P}_{\mathbb{R}_+^{1 \times N_3}}\left(\widetilde{S}_3^{(j)}\right)$

15: **end for**

16: **Output:** $S^{\star} = (S_1^{\star}, S_2^{\star}, S_3^{\star}) = \lim_{j \to +\infty}(S_1^{(j)}, S_2^{(j)}, S_3^{(j)})$

Steps 12–14 describe the projection steps. The projection of a variable $x \in \mathbb{R}^N$ onto a nonempty, closed, convex set $C \subset \mathbb{R}^N$ is given by $\mathcal{P}_C(x) = \min_{y \in C} \|x - y\|_2^2$. This corresponds to the closest point of x belonging to C, using an Euclidean distance [51]. Step 12 is the projection onto $\mathcal{B}_{1,W}^+(\kappa)$, related to the white matter. In accordance with the ℓ_0 approximation, κ is chosen to represent the S_1 sparsity level. Steps 13 and 14 are the projections related to the gray matter and the CSF, respectively, performed onto $\mathbb{R}_+^{1 \times N_2}$ and $\mathbb{R}_+^{1 \times N_3}$.

Let $(S^{(j)})_{j \in N}$ be a sequence generated by Algorithm 1. According to [51], if $0 < \gamma < 2/\|\mathcal{A}\|_S^2$, where $\|\cdot\|_S$ denotes the spectral norm, then $(S^{(j)})_{j \in N}$ converges to a solution $S^{\star}$ of problem (7). In our simulations, we consider that Algorithm 1 has converged when $\|S^{(j+1)} - S^{(j)}\|_2 < \nu \|S^{(j)}\|_2$, where $\nu = 10^{-3}$ is a tolerance, fixed by the user.

2.3.3. Weights Computation and Reweighting Procedure

We propose to use an ℓ_1 reweighting procedure to approximate the ℓ_0 pseudo-norm [29]. To this aim, the matrix of weights W in Step 12 of Algorithm 1, is chosen using a similar approach as described in [24]. The reweighting method is described in Algorithm 2.

Algorithm 2 Reweighting procedure

1: **Input:** Let $S^{(0)} = (S_1^{(0)}, S_2^{(0)}, S_3^{(0)})$, with $S_1^{(0)} \in \mathbb{R}_+^{n \times N_1}$, $S_2^{(0)} \in \mathbb{R}_+^{1 \times N_2}$, and $S_3^{(0)} \in \mathbb{R}_+^{1 \times N_3}$.

2: Let $W^{(0)} = \mathbf{1}_{n \times N_1}$.

3: **Iterations:**

4: **For** $t = 0, 1, \ldots, T - 1$

5: 　　$S^{(t+1)} = $ **Algorithm 1** $\left(S^{(t)}, W^{(t)} \right)$

6: 　　**compute** $W^{(t+1)}$ as per Equations (8) and (9)

7: **end for**

8: **Output:** $S^{(T)} = (S_1^{(T)}, S_2^{(T)}, S_3^{(T)})$

At each cycle $t \in \{0, \ldots, T - 1\}$ of the reweighting scheme, problem (7) is solved using Algorithm 1 for a particular choice of the weighting matrix (see Step 5). The weighting matrix $W^{(t+1)} = (W^{(t+1)})_{1 \leq d \leq n+2, 1 \leq v \leq N_1}$ is updated in Step 6 of Algorithm 2, as follows:

$$W_{d,v}^{(t+1)} = \frac{1}{\tau^{(t+1)} + B_{d,v}^{(t+1)}}, \tag{8}$$

where

$$B_{d,v}^{(t+1)} = \frac{1}{|\mathcal{N}(v)|} \sum_{d'v' \in \mathcal{N}(dv)} |S_{1_{d'v'}}^{(t+1)}|, \tag{9}$$

and

$$\begin{cases} \tau^{(0)} = \mathrm{Var}\left(B^{(0)} \right), \\ \tau^{(t+1)} = \max\left\{ \frac{\tau^{(t)}}{10}, \overline{\tau} \right\}, \end{cases} \tag{10}$$

with $\overline{\tau} > 0$. The parameter $\tau^{(t)}$ is a stability parameter that avoids the weights going to infinity when $B_{d,v}^{(t)} = 0$. Essentially, $B^{(t)}$ is defined as the set of voxels in the support of $S_1^{(t)}$ that gives a blurred version of $S_1^{(t)}$. The *spatial neighborhood* $\mathcal{N}(v)$ is defined as the group of voxels that share either a face, an edge or a vertex with the voxel of interest $v \in \{1, \ldots, N_1\}$. The *angular neighborhood* $\mathcal{N}(d)$ is defined as the set of atoms associated with the directions that fall within a cone of $15°$ with the direction of interest d. The neighborhood of an element individuated by the indices d, and v is then indicated by $\mathcal{N}(dv)$ and corresponds to all the indices in the support of $S_1^{(t)}$ that are simultaneously included in $\mathcal{N}(v)$ and $\mathcal{N}(d)$.

Weights designed in [24] have two-fold effects. Each element $W_{d,v}^{(t)}$ affects the corresponding element $S_{1_{d,v}}^{(t+1)}$ in such a way that large weights progressively force to zero spurious peaks, while small weights favor the presence of the FOD coefficients. Consequently the weighting matrix, associated with the ℓ_1 norm, promotes sparsity. In addition, the smoothness of the variation of the FOD coefficients across spatial neighborhoods is also promoted in (9) by averaging over neighbor voxels. Due to the discretization of the dictionary, fiber contributions are usually spread over a small angular support.

That is the reason why (9) contains a summation (rather than averaging) over neighbor directions. If a spatio-angular neighborhood is reasonably homogeneous, the weight in (8) will be reasonably small and the corresponding FOD coefficient will be preserved at the next iteration. On the contrary, spurious peaks isolated in their spatio-angular neighborhood will be associated with large weights. These large weights will tend to prevent such isolated peaks at the next iteration to not violate the weighted ℓ_1-norm constraint.

Finally, the reweighting process stops when the number of allowed iterations T is reached or there is no substantial relative variation between successive estimates of S_1, i.e., $\|S_1^{(t+1)} - S_1^{(t)}\|_2 < 10^{-3} \|S_1^{(t+1)}\|_2$. Empirical observations suggest to choose $T = 10$ [24].

2.3.4. Data Post-Processing

Once the solution is found, a post-processing procedure is performed along the columns of S_1 in order to extract the directions of the fibers within each voxel. The identification of the highest peaks is performed among all the directions contained within a cone of 30° for each different direction. No more than eight local peaks are assumed per voxel, and peaks smaller than 20% of the maxima are disregarded in order to suppress spurious contributions [22,24].

3. Experimental Setting

In this section, we provide the description of the different experimental settings. In particular, the specifications of the sampling schemes considered in the q- and in the k-spaces are provided in Sections 3.1 and 3.2, respectively. The method proposed to generate the dictionary is described in Section 3.3. Sections 3.4 and 3.5 provide the procedures used to estimate the phase and the coil sensitivities, respectively. A description of the metric used to quantitatively estimate the FOD reconstruction is provided in Section 3.6.

3.1. q-Space Under-Sampling

In this section, we describe the scheme adopted for sampling the signal in q-space. In the case of the synthetic data, the diffusion signal is uniformly sampled over a single shell using 30 diffusion gradients with $b = 1000\,\text{s/mm}^2$. In order to assess the FOD recovery in the presence of various kq-space regimes, the data provided was retrospectively under-sampled by selecting 6, 10, 15 and 20 uniformly distributed q-points, as reported in Figure 2.
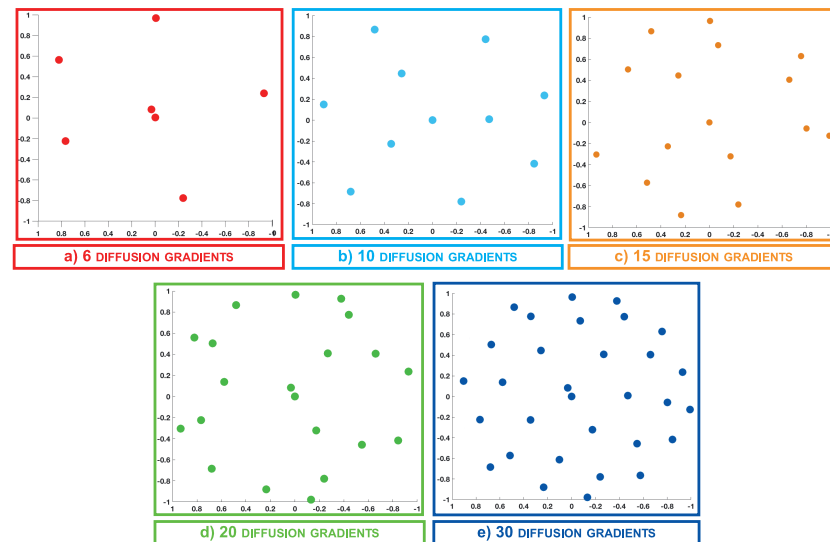

Figure 2. Equatorial plane projection of the q-space under-sampling schemes considered for the synthetic data acquired with 6 (**a**), 10 (**b**), 15 (**c**), 20 (**d**) and 30 (**e**) diffusion gradients.

