



Accelerated Microstructure Imaging via Convex Optimization (AMICO) from diffusion MRI data



Alessandro Daducci^{a,b,*}, Erick J. Canales-Rodríguez^{c,d}, Hui Zhang^f, Tim B. Dyrby^e, Daniel C. Alexander^f, Jean-Philippe Thiran^{a,b}

^a Signal Processing Lab (LTS5), École Polytechnique Fédérale de Lausanne, Switzerland

^b University Hospital Center (CHUV) and University of Lausanne (UNIL), Switzerland

^c FIDMAG Germanes Hospitaláries, Spain

^d Centro de Investigación Biomédica en Red de Salud Mental, CIBERSAM, Spain

^e Danish Research Centre for Magnetic Resonance, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University Hospital Hvidovre, Denmark

^f Department of Computer Science and Centre for Medical Image Computing, University College London, UK

ARTICLE INFO

Article history:

Accepted 12 October 2014

Available online 22 October 2014

Keywords:

Diffusion MRI

Microstructure imaging

Convex optimization

ABSTRACT

Microstructure imaging from diffusion magnetic resonance (MR) data represents an invaluable tool to study non-invasively the morphology of tissues and to provide a biological insight into their microstructural organization. In recent years, a variety of biophysical models have been proposed to associate particular patterns observed in the measured signal with specific microstructural properties of the neuronal tissue, such as axon diameter and fiber density. Despite very appealing results showing that the estimated microstructure indices agree very well with histological examinations, existing techniques require *computationally very expensive non-linear procedures* to fit the models to the data which, in practice, demand the use of powerful computer clusters for large-scale applications. In this work, we present a general framework for *Accelerated Microstructure Imaging via Convex Optimization* (AMICO) and show how to re-formulate this class of techniques as convenient linear systems which, then, can be efficiently solved using very fast algorithms. We demonstrate this linearization of the fitting problem for two specific models, i.e. ActiveAx and NODDI, providing a very attractive alternative for parameter estimation in those techniques; however, the AMICO framework is general and flexible enough to work also for the wider space of microstructure imaging methods. Results demonstrate that AMICO represents an effective means to accelerate the fit of existing techniques drastically (up to four orders of magnitude faster) while preserving accuracy and precision in the estimated model parameters (correlation above 0.9). We believe that the availability of such ultrafast algorithms will help to accelerate the spread of microstructure imaging to larger cohorts of patients and to study a wider spectrum of neurological disorders.

© 2014 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/3.0/>).

Introduction

The brain is the principal organ of the central nervous system that governs all the vital functions of a human being and defines its behavior. The white matter (WM) is the intricate neuronal circuitry responsible for transmitting the information between different cortical regions of the gray matter (GM). Any local disruption to this complex system may lead to an overall malfunctioning of the whole organism, causing a broad spectrum of possible neurological disorders. Diffusion magnetic resonance imaging (dMRI) offers a unique tool to study these pathological conditions as it provides the possibility to assess non-invasively the microstructure of the neuronal tissue by probing the natural thermal

motion of water molecules (Le Bihan et al., 1986; Beaulieu, 2002). Diffusion Tensor Imaging (DTI) (Basser et al., 1994) was the first attempt to formally describe the anisotropy of this random process in biological tissues. DTI can estimate the principal diffusion direction from very few dMRI measurements but, as the model assumes Gaussianity of the diffusion process, more than one fiber population cannot be resolved in the same voxel.

A large number of alternatives have been proposed to circumvent this limitation. A non-comprehensive list includes models that directly extend conventional DTI (Tuch et al., 2002; Schultz and Seidel, 2008; Barmpoutis et al., 2009), Persistent Angular Structure (Jansons and Alexander, 2003), Diffusion Orientation Transform (Özarslan et al., 2006), Diffusion Spectrum Imaging (Wedeen et al., 2005) and all the methods based either on Q-Ball Imaging (Tuch, 2004; Canales-Rodríguez et al., 2009; Aganj et al., 2010) or Spherical Deconvolution (SD) (Tournier et al., 2004; Alexander, 2005; Dell'Acqua et al., 2007; Descoteaux et al., 2009); for a

* Corresponding author at: EPFL STI IEL LTS5, ELD 232, Station 11, CH-1015 Lausanne, Switzerland.

E-mail address: alessandro.daducci@epfl.ch (A. Daducci).

more detailed overview, see Daducci et al. (2014a) and references therein. Above all, SD approaches have attracted a lot of consideration over the past few years and the reasons for their success are twofold. First, they can be formulated as simple linear systems and, thus, solved with very efficient algorithms (Jian and Vemuri, 2007). Second, they have been shown to produce very accurate and robust reconstructions also in the case of complex intra-voxel fiber configurations, using a number of dMRI measurements comparable to DTI (Ramirez-Manzanares et al., 2011; Daducci et al., 2014a). However, all these techniques recover exclusively the directional information of the fiber populations in a voxel, but do not provide any insight into their microstructural organization. The diameter of the axons, for instance, is directly related to the propagation speed of the nerve impulses within WM fascicles (Waxman, 1980); the knowledge of such microstructure properties is crucial to be able to study any pathological condition from a biological perspective.

Another class of techniques, known as microstructure imaging, aims at extracting such precious information by using explicit biophysical modeling of the decay patterns in different tissue compartments, e.g. axons, glial cells and extra-axonal space. These methods can infer not only the orientation of the main fiber populations in a voxel, but also very important microstructural properties of the tissue, such as the average diameter and density of the axons. A variety of approaches have been proposed in the literature; Panagiotaki et al. (2012) provide a comprehensive survey of existing techniques. A simple evolution of DTI was proposed by Behrens et al. (2003) to distinguish between water molecules that diffuse inside and around the axons, which are modeled as ideal cylinders with zero radius, and those moving freely with isotropic diffusion. Assaf and Basser (2005) made a further distinction, in their CHARMED model, between molecules that are restricted within the axons, i.e. intra-axonal space, and those that are hindered in the extra-axonal space around them, assuming no exchange between the two pools. The axons are approximated by parallel cylinders with a fixed radius and the corresponding signal profiles are modeled using the analytical expression of Neuman (1974) that describes particles diffusing within cylindrical geometries; an anisotropic Gaussian process is instead assumed in the extra-axonal space. The AxCaliber model (Assaf et al., 2008) is an extension of CHARMED in which the axon radii are explicitly modeled using Gamma distributions, rather than being fixed to a given size. AxCaliber allows the estimation of the axon diameter from diffusion MRI, in a similar way to the model proposed by Stanisz et al. (1997) using NMR spectroscopy data. However, AxCaliber requires prior knowledge about the orientation of the fascicle to probe, limiting de facto axon diameter mapping to specific brain structures; moreover, the intrinsic long scan times are not suitable for clinical applications.

With this in mind, the ActiveAx technique developed by Alexander et al. (2010) was specifically designed to overcome both limitations, as it allows the estimation of orientationally-invariant indices of axon diameter and density in scan time tolerable by live human subjects, using an optimized acquisition protocol. This method made it possible to extend axon diameter mapping to the whole brain, thus enabling the combination of microstructure indices with tractography (Sherbondy et al., 2010; Daducci et al., 2014b). ActiveAx uses a Minimal Model of White Matter Diffusion (MMWMD) with four compartments to describe the measured dMRI signal (Alexander et al., 2010; Dyrby et al., 2013); besides the restricted and hindered compartments previously considered by Assaf and Basser (2005), the MMWMD accounts also for stationary water trapped within small structures such as glial cells, in a similar way to Stanisz's model (Stanisz et al., 1997), as well as free water characterized by isotropic diffusion. Recently, the MMWMD was extended to improve the estimation in brain regions with orientation dispersion (Zhang et al., 2011b) and crossing fibers (Zhang et al., 2011a). However, the corresponding acquisition protocols require about a 1-hour scan and, thus, are still difficult to be routinely included in clinical studies. To enable the estimation of useful microstructural information also within clinical scan times, e.g. 10–15 min, the MMWMD with orientation

dispersion was later simplified by Zhang et al. (2012). In the resulting technique, termed Neurite Orientation Dispersion and Density Imaging (NODDI), the axon diameter parameter was dropped from the model and the formulation was rather optimized to describe the observed dMRI signal as a function of the volume fraction and orientation dispersion of the axons, as well as the partial volume with cerebrospinal fluid (CSF). On the one hand, all these techniques have demonstrated the practical possibility to estimate microstructural information from dMRI data in addition to just the orientation of the fiber populations in a voxel, and the estimated microstructural indices have been shown to agree very well with known anatomical patterns observed with histology (Alexander et al., 2010; Zhang et al., 2012; Dyrby et al., 2013). On the other hand, however, the non-linear routines usually employed to fit these models, as well as other diffusion modalities (Hernández et al., 2013; Chang et al., 2014), are computationally very intensive and cause practical problems for their application in clinical studies, especially with large cohorts of subjects.

In this paper, we propose to take advantage of the versatility of convex optimization to re-formulate microstructure imaging techniques as equivalent but convenient linear systems that can be solved efficiently using very fast algorithms, thus meeting real application demands; we call this framework AMICO, standing for *Accelerated Microstructure Imaging via Convex Optimization*. As a proof of concept, we demonstrate the effectiveness of the proposed framework for two specific applications, i.e. ActiveAx and NODDI, but its flexibility makes it possible to linearize also other popular microstructure imaging techniques. The AMICO framework is presented in the following section together with the experimental settings used for validation. We report and discuss results with both numerical simulations and real data, highlighting advantages and limitations of the proposed approach. The source code of AMICO is available at <https://github.com/daducci/AMICO/>.

Materials and methods

In this section, we first revise the general framework adopted in classical SD methods to recover the fiber orientations in a voxel, as it will lay the foundations to introduce the AMICO approach. Then, we will show how a straightforward extension of such SD framework allows us to formulate also classical microstructure imaging techniques by means of linear systems of equations and, thus, solve them efficiently using convex optimization techniques.

