COMMIT: Convex Optimization Modeling for Microstructure Informed Tractography

Alessandro Daducci*, Alessandro Dal Palù, Alia Lemkaddem, Jean-Philippe Thiran, *Senior Member, IEEE*

Abstract—Tractography is a class of algorithms aiming at in-vivo mapping the major neuronal pathways in the white matter from diffusion MRI data. These techniques offer a powerful tool to noninvasively investigate at the macroscopic scale the architecture of the neuronal connections of the brain. However, unfortunately, the reconstructions recovered with existing tractography algorithms are not really quantitative even though diffusion MRI is a quantitative modality by nature. As a matter of fact, several techniques have been proposed in recent years to estimate, at the voxel level, intrinsic microstructural features of the tissue, such as axonal density and diameter, by using multicompartment models. In this article, we present a novel framework to reestablish the link between tractography and tissue microstructure. Starting from an input set of candidate fiber-tracts, which are estimated from the data using standard fiber-tracking techniques, we model the diffusion MRI signal in each voxel of the image as a linear combination of the restricted and hindered contributions generated in every location of the brain by these candidate tracts. Then, we seek for the global weight of each of them, i.e. the effective contribution or volume, such that they globally fit the measured signal at best. We demonstrate that these weights can be easily recovered by solving a global convex optimization problem and using efficient algorithms. The effectiveness of our approach has been evaluated both on a realistic phantom with known ground-truth and in-vivo brain data. Results clearly demonstrate the benefits of the proposed formulation, opening new perspectives for a more quantitative and biologically-plausible assessment of the structural connectivity of the brain.

Index Terms—Diffusion MRI, global tractography, tissue microstructure, convex optimization.

I. INTRODUCTION

DIFFUSION MRI (dMRI) is a powerful imaging modality capable of inferring the local axonal structure in each imaging voxel by exploiting the natural random movement of water molecules in biological tissues [1], [2]. This ability is particularly effective for studying the connectivity of the brain, as it permits to noninvasively estimate the major neuronal pathways in the white matter (WM) by means of the so called tractography algorithms (also known as fiber-tracking). In the last decade a variety of approaches have been proposed to tackle the huge dimensionality of the problem, but tractography still represents a tough challenge in this field. For a review of existing techniques see, for instance, [3]–[5] and references

A. Daducci and J.-P. Thiran are with the Signal Processing Lab (LTS5), École Polytechnique Fédérale de Lausanne, Switzerland and the University Hospital Center (CHUV) and University of Lausanne (UNIL), Switzerland.

A. Lemkaddem is with the Signal Processing Lab (LTS5), École Polytechnique Fédérale de Lausanne, Switzerland.

A. Dal Palù is with the Department of Mathematics and Computer Science, University of Parma, Italy.

* indicates the corresponding author.

therein. Among them, line-propagation methods [6], [7] are characterized by greedy deterministic algorithms exploiting solely the local information available in the neighborhood of each voxel, usually in the form of the Orientation Distribution Function (ODF) or Fiber Orientation Distribution (FOD) of the diffusion process. They are very fast, but heavily suffer from the propagation of local estimation inaccuracies along the path. Probabilistic methods [8], [9] introduced the notion of uncertainty in the tracking by repeatedly seeding deterministic tractography in a Monte Carlo-like fashion, but they are very time consuming and they only partially solve the intrinsic issues of these local formulations. To overcome the local nature of previous approaches, front-evolution methods have been introduced [10], [11]. The rationale behind these techniques is to consider the local diffusion profiles as a speed function controlling the propagation of a front from a seed point and to interpret the path with the minimal arrival time between two brain regions (also known as geodesic) as the fiber-tract which is globally optimal. These methods are computationally efficient and very robust to noise. However, they recover the "optimal path" between two given brain locations one by one, disregarding the interactions with all the other "optimal paths" connecting other regions. Recently, global energy minimization approaches [12], [13] have been proposed with the aim to reconstruct at once the full tractogram, i.e. set of fiber-tracts, over the whole brain. Results obtained with these global algorithms outperform any other previous method [3], but the price to pay is a significant increase in the computational burden which is normally not suitable in a clinical perspective. Furthermore, these methods are based on stochastic optimization procedures and consequently they do not guarantee convergence to the global optimal solution.

Unfortunately, despite dMRI actually being a quantitative modality by nature, the tractograms recovered by existing tractography algorithms are not really quantitative [4]. In fact, the structural connectivity between different brain regions is usually quantified by counting the number of recovered pathways or averaging some scalar maps along them [14]. In both cases, these estimates provide only indirect measures of the true underlying neuronal microstructure and are not truly quantitative. On the other hand, a large number of techniques have been recently proposed in the literature to infer, at the voxel level, microstructural features of the neuronal tissue from dMRI data, such as axonal diameter and density. Generally these techniques assume that the tissue is composed of several different "compartments", e.g. axons, glial cells and extra-axonal space, and that the dMRI signal measured in each voxel can be explained in terms of the unique diffusion pattern

of each of them. The following overview is not exhaustive and many other approaches have been proposed; for a comprehensive list, please see [15], [16] and references therein. A very simple and popular technique that distinguishes different compartments is the Ball&Stick model [9]. According to this model, the dMRI signal arises from water molecules diffusing either inside a hypothetical cylinder with zero radius, i.e. the "stick", or around it with isotropic pattern, i.e. the "ball". Although simplistic, this technique allows estimating the relative volume fraction of the fiber compartment other than merely the orientation. The composite hindered and restricted model of diffusion (CHARMED) [17] extended this technique by considering a more complex multicompartment model to yield an accurate physical description of the two water pools. The model assumes parallel cylinders with a fixed radius, with no exchange between compartments, and the water molecules can be either hindered in the extra-axonal space or restricted in the intra-axonal space. In the former case the diffusion process follows a Gaussian distribution and the signal decay is represented by an anisotropic tensor, while in the latter case diffusion is non-Gaussian and the attenuation is modeled by means of exact analytical expressions for particles diffusing in a restricted cylindrical geometry [18].

The method was later improved in [19] by letting the axon radius to follow a Gamma distribution rather than a fixed value. This enhancement to the model made it possible to characterize the size of the axons from dMRI data, but the imaging protocol was not compatible with a clinical setting. The minimal model of white matter diffusion (MMWMD) [20] opened the way for the estimation of orientationally-invariant indices of axonal diameter with clinically feasible acquisitions; tissue modeling was further refined using four compartments and the multishell protocol was optimized to maximize the sensitivity to the parameters of the model. Lastly, the neurite orientation dispersion and density imaging (NODDI) technique [21] relaxed the assumption of parallel cylinders by using Watson distributions and further optimized the acquisition protocol to allow both the axonal density and orientation dispersion to be estimated on clinical scanners. All these approaches represent an effective and practical opportunity for characterizing in-vivo the microstructural organization of neuronal tissue with dMRI. However, they require either long acquisitions or computationally expensive fitting procedures to obtain robust estimates and are thus not suited for clinical applications or large cohorts of subjects. Also, as these techniques are rather recent, further validation studies will be needed before they are routinely adopted in clinical practice.

To date, fiber-tracking and tissue microstructure estimation have been considered as two separate problems. However, recent studies [22]–[26] have demonstrated the feasibility and the potential benefits of combining tractography with local microstructural features for mapping the connectivity. The common denominator for these methods is the estimation of a large set of candidate fiber-tracts using classical fiber-tracking algorithms, e.g streamline, followed by a selection of a subset that best fits the acquired dMRI signal (Fig. 1). They tackle the problem in a *top-down* fashion, whereas all previous tractography approaches were characterized by
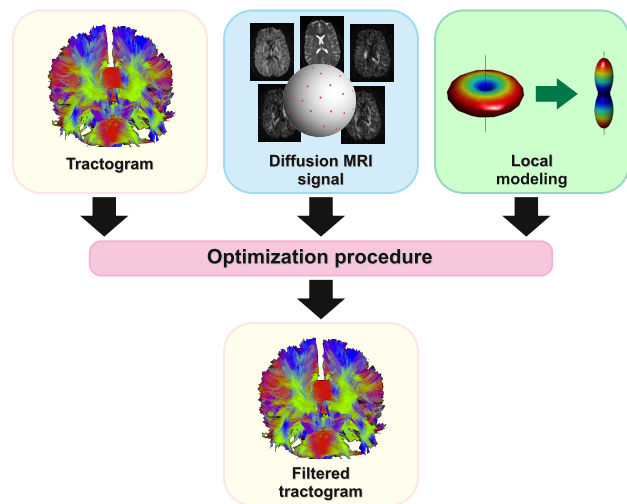

Fig. 1. Block diagram of *top-down strategies* to combine tractography reconstructions with local properties of the tissue.

