NeuroImage 146 (2017) 507–517

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Convexity-constrained and nonnegativity-constrained spherical factorization in diffusion-weighted imaging

Daan Christiaens[a,d,*], Stefan Sunaert[b,c,d], Paul Suetens[a,d], Frederik Maes[a,d]

[a] KU Leuven, Department of Electrical Engineering, ESAT/PSI, Leuven, Belgium
[b] KU Leuven, Department of Imaging and Pathology, Translational MRI, Leuven, Belgium
[c] UZ Leuven, Department of Radiology, Leuven, Belgium
[d] UZ Leuven, Medical Imaging Research Center, Leuven, Belgium

ARTICLE INFO

Keywords:
Diffusion-weighted imaging
Factorization
Spherical deconvolution
Multi-tissue model
Multi-shell HARDI
Blind source separation

ABSTRACT

Diffusion-weighted imaging (DWI) facilitates probing neural tissue structure non-invasively by measuring its hindrance to water diffusion. Analysis of DWI is typically based on generative signal models for given tissue geometry and microstructural properties. In this work, we generalize multi-tissue spherical deconvolution to a blind source separation problem under convexity and nonnegativity constraints. This spherical factorization approach decomposes multi-shell DWI data, represented in the basis of spherical harmonics, into tissue-specific orientation distribution functions and corresponding response functions, without assuming the latter as known thus fully unsupervised. In healthy human brain data, the resulting components are associated with white matter fibres, grey matter, and cerebrospinal fluid. The factorization results are on par with state-of-the-art supervised methods, as demonstrated also in Monte-Carlo simulations evaluating accuracy and precision of the estimated response functions and orientation distribution functions of each component. In animal data and in the presence of oedema, the proposed factorization is able to recover unseen tissue structure, solely relying on DWI. As such, our method broadens the applicability of spherical deconvolution techniques to exploratory analysis of tissue structure in data where priors are uncertain or hard to define.

1. Introduction

Diffusion-weighted imaging (DWI) is a non-invasive magnetic resonance imaging technique with the unique ability to probe tissue microstructure in vivo, by measuring its hindrance to water diffusion (Le Bihan et al., 1986). The water diffusion process is sensitive to the cellular structure of the surrounding tissue, in particular the presence of cell membranes and intracellular organelles (Beaulieu, 2002). DWI is applied in both neuroscientific research and clinical practice, for studying brain organization, detecting pathology, and measuring disease progression.

The DWI signal can be represented in many ways, including the spherical harmonics (SH) basis (Frank, 2002) and the cumulant expansion (Kiselev, 2010) of which diffusion tensor imaging (DTI) (Basser et al., 1994) is a special case. Parameters such as fractional anisotropy (FA) introduced in the context of such signal representations are sensitive to changes in the underlying tissue microstructure. However, their interpretation at the cellular level is less straightforward.

In an effort to provide more specific measures, a myriad of models have been introduced that relate the measured signal to neural tissue structure. These models typically decompose the diffusion signal into cellular compartments, such as intra- and extra-axonal space or free water (Panagiotaki et al., 2012), weighted by their respective volume fractions. Similarly, nonnegativity-constrained spherical deconvolution (CSD) adopts a single fibre compartment of fixed anisotropy, the fibre response function (RF), which contributes linearly and independently to the DWI signal across all fibre orientations in the voxel (Tournier et al., 2004, 2007). Deconvolution then facilitates estimating the orientation distribution function (ODF) of fibres in that voxel, a metric of apparent fibre density in white matter (Raffelt et al., 2012; Dell'Acqua et al., 2013). CSD was later extended to multi-tissue (MT-

Abbreviation: CNSF, convexity- and nonnegativity-constrained spherical factorization; CSD, constrained spherical deconvolution; CSF, cerebrospinal fluid; DTI, diffusion tensor imaging; DWI, diffusion-weighted imaging; FA, fractional anisotropy; GM, grey matter; HARDI, high angular resolution diffusion imaging; MT-CSD, multi-tissue CSD; NMF, nonnegative matrix factorization; ODF, orientation distribution function; PVE, partial volume effect; QP, quadratic programming; RF, response function; RMS, root-mean-square; SH, spherical harmonics; SNR, signal-to-noise ratio; T1, T_1-weighted image; WM, white matter

* Corresponding author at: KU Leuven, Department of Electrical Engineering, ESAT/PSI, Leuven, Belgium.
E-mail address: daan.christiaens@kuleuven.be (D. Christiaens).

http://dx.doi.org/10.1016/j.neuroimage.2016.10.040
Received 19 April 2016; Accepted 25 October 2016
Available online 27 October 2016

)CSD (Jeurissen et al., 2014), which incorporates partial voluming with adjacent tissues that are not adequately modelled by the fibre response function (Parker et al., 2013; Roine et al., 2014). Each tissue compartment is then characterized by a fixed response function, assumed to be known a priori.