In the case of the real data, the signal in q-space is sampled by 60 gradients over eight shells, considering b-values going from 550 to $4800\,\text{s/mm}^2$. In particular, 3, 6, 4, 3, 12, 12, 6 and 13 q-points are considered on sphere 1 ($b = 550\,\text{s/mm}^2$), sphere 2 ($b = 1050\,\text{s/mm}^2$), sphere 3 ($b = 1600\,\text{s/mm}^2$), sphere 4 ($b = 2150\,\text{s/mm}^2$), sphere 5 ($b = 2650\,\text{s/mm}^2$), sphere 6 ($b = 3200\,\text{s/mm}^2$), sphere 7 ($b = 4250\,\text{s/mm}^2$) and sphere 8 ($b = 4800\,\text{s/mm}^2$), respectively.

Three under-sampled datasets were created by retrospectively selecting the signal associated with 15, 30 and 45 q-points in order to compare the FOD recovery in the presence of different q-space under-sampling regimes. Each dataset was created by retrospectively sampling q-points from different shells following a Gaussian distribution centered in the fifth shell, that is to say, q-points with b-values that ranges between $2000\,\text{s/mm}^2$ and

$3000\,\text{s/mm}^2$ are sampled with higher probability [52]. Five different realizations were created for each q-space setting.

For all the synthetic and real datasets, the acquisition of the signal in the absence of diffusion is required in order to obtain $\mathbf{s}_0$. Examples of this image, for both the synthetic and real data, are provided in Figure 3A,D, respectively.
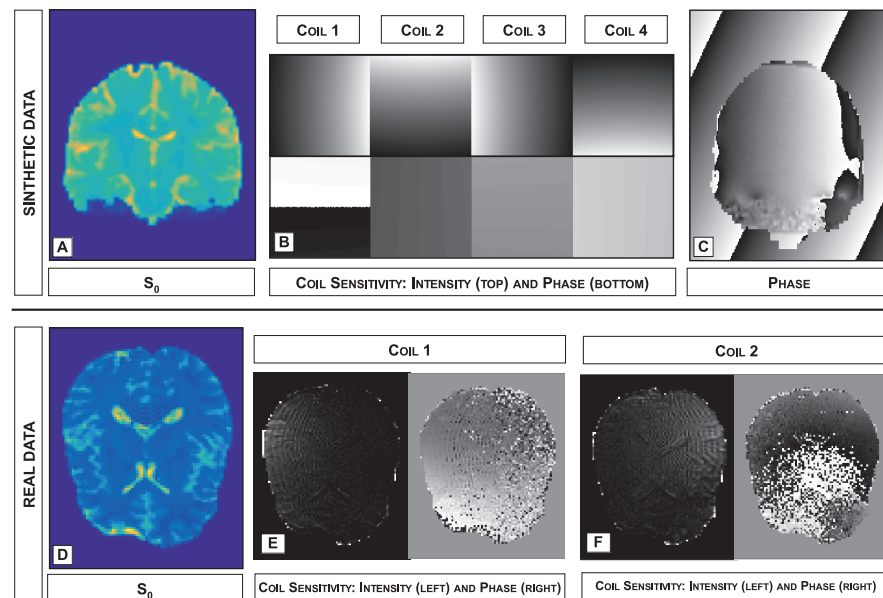

Figure 3. First row: s_0 image (**A**) and coil sensitivity maps (magnitude (**top B**) and phase (**bottom B**)) used in the experiments performed with synthetic data. (**C**) Example of phase map estimated with synthetic data when considering motion and magnetic field inhomogeneities. Second row: s_0 image (**D**) and magnitude and phase estimated for two different coils (**E,F**) for the experiments performed with real data.

3.2. k-Space Under-Sampling

As already explained, the objective of this work is to provide a robust reconstruction method when considering both q- and k-space under-sampling. The description of the schemes used to sample the signal in k-space is provided in this section.

In both synthetic and real data, for each gradient $q \in \{1, \ldots, M\}$, selection masks $M^{(q)}$ are built such that the signal in k-space is sampled within a Cartesian grid by a continuous trajectory. We call *k-space under-sampling factor* the ratio between the total amount of grid points in k-space, for a given field of view and resolution, and the number of samples considered for the reconstruction.

The potential of the proposed method is fully exploited in the 3D setting where the acquisition time can be significantly reduced by limiting both the number of excitations to cover each DW volume and the number of DW sub-volumes. As a proof of concept, the proposed FOD recovery method is evaluated for a specific class of 3D Fourier sampling patterns where the same 2D k-space under-sampling pattern is probed at all spatial frequencies in the third dimension of a selected sub-volume. By a simple inverse Fourier transform in this third dimension, the corresponding data can be cast, without loss of generality, as a slice-by-slice-identical 2D k-space under-sampling. The scheme chosen for the 2D k-space sampling consists in fully sampling the central lines of the k-space while regularly skipping lines at the periphery, with the central zone of the k-space is required to estimate the phase affecting each shot (i.e., each sub-volume).

We argue that a two-phase EPI acquisition protocol could be used to that effect. EPI commonly considers regularly sampled lines to enable a simple calibration of the eddy currents affecting each different k-space line. The proposed k-space under-sampling scheme could result from the combination of two uniform EPI schemes, one for the center

and one for the periphery, with the eddy currents calibration critical only at the interface between the two regions. A schematic representation of the under-sampling scheme used with both synthetic and real data is reported in Figure 4. An experimental evaluation of a precise acquisition protocol is beyond the scope of the present work.
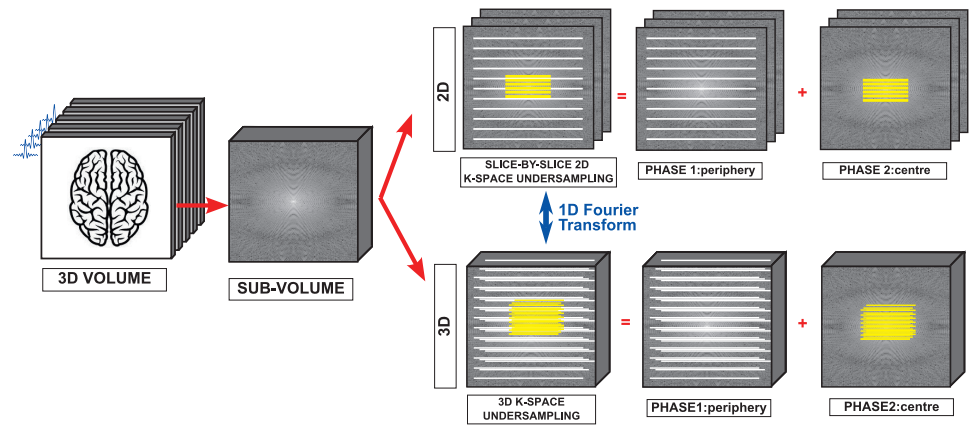

Figure 4. Illustration of the acquisition of the 3D volume containing the brain. Each sub-volume is acquired considering a 3D two-phase EPI scheme (bottom line). The presented results were obtained considering a sequence of identical slice-by-slice 2D k-space under-sampling schemes (top line), which can be brought back to a specific class of 3D k-space under-sampling scheme to which a Fourier transform in the third dimension was applied. The proposed pattern results from the combination of two different regular schemes: the fully sampled central zone (yellow lines) and the regularly under-sampled periphery (white lines).

In the remainder of this work, all our experiments are designed with this k-space sampling scheme. Note, again without loss of generality, that the performance of the proposed approach is evaluated in a setting including a single DW (sub-)volume at each q-space point. Multiple sub-volumes would simply be acquired sequentially and reconstructed separately.

Note that the proposed method requires the complete k-space acquisition of the s_0 signal to implicitly force the FOD coefficients to sum up to one in each voxel. Furthermore, s_0 is used for normalization, segmentation and coil sensitivities calibration purposes as described in detail in Sections 3.4 and 3.5.

3.3. Dictionary Generation

In this section, we provide the description of the method used to create Φ for both the simulated and the real data. The elements of the dictionary Φ are generated by relying on the Gaussian Mixture Model of the q-space signal [8]. More specifically, the dictionary $\Phi = \left[\Phi_{q,d} \right]_{1 \leqslant q \leqslant M, 1 \leqslant d \leqslant (n+2)}$ is defined as $\Phi_{q,d} = \exp(-b\hat{\mathbf{q}}^T \mathbf{D}^{(d)} \hat{\mathbf{q}})$, where q is the index that explores the diffusion gradients $\mathbf{q}$ (with b-value b and orientation $\hat{\mathbf{q}}$) considered for the acquisition, and d is the index associated with the orientations chosen for the discretization of half of the unit sphere. We consider $n = 500$ points for the discretization of the dictionary.