bottom-up strategies. In [22], a BlueGene/L supercomputer with 2048 cores and 500 GB of RAM was employed to search for the optimal combination of fiber-tracts from a huge set of candidates. However, the procedure was too onerous (9 days to process a single brain dataset) to be of real interest. The framework was optimized in MicroTrack [23] and a multicompartment model similar to [20] was adopted to filter out false-positive candidates, i.e. fibers that are not anatomically plausible. The model assumes that the microstructural properties of the tracts, such as axon radius and myelin volume, remain constant along the trajectory. The signal in each voxel is modeled as a combination of the intra-axonal contributions from all the pathways crossing the voxel and further extra-axonal and isotropic compartments considered locally. The results clearly showed the benefits of combining tractography and microstructure, but the stochastic algorithm employed was still computationally heavy (16 hours on 30 2.3 GHz CPUs) and, most importantly, the complexity of the formulation did not guarantee to converge to the optimal solution. Simplified methods that are similar in essence to MicroTrack have been proposed lately. In [24], a particle filter mechanism is employed to improve tractography by exploiting microstructural properties of the tissue, i.e. axon dispersion computed using NODDI [21], during the propagation of the streamlines. The split-and-merge tractography (SMT) [25] splits the tracts at unreliable points according to a reliability measure and produces clusters of short tracts as output, hence providing a tool to explore a tractogram rather than an actual tractography method. The spherical-deconvolution informed filtering of tractograms (SIFT) method [26] uses the FOD reconstructed with constrained spherical deconvolution (CSD) [27] to choose which candidates have to be removed for better fitting the measured signal. Since the algorithm removes candidates iteratively, there is no guarantee to recover the global optimal solution, as the procedure can get trapped in the many local minima of this very high-dimensional space. Notably, despite showing reasonable performances (few hours for a single brain), these simplified methods either

do not implement a real global approach or do not actually have access to the tissue microstructure, as they do not use multiple b-values acquisitions and proper tissue models that are required to be sensitive to such features [15], [20], [21].

In this work, we present a flexible and efficient formulation to reestablish the link between tractography and tissue microstructure. We propose to reformulate tractography in the *framework of convex optimization* with the aim to (i) improve the quality of the reconstructed tractograms by combining them with microstructural properties of the tissue, (ii) reduce the computational cost to accommodate real application demands and (iii) guarantee to recover the global optimal solution. The framework proposed here is an improved and extended version of our work presented in [28]. Our approach is close in spirit to MicroTrack [23], for both use multi-compartment models and global optimization techniques to combine fiber-tracking with microstructure tissue parameters. Nevertheless, our convex formulation drastically reduces the computational complexity hence opening the way for *quantitative tractography* to become practical. To the best of our knowledge, this is the first attempt to use convex optimization for tractography and microstructure estimation. Source code is available at https://github.com/daducci/COMMIT/. The rest of the paper is organized as follows. The proposed formulation is presented in detail in Section II. Results obtained on both realistic phantom and in-vivo data are reported in Section III, followed by a discussion on benefits and limitations of our approach. We conclude with future perspectives and potential implications of this novel formulation.

II. A CONVEX FORMULATION FOR TRACTOGRAPHY

Given a tractogram $\mathcal{F}$, the corresponding measured dMRI image $\boldsymbol{I} \in \mathbb{R}_+^{n_x \times n_y \times n_z \times n_d}$, composed of n_d q-space samples acquired over $n_v = n_x n_y n_z$ voxels, can be modeled as $\boldsymbol{I} = \mathcal{A}(\mathcal{F}) + \eta$, where $\mathcal{A} : \mathcal{F} \rightarrow \boldsymbol{I}$ is an operator modeling the signal contribution of each fiber in all imaging voxels and η is the acquisition noise. Its *inverse formulation* (i.e. given an input dMRI image $\boldsymbol{I}$, finding the set of fibers $\tilde{\mathcal{F}}$ that best describes the data) represents the ultimate goal of every tractography algorithm. Our approach to solve this inverse problem consists of two steps. First, the candidate pathways are estimated from the data using classical fiber-tracking techniques, with the only requisite that these candidates represent a valid superset of the anatomically plausible tracts, i.e. true positives, possibly including also many false positives. Then, we seek for the weight of each of them, i.e. real contribution or volume, by solving a global convex optimization problem that exploits multicompartment models to explain the measured dMRI signal at best. For this reason, we named our framework COMMIT, acronym for *Convex Optimization Modeling for Microstructure Informed Tractography*.

A. Estimation of the candidate tracts

Any tractography algorithm, or combination of them, can be used for the estimation of the candidates. The study of the most appropriate approach to construct this initial set of tracts, i.e. local reconstruction models and tracking

methods, goes beyond the scope of this paper and will be the subject of future research. Yet, to show the flexibility of our framework, in our experiments we tested COMMIT with different inputs; in particular, and without loss of generality, we used a classical line-propagation method based on [6], a front-evolution algorithm similar to [10] and a global approach [13]. In the case of the first two algorithms, we employed the FOD as local diffusion model for the propagation of the tracts in all our experiments, computed by means of the CSD technique [27] using MRtrix[1] and default settings. The global approach of [13] does not construct the tracts by propagation, but all the trajectories are estimated at once by finding the configuration that best describes the measured data, using the Stick [15] as local model. Finally, a binary mask was used as support for all the methods (both for seeding and constraining the tracts in the WM) and the tracking was performed using standard parameters as found in the literature, such as maximum curvature and stopping criteria.

B. Local forward-model

Once the pathways have been estimated with any given tractography algorithm, the dMRI signal contribution of every tract must be mapped to each voxel of the image. To this aim, we adopt a *multicompartment model* similar to [23] to characterize the neuronal tissue, accounting for both restricted (intra-axonal) and hindered (extra-axonal) water pools, as well as partial volume with isotropic diffusion, e.g. CSF. Computationally expensive nonlinear procedures are normally required to fit such models; however, when the response functions are known (or can be determined) a priori, these models can be efficiently estimated by means of systems of linear equations [27], [29]. Hence, in a voxel, the predicted signal $S(\boldsymbol{q})$ at q-space location $\boldsymbol{q} \in \mathbb{R}^3$ can be expressed as a *linear combination* with the following general formulation:

$$S(\boldsymbol{q}) = \underbrace{f^{\text{IC}} R^{\text{IC}}(\boldsymbol{q})}_{\text{restricted}} + \underbrace{f^{\text{EC}} R^{\text{EC}}(\boldsymbol{q})}_{\text{hindered}} + \underbrace{f^{\text{ISO}} R^{\text{ISO}}(\boldsymbol{q})}_{\text{isotropic}}, \qquad (1)$$

where R^{IC}, R^{EC}, R^{ISO} are the predicted signal profiles of the compartments and f^{IC}, f^{EC}, f^{ISO} the corresponding volumes. R^{IC}, R^{EC} and R^{ISO} can be either estimated from the data, as in [27], or synthetically generated with analytic models, as in [15]. Noteworthy, the classical CSD [27] can be seen as a special case of (1) where only one compartment is considered.

Besides, in this context, the restricted contribution of a voxel is given by all the tracts traversing the voxel itself and the hindered space around them also depends on their organization. Thus, Eq. (1) we be rewritten as follows:

$$S(\boldsymbol{q}) = \sum_{\mathcal{F}_i \in \mathcal{F}} \boldsymbol{f}_i^{\text{IC}} R_{(i)}^{\text{IC}}(\boldsymbol{q}) + \sum_{\mathcal{F}_i \in \mathcal{F}} \boldsymbol{f}_i^{\text{EC}} R_{(i)}^{\text{EC}}(\boldsymbol{q}) + f^{\text{ISO}} R^{\text{ISO}}(\boldsymbol{q}),$$
$$(2)$$

where $\boldsymbol{f}_i^{\text{IC}}$ is the *global contribution* of the restricted diffusion arising from fiber $\mathcal{F}_i \in \mathcal{F}$, $\boldsymbol{f}_i^{\text{EC}}$ accounts for the hindered space around the axons in the direction of $\mathcal{F}_i$ and the summation is done over all the candidates. Each $R_{(i)}^{\text{IC}}$ represents a rotated version of the response function R^{IC} to match the local

---

[1] Available online: www.brain.org.au/software/mrtrix

Fig. 2. *The* `COMMIT` *model.* When both the fiber pathways and the response functions of different tissue compartments can be estimated a priori, tractography can be expressed as a convex optimization problem. The mapping $\mathcal{A} : \mathcal{F} \rightarrow \boldsymbol{I}$ between the input fibers $\mathcal{F}$ (controlled by the coefficients $\boldsymbol{x}$) and the dMRI image $\boldsymbol{I}$ (stored in the vector $\boldsymbol{y}$) can be implemented as a linear operator $\boldsymbol{A}$ by a *natural extension of classical deconvolution approaches to the domain of fiber-tracts*. The signal in each voxel is given by a linear combination of the restricted diffusion arising from all the fibers intersecting the voxel, possibly in addition to local extra-axonal and isotropic contributions at the voxel level.

orientation of $\mathcal{F}_i$ and it is scaled by the actual length (in mm) of the portion of $\mathcal{F}_i$ intersecting the voxel, i.e. its intravoxel length. Analogously for the $R_{(i)}^{\text{EC}}$ response functions. Clearly, if a fiber does not cross a voxel there is no signal contribution from it in that voxel.