General framework for fiber orientation reconstruction

In classical spherical deconvolution methods, the dMRI signal $E(\mathbf{q})$ in each voxel is modeled as the convolution of a *fiber orientation distribution* (FOD) function $f: \mathbb{S}^2 \rightarrow \mathbb{R}^+$ with a response function $K(\cdot, \hat{\mathbf{u}})$ corresponding to the signal attenuation of a single fiber with orientation $\hat{\mathbf{u}} \in \mathbb{S}^2$:

$$E(\mathbf{q}) = E_0 \int_{\mathbb{S}^2} K(\mathbf{q}, \hat{\mathbf{u}}) f(\hat{\mathbf{u}}) d\hat{\mathbf{u}}, \quad (1)$$

where E_0 is the signal without diffusion weighting and the integration is performed over the unit sphere \mathbb{S}^2 . The FOD is usually expressed as a linear combination of N_k basis functions, also called atoms, as $f(\hat{\mathbf{u}}) = \sum_{i=1}^{N_k} w_i f_i(\hat{\mathbf{u}})$ and several alternatives have been proposed in the literature, e.g. discrete mixture of Gaussians (Ramirez-Manzanares et al., 2007) or spherical harmonics (Tournier et al., 2007). As we have already mentioned, a key factor for the success of these approaches is that, when the response functions are known (or can be estimated) a priori, the measurement process can be expressed as a system of linear equations, as follows:

$$\mathbf{y} = \Phi \mathbf{x} + \eta, \quad (2)$$

where $\mathbf{x} \in \mathbb{R}_+^{N_k}$ are the coefficients of the FOD to be estimated, $\mathbf{y} \in \mathbb{R}^{N_d}$ is the vector containing the N_d normalized q-space measurements, $\Phi = \{\phi_{ij}\} \in \mathbb{R}^{N_d \times N_k}$ is the linear operator, also called dictionary, that explicitly models the convolution operation in Eq. (1) with $\phi_{ij} = \int_{\mathbb{S}^2} K(\mathbf{q}_i, \hat{\mathbf{u}}) f_j(\hat{\mathbf{u}}) d\hat{\mathbf{u}}$ and η accounts for the acquisition noise. Consequently, the linear problem (2) can be efficiently solved using a variety of algorithms based on convex optimization. Without loss of generality, existing methods can be recast into the following general regularized least-squares formulation:

$$\underset{\mathbf{x} \geq 0}{\operatorname{argmin}} \underbrace{\frac{1}{2} \|\Phi \mathbf{x} - \mathbf{y}\|_2^2}_{\text{data fitness}} + \lambda \underbrace{\Psi(\mathbf{x})}_{\text{regularization}}, \quad (3)$$

where $\|\cdot\|_2$ is the standard ℓ_2 -norm in \mathbb{R}^n , the positivity constraint is explicitly imposed as the coefficients \mathbf{x} correspond to volume fractions, $\Psi(\cdot)$ represents a generic regularization function and the parameter $\lambda > 0$ controls the trade-off between data and regularization terms. The optimal value for λ can be either set empirically, as done in previous studies (Tournier et al., 2007; Descoteaux et al., 2007; Landman et al., 2012), or estimated using ad hoc techniques such as the L-curve (Hansen, 1999) or the generalized cross-validation (Golub et al., 1979). For $\lambda = 0$, Eq. (3) is the standard non-negative least-squares (NNLS). However, as pointed out by Jian and Vemuri (2007), “most deconvolution models used in literature result in extremely ill-conditioned linear systems”; besides, in many practical situations the dictionary Φ can be under-determined, i.e. more unknowns than measurements, for the number of acquired dMRI images is usually as low as possible to reduce the scan time. Hence, a regularization is required in most cases to either improve the stability of the reconstruction problem or to inject prior knowledge. $\Psi = \|\cdot\|_1$ is a popular choice to promote sparsity in the FOD (Ramirez-Manzanares et al., 2007; Michailovich et al., 2011; Landman et al., 2012), when in fact Tikhonov regularization (Tikhonov and Arsenin, 1977) is usually adopted to reduce ill-conditioning in the matrix Φ (Jian and Vemuri, 2007; Tournier et al., 2007).

General framework for tissue microstructure quantification (AMICO)

To express microstructure imaging techniques as linear systems of the form in Eq. (2), we propose to decouple the reconstruction of the intra-voxel fiber geometry, i.e. number and orientation of fiber populations, from the assessment of their microstructure properties, i.e. diameter and density. This allows us to break down the complexity of the original methods into two simpler sub-problems. To achieve that, we first estimate the orientation of the major fiber populations using standard methods, e.g. DTI or SD, as they are well-known for providing very accurate and robust reconstructions (Daducci et al., 2014a). Since in this work we focus on models that assume only one fiber population, i.e. ActiveAx and NODDI, we estimate its direction $\boldsymbol{\mu} \in \mathbb{S}^2$ using DTI (Basser et al., 1994). Once $\boldsymbol{\mu}$ is known, it is possible to extend the linear operator Φ in Eq. (2) to account for the distinct water pools that arise from axons oriented in direction $\boldsymbol{\mu}$.

To simplify the exposition, in the following we will describe how to build the linear operator Φ for a canonical orientation, i.e. z-axis; then, as Φ depends on the particular $\boldsymbol{\mu}$ estimated in a voxel, we will use the shorthand notation $\tilde{\Phi} = \mathcal{R}_{\boldsymbol{\mu}}(\Phi)$ to denote the operation $\mathcal{R}_{\boldsymbol{\mu}} : \mathbb{R}^{N_d \times N_k} \rightarrow \mathbb{R}^{N_d \times N_k}$ that rotates each atom of Φ to match the actual direction of $\boldsymbol{\mu}$ estimated in the voxel. Besides, since ActiveAx and NODDI implement rather different models, we will treat these two cases separately and we will show how to construct ad hoc dictionaries, labeled Φ_A for the former and Φ_N for the latter, in order to express both models by means of the same general formulation of Eq. (3). We refer to this framework as AMICO (*Accelerated Microstructure Imaging via Convex Optimization*).

Linear formulation for ActiveAx

To express the ActiveAx model (Alexander et al., 2010) as a linear system, it is useful to partition the dictionary $\Phi_A \in \mathbb{R}^{N_d \times N_k}$ into three sub-matrices:

$$\Phi_A = [\Phi_A^r | \Phi_A^h | \Phi_A^i], \quad (4)$$

where $\Phi_A^r \in \mathbb{R}^{N_d \times N_r}$, $\Phi_A^h \in \mathbb{R}^{N_d \times N_h}$ and $\Phi_A^i \in \mathbb{R}^{N_d \times N_i}$ model explicitly the intra-axonal, extra-axonal and isotropic contributions to the dMRI signal in the voxel, with $N_k = N_r + N_h + N_i$. The construction of the three sub-dictionaries for the case of the ex-vivo monkey samples used in Alexander et al. (2010) is detailed below; a similar procedure can be used also for in-vivo data.

- Each column in $\Phi_A^r \in \mathbb{R}^{N_d \times N_r}$ corresponds to the signal attenuation of the water molecules restricted within parallel cylinders with a specific diameter. We considered $N_r = 22$ different axon radii, in the range 0.01–10 μm , and the corresponding signal profiles were estimated according to the *cylinder* model, borrowing the taxonomy of Panagiotaki et al. (2012), and assuming longitudinal diffusivity $d_{\parallel} = 0.6 \times 10^{-3}$ mm^2/s , which is typical for ex-vivo fixed samples, as in Alexander et al. (2010).
- Likewise, the atoms in $\Phi_A^h \in \mathbb{R}^{N_d \times N_h}$ aim at describing the hindered space around these axons. We adopted the *zeppelin* model to include $N_h = 7$ different hindered micro-environments characterized by distinct perpendicular diffusivities d_{\perp} between 0.06×10^{-3} and 0.42×10^{-3} mm^2/s , which were calculated via the tortuosity model in Szafer et al. (1995) assuming the same d_{\parallel} and intra-axonal volume fractions in the range 0.3–0.9.
- Finally, a single compartment was considered, i.e. $\Phi_A^i \in \mathbb{R}^{N_d}$, to account for any isotropic diffusion contribution to the voxel. The corresponding response function was generated according to the *ball* model and setting the isotropic diffusivity $d_{iso} = 2.0 \times 10^{-3}$ mm^2/s , usually observed in ex-vivo samples, as in Alexander et al. (2010). Hence, the final dictionary Φ_A consists of $N_k = 30$ atoms in total.

Then, ActiveAx can be formulated as a convex optimization problem as follows:

$$\underset{\mathbf{x} \geq 0}{\operatorname{argmin}} \frac{1}{2} \|\tilde{\Phi}_A \mathbf{x} - \mathbf{y}\|_2^2 + \lambda \frac{1}{2} \|\mathbf{x}\|_2^2, \quad (5)$$

where the classical Tikhonov regularization $\Psi = \frac{1}{2} \|\cdot\|_2^2$ was employed to improve the stability of the problem. If we naturally partition the estimated coefficients \mathbf{x} into $[\mathbf{x}^r | \mathbf{x}^h | \mathbf{x}^i]$, the *intra-axonal volume fraction* ν' and the *indices of axon diameter* a' and *density* ρ' , defined in Alexander et al. (2010), can be expressed as:

$$\nu' = \frac{\sum_{j=1}^{N_r} \mathbf{x}_j^r}{\sum_{j=1}^{N_r} \mathbf{x}_j^r + \sum_{j=1}^{N_h} \mathbf{x}_j^h} \quad (6)$$

$$a' = \frac{\sum_{j=1}^{N_r} 2R_j \mathbf{x}_j^r}{\sum_{j=1}^{N_r} \mathbf{x}_j^r} \quad (7)$$

$$\rho' = \frac{4\nu'}{\pi a'^2} \quad (8)$$

where, for $j \in \{1, \dots, N_r\}$, we denote with R_j the radius of the cylinders corresponding to the j -th atom in Φ_A^r . We will refer to this formulation as ActiveAx_{amico}, as opposed to the original version that here we call ActiveAx_{orig}.

Linear formulation for NODDI

The NODDI model (Zhang et al., 2012) does not attempt to estimate the diameter of the axons, but it aims at explaining the anisotropy observed in the dMRI signal in terms of their density and orientation dispersion. In NODDI, the dMRI signal $E(\mathbf{q}) \in \mathbb{R}_+^{N_d}$ is described with the following hierarchical model:

$$E(\mathbf{q}) = \nu_{iso} E_{iso}(\mathbf{q}) + (1 - \nu_{iso}) [\nu_{ic} E_{ic}(\mathbf{q}) + \nu_{ec} E_{ec}(\mathbf{q})], \quad (9)$$

where $E_{ic} \in \mathbb{R}_+^{N_d}$, $E_{ec} \in \mathbb{R}_+^{N_d}$ and $E_{iso} \in \mathbb{R}_+^{N_d}$ are, respectively, the normalized dMRI signal of the intra-cellular, extra-cellular and isotropic compartments, and ν_{ic} , ν_{ec} and ν_{iso} are the corresponding volume fractions, with $\nu_{ec} = 1 - \nu_{ic}$. E_{ic} and E_{ec} adopt the orientation-dispersed cylinder model introduced in Zhang et al. (2011b) based on Watson distributions, which are real functions on the unit sphere \mathbb{S}^2 that are rotationally symmetric about $\boldsymbol{\mu} \in \mathbb{S}^2$ and whose concentration parameter κ controls the amount of dispersion around it. The latter can be conveniently reported between 0 and 1 using the *orientation dispersion index* (Zhang et al., 2012):

$$OD = \frac{2}{\pi} \arctan\left(\frac{1}{\kappa}\right). \quad (10)$$

NODDI explicitly captures the natural coupling between the intra-cellular and extra-cellular compartments, rather than treating the apparent parallel and perpendicular diffusivities as independent free parameters like previous techniques. In the NODDI model, the anisotropy of the extra-cellular compartment is dictated by the density and the orientation dispersion of the intra-cellular compartment; as a result, E_{ec} depends on both ν_{ic} and κ , which are parameters to be estimated. This dependence makes this model slightly more complicated to formulate as a linear system than ActiveAx. To linearize the NODDI model, then, it is convenient to partition the dictionary $\Phi_N \in \mathbb{R}^{N_d \times N_k}$ into the following two blocks:

$$\Phi_N = \left[\Phi_N^t \mid \Phi_N^i \right], \quad (11)$$

in which the isotropic contributions are modeled through $\Phi_N^i \in \mathbb{R}^{N_d \times N_i}$ as in ActiveAx_{amico} but, in this case, the sub-matrix $\Phi_N^t \in \mathbb{R}^{N_d \times N_t}$ accounts explicitly for the coupled intra- and extra-cellular compartments in the tissue, with $N_k = N_t + N_i$. These sub-dictionaries, for the in-vivo human data used in Zhang et al. (2012), are constructed as follows; a similar procedure can be used also for ex-vivo data.