Moreover, two additional atoms are considered to model the isotropic diffusion in the gray matter and CSF. The d-th atom of the dictionary is indicated by $\left[\Phi_{q,d} \right]_{1 \leq q \leq M}$, and this corresponds to the response of a single fiber, oriented along the d-th direction, subject to M different diffusion gradients. The diffusion tensor $\mathbf{D} = \mathrm{diag}(\lambda_1, \lambda_2, \lambda_3)$ characterizes each fiber population, where λ_1 is the longitudinal diffusivity, and λ_2 and λ_3 are the 2 transverse diffusivity coefficients. In particular, $\mathbf{D}^{(d)}$ is the rotated version of $\mathbf{D}$ along the direction d.

The diffusivity coefficients for the white matter fibers are fixed to the following values: $\lambda_1 = 17 \times 10^{-4} \, \mathrm{mm}^2/\mathrm{s}$ and, $\lambda_2 = \lambda_3 = 3 \times 10^{-4} \, \mathrm{mm}^2/\mathrm{s}$ [18,19]. The diffusion tensors associated with the isotropic atoms are composed by 3 equal values in the diagonal.

In particular, we consider $\lambda_1 = \lambda_2 = \lambda_3 = 17 \times 10^{-4}\,\mathrm{mm^2/s}$ and $\lambda_1 = \lambda_2 = \lambda_3 = 30 \times 10^{-4}\,\mathrm{mm^2/s}$ for the gray matter and CSF diffusivity coefficients, respectively.

3.4. Phase Estimation

In theory, diffusion images are real-valued. However, in practice, they are often contaminated by phase factors during the diffusion encoding process. This phase contamination is mostly due to magnetic field inhomogeneities and biological motion (e.g., physiological and involuntary patient motion).

Usually, methods dealing with signals directly in q-space overcome this difficulty by simply taking the modulus of the complex diffusion signal, in order to obtain real diffusion-weighted images. However this method cannot be used in k-space and kq-space. Indeed, since the phase contamination breaks the Hermitian symmetry of the images, it cannot trivially be removed. In particular, in the linear operator expressed in (2), the diagonal matrix $H^{(q,c)}$ takes into account this phase factor, for a fixed gradient q and coil receiver c.

In the case of synthetic data, the phase contamination produced by the magnetic field inhomogeneities is provided within the contest data set. In addition, effects due to motion are mimicked by adding a different linear phase map to the signal associated with each sub-volume of the phantom and coil receiver. The linear phase maps are generated in such a way that the corresponding k-space shift is constrained within 10 phase-encoding lines. In each voxel $v \in \{1, \ldots, N\}$, the intensity of the signal in q-space is multiplied by $e^{i\theta_v}$, θ_v as the sum of the phases due to both motion and magnetic field inhomogeneities associated with v. In Figure 3C, an example of estimated phase map due to motion and magnetic field inhomogeneities is provided.

For both synthetic and real data, phase-distortions are estimated from the central portion of the k-space data [53,54], which is always fully sampled, (see Section 3.2). By performing the inverse Fourier transform of a zero-padded version of the fully sampled central part of the k-space, we obtain a complex-valued signal whose phase provides the diagonal elements of $H^{(q,c)}$.

3.5. Coil Sensitivity Estimation

In this section, we provide the method to estimate the coil sensitivity coefficients stored in the diagonal matrix $U^{(c)}$, appearing in the linear measurements described in Equations (1) and (2) for both synthetic and real data. Images acquired in absence of diffusion from different coil receivers are combined through the *sum of squares* [55] method in order to obtain a baseline image. Subsequently, the images associated with each receiver coil are divided by the baseline image in order to determine the corresponding sensitivity map.

For synthetic data, the acquisition from four different coil receivers was simulated by considering the toolbox at http://bigwww.epfl.ch/algorithms/mri-reconstruction/ (acessed on 11 August 2021). Examples of the coil sensitivity, when $C = 4$, are shown in Figure 3B.

Examples of the estimated sensitivity maps obtained in the real data experiment are provided in Figure 3E.

3.6. Evaluation Criteria

The quality of the fiber reconstruction was evaluated by using the metrics proposed in [24]. For each single voxel, the fiber reconstruction evaluation takes into account the number of fibers correctly identified and the angular accuracy of the recovered direction.

First, we define the *success rate* (SR) index representing the proportion of voxels in which the number of fibers is correctly identified. More precisely, when all the estimated fibers fall within a tolerance cone of 30° around the true fibers, we have a success. The SR index depends on the *false positive* and *negative* rates that represent the average over all voxels of the number of overestimated and underestimated fibers per voxel.

Secondly, we define the *mean angular error* $(\bar{\theta})$ as the average of the angular errors associated with each true fiber. The angular error is defined as $\theta = \frac{180}{\pi} arccos(|d_{true} \cdot d_{estimated}|)$, where d_{true} and $d_{estimated}$ are the true fiber direction and the direction of the closest estimated fiber, respectively.

4. Results

In this section, we present the results obtained when using the proposed method with both synthetic and real data. Specifically, in Sections 4.1 and 4.2 we discuss the results obtained for synthetic and real data, respectively.

4.1. Synthetic Data

A first analysis of the FOD reconstruction using the proposed kq-space under-sampling method was performed relying on the fiber configuration of the numerical phantom proposed in the *ISMRM Tractography Challenge 2015* [56]. The phantom consists of a volume of $N = 90 \times 90 \times 108$ voxels, acquired by using $M = 32$ diffusion gradients distributed over a single shell (see Section 3.1 for more details).

We consider the acquisition from four coil receivers, where the sensitivity maps are simulated and estimated using the process provided in Section 3.5. Phase contamination due to motion and magnetic field inhomogeneities was generated as described in Section 3.4 and incorporated in the data. The q-space signal was then converted to k-space through the Fourier transform and contaminated with Gaussian noise with zero mean and standard deviation $\sigma = \frac{\langle s_0 \rangle}{\text{SNR}}$. Note that $SNR = 30$ was chosen for the numerical simulations. Lastly, selection masks, built as described in Section 3.2, were applied to the k-space of the diffusion-weighted images to obtain the under-sampled data in kq-space. The dictionary Φ was generated using the procedure described in Section 3.3.

In the following section, the proposed recovery scheme is compared to three different reconstruction frameworks, which represent some of the state-of-the-art approaches for the FOD recovery in the presence of measurements under-sampled both in k and q-spaces separately and in kq-space jointly. Among the state-of-the-art methods investigated, we can distinguish between two types of approaches: the two-step and the one-step approaches.

The two-step approaches are the conventional approach to recover the FOD coefficients in the presence of under-sampled kq-space data, where the DW images are first reconstructed, and the FOD coefficients are subsequently estimated. As already mentioned in Section 3.4, two-step approaches avoid to model the phase contamination due to motion and magnetic field inhomogeneities, by recovering complex images, whose imaginary part is discarded by taking the magnitude. Among all the k-space methods, the TV-regularized approaches were considered for this comparison. The one-step approaches consists on the fiber orientation estimation from the under-sampled kq-space signal directly. Contrary to the two-step approaches, the one-step approaches require the phase contamination to be modeled. In particular, the proposed approach is a one-step method. Below is a summary of the considered methods.

1. $TV - L2_+$ is a two-step approach consisting in a first step where the DW volumes are recovered by using the TV_+ regularization and in a second step where FOD coefficients are recovered by relying on the non-negative least squares problem.

$$STEP\ 1: \quad \min_{I \in \mathbb{C}^{M \times N}} \|\mathcal{F}(I) - \hat{Y}\|_2^2 + \lambda \|I\|_{TV}, \tag{11}$$

where $\hat{Y} \in \mathbb{C}^{MC \times N}$ contains the signal in kq-space. Complex DW volumes are recovered in $I \in \mathbb{C}^{M \times N}$ from the TV regularized problem, where $\mathcal{F}(I) = (\mathcal{F}_{q,c}(I))$, with $\mathcal{F}_{q,c}(I) = IU^{(c)}FM^{(q)}$ for $q \in \{1, \dots, M\}$ and $c \in \{1, \dots, C\}$. To exclude the

remaining phase contamination, the magnitude of the complex images is considered for the FOD recovery (i.e., $\tilde{I} = \frac{|I|}{S_0}$).

$$STEP\ 2: \quad \min_{X \in \mathbb{R}_+^{(n+2) \times N}} \|\mathcal{D}(X) - \tilde{I}\|_2^2 \tag{12}$$

where $\mathcal{D}_q(X) = \Phi_q X$ for all $q \in \{1, \ldots, M\}$. Positivity of the FOD coefficients is imposed in this step.

2. $TV - STR_+$ is a two-step approach consisting in the recovery of both DW volumes and FOD coefficients from regularized problems. In the first step, the TV prior is considered while, in the second step, the *structured sparsity* prior proposed in [24] is exploited.

$$STEP\ 1: \min_{I \in \mathbb{C}^{M \times N}} \|\mathcal{F}(I) - \hat{Y}\|_2^2 + \lambda \|I\|_{TV}$$
$$STEP\ 2: \min_{X \in \mathbb{R}_+^{(n+2) \times N}} \|\mathcal{D}(X) - \tilde{I}\|_2^2 \quad \text{s.t} \quad X \in \mathcal{B}_{1,W}^+(\kappa)$$

3. TV_+ is a one-step approach consisting in the recovery of the FOD coefficients from the kq-space signal applying the TV prior to the images in order to implicitly promote a smooth variation of the FOD coefficients within neighbor voxels.

$$\min_{X \in \mathbb{R}_+^{(n+2) \times N}} \|\mathcal{A}(X) - \hat{Y}\|_2^2 + \lambda \|\Phi X\|_{TV}$$

4. $STR_+ - TISS$ is the proposed one-step approach consisting in the recovery of the FOD coefficients by solving the problem proposed in (7) where the *structured sparsity* prior and the spatial distribution of the different tissues is taken into account.