C. The `COMMIT` model

In this context, both the trajectories of the fibers and the response functions of the compartments are known or can be estimated a priori. Consequently, the mapping $\mathcal{A} : \mathcal{F} \rightarrow \boldsymbol{I}$ can be implemented as a *linear operator* by a natural extension of Eq. (2) to the space of fiber-tracts which accounts for all voxels of the image $\boldsymbol{I}$, thus allowing us to express tractography as a convex optimization problem. A schematic representation of our formulation is illustrated in Fig. 2. The observation model can be written in matrix form as

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \eta, \qquad (3)$$

where $\boldsymbol{y} \in \mathbb{R}_+^{n_d n_v}$ is the vector containing the n_d q-space samples acquired in all n_v voxels, η accounts for both acquisition noise and modeling errors, $\boldsymbol{A} \in \mathbb{R}^{n_d n_v \times n_c}$ is the observation matrix (i.e. dictionary) modeling explicitly the multicompartment model of Eq. (2) in every voxel and the positive weights $\boldsymbol{x} \in \mathbb{R}_+^{n_c}$ are the contributions of the n_c basis functions in $\boldsymbol{A}$ (i.e. atoms). Please note that even if the size of $\boldsymbol{A}$ may appear prohibitive, in practice it contains many zeros as each fiber-tract traverses only a very small portion of the voxels of the image.

The linear operator $\boldsymbol{A}$ is a *block matrix*:

$$\boldsymbol{A} = \left[\begin{array}{c|c|c} \boldsymbol{A}^{\text{IC}} & \boldsymbol{A}^{\text{EC}} & \boldsymbol{A}^{\text{ISO}} \end{array} \right], \qquad (4)$$

in which $n_c = n_r + n_h + n_i$ and the three sub-matrices $\boldsymbol{A}^{\text{IC}} \in \mathbb{R}^{n_d n_v \times n_r}$, $\boldsymbol{A}^{\text{EC}} \in \mathbb{R}^{n_d n_v \times n_h}$ and $\boldsymbol{A}^{\text{ISO}} \in \mathbb{R}^{n_d n_v \times n_i}$ encode, respectively, the n_r restricted, n_h hindered and n_i isotropic contributions to the image, defined as follows.

1) Restricted: The matrix $\boldsymbol{A}^{\text{IC}}$ has one column (or more) for each candidate fiber $\mathcal{F}_i \in \mathcal{F}$ and its rows correspond to the restricted signal contribution of all the fibers $\mathcal{F}$ for a given voxel and diffusion gradient. As in previous approaches [22], [23], [26], we assume as well in our formulation that, at the scale measurable with dMRI, the microstructural properties of the tracts remain constant along their trajectories. The restricted contribution in every voxel due to fiber $\mathcal{F}_i$ is computed by rotating the response function $R^{\text{IC}} \in \mathbb{R}^{n_d}$ to match the local orientation of $\mathcal{F}_i$. In addition, this contribution is scaled by the relative length of fiber $\mathcal{F}_i$ in each voxel; no signal contribution is generated if a fiber does not traverse a voxel. Finally, multiple contributions can be specified for a fiber, e.g. considering each pathway as consisting of distinct populations of axons with different diameters; in that case, multiple columns in $\boldsymbol{A}^{\text{IC}}$ will be associated with each fiber.

2) Hindered: Several strategies are possible to model hindered diffusion in a voxel as, quoting [15], the effect of multiple fiber populations on the extra-axonal space is not very clear. For the sake of generality, in this work we allow the possibility to specify multiple anisotropic contributions in every voxel. Yet, to avoid the redundancy caused by the fact that multiple pathways can follow locally the same direction, only one extra-axonal compartment is considered for every unique fiber population in each voxel $v \in \{1, \ldots, n_v\}$. In particular, every atom in $\boldsymbol{A}^{\text{EC}}$ is associated with one of the principal diffusion directions estimated in a given voxel by

means of any local reconstruction technique, CSD in this work. It consists of the anisotropic response function $R^{\text{EC}} \in \mathbb{R}^{n_d}$ reoriented to match the corresponding unique fiber population in the voxel; all remaining $(n_v - 1)n_d$ elements of the column are null.

3) Isotropic restriction: An independent contribution is added to each voxel to model any partial volume with brain and non-brain tissues that are characterized by isotropic restriction. Each atom in $\boldsymbol{A}^{\text{ISO}}$ contributes to a single voxel $v \in \{1, \dots, n_v\}$ with the predicted signal $R^{\text{ISO}} \in \mathbb{R}^{n_d}$; all remaining $(n_v - 1)n_d$ elements of the column are null.

D. Model fitting

The weights $\boldsymbol{x}$ can be estimated by solving the following *non-negative least-squares* (NNLS) problem:

$$\operatorname*{argmin}_{\boldsymbol{x} \geq 0} ||\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}||_2^2, \qquad (5)$$

where $|| \cdot ||_2$ is the usual ℓ_2 norm in $\mathbb{R}^n$. Unfortunately, the size of $\boldsymbol{A}$ prevents to solve (5) using direct methods, as the computation of the corresponding pseudo-inverse or Cholesky decomposition is intractable. However, several approaches have been developed to iteratively find the least-squares solution of such large-scale problems [30], [31].

In many practical situations the system of equations (3) might be under-determined, i.e. more unknowns than measurements; in this context, that condition can arise if the number of q-space measurements in the input image is small (under-sampling) or the set of candidate fibers is very large. In these situations the inverse problem in (5) does not have, in general, a unique solution. However, if the coefficients to be recovered are known to be sparse, the use of ℓ_1-regularization is a very popular and effective choice to find a solution to these ill-posed problems [32], [33]. A solution to (3) can then be recovered by solving, for instance, the following ℓ_1-*minimization* problem:

$$\operatorname*{argmin}_{\boldsymbol{x} \geq 0} ||\boldsymbol{x}||_1 \text{ subject to } ||\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}||_2 \leq \epsilon, \qquad (6)$$

where the $|| \cdot ||_1$ is the usual ℓ_1 norm in $\mathbb{R}^n$, which is used to promote sparsity in the solution $\boldsymbol{x}$, and the parameter ϵ is a bound on the noise level and modeling errors. This formulation is known as *basis pursuit de-noise* (BPDN). In the following experiments we will use the most appropriate formulation for each specific experimental condition.

E. Implementation details

In this work we have used off-the-shelf software that is publicly available to solve the inverse problems (5) and (6); specifically, we used the `sbb`[2] [34] algorithm to solve NNLS problems and `spgl1`[3] [35] for BPDN. Both solvers are implemented in Matlab and are specifically designed to solve large-scale problems. In particular, they iteratively search for the optimal solution using canonical gradient-projection optimization approaches that rely only on matrix-vector operations, as the inversion of $\boldsymbol{A}$ would be too expensive. This property is of paramount importance in our framework; in fact, the matrix $\boldsymbol{A}$ does not need to be stored explicitly and it can rather be represented by means of sparse data structures. Moreover, as tractography is an intrinsic WM reconstruction technique, only WM voxels must be kept in memory. Our implementation internally uses precomputed lookup-tables for the response functions of the compartments and thus the two operations $\boldsymbol{A}\boldsymbol{x}$ and $\boldsymbol{A}^\dagger \boldsymbol{y}$ can be efficiently calculated at runtime. The response functions have been generated using the tool `datasynth` available in the Camino toolbox[4] which offers a wide gamma of analytic models [15] and different alternatives have been tested in this work; more details will be provided regarding the specific models used in each experiment.

F. Data and experiments

To show the flexibility and effectiveness of our formulation, we tested `COMMIT` in multiple scenarios using both synthetic and in-vivo data experiments. On one hand, we quantitatively assessed its performance using the popular FiberCup phantom [36] as its ground-truth is known. Fig. 3 shows the 12 regions of interest (ROIs) delineating the 7 true bundles and the WM mask used for tracking. For compactness, the 3 orange sub-bundles are shown as one. For the scope of this study, we used the dataset acquired with 64 diffusion directions on a shell at $b = 1500 \, \text{s/mm}^2$ and $3 \, \text{mm}$ isotropic resolution.