- Each column in $\Phi_N^t \in \mathbb{R}^{N_d \times N_t}$ corresponds to the signal attenuation that arises from a micro-environment characterized by a specific density and orientation dispersion of the axons. We considered $N_t = 144$ distinct combinations, with 12 values for $\nu_{ic} \in \{0.1, \dots, 1\}$ and 12 different $\kappa \in \{0, \dots, 20\}$. The corresponding signal profiles were estimated according to the model of dispersed WM based on the Watson distributions (Zhang et al., 2011b) and assuming longitudinal diffusivity $d_{||} = 1.7 \times 10^{-3}$ mm²/s, which is commonly observed in in-vivo human data, as in Zhang et al. (2012).
- Isotropic contributions are modeled as in Eq. (4), i.e. $\Phi_N^i \equiv \Phi_N^a$, but setting the isotropic diffusivity $d_{iso} = 3.0 \times 10^{-3}$ mm²/s, which is typical in in-vivo human data, as in Zhang et al. (2012). Therefore, the final dictionary Φ_N consists of $N_k = 145$ atoms in total.

Hence, NODDI can be formulated as a convex optimization problem as:

$$\underset{\mathbf{x} \geq 0}{\operatorname{argmin}} \frac{1}{2} \|\tilde{\Phi}_N \mathbf{x} - \mathbf{y}\|_2^2 + \lambda \frac{1}{2} \|\mathbf{x}\|_2^2 + \gamma \|\mathbf{x}\|_1, \quad (12)$$

where, besides the Tikhonov prior $\Psi_1 = \frac{1}{2} \|\cdot\|_2^2$ used in Eq. (5), we employed an additional regularization function $\Psi_2 = \|\cdot\|_1$ to promote sparsity in the recovered coefficients, with $\|\cdot\|_1$ the standard ℓ_1 -norm

in \mathbb{R}^n . It is worth noting, though, that the solution is sparse only in the N_t atoms of Φ_N^t and that the isotropic contribution must be free to take any value without restriction. Unfortunately, however, publicly available solvers (such as the one used in this work, see later) offer only the possibility to enforce the sparsity prior on the whole coefficient vector \mathbf{x} . We adopted the following strategy to overcome this limitation:

1. The volume fraction of the isotropic compartment ν_{iso} , which is actually not affected by the sparsity prior, is first estimated by solving Eq. (12) without regularization, i.e. $\lambda = \gamma = 0$.
2. We then remove this isotropic contribution from the dMRI signal, i.e. $\mathbf{y} = \mathbf{y} - \tilde{\Phi}_N^i \nu_{iso}$, and solve again Eq. (12) but enforcing this time the sparsity prior on the remaining coefficients. This step identifies the support of the solution, i.e. the smallest subset of atoms needed to explain the signal \mathbf{y} , but solutions are known to be biased because the ℓ_1 norm tends to under-estimate the true value of the coefficients (Figueiredo et al., 2007).
3. For this reason, we finally apply a debiasing step to correct the magnitude of the recovered coefficients, by solving once more Eq. (12), without regularization, over the previously identified support set of the solution, as in Figueiredo et al. (2007).

Let $[\mathbf{x}^t \mid \mathbf{x}^i]$ be the partition of the coefficients \mathbf{x} according to Eq. (11); the parameters of the NODDI model, defined in Zhang et al. (2012), can then be computed as:

$$\nu_{ic} = \frac{\sum_{j=1}^{N_t} f_j \mathbf{x}_j^t}{\sum_{j=1}^{N_t} \mathbf{x}_j^t} \quad (13)$$

$$\kappa = \frac{\sum_{j=1}^{N_t} k_j \mathbf{x}_j^t}{\sum_{j=1}^{N_t} \mathbf{x}_j^t} \quad (14)$$

$$\nu_{iso} = \sum_{j=1}^{N_i} \mathbf{x}_j^i \quad (15)$$

where, for $j \in \{1, \dots, N_t\}$, we denote with f_j and k_j , respectively, the intra-cellular volume fraction and the concentration parameter corresponding to the j -th atom in Φ_N^t . We will refer to this formulation as NODDI_{amico}, as opposed to the original version here termed NODDI_{orig}.

Datasets

We tested AMICO using data and experimental settings used in the original publications, respectively (Alexander et al., 2010) for ActiveAx and (Zhang et al., 2012) for NODDI. For convenience, in the following we summarize a few details about the data.

ActiveAx data

In Alexander et al. (2010), ex-vivo dMRI images¹ of a fixed monkey brain, prepared as in Dyrby et al. (2011), were acquired on a 4.7 T Varian system. Three slices centered on the mid-sagittal plane of the corpus callosum (CC) were acquired with a 0.4 mm isotropic spatial resolution and an interslice gap of 2 mm. The imaging protocol was optimized for a maximum gradient strength $G_{max} = 140$ mT/m using the procedure detailed in Alexander (2008). A total of 360 measurements were divided into 4 shells with b-values {1930, 1930, 3090, 13190} S/mm² corresponding, respectively, to gradient amplitudes $G = \{140, 140, 131, 140\}$ mT/m, $\delta = \{10.2, 10.2, 7.6, 17.7\}$ ms, $\Delta = \{16.7, 16.7, 45.9, 35.8\}$ ms and same $TR/TE = 5000/60$ ms for all images. No spatial smoothing was applied to the images. In addition, to quantitatively validate the method, synthetic data was generated for the same imaging protocol using the diffusion simulator system in the Camino toolkit² (Hall and Alexander, 2009), by

¹ <http://dig.drcmr.dk/activeax-dataset/>.

² www.camino.org.uk.

simulating random walks of water molecules in 44 different substrates of WM microstructure with known geometry, as in Alexander et al. (2010). Each substrate consists of non-abutting and impermeable cylinders with realistic axon diameter distributions and fiber densities; a total of 22 different axon radii histograms were considered, each simulated for 2 distinct packing densities. Finally, the signal was contaminated with Rician noise (Gudbjartsson and Patz, 1995) to reproduce a signal-to-noise ratio (SNR) in the E_0 signal with values ranging from 10 to 50; for each substrate, 100 noise realizations have been tested, for a total of 4400 experiments for each SNR.

NODDI data

In Zhang et al. (2012), dMRI data was synthesized according to the model of dispersed WM based on the Watson distributions (Zhang et al., 2011b). A total of 80 different substrates were tested as combinations of the following parameters typically observed in human brain tissue: intra-cellular volume fraction $\nu_{ic} \in \{0.2, 0.4, 0.6, 0.8\}$, isotropic volume fraction $\nu_{iso} \in \{0.0\}$, average axon diameter $a \in \{0.5, 1, 2, 4\}$ μm and concentration of the Watson distribution $\kappa \in \{0, 0.25, 1, 4, 16\}$. Each configuration was evaluated along 250 directions evenly distributed over the sphere in order to assess any potential dependence on the mean orientation μ ; as we did in the case of ActiveAx, the corresponding signal was contaminated with Rician noise using the same SNR range. Among all imaging protocol variants investigated in Zhang et al. (2012), in our current experiments we used the following one as it matches the in-vivo human dataset released with the NODDI toolbox: two shells with 24 measurements at b-value 700 s/mm^2 and 48 at 2000 s/mm^2 , maximum gradient strength $G_{max} = 40$ mT/m , spatial resolution $1.875 \times 1.875 \times 2.5$ mm^3 and using the same $\delta/\Delta/\text{TR}/\text{TE} = 27.7/32.2/12400/86.6$ ms for all images.

Experiments

The free parameters λ and γ controlling the degree of regularization in the AMICO formulations (5) and (12) must be determined. As the goal of this work was to reproduce the microstructure estimates of the original models using faster algorithms, we have searched empirically for values that allowed us to obtain similar (or better) estimates. Specifically, we restricted the search in the range 0–0.5 and computed the average absolute error of the estimated model parameters in the synthetic data with SNR = 30, as it corresponds to the typical noise level in real acquisitions. The optimal values identified this way were then used consistently throughout the rest of the article, i.e. all synthetic and real data, in order to avoid computing the optimal regularization coefficients for each single experiment, which would unnecessarily overload the exposition. For comparison, reconstructions without regularization ($\lambda = 0$ and $\gamma = 0$) will be also reported.

To quantitatively evaluate the effectiveness of the proposed convex formulations, we have compared the microstructure indices estimated with AMICO on the synthetic substrates, for which the ground-truth is known, to those computed using the Matlab implementations³ of the original algorithms used in Alexander et al. (2010) and Zhang et al. (2012). The quality of the reconstructions was assessed by means of:

- Absolute error between the true, i.e. p_j , and the estimated microstructure indices, i.e. \tilde{p}_j , defined as $\Delta p = \frac{1}{N} \sum_{j=1}^N |p_j - \tilde{p}_j|$, where N is the number of voxels. The metric Δp has the same unit as its values p_j . Both bias and variance in the estimation are analyzed.
- Correlation between the estimated and the ground-truth values, expressed as the Pearson correlation coefficient r .

³ NODDI_{orig} is available online via the NODDI Matlab toolbox at <http://cmic.cs.ucl.ac.uk/mig/index.php?n=Tutorial.NODDIatlab>. The Matlab implementation of ActiveAx_{orig} is not publicly available, but a Java implementation is distributed as part of the Camino toolkit at www.camino.org.uk; the results of the two versions differ, but provide very similar estimates.

In the case of real acquisitions, the ground-truth was clearly unknown, both in-vivo and ex-vivo; hence, only a qualitative comparison between the estimated parameters was possible.

We have also compared the two approaches in terms of the time required to fit the models to the data. If on the one hand the linearization of the models with AMICO allows us to effectively break down the complexity of the original problem into two simpler sub-problems, on the other hand a valid concern about this strategy is that it might be sensitive to the initial estimation of the fiber orientation. For this reason, we have explicitly tested the effect of inaccuracies in the fiber orientation estimation on the quality of the reconstructions with AMICO.

Implementation details

In all our experiments, we used standard DTI (Basser et al., 1994) to estimate the orientation of the main fiber population $\mu \in \mathbb{S}^2$ in each voxel. In particular, we fit a diffusion tensor to the log-transformed dMRI measurements using the ordinary least-squares (LS) procedure implemented in Matlab and we set μ to the estimated principal eigenvector. Then, as the dictionaries Φ_A and Φ_N in a voxel depend on the specific μ estimated in it, we precomputed rotated versions of all the atoms with an angular resolution of 1° and used look-up tables to accelerate the construction of the ad hoc dictionaries in each voxel. In the case of ActiveAx, the response functions were generated using the tool *datasynth* that is available in the Camino toolkit, as it provides analytic models (Panagiotaki et al., 2012) to efficiently synthesize the dMRI measurements corresponding to all necessary compartments. The response functions for the NODDI dictionaries were generated according to the model of dispersed WM using the Watson distributions implemented in the NODDI toolbox (Zhang et al., 2011b). Lastly, to solve Eqs. (5) and (12), we used the SPArse Modeling Software (SPAMS) optimization toolbox.⁴ This open-source library provides a very flexible implementation in C++ of the proximal splitting method of Beck and Teboulle (2009), which allows the resolution of a large class of linear problems using a wide range of regularization functions. All experiments have been conducted on a standard workstation (Intel Core i7, 2.80 GHz, 6 GB ram) without multi-threading or parallel computing.

Results and discussion

Sensitivity to the regularization parameter

The left plots in Fig. 1 report the average absolute error of ActiveAx_{amico} in the estimation of ν' (top) and a' (bottom) as a function of the regularization parameter λ . Results correspond to the reconstructions on the 44 substrates at SNR = 30 and the green line indicates the performance of the original algorithm. As can be noticed, with sufficient regularization, i.e. $\lambda \geq 0.1$, ActiveAx_{amico} estimates both microstructure indices more accurately than the original algorithm; on the other hand, with no regularization, i.e. $\lambda = 0$, the average errors are much higher. The intra-axonal volume fraction ν' shows very similar performances starting from a regularization $\lambda \approx 0.02$ and its estimates appear very stable over the rest of the range. The axon diameter index a' presents a minimum error at $\lambda \approx 0.1$ and then slowly deteriorates, but it always stays below the corresponding error of ActiveAx_{orig} (≈ 0.7 μm , green line). For this reason, we set $\lambda = 0.1$ for the rest of the article and report the results only for this level of regularization. A similar analysis (right plots) led to the identification of $\lambda = 0.001$ and $\gamma = 0.5$ as optimal regularization parameters in the case of NODDI_{amico}.