The Primal-Dual [57] and the FB algorithms were used to, respectively, solve the first and the second step of the two-step approaches. More specifically, in the second step of $TV - STR_+$, the FB algorithm was used multiple times to solve 10 weighted-ℓ_1 problems. For all the frameworks, the optimal parameters were found using a grid-search approach. The evaluation of the image reconstruction in the case of the two-step approaches was performed relying on the SNR index. On the other hand, the SR index was chosen to evaluate the quality of the fiber orientation reconstruction and determine the best parameters for both the second step of the two-step approaches and the one-step approaches.

The parameters obtained in such a way are as follows: $1 \times 10^{+3} \leq \lambda \leq 2.1 \times 10^{+3}$ and $3.9N \leq \kappa \leq 4.2N$. TV_+ is solved using the Primal-Dual algorithm, considering $4 \times 10^{+5} \leq \lambda \leq 9.8 \times 10^{+5}$. On the other hand, $STR_+ - TISS$ was solved using the Algorithm 2, with $\kappa = 4N_{wm}$, with N_{wm} as the number of white matter voxels. However, in the case of the synthetic data, where only four coil receivers are considered, the gradient of the data fidelity term is fully computed and not approximated. Hence, at each iteration $j \in \mathbb{N}$, we choose $\mathcal{S}_c^{(j)} = \{1, 2, 3, 4\}$ in Algorithm 1, which reduces to the FB algorithm.

We emphasize that additional tests were performed for $TV - L2_+$ and $TV - STR_+$ with $\lambda = 0$, corresponding to remove the TV regularization for the DW images in the least-squares minimization problem. These tests have shown that the use of the TV prior provides better results.

In Figure 5, the quantitative evaluation of the fiber reconstruction is provided in the case of various kq-space under-sampling regimes for the four different methods described above. Performances obtained considering different q-space under-sampling ratios are reported in different colors, while performances obtained considering different *k-space under-sampling factors* are reported along the x-axis. The quality of the reconstruction was evaluated considering noisy data with SNR = 30, and 5 different q-space under-sampling factors (i.e., $M \in \{6, 10, 15, 20, 30\}$), and 6 different *k-space under-sampling factors* (i.e., $N/K \in \{1,2,4,6,8,10\}$).

The obtained SR index, mean angular error, and the rate of *false positives* and *negatives* are shown in Figure 5, from the first to the last row, respectively. The error bars correspond to the standard deviation of the evaluation indices obtained from k-space data corrupted by different noise realisations.
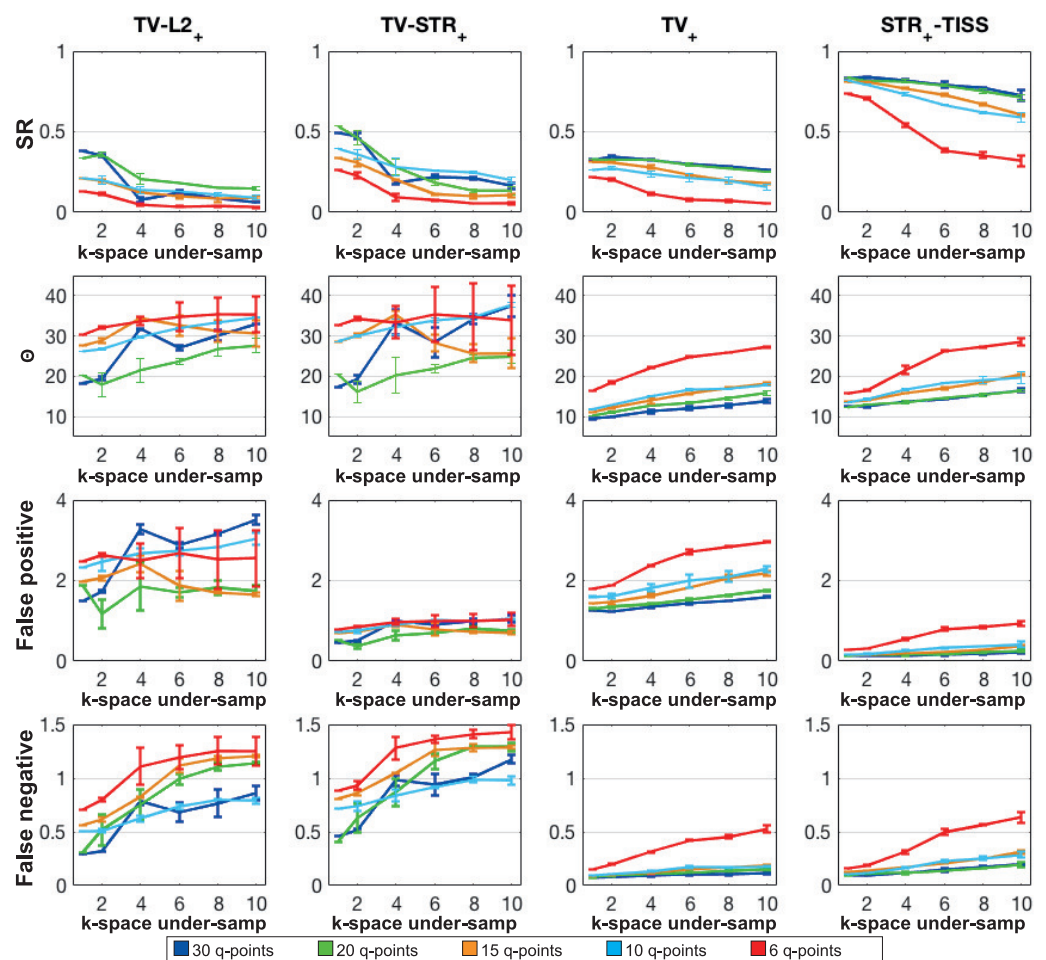
Figure 5. Quantitative evaluation of the fiber orientation estimated by using four different methods for the synthetic data from the ISMRM Tractography Challenge 2015. For the evaluation, the SR (**first row**), the mean angular error (**second row**), the number of *false positives* (**third row**) and *false negatives* (**fourth row**) are shown for each different method along each column. Reconstructions were performed considering different number of q-points (i.e., $M \in \{6, 10, 15, 20, 30\}$), shown using different colors and for different *k-space under-sampling factors* (i.e., $N/K \in \{1, \ldots, 10\}$).

Performances obtained with the two-step approaches are reported in the graphs in the first two columns of Figure 5. In the presence of k-space under-sampling regimes, both $TV - L2_+$ and $TV - STR_+$ are characterized by SR rates lower that the proposed $STR_+ - TISS$ method. For $TV - L2_+$, the low performances result from the high number of spurious peaks that is estimated in each voxel as can be observed in the diagram in the first column of Figure 5. When using the *structured sparsity* prior to regularize the FOD coefficients in $TV - STR_+$ the number of *false positive* drastically decreases, as shown in the corresponding graph in the second column of Figure 5. However, the number of *false negative* moderately rises up, preventing the SR index to increase. In general, the direct regularization of the FOD provides better performances than $TV - L2_+$.

The last two columns of Figure 5 give the graphs of the two one-step approaches: the TV_+ method (third column), and the proposed method $STR_+ - TISS$ (fourth column). TV_+ does not reach high performances for any of the kq-space under-sampling regime. The SR index remains below 0.31, that is reached considering 30 q-points with no k-space under-sampling. These low SR values are due to the fact the TV regularization is incapable of excluding the numerous spurious peaks that affect the FOD reconstruction.

However, the performances obtained with the one-step approach TV_+, are shown to be quite robust to the k-space under-sampling. Indeed, the SR decreases only of 0.15 (resp. 0.15) and θ increases of $4°$ (resp. $5°$) going from no k-space under-sampling to a *k-space under-sampling factor* of 10 with 30 q-points (resp. 10 q-points). As discussed in Section 2.3.1, and shown in [22], since fiber coefficients are implicitly forced to sum up to one in each voxel, imposing ℓ_1 norm in this context is ineffective. Additional simulations confirmed this behavior showing that using ℓ_1 norm to regularize the FOD coefficients in addition to TV did not lead to a significant difference in the results.

Concerning the proposed $STR_+ - TISS$ method, $STR_+ - TISS$ outperforms both the two-step and the one-step approaches here above presented. When compared to TV_+, $STR_+ - TISS$ shows to substantially benefit from the direct regularization of the FOD coefficients, through the *structured sparsity* prior. Considering no k-space under-sampling with 6-points (resp. 30 q-points), SR = 0.31 for TV_+ (resp. 0.48), against SR = 0.84 for $STR_+ - TISS$ (resp. 0.86). For a *k-space under-sampling factor* of 10 with 6-points (resp. 30 q-points), SR = 0.06 for TV_+ (resp. 0.34), against SR = 0.62 for $STR_+ - TISS$ (resp. 0.75). Thus $STR_+ - TISS$ exhibits much higher performances when compared to TV_+. However, both methods are characterized by a similar amount of *false negative*.