Fig. 3.   *The FiberCup data.* Ground-truth, tracking mask and regions of interest (1–12) used throughout the experiments.

On the other hand, we demonstrated the efficacy and the benefits of `COMMIT` also in case of in-vivo human brain data using a publicly-available dataset[5]. This data was acquired from a healthy subject on a 3T MRI system (GE Signa Excite HDx, Milwaukee, WI) with maximum gradient strength $G_{max} = 40 \, \text{mT/m}$ and using the following imaging protocol: 24 images at $b = 700 \, \text{s/mm}^2$ with $G = 23.7 \, \text{mT/m}$, 48 at $b = 2000 \, \text{s/mm}^2$ with $G = 40 \, \text{mT/m}$, 9 with no diffusion weighting (b=0), spatial resolution $1.875 \times 1.875 \times 2.5 \, \text{mm}^3$ and using the same $TR/TE = 12\,400/86.6 \, \text{ms}$ and $\delta/\Delta = 27.7/32.2 \, \text{ms}$ for all images. The data was not normalized to the reference b=0 image. Therefore, all images have been corrected for nonuniform intensity using the N4 algorithm [37]; the bias field was estimated from the average b=0 volume and, subsequently, all diffusion images were corrected using

---

[2]Available online: http://suvrit.de/work/progs/nnls.html
[3]Available online: http://www.cs.ubc.ca/~mpf/spgl1

[4]Available online: www.camino.org.uk
[5]Available online: www.nitrc.org/frs/?group_id=716

Fig. 4. *Qualitative evaluation on the FiberCup phantom.* The tracts reconstructed with our approach (bottom plots) are visually compared to those recovered by the state-of-the-art global tracking algorithm, i.e. GIBBS (top plots). Each true bundle of the ground-truth (Fig. 3) is analyzed separately; the bundle 5 is actually composed of three separate sub-bundles which are here reported as one for compactness. The blue arrows highlight the spurious tracts reconstructed.

this information. Finally, a WM mask was obtained from a high-resolution T1-weighted image using Freesurfer[6].

III. RESULTS AND DISCUSSION

A. Quantitative evaluation on realistic phantom

COMMIT is a framework that combines an existing tractogram with the microstructure features of the tissue and, as such, it cannot be considered as a proper tractography algorithm. Nevertheless, as a first experiment, we investigated whether this framework could be exploited to achieve the tracking quality of advanced global techniques using, instead, simpler and faster algorithms in combination with COMMIT as post-processing. To this aim, we compared the performance of this approach with the state-of-the-art global algorithm of [13], hereafter termed GIBBS, as this latter was ranked first in a recent study evaluating several tractography algorithms [3]. To estimate the candidate tracts we implemented a front-evolution algorithm similar to [10], from now on termed GEODESIC (see Appendix); this choice was motivated by three main reasons. First, this family of methods are very efficient, versatile and robust to noise. Second, the reconstructed pathways satisfy the important anatomical constraint to originate in the gray matter (GM) and develop in the WM. Third, it is virtually possible to recover a unique pathway for any pair of locations by properly tuning the parameters of the algorithm. Finally, in this case the local forward-model consisted of a single anisotropic tensor for the contribution of the fibers, estimated from the data as in [27], no hindered compartment and an independent isotropic contribution at the voxel level.

Fig. 4 presents the results of the *qualitative comparison* of the tractograms reconstructed with GIBBS (top row) and using COMMIT in combination with GEODESIC (bottom row), i.e. GEODESIC + COMMIT. For each of the true fiber bundles

[6]Available online: http://surfer.nmr.mgh.harvard.edu

present in the ground-truth we reported all the reconstructed tracts which originate from one of the ROIs identifying each bundle (reported in yellow in each plot). No post-processing such as filtering or clustering has been done on the resulting tracts. In case of COMMIT, pathways that were assigned a null weight are not displayed. Looking at the plots we can easily appreciate how GIBBS estimates lots of false positives fibers which are not in the ground-truth (pointed by blue arrows). On the contrary, the tractograms reconstructed with the combination GEODESIC + COMMIT look much cleaner and contain less spurious fiber-tracts. This is more evident looking at the plots corresponding to the bundles 3 and 5. We observe that both methods still have troubles disentangling the kissing configuration of the bundles 1 and 3, which might probably require optimized acquisitions or more anatomical priors for their proper characterization as the current data does not allow to separate them. Inspecting carefully the plots, the tractograms reconstructed with our approach may appear more scattered as compared to GIBBS. However, it is important to remember that COMMIT assigns a *global weight* to each tract in the candidate set, i.e. x_i, which is not reported in the figures. As GIBBS implicitly assigns a constant weight to each pathway, the algorithm has to recover multiple instances of the same tract to obtain the same effect and so the tractograms appear denser.

These results indeed demonstrate that, using COMMIT, it is possible to attain the tracking quality of global tractography using simpler non-global algorithms. Concerning the *computational cost* of the two approaches, it is worth noting that, to obtain the tractograms in Fig. 4, GIBBS required about 44 minutes using optimal parameters whereas our approach took less than a minute (≈ 25 seconds), including the estimation of the candidate fibers with GEODESIC and the subsequent optimization stage with COMMIT. All the experiments were conducted on a standard laptop equipped with 8 GB of RAM

without using any kind of parallel computing. Our findings can thus have several implications in a clinical perspective, where faster but less accurate tracking algorithms are usually preferred due to stringent time constraints.

The following *quantitative analysis* might shed some light to understand the reasons for this gap. Table I reports the scores computed using the Tractometer methodology proposed in [38], which is an evaluation system for tractography pipelines with a particular emphasis on global connectivity. Results are reported both for the tracking algorithms alone, i.e. `GIBBS` and `GEODESIC`, as well as when applying `COMMIT` afterwards, i.e. `GIBBS + COMMIT` and `GEODESIC + COMMIT`. As `GIBBS` implements a stochastic process, its scores are reported as an average over a series of 10 runs. All methods successfully recover all 7 Valid Bundles (`VB`); however, a high number of Invalid Bundles (`IB`), a.k.a. false positives that connect unexpected ROIs, are also reconstructed by the two tracking methods: on average 12.5 with `GIBBS` and 6 with `GEODESIC`. On the contrary, the number of false positives is drastically reduced when applying `COMMIT`, respectively 1.7 and 1. In addition, it is worth to observe that of all the tracts recovered by `GIBBS`, 76.5% do not connect two ROIs and are therefore discarded as anatomically impossible; as a consequence, the support is represented only by the remaining 23.5% of the tracts and the risk of false negatives might be rather high. On the other hand, with `GEODESIC` no such No Connections (`NC`) are recovered. This analysis could explain the big gap in the execution times of the two methods, for `GIBBS` putting a lot of effort in constructing lots of fibers that later will have to be discarded. This can be further emphasized by comparing the percentage of Valid Connections (`VC`) and Invalid Connections (`IC`), which are the number of tracts connecting expected and unexpected ROIs, respectively. The combination `GEODESIC + COMMIT` recovers 97.9% of

`VC` and 2.1% of `IC`, while `GIBBS` obtains only 19.8% of `VC` and 3.7% of `IC`. For a fair comparison of the actual performances, though, we ought to normalize the scores by the total number of connecting fibers, i.e. $1 - $ `NC`. By doing this it turns out that `GIBBS` actually recovers 84.3% of `VC` and 15.7% of `IC` and its performances are closer to `COMMIT`. Per contra, our approach requires much less time and, as opposed to stochastic methods, guarantees convergence to the optimal configuration.

	VB (num)	IB (num)	NC (%)	VC (%)	IC (%)	Time (min)
GIBBS	7	12.5	76.5	19.8	3.7	43.6
GIBBS + COMMIT	7	1.7	76.5	21.9	1.6	43.8
GEODESIC	7	6	0	84.2	15.8	0.2
GEODESIC + COMMIT	7	**1**	**0**	**97.9**	**2.1**	0.4

TABLE I
QUANTITATIVE COMPARISON OF `COMMIT` AND `GIBBS` USING THE SCORES PROPOSED IN [38] : VALID BUNDLES (`VB`), INVALID BUNDLES (`IB`), NO CONNECTIONS (`NC`), VALID CONNECTIONS (`VC`), INVALID CONNECTIONS (`IC`). RESULTS FOR `GIBBS` ARE REPORTED AS AN AVERAGE OVER A SERIES OF 10 RUNS. BEST SCORES ARE MARKED IN BOLD.