⁴ <http://spams-devel.gforge.inria.fr>.

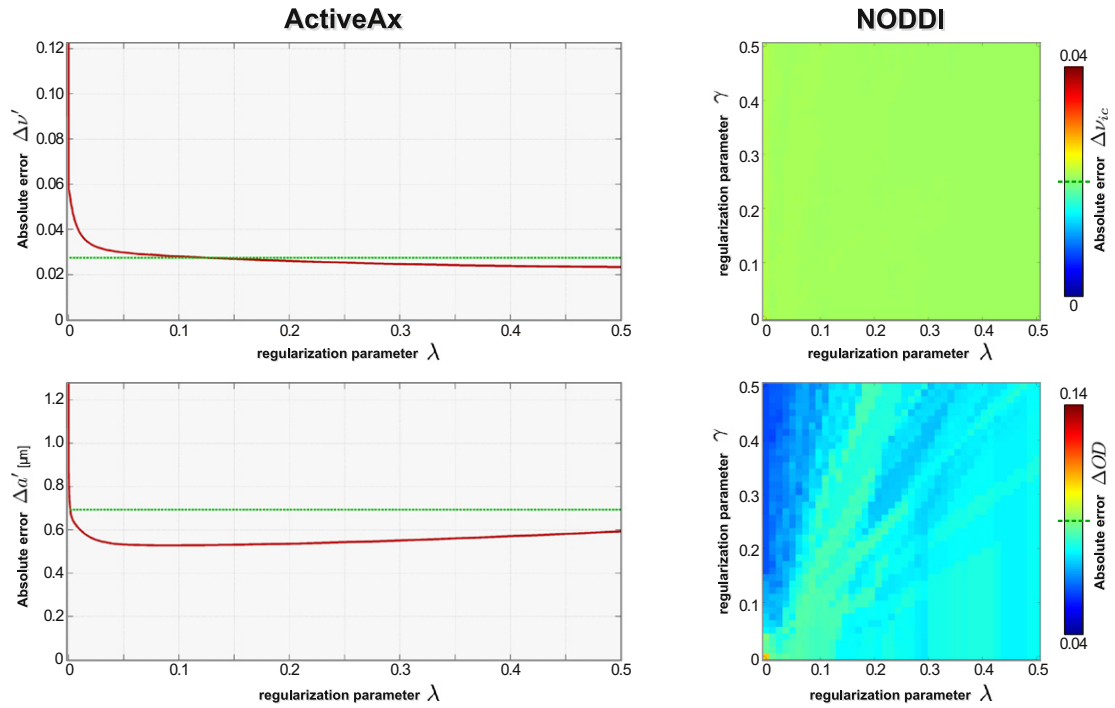


Fig. 1. Sensitivity to the regularization parameters. *Left:* the average absolute error in the estimation of ν' (top) and a' (bottom) is plotted as a function of the regularization parameter λ used in ActiveAx_{amico}. Results correspond to the reconstructions on the synthetic substrates at SNR = 30; for reference, the performance of ActiveAx_{orig} is marked in green. *Right:* in the case of NODDI_{amico}, as two regularization parameters λ and γ are used, the errors $\Delta\nu'_{ic}$ (top) and ΔOD (bottom) are plotted as 2D maps. The color bar is centered at the average performance of NODDI_{orig}, which is highlighted by the green dashed line for reference.

ActiveAx

Fig. 2 compares, as a function of the SNR, the intra-axonal volume fraction ν' (left) and axon diameter a' (right) indices estimated by ActiveAx_{orig} and ActiveAx_{amico} on the 44 synthetic substrates. AMICO is presented both without regularization and for $\lambda = 0.1$. The top plots report the absolute error in estimating ν' and a' (mean and standard

deviation), while in the bottom we compare their correlation with the ground-truth values. The first result that can be clearly observed is that, without regularization, the convex approach does not provide plausible estimates. This is revealed in the much higher average errors for both parameters as well as in their low correlation with the ground-truth. Conversely, if the problem is appropriately regularized, ActiveAx_{amico} provides very accurate, i.e. low average errors, and robust estimates, i.e.

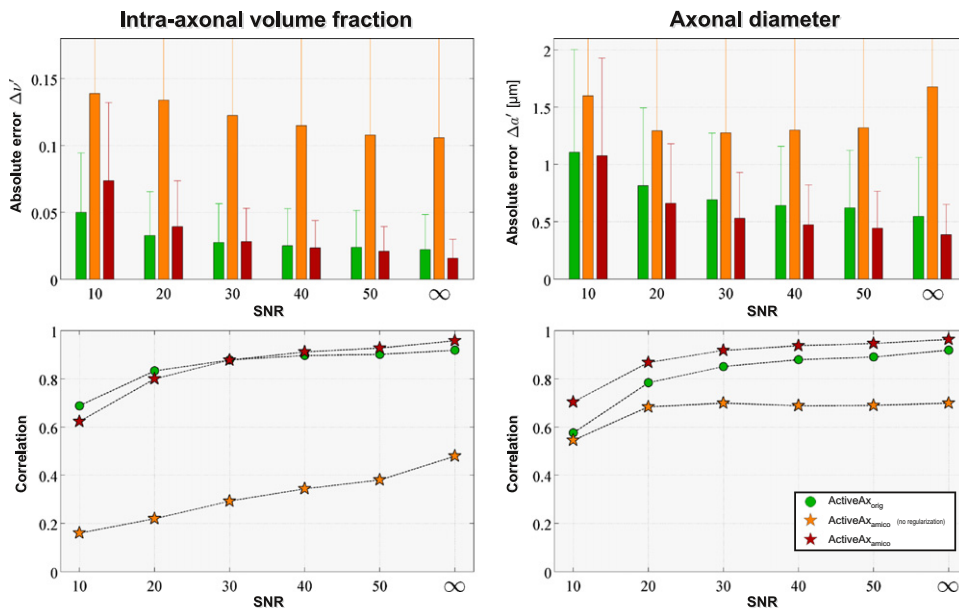


Fig. 2. ActiveAx performances on the synthetic substrates as a function of SNR. ActiveAx_{amico} with $\lambda = 0$ (orange) and $\lambda = 0.1$ (red) is compared to ActiveAx_{orig} (green) by means of the average absolute error (top) and correlation (bottom) of the estimated parameters ν' (left) and a' (right) with respect to the ground-truth.

low standard deviations, of both microstructure parameters, almost indistinguishable from those estimated with the original algorithm. The axon diameter index a' appears to be estimated better with ActiveAx_{amico} than ActiveAx_{orig}, both in terms of absolute errors and correlation to the ground-truth. Concerning the estimation of the intra-axonal volume fraction ν' , the two approaches show similar performances, but the estimates with ActiveAx_{amico} progressively deteriorate for $\text{SNR} \leq 20$, and exhibit slightly higher absolute errors and lower correlations than ActiveAx_{orig}. However, this was expected because, at low SNR regimes, the measurement noise cannot be approximated as a Gaussian (Gudbjartsson and Patz, 1995) and thus the data fidelity term, which is based on the ℓ_2 -norm of the residual, does not represent the optimal log-likelihood model for the noise in Rician contaminated data. ActiveAx_{orig}, on the other hand, implements the proper noise model and thus it is expected to be more accurate in high noise conditions.

Table 1 reports the time required by the fit with two algorithms; the first column corresponds to the results presented in Fig. 2. The original ActiveAx_{orig} algorithm took about 11 days to fit the model to all the voxels, i.e. 44 substrates \times 100 noise trials \times 6 SNR, which corresponds to ≈ 38 s/voxel. On the other hand, this burdensome computation time can be drastically reduced with the proposed convex formulation by several orders of magnitude as indeed, on the same data, ActiveAx_{amico} required less than 20 s, approximately 0.7 ms/voxel.

Fig. 3 zooms in on the individual reconstructions obtained on each substrate, in the noiseless case and at $\text{SNR} = 20$. The estimated parameters ν' and a' are plotted in red against their corresponding ground-truth values ν and α , and their averages over the 100 noise runs are reported as blue crosses. These scatter plots clearly confirm that the parameters ν' and a' estimated by ActiveAx_{amico} without regularization (middle) have very little correlation with the ground-truth, and that their values can depart considerably from the true ones. Nonetheless, the situation changes substantially when the problem is properly regularized (bottom), as indeed the estimated parameters look actually very accurate and concentrated around the true values; for this reason, from now on we will report only the results corresponding to the regularized version of AMICO. Comparing then directly the reconstructions of ActiveAx_{orig} (top) with ActiveAx_{amico} (bottom), we can recognize the presence of the same “lower bound” at about 2 μm observed in Dyrby et al. (2013) for the sensitivity to detection of different axon diameters for any G_{max} . Similarly, ActiveAx_{orig} exhibits a tendency, highlighted in the same publication, to under-estimate the a' index as α increases, whereas ActiveAx_{amico} seems slightly more robust to this effect. In the case of noiseless data, the a' index estimated with ActiveAx_{orig} in large axon configurations ($\approx 6 - 7 \mu\text{m}$) is not affected by the packing density of the substrates, as expected, while ActiveAx_{amico} returns slightly different values. On the other hand, ActiveAx_{orig} seems to be much more sensitive to the packing density than ActiveAx_{amico} for medium axons ($\approx 3 - 4 \mu\text{m}$). The two approaches showed opposite behaviors for the estimation of the intra-axonal volume fraction: in fact, ActiveAx_{orig} tends to under-estimate ν' whereas ActiveAx_{amico} slightly over-estimates it. Another point worth mentioning is that, while in ActiveAx_{orig} the contributions of the compartments sum to unity by definition, this constraint was not explicitly imposed in our convex formulation (5). Nonetheless, results showed that this physical constraint is naturally met also in the case of ActiveAx_{amico}: in fact, the sum of the coefficients is 0.997 ± 0.001 (mean \pm standard deviation)

on the noiseless data and 1.010 ± 0.055 at $\text{SNR} = 20$. Future experiments will be conducted to assess any possible benefit from an explicit incorporation of such constraint in every voxel. Lastly, it can be noticed that the original algorithm appears to suffer from some outliers in the estimates, especially in the case of a' , whereas the AMICO approach seems more robust. A possible explanation stems from the fact that ActiveAx_{orig} uses optimization procedures, notably a final Markov Chain Monte Carlo (MCMC) stage, that can get trapped in one of the local minima of the objective function; on the other hand, ActiveAx_{amico} naturally ensures convergence to the global minimum, as the formulation is convex.

One shortcoming of the proposed approach is that it assumes the main fiber population in the voxel to be accurately estimated beforehand, whereas the original ActiveAx method considers also its orientation as an unknown in the fit. In Fig. 4 we explicitly tested the influence of inaccuracies in the assessment of μ on the performances of ActiveAx_{amico}. The absolute errors (mean and standard deviation) in the estimation of ν' and a' are plotted as a function of the angular separation between the estimated μ and the actual orientation of the fiber population in the voxel. Results correspond to the reconstructions with ActiveAx_{amico} ($\lambda = 0.1$) on the 44 synthetic substrates at $\text{SNR} = 30$; for reference, the performances of ActiveAx_{orig} are also reported in green. No significant deterioration in the assessment of the intra-axonal volume fraction ν' can be observed even with inaccuracies in the estimation of μ up to 15° . The mean absolute error for ActiveAx_{amico} progressively degrades, but it never departs significantly from the corresponding average performance of ActiveAx_{orig} (green dotted line, ≈ 0.03) on the same data. For the estimation of the axonal diameter index a' , ActiveAx_{amico} exhibits stable reconstructions and smaller errors than ActiveAx_{orig} up to 7° , but a rapid deterioration in accuracy can be observed for inaccuracies in the assessment of μ above $9 - 10^\circ$. However, the central pillar this work is based upon is that, actually, classical methods for fiber orientation recovery have been shown to produce quite accurate reconstructions; in particular, in the case of voxels with only one fiber population, as assumed in the ActiveAx model, DTI can estimate its orientation with very high accuracy (Daducci et al., 2014a).