For the sake of brevity, we do not provide the comparison considering or ignoring the tissue spatial distribution. Additional simulations have shown that, when the regularization parameter κ is known, FOD configurations recovered by the two methods are comparable. However, in practice κ needs to be tuned, and using prior knowledge of the tissue distribution improves the reconstruction. This fact can be understood as follows. In the presence of under-estimated κ, coefficients associated with isotropic tissues compartments take precedence over the fiber compartments (since isotropic behavior can more easily fit the data). When the tissue distribution is taken into account, the coefficients associated with the fiber compartments are processed separately thus preventing them from being annihilated by the dominant presence of the isotropic compartment coefficients.

In Figure 6, we provide an illustrative example, considering 15 diffusion gradients with a *k-space under-sampling factor* of 4, of the performance obtained from the four considered methods. In particular, we show the mean angular error (first row), the SR (second row), the *false positive* (third row) and the *false negative* (fourth row) maps. These results are confirming the quantitative results described in Figure 5. In particular, the maps evaluating the fiber configurations recovered from $TV - L2_+$ and $TV - STR_+$ show the worst performances. The SR maps in Figure 6(1B,2B) mostly show red pixels, indicating the presence of voxels with an uncorrected number of estimated fibers. For the $TV - L2_+$ method, the majority of such voxels is characterized by over-estimated fiber populations, as it is evidenced by the map in Figure 6(1C).

By comparing the maps in Figure 6(1C,2C), we can observe that the use of the *structured sparsity* prior significantly decreases the amount of *false positive*. Nevertheless, the map in Figure 6(2D) highlights the presence of higher *false negatives* rates. For the TV_+ method, the SR remains low, as it can be observed from the large areas with red pixels in Figure 6(3B). However, the angular accuracy of the estimated fiber configurations is significantly improved when compared to the two-step approaches (see maps on the first row of Figure 6). In Figure 6(4B), the SR map obtained with the $STR - TISS_+$ method shows the higher rate of yellow pixels indicating the presence of a large number of voxels with the correct number of estimated fibers. In addition, the angular accuracy achieved with $STR - TISS_+$ remains low, as displayed by the map in Figure 6(4A).
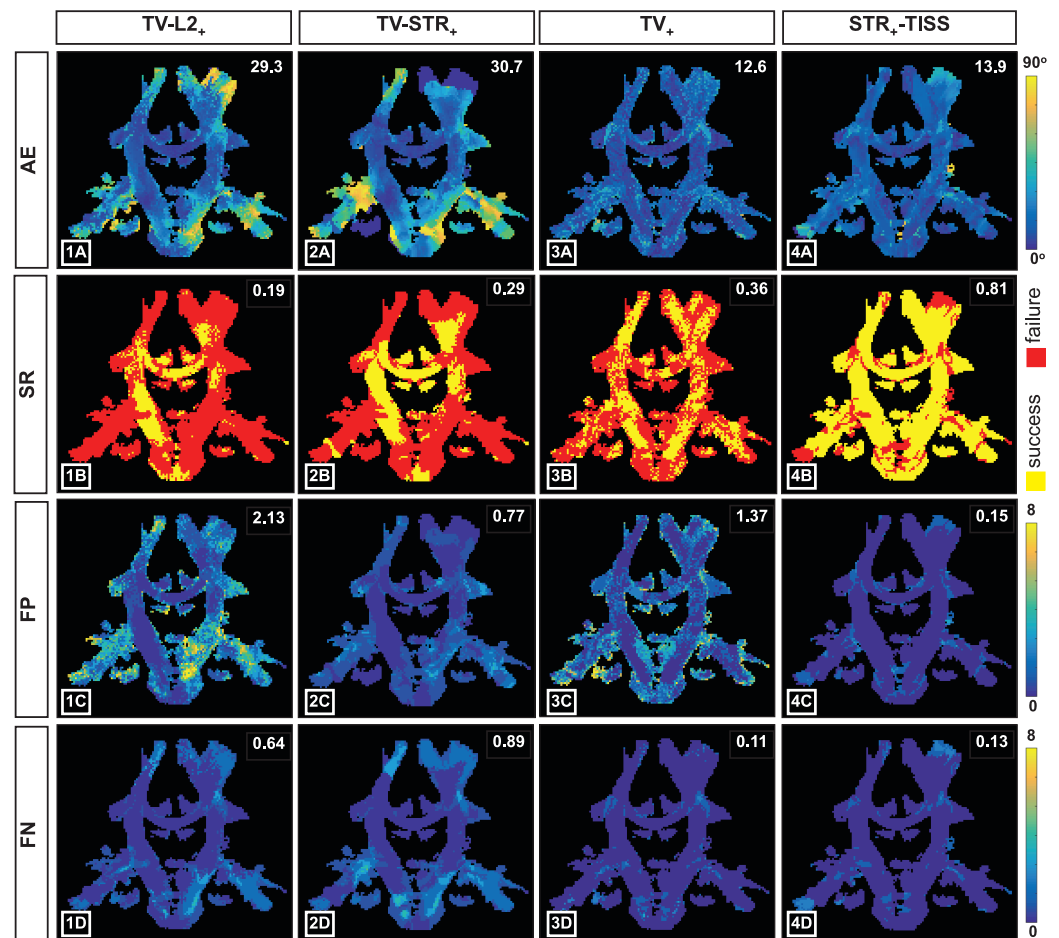
Figure 6. Visualization of the quantitative results obtained by using the four different methods described in Section 4.1, applied to the synthetic data of the ISMRM Tractography Challenge 2015. Mean angular error (**1A–4A**), success rate (**1B–4B**), *false positive* (**1C–4C**) and *false negative* (**1D–4D**) maps obtained from the reconstructions considering M = 15 q-points, with a *k-space under-sampling factor* N/K = 4. On the bottom right corner, values of the corresponding quantitative index are reported for the displayed slice.

4.2. Real Data

The kq-space under-sampling framework was tested for the recovery of the fiber orientation from real data. The real data set consists of a volume of $100 \times 100 \times 60$ voxels acquired on a 3T Magnetom Trio system (Siemens, Germany) with a pixel size of $2.2 \times 2.2 \times 2.2$ mm^3. The dataset was acquired considering 17 different coil receivers, and 60 diffusion gradients distributed over multiple shells, as described in Section 3.1. The complete signal was retrospectively under-sampled both in q- and k-space (as described

in Sections 3.1 and 3.2, respectively) in order to compare the reconstruction performances considering different under-sampling regimes.

In the case of real data, Φ was generated by considering $n = 500$ for the discretization of the sphere using the procedure described in Section 3.3. The tissue segmentation maps were estimated from the segmentation of the s_0 signal and $H^{(q,c)}$ was estimated from the low resolution image as described in Section 3.4. The diagonal of the matrix S_0 is filled with the s_0 values, which were recovered from the multi-coil signals while correcting for the estimated phase through a least squares approach. Ultimately, $U^{(c)}$ was estimated from the s_0 signal as described in Section 3.5. Note that the real data was corrected for the EPI Nyquist ghosting prior to the reconstruction.

Results were obtained by solving problem (7) using Algorithm 2, where the reweighing process is performed 10 times at most, and $\kappa = 5 \times N_{wm}$. The computation of the gradient in Algorithm 1 is approximated choosing 12 out of 17 coils at each iteration, i.e., $|\mathcal{S}_c| = 12$. In particular, four of these coils are fixed (chosen at each iteration), and eight are randomly selected. In addition, a Nesterov acceleration [58] is considered in Algorithm 1. Although no theoretical result ensures the convergence of the resulting method, additional simulations have shown that Nesterov acceleration leads to much faster convergence while the same minimum is achieved using both the methods (with or without Nesterov acceleration).

Considering the fiber configuration recovered from 60 diffusion gradients with full k-space as reference, quantitative evaluation maps are computed and provided in Figure 7, for the fiber configurations obtained from four different kq-space under-sampling settings. For each different setting, the FOD reconstruction was performed 5 times, each considering a different realisation of q-space under-sampling. The results presented here corresponds to the mean and the standard deviation of the evaluation indices obtained considering different q-space under-sampling realizations.

In the first column of Figure 7 the performances obtained considering 15 diffusion gradients with full k-space is reported. The performances achieved with 30 diffusion gradients considering a *k-space under-sampling factor* of 2 and 45 diffusion gradients considering a *k-space under-sampling factor* of 3 are shown in the second and third columns of Figure 7, respectively. Ultimately, the evaluation maps of the fiber geometries recovered from 60 diffusion gradients with *k-space under-sampling factor* of 4 are shown in the fourth column of Figure 7.