B. Qualitative evaluation on in-vivo data

In the previous section we have shown that, indeed, the quality of the tractograms can be substantially improved by considering the microstructural properties of the tissue with `COMMIT`. In this section, we evaluated the effectiveness of our framework for microstructure informed tractography also in case of in-vivo data. However, as the ground-truth is unknown, we investigated the *biological plausibility of the tractograms* with respect to the actual underlying tissue structure, before and after optimization with `COMMIT`. In our experiments, we found that the front-evolution algorithm used in the previous
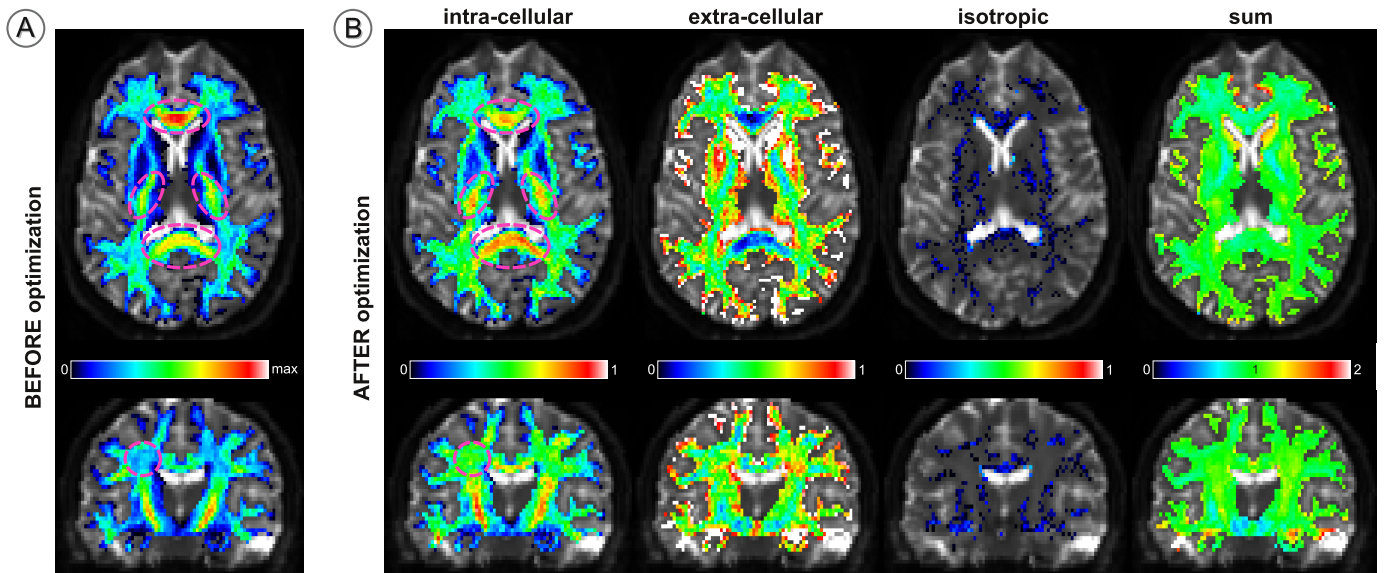

Fig. 5. *Parameter maps estimated using* `COMMIT`. Subplot **A** reports the voxelwise density of tracts recovered with `GIBBS` in two representative slices of the brain before optimization, whereas in **B** we show the maps for the intra-cellular, extra-cellular and isotropic compartments after optimization with `COMMIT` using **A** as input. Interestingly, the sum of these three compartments (rightmost plot) consistently converges to unity in all voxels although it was not imposed in the optimization. The areas highlighted in magenta point to the regions where the improvements with `COMMIT` are more evident.

section does not provide a suitable set of tracts to build the linear operator $\boldsymbol{A}$ also in case of in-vivo data, probably due to the higher complexity of real brain connections; on the other hand, GIBBS seemed to offer a valid support for COMMIT. For this reason, in this first experiment with in-vivo data we have adopted GIBBS to estimate the input set of tracts and we will investigate in detail the effect of using other tracking algorithms in Section III-D. For the local forward-model we considered two restricted contributions for each tract and we used the Gaussian Phase Distribution approximation assuming a longitudinal diffusivity $d_\parallel = 1.7 \times 10^{-3}\,\mathrm{mm}^2/\mathrm{s}$ and two distinct axon radii $\mathrm{R} \in \{0.5, 5\}\,\mu\mathrm{m}$. In addition, in each voxel, we included one hindered contribution for every unique fiber population and we adopted the Zeppelin model assuming a perpendicular diffusivity $d_\perp = 0.5 \times 10^{-3}\,\mathrm{mm}^2/\mathrm{s}$ and same $d_\parallel$. Lastly, two independent compartments with isotropic diffusivity $d \in \{1.7, 3.0\} \times 10^{-3}\,\mathrm{mm}^2/\mathrm{s}$ were added to account for partial volume with GM and CSF, respectively.

In Fig. 5A we computed the density of the tracts in two representative slices of the brain, in a similar way to the track-density imaging (TDI) technique [39], assigning a fixed volume per unit length of the tracts; this is in fact an implicit assumption in all tractography algorithms. Analogously, Fig. 5B reports the intra-cellular, extra-cellular and isotropic maps after computing the actual contributions of all the compartments with COMMIT and then summing the corresponding weights in each voxel, i.e. $\boldsymbol{f}^{\text{IC}}$, $\boldsymbol{f}^{\text{EC}}$ and f^{ISO}. The tracking time with GIBBS was ≈ 12 hours and the optimization time with COMMIT ≈ 13 minutes. Comparing the images before and after optimization, we can notice that the spatial distribution of the raw fiber-tracts does not follow the expected pattern of neuronal density as found in previous studies [40], [41]. On the other hand, the maps estimated after COMMIT indicate a spatial distribution which appears *more biologically plausible* and *in agreement with the known brain anatomy*. In fact, the intra-cellular map shows the highest values in the major WM bundles, e.g. corpus callosum and corticospinal tract highlighted in magenta in upper plots, an homogeneous distribution in crossing areas, e.g. corona radiata indicated in lower plots, and a decreased contribution approaching the GM. Conversely, the extra-cellular map reveals the lowest contributions in areas of highly-packed bundles, increased values at the boundary with GM and an homogeneous pattern elsewhere. Isotropic contributions occur mostly at interfaces with GM and CSF regions, as expected.

Finally, it is worth noting that the sum of the compartments appears to consistently converge to unity (1.001 ± 0.155, mean $\pm$ standard deviation) even though this constraint was not imposed in the optimization. This result represents a *reassuring indication that the model is appropriate* and it leaves, as well, significant room for improvement in the future through the explicit incorporation of such physical constraint in every voxel. Besides, the careful reader might have noticed that this physical property is not perfectly met in few voxels close to GM or CSF. In these regions, inaccuracies in the segmentation of WM can cause fibers to stop prematurely or extend inside GM. In the first case, WM voxels that are not properly covered by the tracts do not possess a valid support

for COMMIT to accurately model the signal in the voxel. Conversely, those GM voxels that are (incorrectly) traversed by the tracts are included in the optimization but COMMIT is unable to fit correctly the model as it is intrinsically a pure WM model.

C. Local vs global fit

As previously stated, nonlinear procedures are normally required to fit multicompartment models to the data and extract microstructure properties of the tissue. Nonlinear approaches offer in fact a high degree of control over the functions to fit and produce very accurate estimates but, per contra, they are computationally very intensive. In this section, we tested the efficacy of using linear methods to fit these models.


Fig. 6. *Local vs global fit.* The intra-cellular volume fraction map (top row) is reported in a representative slice as estimated by: NODDI (left column), COMMIT (right) and a voxel-by-voxel NNLS fit (middle) using the same local forward-model as COMMIT but solving each voxel independently, i.e. without using the regularization provided by the input fiber-tracts. The accuracy of the fit is also reported (bottom row) by means of the NMSE between the measured and estimated dMRI signal in each voxel.

Fig. 6 compares the maps of the intra-cellular compartment as estimated by COMMIT and by a voxel-by-voxel NNLS fit that uses the same local forward-model but solves each voxel independently, i.e. without exploiting the regularization provided by the tracts. For reference, we adopted the voxelwise fitting results from [21] as pseudo ground-truth, hereafter referred as NODDI, as this model has been shown to produce parameter maps consistent with brain anatomy [40], [41]. The accuracy of the fit in each voxel was quantified by means of the normalized mean-squared error (NMSE) between the measured and the reconstructed signal, defined as

$$\text{NMSE} = \frac{||S(\boldsymbol{q}) - \tilde{S}(\boldsymbol{q})||_2^2}{||S(\boldsymbol{q})||_2^2} = \frac{\sum_i [S(\boldsymbol{q}_i) - \tilde{S}(\boldsymbol{q}_i)]^2}{\sum_i S(\boldsymbol{q}_i)^2}, \quad (7)$$

where $S(\boldsymbol{q}_i)$ and $\tilde{S}(\boldsymbol{q}_i)$ are, respectively, the measured and estimated signal at q-space position $\boldsymbol{q}_i$ for $i \in \{1, \ldots, n_d\}$.