Lastly, we tested ActiveAx_{amico} also using measured brain data from the original publication. Fig. 5 compares the maps of the axonal diameter a' (top) and density ρ' (bottom) indices estimated by ActiveAx_{orig} (left) and ActiveAx_{amico} (middle) in the mid-sagittal slice of the CC of the ex-vivo monkey dataset, whose SNR was around 20. The difference maps between the two approaches are also shown. From visual inspection, it can be appreciated that the maps estimated with the proposed convex approach are in very good agreement with those of ActiveAx_{orig}. Both methods exhibit the expected pattern of axonal diameter and density variation across the CC as seen in previous ActiveAx studies (Alexander et al., 2010; Zhang et al., 2011a; Dyrby et al., 2013) as well as with electron microscopy (Lamantia and Rakic, 1990): smaller axons densely packed in the genu and the splenium (low a' and high ρ') while the mid-body composed of bigger axons with lower density (high a' and low ρ'). ActiveAx_{amico} maps are less speckly than those from the original fitting, most likely as a result of the same local minimum issue previously highlighted that causes outliers in the estimates of the original fitting. Moreover, ActiveAx_{amico} shows slightly higher estimates of a' and, correspondingly, smaller ρ' values. It is also worth noting that, for this ex-vivo data, we have used the same regularization previously estimated on the synthetic substrates, thus confirming the robustness to the choice of the regularization parameter highlighted in Fig. 1.

Table 1
Computation time required by the original algorithms and the corresponding convex formulations for each full experiment performed in this study.

	ActiveAx		NODDI	
	Substrates	CC slice	Substrates	Whole brain
Original method	11 days 14 h 40 min	4 h 30 min	40 h 20 min	65 h 10 min
AMICO approach	18 s	0.3 s	6 min 20 s	9 min 30 s

NODDI

Similarly, in Fig. 6 we provide an overall comparison, as a function of the SNR, between NODDI_{orig} and NODDI_{amico} on the 80 synthetic substrates. As before, AMICO results are presented both with and without regularization. The performances have been assessed by means of the absolute error (mean and standard deviation) in the estimated parameters

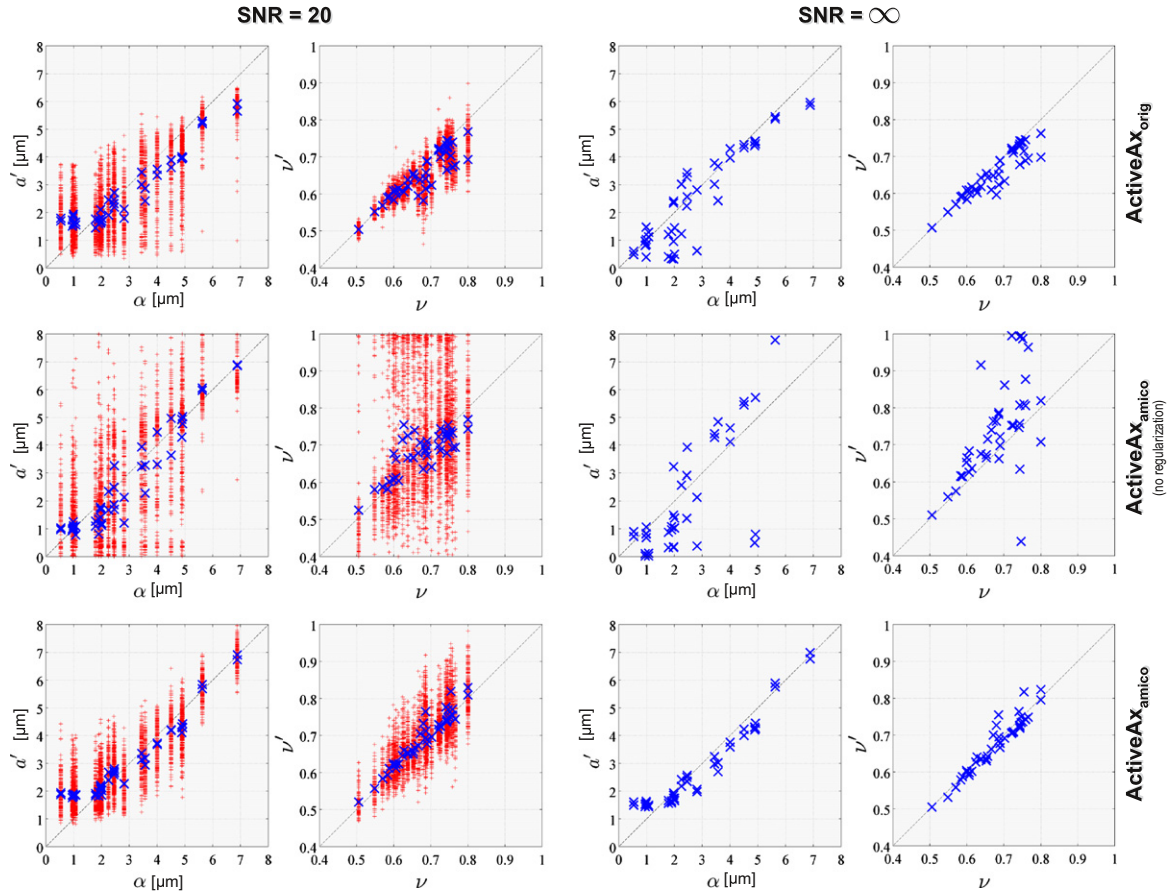


Fig. 3. ActiveAx detailed performances on the synthetic substrates. The parameters ν' and a' estimated by ActiveAx_{orig} (top) and ActiveAx_{amicco} with two levels of regularization ($\lambda = 0$ in the middle and $\lambda = 0.1$ at the bottom) on each individual substrate are plotted against their corresponding idealized indices of intra-axonal volume fraction ν and axon diameter α . Red marks show the individual estimates for each of the 100 noise realizations of every substrate, whereas the corresponding mean values over the trials are marked in blue. Results are reported for SNR = 20 (left) and noiseless data (right).

ν_{ic} and OD and their correlation with the ground-truth values. Results confirm that our proposed convex formulation can provide very accurate estimates also in the case of the NODDI model. If properly regularized ($\lambda = 0.001$ and $\gamma = 0.5$), no appreciable differences in the average absolute error could be observed for both microstructure indices with respect to the original model (compare green and red); in particular, the correlation with ground-truth values appears almost indistinguishable between the two approaches. On the other hand, higher absolute errors and lower correlation are evident if no regularization is employed (orange). As shown in the third column of Table 1, the time required by

the original NODDI_{orig} algorithm to fit the model to all voxels, i.e. 80 substrates \times 250 orientations of $\mu \times 6$ SNR, was about 40 h; on the same data, our convex NODDI_{amicco} approach took only 6 min.

Although it does not use the proper noise model, NODDI_{amicco} exhibits better performance than NODDI_{orig}, i.e. smaller average error, as the noise level increases. Fig. 7 provides a possible explanation for this behavior. The microstructure indices OD (left), ν_{ic} (middle) and ν_{iso} (right) estimated with NODDI_{orig} and NODDI_{amicco} are plotted, as mean and standard deviation, against the corresponding ground-truth values of each substrate. With no noise (top plots), the three algorithms

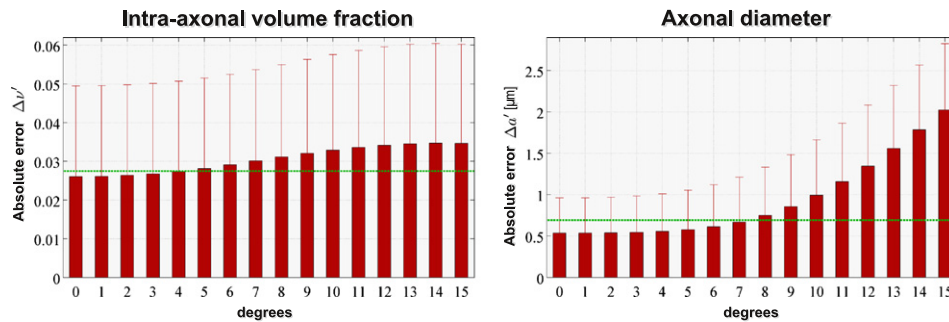


Fig. 4. Robustness to inaccuracies in the estimation of μ . The absolute errors (mean and standard deviation) in the estimation of ν' (left) and a' (right) are plotted as a function of the angular deviation between the true fiber orientation and the estimated μ . Results correspond to the reconstructions with ActiveAx_{amicco} ($\lambda = 0.1$) on the 44 substrates at SNR = 30; for reference, the green line indicates the performance of ActiveAx_{orig}.

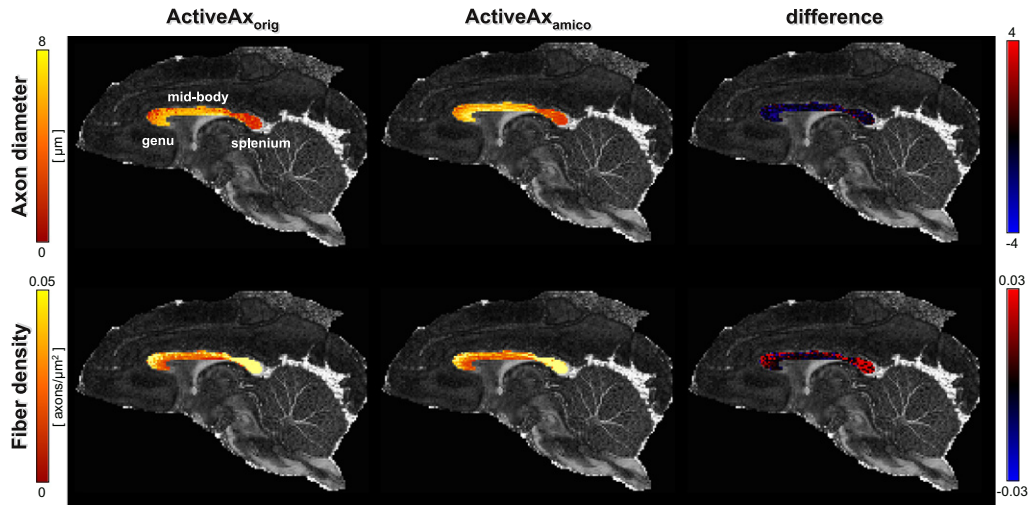


Fig. 5. ActiveAx evaluation on the ex-vivo monkey dataset. Axonal diameter a' (top) and density ρ' (bottom) maps estimated by ActiveAx_{orig} (left) and ActiveAx_{amico} (middle) in the mid-sagittal slice of the corpus callosum. The last column shows the difference between the corresponding ActiveAx_{orig} and ActiveAx_{amico} maps. The SNR in the WM was around 20.

are essentially equivalent, showing very accurate estimates and with very low variability; curiously, though, NODDI_{orig} exhibits a higher variance in the estimation of OD for configurations with an orientation dispersion of 0.84. In the case of noisy data (bottom plots), ν_{ic} is still estimated very accurately, but we could see that all the methods tend to under-estimate the axonal dispersion for very dispersed configurations, as already highlighted in Zhang et al. (2012). At the same time, an over-estimation of ν_{iso} can be observed as the intra-cellular volume fraction of the substrates decreases; please note that no isotropic compartment was simulated in these substrates, i.e. the ground-truth $\nu_{iso} = 0$. These two phenomena correspond to voxel configurations characterized by a diffusion process which is progressively becoming isotropic; in these situations, then, the predicted signals for the different

compartments present in the voxel are very close and it is difficult to distinguish among them. Hence, their contribution is arbitrarily distributed among them, a condition that we picture as a kind of “exchange of energy” between the compartments.