The comparison between the evaluation maps provided in Figure 7 shows that, considering the same overall kq-space under-sampling factor, the recovery from data that were under-sampled only in q-space results in lower quality reconstruction (SR $= 0.18 \pm 0.06$, $\bar{\theta} = 24.5° \pm 5.1$), as highlighted by the larger amount of red pixels and extended areas with high angular error in Figure 7A,E, respectively. On the other hand, the fiber configuration recovered from 60 diffusion gradients with 4-fold k-space under-sampling exhibits the highest reconstruction quality (SR $= 0.59 \pm 0.01$, $\bar{\theta} = 7.5° \pm 0.2$) as can be observed in the fourth column of Figure 7.

In Figure 8, we provide a zoomed view of the fiber configurations recovered from the real data in the presence of the same kq-space under-sampling settings described above. The fiber configuration obtained considering 60 diffusion gradients with full k-space is provided in Figure 8A as golden reference. Fiber geometries reported in Figure 8B–D appear to be very close to the golden reference. In contrast, the fiber configuration estimated in Figure 8E fails to represent the main fiber crossing revealed in the golden reference.

The analysis of the quantitative and qualitative results strongly suggests that strategies with combined kq-space sampling are advisable, when compared to q-space only strategies.
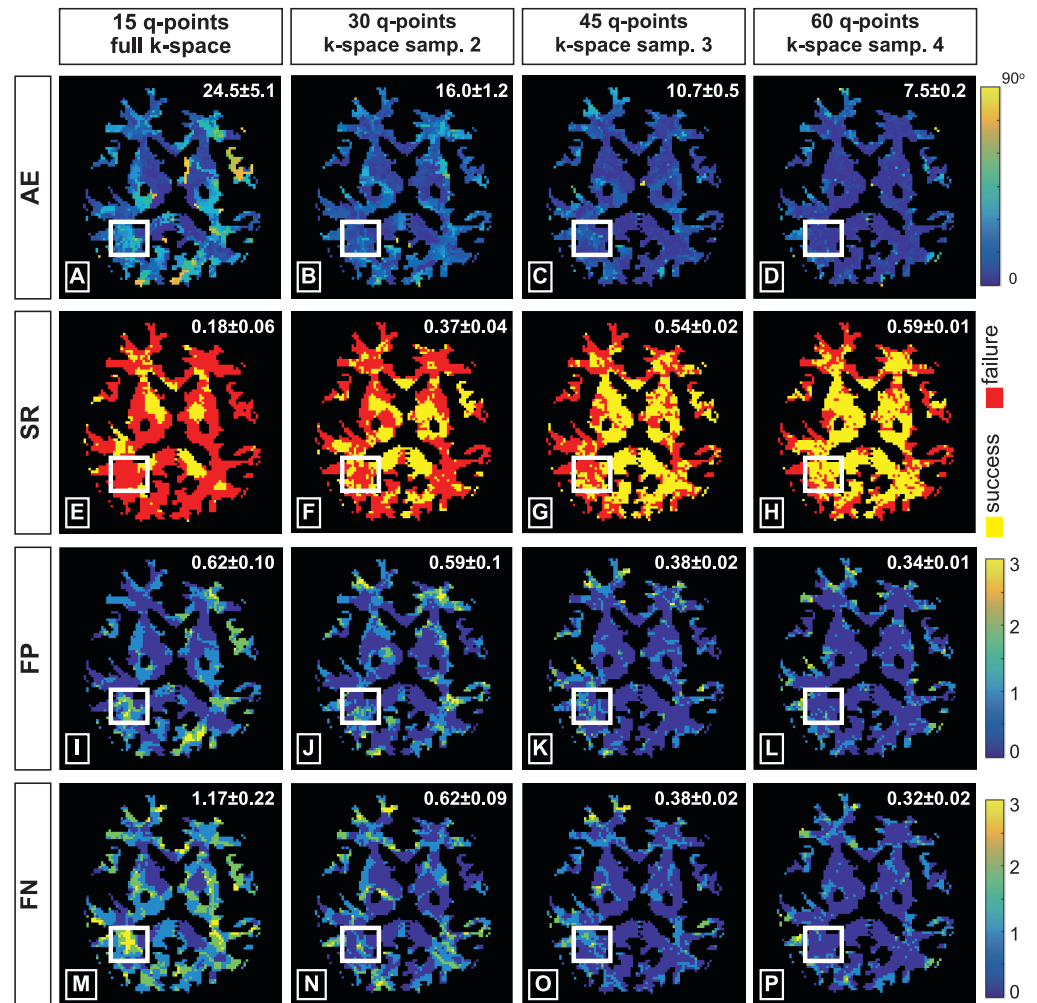
Figure 7. Quantitative evaluation of the fiber orientation reconstruction performed on the real data by using the proposed kq-space under-sampling framework. Reconstructions considering 15 diffusion gradients with full k-space (**first column**), 30 diffusion gradients with *k-space under-sampling factor* of 2 (**second column**), 45 diffusion gradients with *k-space under-sampling factor* of 3 (**third column**) and 60 gradient with *k-space under-sampling* factor of 4 (**fourth column**) were compared with the fiber configuration obtained using 60 q-points and full k-space. The maps of the mean angular error (**first row**), success rate (**second row**), *false positives* (**third row**) and negatives (**fourth row**) are considered for the evaluation.

In Table 1, we compare the computational time required to process the real data with or without stochastic approach. Specifically, we investigate the computational time per iteration either when all the 17 coils are considered at each iteration (deterministic approach), or when 12 of the 17 coils are selected randomly at each iteration (stochastic approach). On the one hand, we can observe that the computational time per iteration when considering 12 coils per iteration corresponds (approximately) to 12/17 times the computational time per iteration when considering 17 coils. This suggests that the computation of the gradient step in Algorithm 1 is much heavier than the computation of the projection step.
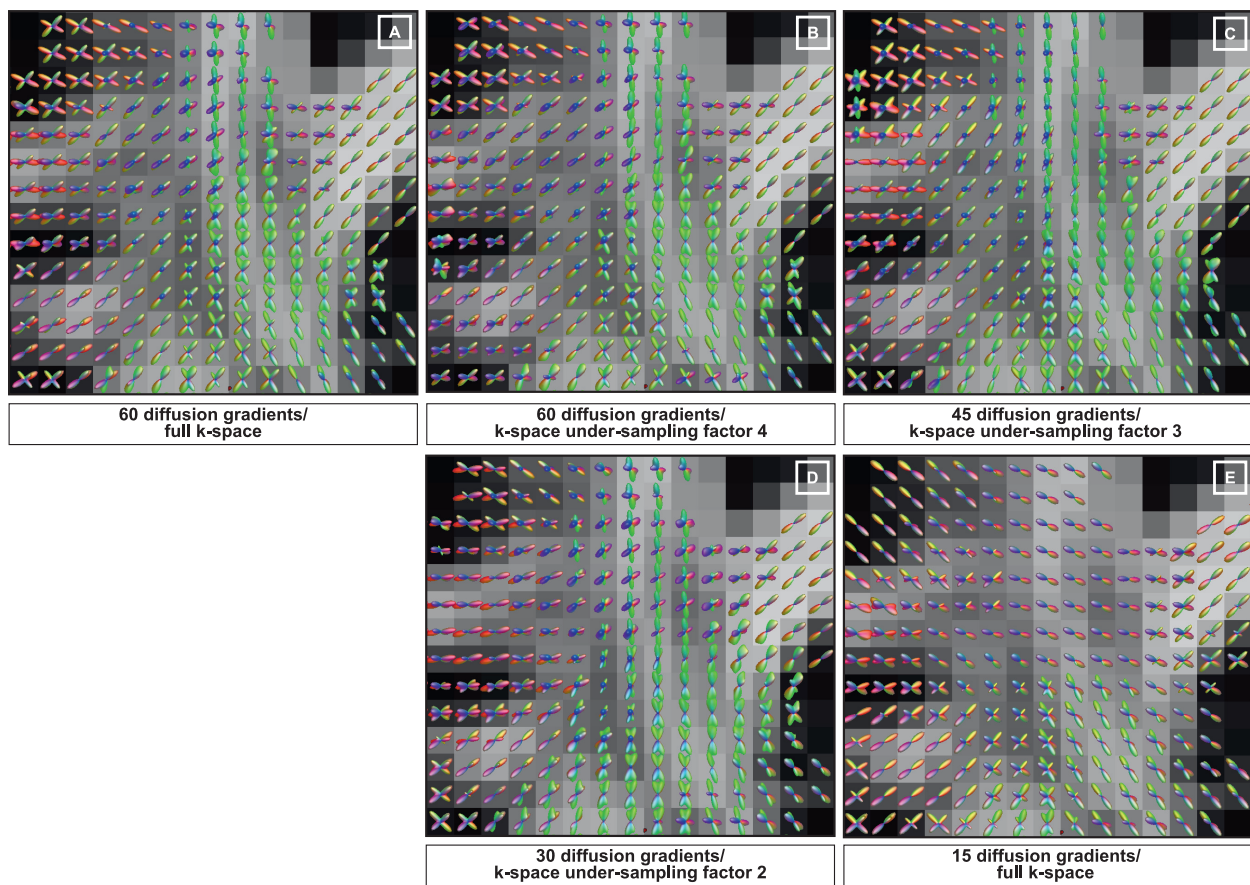
Figure 8. Qualitative evaluation of the fiber orientation reconstruction performed on the real data by using the proposed kq-space under-sampling framework. The fiber configurations displayed correspond to the region framed highlighted by the white frame in Figure 7. Reconstructions were performed considering 60 diffusion with full k-space (**A**), and with *k-space under-sampling factor* of 4 (**B**), considering 45 diffusion gradients with *k-space under-sampling factor* of 3 (**C**), 30 diffusion gradients with with *k-space under-sampling factor* of 2 (**D**) and 15 diffusion gradients with full k-space (**E**).