Results clearly indicate that the simple voxelwise linear-fit is not suitable to produce reasonable estimates. The intra-cellular compartment appears unnaturally high (≈ 0.9–1.0)

and its spatial distribution does not resemble the anatomically-correct pattern appreciable in the ν_{ic} map of NODDI (left). It is interesting to observe that fitting errors are definitely smaller than in NODDI, but this result clearly points to *overfitting*. Actually, it turns out that in the majority of the voxels, the local linear system to be solved is under-determined and hence the problem has too many independent parameters for the data available locally. On the contrary, solving the problem globally and exploiting the input tracts as prior knowledge on brain structure permits to recover biologically plausible estimates also using a linear approach. The trajectories of the tracts act as a *spatial regularization* term that effectively reduces the degrees of freedom of the linear problem. A close inspection of the figure reveals that COMMIT exhibits the expected pattern of *neurite density* (borrowing the terminology in [21]) as shown by NODDI and previous studies [40], [41], i.e. higher contributions in major WM bundles, decreased values approaching the GM and homogeneous elsewhere. Interestingly, the fitting errors in COMMIT are comparable to NODDI, i.e. $0.92\% \pm 0.53\%$ for the former and $0.89\% \pm 0.53\%$ for the latter, suggesting that the spatial regularization provided by the tracts successfully prevented overfitting. NODDI showed superior fitting results (lower NMSE) in areas with highly-packed axons, e.g. corpus callosum, and in proximity of GM, whereas COMMIT seemed to perform better in locations with crossing fibers. This result finds a natural explanation in the intrinsic differences between the two approaches: NODDI models both WM and GM but assumes a single fiber compartment in each voxel, when in fact COMMIT is intrinsically a WM model but can successfully handle multiple fiber populations. Yet, despite these differences, NODDI served the purpose of an independent technique to highlight and compare known anatomical patterns of neurite density in the final maps and point out overfitting situations.

Notably, our findings might have important implications for the next generation of algorithms for *voxelwise microstructure imaging*. In fact, the time required to estimate the maps in Fig. 6 on the whole brain was around 9–10 hours (with 4 cores) with NODDI, whereas only 13 minutes (1 core) were needed with COMMIT (excluding the time for the tracking) and only 2 minutes (1 core) with the voxelwise NNLS approach. NODDI uses a complex fitting procedure based on Gauss-Newton nonlinear optimization (implemented in Matlab) that produces very accurate results at the price of high computational time. Our results clearly show that convex optimization might be used as a means to drastically reduce the complexity of this class of reconstruction algorithms, but regularization seems mandatory to prevent overfitting. Hence, future research will be conducted to investigate more local forms of regularization, i.e. using a neighborhood instead of the whole trajectories, that might improve the stability of the fit at the voxel level and avoid overfitting.

D. Importance of the input set of candidates

The previous sections have shown that GIBBS offers a valid support to regularize the problem (3) and that the individual contributions of the tracts estimated with COMMIT


Fig. 7.   *Dependence on the input tracts.* The map of the restricted compartment is shown, before and after optimization with COMMIT, for candidate tracts estimated with GIBBS (left column), a classical streamline algorithm (middle) and the front-evolution approach used in Section III-A (right).

are distributed accordingly to known brain anatomy. In this section, we evaluated the efficacy of using other tractography algorithms to build the input set of candidates. A large number of different techniques have been proposed to date in the literature [3]–[5] and an exhaustive evaluation would be prohibitive. Hence, without loss of generality, we selected three representative algorithms with the aim to *highlight specific characteristics of each tractograhy class*: the global approach used in the previous experiments, i.e. GIBBS, the front-evolution algorithm we employed in Section III-A, i.e. GEODESIC, and a classical line-propagation method, hereafter STREAMLINE. This latter is an in-house implementation [42] of the deterministic algorithm described in [43], which considers multiple local diffusion directions and multiple random seed points within each voxel and a maximum turning angle of 60 degrees. A summary of the three algorithms can be found in Table II.

	GIBBS	STREAMLINE	GEODESIC
Tracking Time	12 hours	3 minutes	25 minutes
Number of Tracts	85×10^3	1061×10^3	540×10^3
Optimization Time	13 minutes	28 minutes	39 minutes
Anatomical Plausibility	High	Low	Medium

TABLE II
SUMMARY OF THE THREE EVALUATED TRACTOGRAPHY ALGORITHMS FOR CONSTRUCTING THE INITIAL SET OF TRACTS AS INPUT TO COMMIT.

Fig. 7 compares the three tracking methods with respect to their ability to provide a valid support for combining tissue microstructure estimation with tractography. The map of the intra-cellular compartment is shown before and after optimization with COMMIT. In this experiment, we used the formulation (6) to solve the inverse problem, as STREAMLINE and GEODESIC recover a higher number of tracts than GIBBS and using (5) turned out to be inappropriate. As one can see, neither STREAMLINE nor GEODESIC seem to provide

Fig. 8. *COMMIT as a means to study the effects of under-sampled data.* The map of the restricted compartment is shown in the top row of A as a function of the number of under-sampled measurements considered in the optimization. The accuracy of the fit is also reported (bottom row) as the NMSE between the measured and estimated dMRI signal. Subplots B and C summarize, respectively, the time required for COMMIT to converge and the accuracy of the fit.

a set of candidates as good as GIBBS. This latter estimates the tracts using an energy-minimization procedure which, albeit being computationally expensive ($\approx$ 11–12 hours), produces streamlines that have been placed to fit the signal. STREAMLINE, on the contrary, constructs the pathways using a greedy approach. The input set is estimated much faster ($\approx$ 3 minutes) but, despite showing a pretty uniform coverage of the WM, the spatial distribution of the neurite density after optimization does not match known anatomy. A possible reason could be due to the well-known tendency of these local methods for retrieving a lot of similar pathways and, at the same time, for missing many existing ones, i.e. false negatives. As a consequence, this input set of trajectories cannot represent a valid support that includes all the true bundles and, consequently, COMMIT is unable to accommodate the weights accordingly to the underlying axonal structure. On the other hand, front-evolution approaches can mitigate this limitation of local methods, as they can virtually construct a unique pathway for any pair of locations in the brain. Results show indeed that GEODESIC can estimate the input tracts in reasonable time ($\approx$ 20–25 minutes) and that the spatial distribution of neurite density after optimization is more biologically plausible than STREAMLINE, even though not as good as when using GIBBS as support. A possible explanation stems from a peculiarity of geodesic approaches, which can be seen in Fig. 7, namely the fact that they tend to collapse different tracts onto the same "optimal path" in some areas of the brain. As a result, some WM voxels are not properly covered and bottlenecks can be created in terms of neurite density, thus failing to provide a suitable support for COMMIT.

An *interesting implication of our findings* is that the number of tracts is not crucial to obtain a good support for combining tractography and tissue microstructure estimation, but rather the fact that the tracts should account for all true anatomical bundles and cover homogeneously the WM. Our results also suggest that probably the optimal initial set of tracts is represented by *a combination of tracking methods*; interestingly,

this is also the strategy adopted in [22], in such a way to borrow the advantages of each technique. The characterization of the optimal candidate set represents a very important and interesting study by itself, but it goes beyond the scope of this paper and will be the subject of future investigation.

E. Studying the effects of under-sampled data in q-space

In the last experiment, we present a potential application of this new framework as a means to study and optimize dMRI acquisition protocols and accelerate the acquisitions. To date, the effects of under-sampled data have been widely investigated, using the theory of compressed sensing [32], [33], for diffusion modeling at the voxel level (see [44] and references therein). On the other hand, compressed sensing has never been explored so far with the aim to reduce the long scan time required for mapping the microstructure. Also, the impact of under-sampled data on the tractograms has never been addressed. An added value of our formulation is that, in its basis, it is intrinsically suitable for this kind of investigations. Fig. 8 gives an example of such an analysis. We used the same experimental setup of Section III-B but considering subsets of q-space measurements in the optimization with COMMIT. Each subset was created using the tool subsetpoints of Camino in order to obtain gradient directions as evenly spaced as possible; the same under-sampling rate was used for each shell. As the problem becomes quickly under-determined, the BPDN formulation (6) was used to regularize the solutions.