Lastly, we tested our convex formulation for the NODDI model also in the case of in-vivo human data. Fig. 8 shows the maps of the microstructure parameters OD , ν_{ic} and ν_{iso} estimated with NODDI_{orig} and NODDI_{amico}, the latter both with and without regularization, in two representative slices of the brain. The difference maps between the original and the convex approach (with regularization) are also shown. For reference, the Fractional Anisotropy (FA) and Apparent Diffusion Coefficient (ADC) values extracted from standard DTI analysis are also shown. The time required to fit the model in the whole brain

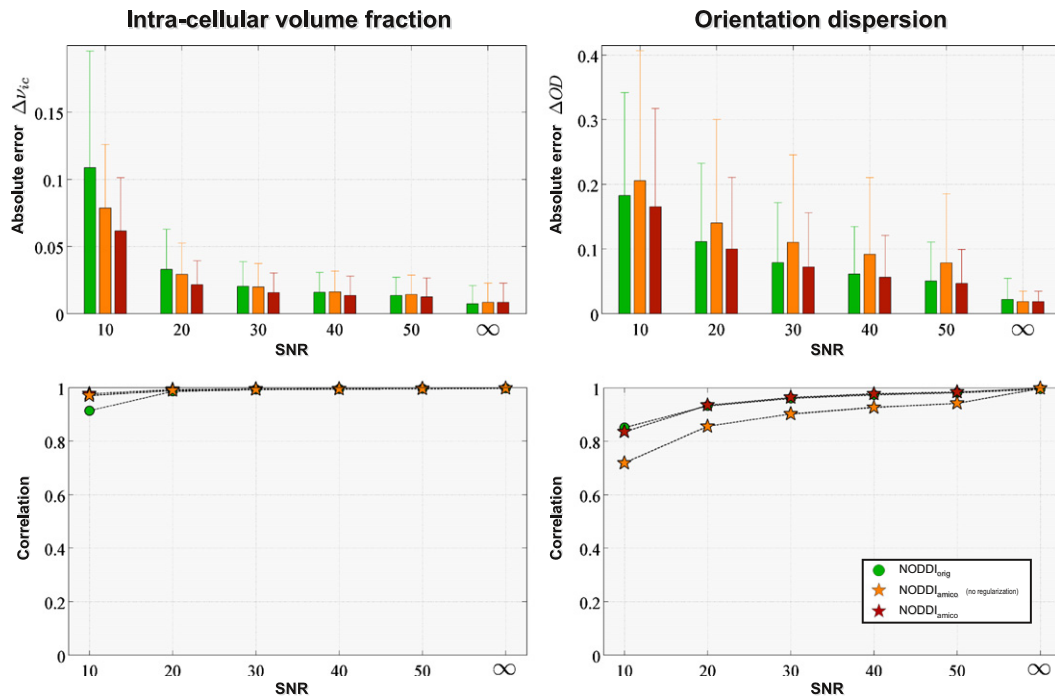


Fig. 6. NODDI performances on the synthetic substrates as a function of SNR. The reconstructions of NODDI_{amico} obtained with (red) and without regularization (orange) are compared to NODDI_{orig} (green) by means of the average absolute error (top) and correlation (bottom) of the estimated parameters ν_{ic} (left) and OD (right) with respect to the ground-truth.

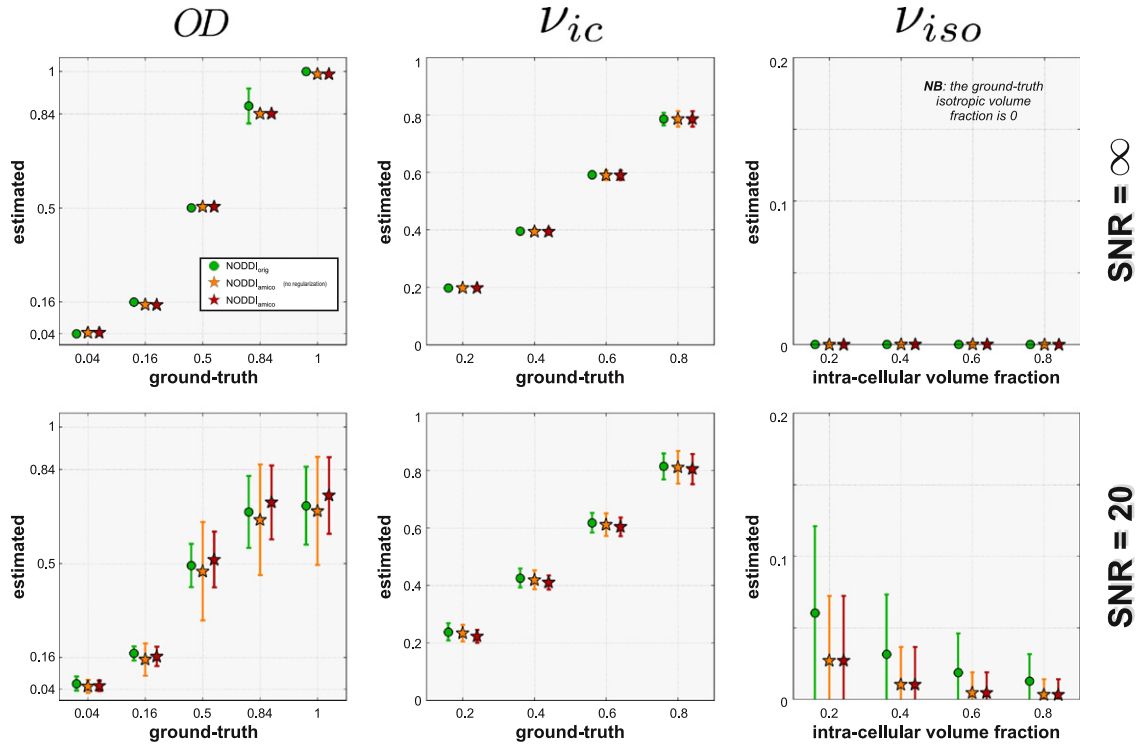


Fig. 7. Detailed evaluation of NODDI on the synthetic substrates. The mean and standard deviation of the microstructure parameters OD (left), ν_{ic} (middle) and ν_{iso} (right) estimated with $NODDI_{orig}$ (green) and $NODDI_{amico}$, the latter reported both with (red) and without (orange) regularization. Two noise levels are presented. To prevent the error bars of the three methods from lying on top of each other, we slightly shifted them with respect to the corresponding x-axis marker.

was about 65 h with the original algorithm, whereas less than 10 min was required using our convex formulation. Despite this big gap in computation time, the ν_{iso} maps estimated by the three approaches appear almost identical; in fact, the average absolute difference between the original NODDI and our AMICO approach over the whole brain is 0.004 ± 0.009 (mean and standard deviation). Also, no significant discrepancies can be observed in the maps of the intra-cellular volume fraction ν_{ic} , with the exception of few voxels, especially at the boundary with CSF, where we can clearly observe the aforementioned “exchange of energy” in locations with almost isotropic diffusion. The average absolute difference is a bit higher in this case, i.e. 0.032 ± 0.119 ; however, if we exclude these voxels from the analysis, we obtain a more indicative measure of the actual discrepancy between the two methods in the brain, i.e. 0.015 ± 0.028 . Finally, it is worth noting that the $NODDI_{orig}$ algorithm provides a heterogeneous classification for such voxels and, in general, they are identified as being characterized by unusually high neurite density (very bright voxels in ν_{ic} maps). On the other hand, with $NODDI_{amico}$ they are consistently classified as containing very few axons, i.e. low ν_{ic} , and having high dispersion, i.e. high OD . In all remaining voxels, the OD maps estimated with the two algorithms are almost indistinguishable, with an average absolute difference of 0.018 ± 0.022 . Again, if no regularization is employed, the resulting OD maps are extremely noisy and not in agreement with the expected anatomical patterns.

Advantages and limitations

The principal benefit of the proposed AMICO framework is to provide an *acceleration factor of several orders of magnitude* in the intrinsic fitting time required by existing microstructure imaging techniques (see Table 1). ActiveAx uses a three-stage routine to fit the model to the data (Alexander et al., 2010), where a grid search of the parameters is followed, in turn, by a non-linear gradient descent stage and a final

MCMC procedure; NODDI uses a similar routine but omits the last MCMC step. These non-linear procedures are very computationally intensive and pose practical difficulties for large scale applications. Microstructure imaging techniques have already been employed successfully in a number of clinical studies (Winston et al., 2014; Kunz et al., 2014; Lemkaddem et al., 2014); the availability of faster algorithms might favor their adoption on a larger scale without the need for powerful computer clusters. The tremendous acceleration enabled with AMICO is mainly due to the linearization of the models achievable by splitting the estimation of the intra-voxel fiber orientation and of its microstructure properties into two simpler sub-problems. A more elegant and robust approach to solve such separable problems would be to develop an alternating minimization algorithm (Csisz et al., 1984) which repeatedly optimizes, in turn, with respect to the orientation of the fiber population and then over its microstructure properties, until convergence.

Besides accelerating the fit, results also showed that AMICO can achieve slightly more accurate and robust parameter estimates than the original techniques. The latter, in particular, appeared to suffer from some outliers in the estimates, which are likely to be caused by the non-linear procedures that are used to fit the models. In fact, besides being computationally demanding, non-linear algorithms are susceptible to get trapped in the local minima of the objective function; to minimize this probability, both ActiveAx and NODDI use in fact an initial grid search over the parameter space to identify a good starting point for the subsequent optimization. By comparison, AMICO formulations are convex and, consequently, they do not need any initialization procedure and, especially, always guarantee to converge to the global minimum.