Table 1. Comparison of the computational time required for the FOD recovery from the real data, considering four different under-sampling settings, using either 17 coils (deterministic approach) or 12 coils selected randomly per iteration (stochastic approach).

Samp.		17 Coils (Deterministic)		12 Coils (Stochastic)	
q	k	Time per Iter. (s)	Iter.	Time per Iter. (s)	Iter.
60	4	21.5 ± 0.13	3520	16.4 ± 0.15	3536
45	3	17.9 ± 0.08	3603	13.5 ± 0.12	3618
30	2	12.8 ± 0.07	3445	9.9 ± 0.13	3445
15	1	7.7 ± 0.25	1856	6.1 ± 0.21	1854

Thus, the computational time per iteration is lower using the stochastic approach than the deterministic approach. On the other hand, it can be noticed that both the deterministic and the stochastic approaches necessitate approximately the same number of iterations to reach the stopping criterion. Consequently, the use of a stochastic approach to recover the fiber configurations considering 12 coils per iteration enables to reduce both the memory requirement and the total computational time.

5. Discussion

We developed a method to accelerate high angular and spatial resolution dMRI acquisition relying on a 3D kq-space under-sampling scheme. We provided a novel formulation

to estimate the FOD coefficients from highly under-sampled sub-volumes. The proposed approach used two types of anatomical priors. On the one hand, a brain tissue segmentation constraint, which does not required additional acquisitions and can be inferred from the no-diffusion-weighted image, was explicitly imposed for the FOD recovery. On the other hand, *structured sparsity*, developed previously in [24], was leveraged to promote simultaneously voxel-wise sparsity and spatial smoothness of fiber orientation.

The resulting minimization problem was approached via a reweighting scheme to solve a sequence of convex minimization problems, using a stochastic FB algorithm structure. By considering the stochastic variation, multi-coil data can be processed while minimizing both the memory requirement and the reconstruction time. Through synthetic and real data experiments, we demonstrated that the proposed recovery framework outperformed the existing method in kq-space and the traditional two-step reconstruction approaches recovering sequentially the diffusion weighted images and the FODs. Furthermore, we observed that, for equal overall under-sampling ratios, the proposed kq-space approach performed better when the k-space under-sampling was exploited rather than heavily under-sampling in q-space only.

Some limitations of our work are to be highlighted. First, a two-phase uniform EPI scheme was proposed to simultaneously fully sample the center and under-sample the periphery of the k-space data in order to minimize the effects arising from the irregular sampling. In this context, the use of more advanced models, taking into account the correction of the geometrical distortions typically affecting the EPI data and the use of acquisition schemes less prone to artifacts, would certainly contribute to improving the performance of the proposed method.

Secondly, one of the critical steps of the proposed kq-space under-sampling method is the motion-induced phase estimation. We considered the phase estimated from the low resolution image associated with a fully sampled central k-space region. However, the use of more sophisticated methods to calibrate the motion-induced phase might further improve the performances of the kq-space approach and will be the focus of future investigations.

Finally, future work should address the validation of our approach when additional microstructure parameters are considered, where the necessity for accelerating the acquisition is even stronger. The proposed approach generalizes naturally to this context, with the main difference lying in the definition of a larger and more general dictionary Φ, accounting not only for fiber orientation, but also the diameter, etc.

References

1. Jara, H. *Theory of Quantitative Magnetic Resonance Imaging*; World Scientific: Singapore, 2013.
2. Le Bihan, D. Looking into the functional architecture of the brain with diffusion MRI. *Nat. Rev. Neurosci.* **2003**, *4*, 469–480. [CrossRef] [PubMed]
3. Sporns, O.; Tononi, G.; Kötter, R. The human connectome: A structural description of the human brain. *PLoS Comput. Biol.* **2005**, *1*, 0245–0251. [CrossRef] [PubMed]
4. Zhang, Y.; Schuff, N.; Du, A.T.; Rosen, H.J.; Kramer, J.H.; Gorno-Tempini, M.L.; Miller, B.L.; Weiner, M.W. White matter damage in frontotemporal dementia and Alzheimer's disease measured by diffusion MRI. *Brain* **2009**, *132*, 2579–2592. [CrossRef] [PubMed]
5. Park, H.J.; Westin, C.F.; Kubicki, M.; Maier, S.E.; Niznikiewicz, M.; Baer, A.; Frumin, M.; Kikinis, R.; Jolesz, F.A.; McCarley, R.W.; et al. White matter hemisphere asymmetries in healthy subjects and in schizophrenia: A diffusion tensor MRI study. *NeuroImage* **2004**, *23*, 213–223. [CrossRef]
6. Callaghan, P.T.; Eccles, C.D.; Xia, Y. NMR microscopy of dynamic displacements: K-space and q-space imaging. *J. Phys. E Sci. Instrum.* **2000**, *21*, 820–822. [CrossRef]
7. Basser, P.J.; Mattiello, J.; Le Bihan, D. MR Diffusion Tensor Spectroscopy and Imaging. *Biophys. J.* **1994**, *66*, 259–267. [CrossRef]
8. Tuch, D.S.; Reese, T.G.; Wiegell, M.R.; Makris, N.; Belliveau, J.W.; Wedeen, V.J. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn. Reson. Med.* **2002**, *48*, 577–582. [CrossRef] [PubMed]
9. Wedeen, V.J.; Hagmann, P.; Tseng, W.Y.I.; Reese, T.G.; Weisskoff, R.M. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magn. Reson. Med.* **2005**, *54*, 1377–1386. [CrossRef] [PubMed]
10. Tuch, D.S. Q-ball imaging. *Magn. Reson. Med.* **2004**, *52*, 1358–1372. [CrossRef]
11. Zhang, H.; Hubbard, P.L.; Parker, G.J.M.; Alexander, D.C. Axon diameter mapping in the presence of orientation dispersion with diffusion MRI. *NeuroImage* **2011**, *56*, 1301–1315. [CrossRef]
12. Zhang, H.; Schneider, T.; Wheeler-Kingshott, C.A.; Alexander, D.C. NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage* **2012**, *61*, 1000–1016. [CrossRef]
13. Daducci, A.; Canales-Rodriguez, E.J.; Descoteaux, M.; Garyfallidis, E.; Gur, Y.; Lin, Y.C.; Mani, M.; Merlet, S.; Paquette, M.; Ramirez-Manzanares, A.; et al. Quantitative comparison of reconstruction methods for intra-voxel fiber recovery from diffusion MRI. *IEEE Trans. Med. Imaging* **2014**, *33*, 384–399. [CrossRef]
14. Ning, L.; Laun, F.; Gur, Y.; DiBella, E.V.; Deslauriers-Gauthier, S.; Megherbi, T.; Ghosh, A.; Zucchelli, M.; Menegaz, G.; Fick, R.; et al. Sparse Reconstruction Challenge for diffusion MRI: Validation on a physical phantom to determine which acquisition scheme and analysis method to use? *Med. Image Anal.* **2015**, *26*, 316–331. [CrossRef] [PubMed]
15. Tournier, J.D.; Calamante, F.; Gadian, D.G.; Connelly, A. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage* **2004**, *23*, 1176–1185. [CrossRef]
16. Alexander, D.C. Maximum entropy spherical deconvolution for diffusion MRI. *Inf. Process. Med. Imaging* **2005**, *19*, 76–87.
17. Acqua, F.D.; Rizzo, G.; Scifo, P.; Clarke, R.A.; Scotti, G.; Fazio, F. A Model-Based Deconvolution Approach to Solve Fiber Crossing in Diffusion-Weighted MR Imaging. *IEEE Trans. Biomed. Eng.* **2007**, *54*, 462–472. [CrossRef] [PubMed]
18. Jian, B.; Vemuri, B.C. A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. *IEEE Trans. Med. Imaging* **2007**, *26*, 1464–1471. [CrossRef]
19. Ramirez-Manzanares, A.; Rivera, M.; Vemuri, B.C.; Carney, P.; Mareci, T. Diffusion basis functions decomposition for estimating white matter intravoxel fiber geometry. *IEEE Trans. Med. Imaging* **2007**, *26*, 1091–1102. [CrossRef]
20. Tristán-Vega, A.; Westin, C.F. Probabilistic ODF estimation from reduced HARDI data with sparse regularization. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*; Springer: Berlin/Heidelberg, Germany, 2011; Volume 6892 LNCS, pp. 182–190.
21. Tournier, J.D.; Calamante, F.; Connelly, A. Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* **2007**, *35*, 1459–1472. [CrossRef] [PubMed]
22. Daducci, A.; Van De Ville, D.; Thiran, J.P.; Wiaux, Y. Sparse regularization for fiber ODF reconstruction: From the suboptimality of ℓ_2 and ℓ_1 priors to ℓ_0. *Med. Image Anal.* **2014**, *18*, 820–833. [CrossRef]
23. Landman, B.A.; Bogovic, J.A.; Wan, H.; ElShahaby, F.E.Z.; Bazin, P.L.; Prince, J.L. Resolution of crossing fibers with constrained compressed sensing using diffusion tensor MRI. *NeuroImage* **2012**, *59*, 2175–2186. [CrossRef] [PubMed]
24. Auría, A.; Daducci, A.; Thiran, J.P.; Wiaux, Y. Structured sparsity for spatially coherent fibre orientation estimation in diffusion MRI. *NeuroImage* **2015**, *115*, 245–255. [CrossRef]
25. Michailovich, O.; Rathi, Y.; Dolui, S. Spatially Regularized Compressed Sensing for High Angular Resolution Diffusion Imaging. *IEEE Trans. Med. Imaging* **2011**, *30*, 1100–1115. [CrossRef]
26. 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.* **2015**, *73*, 126–138. [CrossRef] [PubMed]
27. Donoho, D.L. Compressed sensing. *IEEE Trans. Inf. Theory* **2006**, *52*, 1289–1306. [CrossRef]
28. Candès, E.J. Compressive sampling. In Proceedings of the International Congress of Mathematicians, Madrid, Spain, 22–30 August 2006; pp. 1433–1452.
29. Candes, E.J.; Wakin, M.B.; Boyd, S.P. Enhancing sparsity by reweighted L1 minimisation. *J. Fourier Anal. Appl.* **2008**, *14*, 877–905. [CrossRef]