Fig. 8A shows the map of restricted contributions as estimated using different under-sampling rates, i.e. from using all the input measurements (72 samples) to the inclusion of only 6.25% of them (5 samples), along with the NMSE of the reconstructed signal. The execution time required by COMMIT clearly decreases by reducing the number of measurements considered in the optimization; per contra, as expected, also the quality of the reconstructions is progressively reduced. However, no apparent sign of degradation can be observed in

the reconstructions when discarding 50% of the data (36 samples), whereas the quality quickly deteriorated with stronger under-sampling regimes, i.e. higher NMSE and distribution of the intra-cellular compartment not conform to underlying anatomy. This finding suggests that, despite the protocol was already optimized in [21] with the aim to improve the fit of the `NODDI` model, the *global structure provided by the fibers could be exploited* to either accelerate further the acquisitions or enhance the quality of the reconstructions. Surely, future work will be necessary to characterize the notion of sparse representations of the tractograms and to study more systematically the effectiveness of the ℓ_1 norm for promoting sparsity in the fiber-tracts domain. It will be interesting, as well, to assess the lower bound for data under-sampling that still allows the estimation of useful characteristics of WM tissue, both in terms of trajectory and microstructure properties of the tracts.

F. General considerations and future perspectives

The originality of our formulation lies in the possibility to express tractography and tissue microstructure estimation in a *unified convex problem*, which makes it possible for the first time to deal efficiently with the tremendous computational complexity of these very high-dimensional problems. As a proof of concept, we suggested four possible applications of this novel framework: (i) as a standalone fiber-tracking algorithm with superior performances than state-of-the-art global approaches, (ii) as a means to combine tractography and tissue microstructure estimation to recover the effective contribution of each tract, (iii) as a way for comparing existing tractography algorithms with respect to the biological consistency of the estimated tractograms to the underlying neuronal tissue, and (iv) as a tool to further optimize the acquisition protocols and improve the reconstructions by exploiting the global structure in the data imposed by the fibers. Since this new formulation provides the means to *deal with tractography and tissue microstructure from a different perspective*, we genuinely believe that the range of potential applications of this framework will not be limited, in the future, to the few described here.

A strength of `COMMIT` is that it already showed convincing results even if we used a simple approach to formulate the problem and employed non-specialized software that was publicly available to retrieve the solutions. Yet, we expect that the framework will perform better with a *more appropriate incorporation of prior knowledge* into the model. For example, in our experiments we used a very simple form of regularization that enforces the ℓ_1 norm on all the coefficients; despite being very crude and rudimentary, results already showed the effectiveness of the current implementation. However, ℓ_1 should be more effective if applied only to the coefficients controlling the tracts, i.e. $\boldsymbol{A}^{\mathrm{IC}}$ sub-matrix, whereas Total Variation (TV) regularization seems to be more adequate for imposing spatial smoothness of extra-cellular and isotropic contributions across voxels. We could also set specific bounds on the parameters or explicitly impose, in each voxel, the physical constraint of unit sum of all compartments. All these refinements to the model will surely lead to improved and more robust reconstructions;

from this point of view, thus, our framework has still *a lot of room for improvement*.

Finally, as `COMMIT` relies only on the two operations $\boldsymbol{Ax}$ and $\boldsymbol{A}^{\dagger}\boldsymbol{y}$, we expect also a substantial improvement in terms of execution time with an *implementation on graphics processing units* (GPUs). In fact, sparse matrix-vector multiplication is particularly suitable to benefit from the massive parallelism offered by GPUs and many solutions have already been proposed [45], [46]. GPU implementations will enable hybrid procedures to be developed that alternate `COMMIT` to short runs of classical energy-minimization techniques [12], [13], thus paving the way to obtain a real standalone tractography algorithm which is truly global and quantitative.

The *main limitations* of our approach are twofold. In the first place, the input tracts are static and their position cannot be adapted and, as a consequence, our formulation is sensitive to the choice of the algorithm used to estimate the candidates, as shown in the manuscript. An investigation of the most appropriate local models and tracking methods to obtain plausible and robust estimates is indeed required and it will be the subject of future research. Second, the model assumes that the microstructural properties of the tracts remain constant along their trajectories. Our formulation, as well as all previous approaches, is thus not (directly) suitable in case of pathologies that might locally alter the properties of the axons. However, in those situations, we speculate that any local axonal injury would be rather captured (indirectly) as altered extra-cellular or isotropic contributions in the voxel affected by the disease, as in our model these contributions are spatially independent. Clearly, a comprehensive evaluation of `COMMIT` in such pathological conditions is necessary to assess the applicability of this framework in clinical applications.

IV. CONCLUSION

We have proposed a novel formulation for diffusion MRI tractography that combines the estimation of local microstructure properties of the tissue with the versatility of classical fiber-tracking algorithms in a *unified and efficient framework*. Because of its convex formulation, our approach is fast, accurate and can be used in a wide range of applications. We demonstrated the feasibility and the effectiveness of our framework both on synthetic and in-vivo data. Our findings clearly show that, after optimization with our approach, the tractograms agree more closely with the known brain anatomy and suggest that our proposed formulation represents a viable approach towards *quantitative tractography*. We believe that the framework presented in this work could be received with interest by the diffusion MRI community as it opens new perspectives for the quantification of brain connectivity.

APPENDIX
SHORTEST-PATH TRACKING

Our front-evolution algorithm is based on the *classical shortest-path algorithm* [47]. Fibers are described by paths of nodes that are arranged on regular distributions of points on each face of the voxels within WM. The number of points per face is rather low (e.g. 16 or 25) in order to keep a

limited memory usage and fast execution. The cost function for traversing a segment depends on several parameters that ensure a realistic pathway reconstruction. We account for fiber bending, deviation from local FOD maxima and segment length. Instead of adding the segment costs, as in the original algorithm, we multiply them to penalize paths containing even a single non optimal segment. Multiple source shortest-path is used for efficiency purposes (e.g. nodes lying between WM and a specific ROI) and to recover coherent bundles of fibers. A common drawback of shortest-path is that, among all valid paths that end over a ROI voxel, they often share an initial common part, while other regions are not traversed at all. This causes a slow convergence in the optimization, since the problem is severely under-determined. The algorithm actually traverses every voxel; however many non-optimal paths stop in the middle of WM and they can be exploited in post-processing. We then introduce the *tree widening* procedure to post-process non-optimal results and to produce a set of paths that are better spatially distributed. In practice, the algorithm exploits partial paths and combines them with close-by fibers, only if this junction causes no strong penalties in terms of shortest-path costs. The resulting fibers are more uniformly distributed and there are less shared sub-paths among fibers.

Acknowledgments

References

[1] D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, M. Laval-Jeantet *et al.*, "MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders," *Radiology*, vol. 161, no. 2, pp. 401–7, 1986.

[2] P. J. Basser and C. Pierpaoli, "Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI," *Journal of Magnetic Resonance, Series B*, vol. 111, no. 3, pp. 209–19, 1996.

[3] P. Fillard, M. Descoteaux, A. Goh, S. Gouttard, B. Jeurissen, J. Malcolm, A. Ramirez-Manzanares, M. Reisert, K. Sakaie, F. Tensaouti, T. Yo, J.-F. Mangin, and C. Poupon, "Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom," *NeuroImage*, vol. 56, pp. 220–34, 2011.

[4] S. Jbabdi and H. Johansen-Berg, "Tractography: where do we go from here?" *Brain Connect*, vol. 1, no. 3, pp. 169–83, 2011.

[5] J.-F. Mangin, P. Fillard, Y. Cointepas, D. Le Bihan, V. Frouin, and C. Poupon, "Towards global tractography," *NeuroImage*, vol. 80, pp. 290–6, 2013.

[6] S. Mori, B. Crain, V. Chacko, and P. Van Zijl, "Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging," *Ann Neurol*, vol. 45, pp. 265–9, 1999.

[7] P. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi, "In vivo fiber tractography using DT-MRI data," *Magn Reson Med*, vol. 44, pp. 625–32, 2000.

[8] G. Parker, H. Haroon, and C. Wheeler-Kingshott, "A framework for a streamline-based probabilistic index of connectivity (PICo) using a structural interpretation of MRI diffusion measurements," *J Magn Reson Imaging*, vol. 18, pp. 242–54, 2003.

[9] T. Behrens, M. Woolrich, M. Jenkinson, H. Johansen-Berg, R. Nunes, S. Clare, P. Matthews, J. Brady, and S. Smith, "Characterization and propagation of uncertainty in diffusion-weighted MR imaging," *Magn Reson Med*, vol. 50, pp. 1077–88, 2003.

[10] G. Parker, C. Wheeler-Kingshott, and G. Barker, "Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging," *IEEE Trans Med Imaging*, vol. 21, pp. 505–12, 2002.