In this work, the regularization parameters in Eqs. (5) and (12) have been set by searching empirically for their values in order to reproduce at best the microstructure indices estimated by the original algorithms. Actually, results have shown that the proposed convex formulations are rather insensitive to the specific choice of λ and γ , and that the estimated indices are almost indistinguishable from those obtained using the

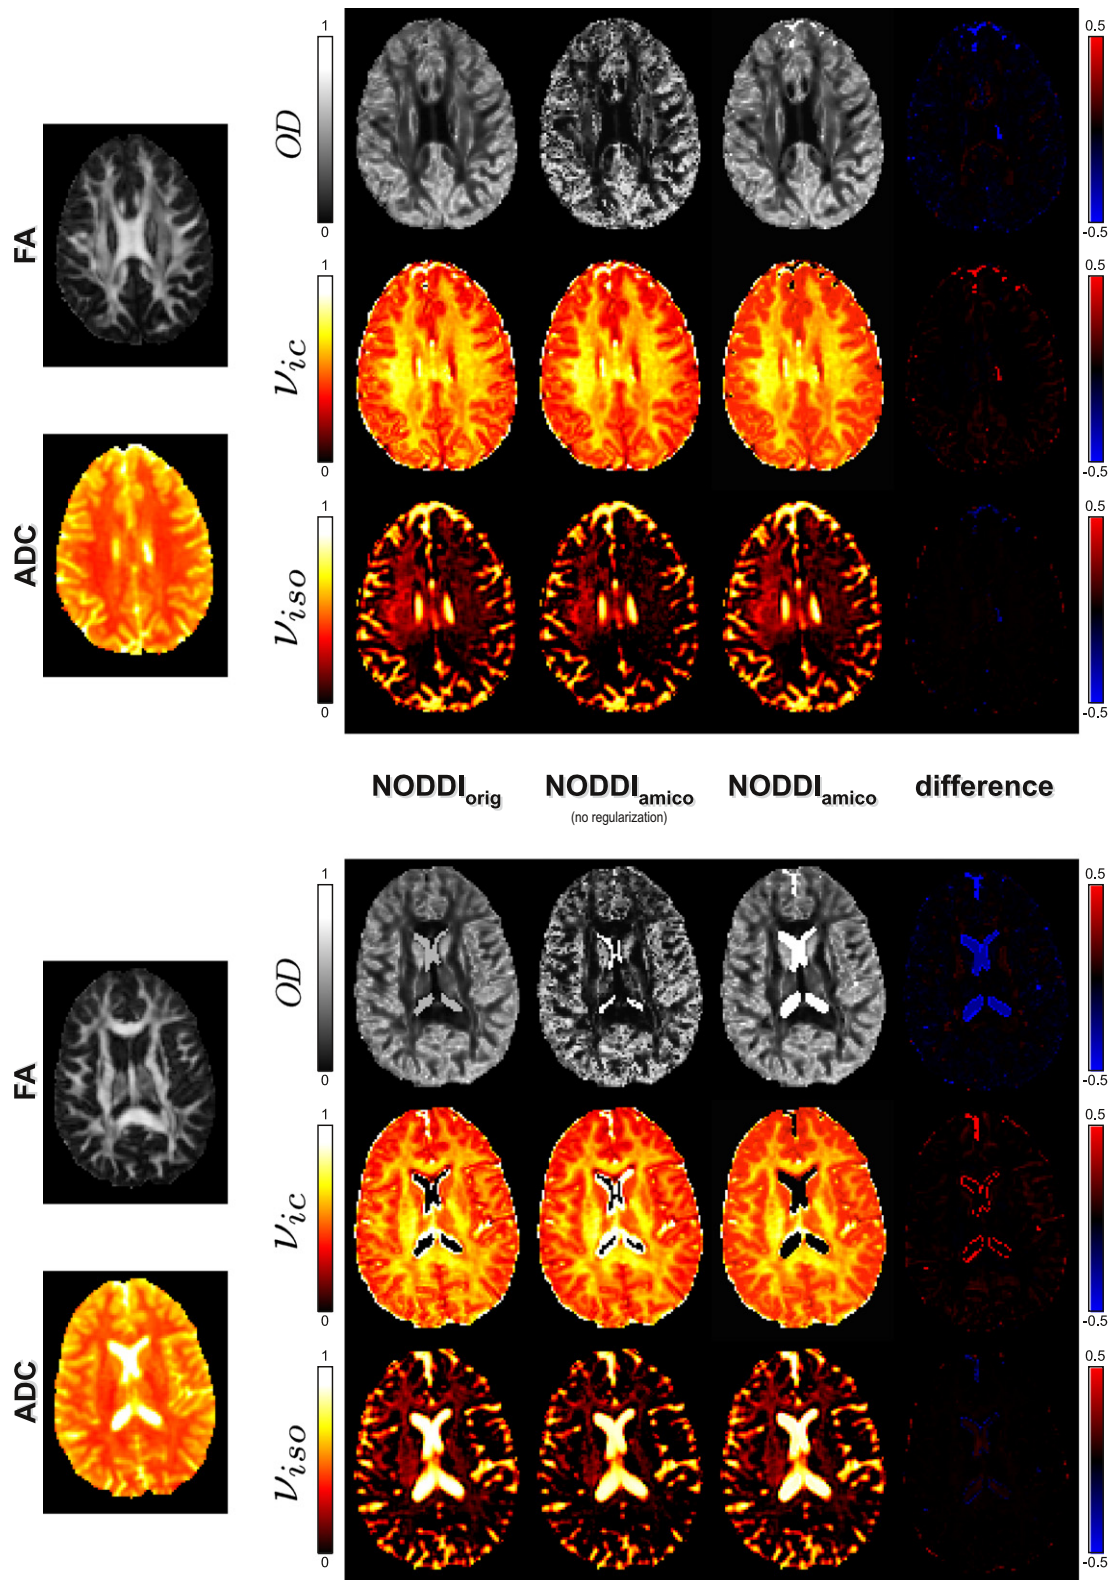


Fig. 8. NODDI evaluation on the in-vivo human dataset. The microstructure parameters OD , V_{ic} and V_{iso} estimated with NODDI_{orig} and NODDI_{amico}, both with and without regularization, are reported in two representative slices of the brain. The last column shows the difference between the corresponding NODDI_{orig} and NODDI_{amico} maps. FA and ADC maps extracted from standard DTI analysis are reported as reference.

original algorithms for a broad range of regularization parameters (Fig. 1). Empirical procedures are commonly employed in the field for setting the regularization trade-off of similar ill-posed problems; see for instance

classic methods like Tournier et al. (2007); Descoteaux et al. (2007); and Landman et al. (2012). Nonetheless, ad hoc techniques exist to estimate the optimal values for these parameters, using for instance the

L-curve method (Hansen, 1999) or the generalized cross-validation (Golub et al., 1979); these strategies will be explored in the future to further improve the estimates with AMICO.

General observations

The results presented in this paper served as a proof of concept to show how microstructure imaging methods can be accelerated using convex optimization. The linear operators introduced in Eqs. (4) and (11) indeed showed to provide excellent agreement with the original models in all synthetic and real data experiments, thus fulfilling the aim of the present work. It would be very interesting, though, to study the performances of the proposed formulations for different choices of the linear operators, e.g. number of the atoms and considered ranges for the parameters; despite its importance, it goes beyond the purpose of the present work. Yet, as the current formulations do not fully exploit all the versatility of the framework, we foresee that AMICO can be further enhanced in a number of ways. For instance, it is reasonable to expect similar contributions from each compartment in spatially adjacent voxels; to this end, Total Variation (TV) regularization is a popular choice for promoting spatial smoothness of the coefficients across voxels. Due to magnitude improvement in the fitting speed with AMICO, one could also combine, in the search for finer anatomical details, such regularization approaches with interpolation of the raw dMRI data (Dyrby et al., 2014). Additional priors might be also imposed in each voxel among the compartments, to model more in detail such complex micro-environments. In the case of NODDI, for example, there exists a coupling between the intra-cellular volume fraction ν_{ic} and the shape of the extra-cellular response functions, which depend on both ν_{ic} and κ . This relationship implicitly imposes a structure on the coefficients, partitioning them into disjoint groups such that coefficients in the same group tend to be zeros all together; in these situations, a regularization based on group sparsity (Yuan and Lin, 2006) might improve significantly the reconstructions, as this structure is adequately captured by means of a group ℓ_1 -norm penalty. Future research will be devoted to investigate all these additional forms of regularization.

In this study, we have focused our attention on two specific techniques, i.e. ActiveAx and NODDI, that actually assume only one fiber population in each voxel. Recently, though, some efforts have been made to extend these models and account for more complex intra-voxel fiber configurations; ActiveAx, for example, has been extended to allow axon diameter mapping also in the presence of axonal dispersion (Zhang et al., 2011b) and crossing fibers (Zhang et al., 2011a), whereas Bingham distributions have been introduced in NODDI to account for anisotropic rather than isotropic only dispersion (Tariq et al., 2014), as also done in Kaden et al. (2007) and Sotiropoulos et al. (2012). All these enhancements can be as well integrated into the linear systems (5) and (12), as the flexibility and generality of the proposed framework allow us to conveniently inject additional prior knowledge. Nevertheless, for the sake of clarity and to avoid overloading the exposition, in this work we preferred to concentrate only on the original formulations and leave the inclusion of these extended models to future research, where we will also investigate the effectiveness of using AMICO to linearize other popular microstructure imaging techniques, e.g. CHARMED (Assaf and Basser, 2005), and the wide range of similar compartment models explored in Panagiotaki et al. (2012) and Ferizi et al. (2013).

Another possible area for improvement of the current formulations is the direct incorporation of the proper noise model into the data fidelity term. The implementations presented in this work are in fact based on the ℓ_2 -norm of the residual, i.e. $\|\Phi\mathbf{x} - \mathbf{y}\|_2$, and therefore they implicitly assume a zero-mean Gaussian distribution for the noise. However, it has been demonstrated that this assumption is never met in most practical situations, especially in the case of multichannel acquisitions where the noise has been shown to rather follow Rician or Non-

central χ distributions (Gudbjartsson and Patz, 1995; Dietrich et al., 2008). Recent studies have incorporated these proper noise models in classical SD methods for fiber orientation recovery (Dolui et al., 2012; Canales-Rodríguez et al., under review); future work will be required to adapt these expressions to the microstructure imaging algorithms considered in this work.

Conclusion

The aim of this study was to improve the burdensome fitting time required by existing microstructure imaging techniques in diffusion MRI. We have shown how to conveniently re-formulate these models as simple linear systems that can be solved very efficiently using convex optimization, thus enabling a drastic reduction in the computation time by orders of magnitude, while still preserving the accuracy and robustness of the estimated parameters. We have demonstrated the effectiveness of the AMICO framework for two well-known methods, i.e. ActiveAx and NODDI, but the flexibility of the proposed formulation will allow its extension to the broader range of existing techniques. The availability of such ultrafast fitting algorithms will help the practical application of these models in larger cohorts of patients.

Acknowledgments

This work is supported by the Center for Biomedical Imaging (CIBM) of the Geneva-Lausanne Universities and the EPFL, as well as the foundations Leenaards and Louis-Jeantet.