30. Calabrese, E.; Badea, A.; Coe, C.L.; Lubach, G.R.; Styner, M.A.; Johnson, G.A. Investigating the tradeoffs between spatial resolution and diffusion sampling for brain mapping with diffusion tractography: Time well spent? *Hum. Brain Mapp.* **2014**, *35*, 5667–5685. [CrossRef]

31. Zhan, L.; Jahanshad, N.; Ennis, D.B.; Jin, Y.; Bernstein, M.A.; Borowski, B.J.; Jack, C.R.; Toga, A.W.; Leow, A.D.; Thompson, P.M. Angular Versus Spatial Resolution Trade-Offs for Diffusion Imaging Under Time Constraints. *Hum. Brain Mapp.* **2013**, *32*, 2688–2706. [CrossRef]

32. Vos, S.B.; Aksoy, M.; Han, Z.; Holdsworth, S.J.; Maclaren, J.; Viergever, M.A.; Leemans, A.; Bammer, R. NeuroImage Trade-off between angular and spatial resolutions in vivo fiber tractography. *NeuroImage* **2016**, *129*, 117–132. [CrossRef] [PubMed]

33. Scherrer, B.; Afacan, O.; Taquet, M.; Prabhu, S.P.; Gholipour, A.; Warfield, K. Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisition. *HHS Pubblic Access* **2012**, *16* , 1465–1476. [CrossRef] [PubMed]

34. 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. [CrossRef] [PubMed]

35. Haldar, J.P.; Wedeen, V.J.; Nezamzadeh, M.; Dai, G.; Weiner, M.W.; Schuff, N.; Liang, Z.P. Improved diffusion imaging through SNR-enhancing joint reconstruction. *Magn. Reson. Med.* **2013**, *69*, 277–289. [CrossRef]

36. Jeong, E.K.; Kim, S.E.; Parker, D.L. 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. [CrossRef]

37. Cheng, J.; Shen, D.; Basser, P.J.; Yap, P.T. Joint 6D k-q Space Compressed Sensing for Accelerated High Angular Resolution Diffusion MRI. In *International Conference on Information Processing in Medical Imaging*; Springer: Cham, Switzerland, 2015; Volume 9123, pp. 782–793.

38. Mani, M.M.M.; Jacob, M. qModeL: A plug-and-play model-based reconstruction for highly accelerated multi-shot diffusion MRI using learned priors. *Magn. Reson. Med.* **2021**, *86*, 835–851. [CrossRef]

39. Ramos-Llordén, G.; Ning, L.; Liao, R.M.C.; Michailovich, O.; Setsompop, Y.R.K. High-fidelity, accelerated whole-brain submil-limeter in vivo. *Magn. Reson. Med.* **2020**, *84*, 1781–1795. [CrossRef]

40. Sun, J.; Sakhaee, E.; Entezari, A.; Vemuri, B.C. Leveraging EAP-Sparsity for Compressed Sensing of MS-HARDI in (k,q)-Space. In *International Conference on Information Processing in Medical Imaging*; Springer: Cham, Switzerland, 2015; Volume 24, pp. 375–386.

41. Awate, S.P.; Dibella, V.R.E. Compressed sensing HARDI via rotation-invariant concise dictionaries, flexible k-space under-sampling, and multiscale spatial regularity. In Proceedings of the 2013 IEEE 10th International Symposium on Biomedical Imaging, San Francisco, CA, USA, 7–11 April 2013; pp. 9–12. [CrossRef]

42. Combettes, P.L.; Pesquet, J.C. Stochastic forward-backward and primal-dual approximation algorithms with application to online image restoration. In Proceedings of the 2016 24th European Signal Processing Conference (EUSIPCO), Budapest, Hungary, 29 August–2 September 2016; Volume 1, pp. 1813–1817.

43. Gudbajartsson, S.P.H. The Rician Distribution of Noisy MRI Data. *Changes* **1995**, *29*, 997–1003. [CrossRef]

44. Henkelman, R.M. Measurement of signal intensities in the presence of noise in MR images. *Med. Phys.* **1985**, *12*, 232–233. [CrossRef]

45. Chen, S.S.; Donoho, D.L.; Saunders, M.A. Atomic Decomposition by Basis Pursuit. *SIAM J. Sci. Comput.* **2001**, *43*, 129–159. [CrossRef]

46. Donoho, D.L. De-Noising by Soft-Thresholding. *IEEE Trans. Inf. Theory* **1995**, *41*, 613–627. [CrossRef]

47. Ochs, P.; Dosovitskiy, A.; Brox, T.; Pock, T. On iteratively reweighted algorithms for non-smooth nonconvex optimization in computer vision. *SIAM J. Imaging Sci.* **2015**, *8*, 331–372. [CrossRef]

48. Geiping, J.; Moeller, M. Composite optimization by nonconvex majorization-minimization. *SIAM J. Imaging Sci.* **2018**, *11*, 2494–2598. [CrossRef]

49. Ochs, P.; Fadili, J.; Brox, T. Non-smooth non-convex Bregman minimization: Unification and new algorithms. *J. Optim. Theory Appl.* **2019**, *181*, 244–278. [CrossRef]

50. Repetti, A.; Wiaux, Y. Variable Metric Forward-Backward Algorithm for Composite Minimization Problems. *SIAM J. Optim.* **2021**, *31*, 1215–1241. [CrossRef]

51. Combettes, P.L.; Wajs, V.R. Signal Recovery by Proximal Forward-Backward Splitting. *Multiscale Model. Simul.* **2005**, *4*, 1168–1200. [CrossRef]

52. Caruyer, E.; Lenglet, C.; Sapiro, G.; Deriche, R. Design of multishell sampling schemes with uniform coverage in diffusion MRI. *Magn. Reson. Med.* **2013**, *69*, 1534–1540. [CrossRef]

53. Miller, K.L.; Pauly, J.M. Nonlinear phase correction for navigated diffusion imaging. *Magn. Reson. Med.* **2003**, *50*, 343–353. [CrossRef] [PubMed]

54. Pipe, J.G.; Farthing, V.G.; Forbes, K.P. Multishot diffusion-weighted FSE using PROPELLER MRI. *Magn. Reson. Med.* **2002**, *47*, 42–52. [CrossRef] [PubMed]

55. Roemer, P.B.; Edelstein, W.A.; Hayes, C.E.; Souza, S.P.; Muller, O.M. The NMR Phase Array. *Magn. Reson. Med.* **1990**, *16*, 192–225. [CrossRef]

56. Maier-Hein, K.H.; Neher, P.F.; Houde, J.-C.; Côté, M.-A.; Garyfallidis, E.; Zhong, J.; Chamberland, M.; Yeh, F.-C.; Lin, Y.-C.; Ji, Q.; et al. The challenge of mapping the human connectome based on diffusion tractography. *Nat. Commun.* **2017**, *8*, 1–23. [CrossRef] [PubMed]

57. Komodakis, N.; Pesquet, J.C. Playing with Duality: An Overview of Recent Primal-Dual Approaches for Solving Large-Scale Optimization Problems. *IEEE Signal Process. Mag.* **2015**, *32*, 31–54. [CrossRef]
58. Beck, A.; Teboulle, M. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM J. Imaging Sci.* **2009**, *2*, 183–202. [CrossRef]