[11] S. Jbabdi, P. Bellec, R. Toro, J. Daunizeau, M. Pélégrini-Issac, and H. Benali, "Accurate anisotropic fast marching for diffusion-based geodesic tractography," *Int J Biomed Imaging*, pp. 1–13, 2008.

[12] P. Fillard, C. Poupon, and J.-F. Mangin, "A novel global tractography algorithm based on an adaptive spin glass model," in *Proc. MICCAI*, 2009, pp. 927–34.

[13] M. Reisert, I. Mader, C. Anastasopoulos, M. Weigel, S. Schnell, and V. Kiselev, "Global fiber reconstruction becomes practical," *NeuroImage*, vol. 54, pp. 955–62, 2011.

[14] D. K. Jones, T. R. Knösche, and R. Turner, "White matter integrity, fiber count, and other fallacies: the do's and don'ts of diffusion MRI," *NeuroImage*, vol. 73, pp. 239–54, 2012.

[15] E. Panagiotaki, T. Schneider, B. Siow, M. G. Hall, M. F. Lythgoe, and D. C. Alexander, "Compartment models of the diffusion MR signal in brain white matter: a taxonomy and comparison," *NeuroImage*, vol. 59, no. 3, pp. 2241–54, 2012.

[16] G. P. Winston, "The physical and biological basis of quantitative parameters derived from diffusion MRI," *Quant Imaging Med Surg*, vol. 2, no. 4, pp. 254–65, 2012.

[17] Y. Assaf and P. J. Basser, "Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain," *NeuroImage*, vol. 27, no. 1, pp. 48–58, 2005.

[18] C. Neuman, "Spin echo of spins diffusing in a bounded medium," *J Chem Phys*, vol. 60, p. 4508, 1974.

[19] Y. Assaf, T. Blumenfeld-Katzir, Y. Yovel, and P. J. Basser, "AxCaliber: a method for measuring axon diameter distribution from diffusion MRI," *Magn Reson Med*, vol. 59, no. 6, pp. 1347–54, 2008.

[20] D. C. Alexander, P. L. Hubbard, M. G. Hall, E. A. Moore, M. Ptito, G. J. Parker, and T. B. Dyrby, "Orientationally invariant indices of axon diameter and density from diffusion MRI," *NeuroImage*, vol. 52, no. 4, pp. 1374–89, 2010.

[21] H. Zhang, T. Schneider, C. A. Wheeler-Kingshott, and D. C. Alexander, "NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain," *NeuroImage*, vol. 61, pp. 1000–16, 2012.

[22] A. Sherbondy, R. Dougherty, R. Ananthanarayanan, D. Modha, and B. Wandell, "Think global, act local; projectome estimation with Blue-Matter," in *Proc. MICCAI*, 2009, pp. 861–8.

[23] A. Sherbondy, M. Rowe, and D. Alexander, "Microtrack: An algorithm for concurrent projectome and microstructure estimation," in *Proc. MICCAI*, 2010, pp. 183–90.

[24] M. Rowe, H. G. Zhang, N. Oxtoby, and D. C. Alexander, "Beyond crossing fibers: tractography exploiting sub-voxel fibre dispersion and neighbourhood structure," in *Proc. IPMI*, vol. 23, 2013, pp. 402–13.

[25] B. Yoldemir, B. Acar, Z. Firat, and O. Kilickesmez, "SMT: A Reliability Based Interactive DTI Tractography Algorithm," *IEEE Trans Med Imaging*, vol. 31, no. 10, pp. 1929–40, 2012.

[26] R. E. Smith, J.-D. Tournier, F. Calamante, and A. Connelly, "SIFT: Spherical-deconvolution informed filtering of tractograms," *NeuroImage*, vol. 67, pp. 298–312, 2012.

[27] J.-D. Tournier, F. Calamante, and A. Connelly, "Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution," *NeuroImage*, vol. 35, pp. 1459–72, 2007.

[28] A. Daducci, A. Dal Palú, A. Lemkaddem, and J.-P. Thiran, "A convex optimization framework for global tractography," in *Proc. IEEE ISBI*, 2013, pp. 524–7.

[29] N. S. White, T. B. Leergaard, H. D'Arceuil, J. G. Bjaalie, and A. M. Dale, "Probing tissue microstructure with restriction spectrum imaging: histological and theoretical validation," *Hum Brain Mapp*, vol. 34, no. 2, pp. 327–46, 2013.

[30] C. L. Lawson and R. J. Hanson, *Solving least squares problems*. Society for Industrial and Applied Mathematics, 1974, vol. 161.

[31] Y. Saad, *Iterative methods for sparse linear systems*. Society for Industrial and Applied Mathematics, 2003.

[32] D. Donoho, "Compressed sensing," *IEEE Trans Inf Theory*, vol. 52, no. 4, pp. 1289–306, 2006.

[33] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans Inf Theory*, vol. 52, no. 2, pp. 489–509, 2006.

[34] D. Kim, S. Sra, and I. S. Dhillon, "A non-monotonic method for large-scale non-negative least squares," *Optim Methods Softw*, pp. 1–28, 2012.

[35] E. van den Berg and M. Friedlander, "Probing the pareto frontier for basis pursuit solutions," *SIAM J Sci Comput*, vol. 31, pp. 890–912, 2008.

[36] C. Poupon, B. Rieul, I. Kezele, M. Perrin, F. Poupon, and J.-F. Mangin, "New diffusion phantoms dedicated to the study and validation of high-angular-resolution diffusion imaging (hardi) models," *Magn Reson Med*, vol. 60, pp. 1276–83, 2008.

[37] N. J. Tustison, B. B. Avants, P. A. Cook, Y. Zheng, A. Egan, P. A. Yushkevich, and J. C. Gee, "N4ITK: improved N3 bias correction," *IEEE Trans Med Imaging*, vol. 29, no. 6, pp. 1310–20, 2010.

[38] M.-A. Côté, G. Girard, A. Boré, E. Garyfallidis, J.-C. Houde, and M. Descoteaux, "Tractometer: Towards validation of tractography pipelines," *Med Image Anal*, vol. 17, no. 7, pp. 844–57, 2013.

[39] F. Calamante, J.-D. Tournier, G. D. Jackson, and A. Connelly, "Track-density imaging (TDI): super-resolution white matter imaging using whole-brain track-density mapping," *NeuroImage*, vol. 53, no. 4, pp. 1233–43, 2010.

[40] S. Jespersen, C. Bjarkam, J. Nyengaard, M. Chakravarty, B. Hansen, T. Vosegaard, L. Ostergaard, D. Yablonskiy, N. Nielsen, and P. Vestergaard-Poulsen, "Neurite density from magnetic resonance diffusion measurements at ultrahigh field: comparison with light microscopy and electron microscopy," *NeuroImage*, vol. 49, no. 1, pp. 205–16, 2010.

[41] S. Jespersen, L. A. Leigland, A. Cornea, and C. D. Kroenke, "Determination of axonal and dendritic orientation distributions within the developing cerebral cortex by diffusion tensor imaging," *IEEE Trans Med Imaging*, vol. 31, no. 1, pp. 16–32, 2012.

[42] A. Daducci, S. Gerhard, A. Griffa, A. Lemkaddem, L. Cammoun, X. Gigandet, R. Meuli, P. Hagmann, and J.-P. Thiran, ""the connectome mapper: an open-source processing pipeline to map connectomes with mri","" *PloS one*, vol. 7, no. 12, p. e48121, 2012.

[43] P. Hagmann, M. Kurant, X. Gigandet, P. Thiran, V. J. Wedeen, R. Meuli, and J.-P. Thiran, ""mapping human whole-brain structural networks with diffusion mri","" *PloS one*, vol. 2, no. 7, p. e597, 2007.

[44] A. Daducci, E. J. Canales-Rodríguez, M. Descoteaux, E. Garyfallidis, Y. Gur, Y.-C. Lin, M. Mani, S. Merlet, M. Paquette, A. Ramirez-Manzanares *et al.*, "Quantitative comparison of reconstruction methods for intra-voxel fiber recovery from diffusion MRI," *IEEE Trans Med Imaging*, vol. 33, no. 2, pp. 384–399, 2014.

[45] N. Bell and M. Garland, "Implementing sparse matrix-vector multiplication on throughput-oriented processors," in *Conference on High Performance Computing Networking, Storage and Analysis*, 2009, pp. 18:1–18:11.

[46] H. H. B. Sørensen, "High-performance matrix-vector multiplication on the GPU," in *Proc. International Conference on Parallel Processing*, 2012, pp. 377–86.

[47] E. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, pp. 269–71, 1959.