References

- Aganj, I., Lenglet, C., Sapiro, G., Yacoub, E., Ugurbil, K., Harel, N., 2010. Reconstruction of the orientation distribution function in single- and multiple-shell Q-ball imaging within constant solid angle. *Magn. Reson. Med.* 64, 554–566.
- Alexander, D.C., 2005. Maximum entropy spherical deconvolution for diffusion MRI. *Proc. IPMI* vol. 19, pp. 76–87.
- Alexander, D.C., 2008. A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. *Magn. Reson. Med.* 60 (2), 439–448.
- Alexander, D.C., Hubbard, P.L., Hall, M.G., Moore, E.A., Pitro, M., Parker, G.J.M., Dyrby, T.B., 2010. Orientationally invariant indices of axon diameter and density from diffusion MRI. *NeuroImage* 52 (4), 1374–1389.
- Assaf, Y., Basser, P.J., 2005. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. *NeuroImage* 27 (1), 48–58.
- Assaf, Y., Blumenfeld-Katzir, T., Yovel, Y., Basser, P.J., 2008. AxCaliber: a method for measuring axon diameter distribution from diffusion MRI. *Magn. Reson. Med.* 59 (6), 1347–1354.
- Barmppoutis, A., Hwang, M.S., Howland, D., Forder, J.R., Vemuri, B.C., 2009. Regularized positive-definite fourth order tensor field estimation from DW-MRI. *NeuroImage* 45, 153–162.
- Basser, P., Mattiello, J., LeBihan, D., 1994. MR diffusion tensor spectroscopy and imaging. *Biophys. J.* 66, 259–267.
- Beaulieu, C., 2002. The basis of anisotropic water diffusion in the nervous system — a technical review. *NMR Biomed.* 15 (7–8), 435–455.
- Beck, A., Teboulle, M., 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.* 2 (1), 183–202.
- Behrens, T.E.J., Woolrich, M.W., Jenkinson, M., Johansen-Berg, H., Nunes, R.G., Clare, S., Matthews, P.M., Brady, J.M., Smith, S.M., 2003. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn. Reson. Med.* 50, 1077–1088.
- Canales-Rodríguez, E., Melie-García, L., Iturria-Medina, Y., 2009. Mathematical description of q-space in spherical coordinates: exact q-ball imaging. *Magn. Reson. Med.* 61, 1350–1367.
- Canales-Rodríguez, E.J., Daducci, A., Sotiropoulos, S.N., Caruyer, E., Aja-Fernández, S., Radua, J., Iturria-Medina, Y., Melie-García, L., Alemán-Gómez, Y., Thiran, J.-P., Sarró, S., Pomarol-Clotet, E., Salvador, R., 2014w. Spherical Deconvolution of Multichannel Diffusion MRI Data with Non-Gaussian Noise Models and Total Variation Spatial Regularization, (under review).
- Chang, L.-C., El-Araby, E., Dang, V.Q., Dao, L.H., 2014. GPU acceleration of nonlinear diffusion tensor estimation using CUDA and MPI. *Neurocomputing* 135, 328–338.
- Csisz, I., Tusnády, G., et al., 1984. Information geometry and alternating minimization procedures. *Stat. Decis.* 205–237.
- Daducci, A., Canales-Rodríguez, E.J., Descoteaux, M., Garyfallidis, E., Gur, Y., Lin, Y.-C., Mani, M., Merlet, S., Paquette, M., Ramirez-Manzanares, A., et al., 2014a. Quantitative comparison of reconstruction methods for intra-voxel fiber recovery from diffusion MRI. *IEEE Trans. Med. Imaging* 33 (2), 384–399.
- Daducci, A., Dal Palú, A., Lemkaddem, A., Thiran, J.-P., 2014b. COMMIT: convex optimization modeling for microstructure informed tractography. *IEEE Trans. Med. Imaging* <http://dx.doi.org/10.1109/TMI.2014.2352414> (to appear).

- Dell'Acqua, F., Rizzo, G., Scifo, P., Clarke, R.A., Scotti, G., Fazio, F., 2007. A model-based deconvolution approach to solve fiber crossing in diffusion-weighted MR imaging. *IEEE Trans. Biomed. Eng.* 54, 462–472.
- Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R., 2007. Regularized, fast, and robust analytical Q-ball imaging. *Magn. Reson. Med.* 58 (3), 497–510.
- Descoteaux, M., Deriche, R., Knosche, T., Anwander, A., 2009. Deterministic and probabilistic tractography based on complex fibre orientation distributions. *IEEE Trans. Med. Imaging* 28, 269–286.
- Dietrich, O., Raya, J.G., Reeder, S.B., Ingrisch, M., Reiser, M.F., Schoenberg, S.O., 2008. Influence of multichannel combination, parallel imaging and other reconstruction techniques on MRI noise characteristics. *Magn. Reson. Imaging* 26 (6), 754–762.
- Dolui, S., Kuurstra, A., Michailovich, O.V., 2012. Rician compressed sensing for fast and stable signal reconstruction in diffusion MRI. *Proc. SPIE*, p. 83144Q.
- Dyrby, T.B., Baaré, W.F.C., Alexander, D.C., Jelsing, J., Garde, E., Søgaard, L.V., 2011. An ex vivo imaging pipeline for producing high-quality and high-resolution diffusion-weighted imaging datasets. *Hum. Brain Mapp.* 32 (4), 544–563.
- Dyrby, T.B., Søgaard, L.V., Hall, M.G., Pfito, M., Alexander, D.C., 2013. Contrast and stability of the axon diameter index from microstructure imaging with diffusion MRI. *Magn. Reson. Med.* 70, 711–721.
- Dyrby, T.B., Lundell, H., Burke, M.W., Reislev, N.L., Paulson, O.B., Pfito, M., Siebner, H.R., 2014. Interpolation of diffusion weighted imaging datasets. *NeuroImage* 103, 202–213.
- Ferizi, U., Schneider, T., Panagiotaki, E., Nedjati-Gilani, G., Zhang, H., Wheeler-Kingshott, C.A.M., Alexander, D.C., 2013. A ranking of diffusion MRI compartment models with in vivo human brain data. *Magn. Reson. Med.* <http://dx.doi.org/10.1002/mrm.25080>.
- Figueiredo, M., Nowak, R., Wright, S., 2007. Gradient projection for sparse reconstruction: application to compressed sensing and other inverse problems. *IEEE J. Sel. Top. Sign. Process.* 1 (4), 586–597.
- Golub, G.H., Heath, M., Wahba, G., 1979. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics* 21 (2), 215–223.
- Gudbjartsson, H., Patz, S., 1995. The Rician distribution of noisy MRI data. *Magn. Reson. Med.* 34 (6), 910–914.
- Hall, M.G., Alexander, D.C., 2009. Convergence and parameter choice for Monte-Carlo simulations of diffusion MRI. *IEEE Trans. Med. Imaging* 28 (9), 1354–1364.
- Hansen, P.C., 1999. The L-curve and Its Use in the Numerical Treatment of Inverse Problems. IMM, Department of Mathematical Modelling, Technical University of Denmark.
- Hernández, M., Guerrero, G.D., Cecilia, J.M., García, J.M., Inuggi, A., Jbabdi, S., Behrens, T.E., Sotiropoulos, S.N., 2013. Accelerating fibre orientation estimation from diffusion weighted magnetic resonance imaging using GPUs. *PLoS ONE* 8 (4), e61892.
- Jansons, K., Alexander, D.C., 2003. Persistent angular structure: new insights from diffusion magnetic resonance imaging data. *Inverse Probl.* 19, 1031–1046.
- Jian, B., Vemuri, B.C., 2007. A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. *IEEE Trans. Med. Imaging* 26 (11), 1464–1471.
- Kaden, E., Knösche, T.R., Anwander, A., 2007. Parametric spherical deconvolution: inferring anatomical connectivity using diffusion MR imaging. *NeuroImage* 37 (2), 474–488.
- Kunz, N., Zhang, H., Vasung, L., O'Brien, K.R., Assaf, Y., Lazeyras, F., Alexander, D.C., Hüppi, P.S., 2014. Assessing white matter microstructure of the newborn with multi-shell diffusion MRI and biophysical compartment models. *NeuroImage* 96, 288–299.
- Lamantia, A.S., Rakic, P., 1990. Cytological and quantitative characteristics of four cerebral commissures in the rhesus monkey. *J. Comp. Neurol.* 291 (4), 520–537.
- Landman, B.A., Bogovic, J.A., Wan, H., Elshahaby, F.E.Z., Bazin, P.-L., Prince, J.L., 2012. Resolution of crossing fibers with constrained compressed sensing using diffusion tensor MRI. *NeuroImage* 59 (3), 2175–2186.
- Le Bihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., Laval-Jeantet, M., et al., 1986. MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology* 161 (2), 401–407.
- Lemkaddem, A., Daducci, A., Kunz, N., Lazeyras, F., Seeck, M., Thiran, J.-P., Vulliémot, S., 2014. Connectivity and tissue microstructural alterations in right and left temporal lobe epilepsy revealed by diffusion spectrum imaging. *NeuroImage Clin.* 5 (349–58).
- Michailovich, O., Rathi, Y., Dolui, S., 2011. Spatially regularized compressed sensing for high angular resolution diffusion imaging. *IEEE Trans. Med. Imaging* 30 (5), 1100–1115.
- Neuman, C., 1974. Spin echo of spins diffusing in a bounded medium. *J. Chem. Phys.* 60, 4508.
- Özarslan, E., Shepherd, T., Vemuri, B., Blackband, S., Mareci, T., 2006. Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT). *NeuroImage* 31, 1086–1103.
- Panagiotaki, E., Schneider, T., Siow, B., Hall, M.G., Lythgoe, M.F., Alexander, D.C., 2012. Compartment models of the diffusion MR signal in brain white matter: a taxonomy and comparison. *NeuroImage* 59 (3), 2241–2254.
- Ramirez-Manzanares, A., Rivera, M., Vemuri, B., Carney, P., Mareci, T., 2007. Diffusion basis functions decomposition for estimating white matter intravoxel fiber geometry. *IEEE Trans. Med. Imaging* 26, 1091–1102.
- Ramirez-Manzanares, A., Cook, P.A., Hall, M., Ashtari, M., Gee, J.C., 2011. Resolving axon fiber crossings at clinical b-values: an evaluation study. *Med. Phys.* 38 (9), 5239–5253.
- Schultz, T., Seidel, H.P., 2008. Estimating crossing fibers: a tensor decomposition approach. *IEEE Trans. Vis. Comput. Graph.* 14, 1635–1642.
- Sherbondy, A., Rowe, M., Alexander, D.C., 2010. Microtrack: an algorithm for concurrent projectome and microstructure estimation. *Proc. MICCAI*, pp. 183–190.
- Sotiropoulos, S.N., Behrens, T.E.J., Jbabdi, S., 2012. Ball and rackets: inferring fiber fanning from diffusion-weighted MRI. *NeuroImage* 60 (2), 1412–1425.
- Stanisz, G.J., Szafer, A., Wright, G.A., Henkelman, R.M., 1997. An analytical model of restricted diffusion in bovine optic nerve. *Magn. Reson. Med.* 37 (1), 103–111.
- Szafer, A., Zhong, J., Gore, J.C., 1995. Theoretical model for water diffusion in tissues. *Magn. Reson. Med.* 33 (5), 697–712.
- Tariq, M., Schneider, T., Alexander, D.C., Wheeler-Kingshott, C.A., Zhang, H., 2014. In vivo estimation of dispersion anisotropy of neurites using diffusion MRI. *Proc. MICCAI*, pp. 241–248.
- Tikhonov, A.N., Arsenin, V.Y., 1977. *Solutions of Ill-posed Problems*. Winston & Sons, New York.
- Tournier, J.D., Calamante, F., Gadian, D.G., Connelly, A., 2004. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage* 23, 1176–1185.
- Tournier, J.D., Calamante, F., Connelly, A., 2007. Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* 35, 1459–1472.
- Tuch, D., 2004. Q-ball imaging. *Magn. Reson. Med.* 52, 1358–1372.
- Tuch, D., Reese, T., Wiegell, M., Makris, N., Belliveau, J., Wedeen, V., 2002. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn. Reson. Med.* 48, 577–582.
- Waxman, S.G., 1980. Determinants of conduction velocity in myelinated nerve fibers. *Muscle Nerve* 3 (2), 141–150.
- Wedeen, V., Hagmann, P., Tseng, W.-Y., Reese, T., Weisskoff, R., 2005. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magn. Reson. Med.* 54, 1377–1386.
- Winston, G.P., Micallef, C., Symms, M.R., Alexander, D.C., Duncan, J.S., Zhang, H., 2014. Advanced diffusion imaging sequences could aid assessing patients with focal cortical dysplasia and epilepsy. *Epilepsy Res.* 108 (2), 336–339.
- Yuan, M., Lin, Y., 2006. Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 68 (1), 49–67.
- Zhang, H., Dyrby, T.B., Alexander, D.C., 2011a. Axon diameter mapping in crossing fibers with diffusion MRI. *Proc. MICCAI*, pp. 82–89.
- Zhang, H., Hubbard, P.L., Parker, G.J.M., Alexander, D.C., 2011b. Axon diameter mapping in the presence of orientation dispersion with diffusion MRI. *NeuroImage* 56 (3), 1301–1315.
- Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C., 2012. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage* 61, 1000–1016.