This work generalizes MT-CSD to a blind source separation problem, akin to nonnegative matrix factorization (NMF) (Paatero and Tapper, 1994; Lee and Seung, 1999; Wang and Zhang, 2013). NMF decomposes each input vector as a nonnegative linear combination of unknown source vectors. Similarly, our approach expands the diffusion signal in a basis of response functions, adapted to the tissue structure and to the DWI data at hand. The resulting components can be associated with different normal tissue types and certain types of pathology. As such, our method strikes a balance between signal representation and tissue modelling: it seeks a decomposition that closely represents the data, subject to minimal constraints that give structural interpretation to the component basis functions.

In addition, this method addresses a very practical problem regarding multi-tissue CSD, namely estimating response functions from the data at hand. Originally, white matter (WM) fibre response functions were fitted to the DWI data in a single-fibre mask of high FA, after reorientation of the diffusion tensor eigenvectors (Tournier et al., 2004, 2007). Alternative recursive approaches have been introduced, which segment single-fibre voxels and reorient the data based on the peaks of the fibre ODFs iteratively (Tournier et al., 2013; Tax et al., 2014), or which calibrate the kernel anisotropy in each voxel separately under sparsity constraints (Schultz and Groeschel, 2013). However, these techniques do not directly generalize to other tissue types, such as grey matter (GM) and cerebrospinal fluid (CSF). Current literature therefore relies on tissue segmentation of T_1-weighted images (T1) to define GM and CSF kernels, which requires the T1 to be aligned to the DWI data (Jeurissen et al., 2014). As this is rarely the case in practice, direct DWI tissue segmentation methods have been introduced independently and simultaneously, based on sparsity-constrained NMF (Jeurissen et al., 2015) or convexity-constrained NMF (Christiaens et al., 2015b, Appendix A) of the isotropic mean DWI signal per shell. These methods circumvent T1 requirement and are thus applicable in any reference frame without external input, but still rely on the diffusion tensor model for reorienting the DWI data in each single-fibre voxel. Here, we account for the full anisotropy of the DWI signal by extending NMF to convolution in spherical harmonics.

In related work, Xie et al. (2011) applied NMF to single-shell diffusion tensor data. Reisert et al. (2014) have introduced a more general dictionary learning method that imposes sparsity on the tissue ODFs. In contrast to their approach, we do not impose any constraints on the ODFs except for nonnegativity. Instead, we constrain the tissue RFs to be convex combinations of the data voxels. As such, physical plausibility of the tissue responses is ensured in a purely data-driven manner.

Extending our previous conference paper (Christiaens et al., 2015a), we made improvements to the initialization, the optimization, and the convergence criterion, improving the overall performance and speed of the algorithm. The accuracy and precision of our convexity- and nonnegativity-constrained spherical factorization (CNSF) technique are evaluated in Monte Carlo simulations at various noise levels. In addition, we include results on healthy brain data, both in vivo and ex vivo, and in the presence of pathology, and show that the decomposition can be associated to known anatomy.

2. Method

2.1. Multi-tissue spherical deconvolution

Multi-tissue spherical deconvolution (Tournier et al., 2007; Jeurissen et al., 2014) assumes linear partial volume effect (PVE) to decompose the DWI signal into n tissue components, each of which is the spherical convolution of a response function (RF) and an orientation distribution function (ODF). The response function is an axially symmetric function $H_{t,b}(\theta)$ that characterizes the signal anisotropy and attenuation across b-values for each component t. Each RF is assumed to be spatially-invariant. The ODF $F_t(\theta, \phi)$ is a nonnegative function on the sphere that determines the local directionality and density of that particular component in the voxel. As such, the diffusion signal $S_b(\mathbf{g})$ in each voxel, for gradient direction $\mathbf{g}$ and given b-value, becomes

$$S_b(\mathbf{g}) \approx \sum_{t=1}^{n} (H_{t,b} * F_t)(\mathbf{g}). \tag{1}$$

All functions are commonly represented in the basis of real, symmetric spherical harmonics (SH) of maximum order $\ell_{\max}$ (Tournier et al., 2007; Descoteaux et al., 2009; Jeurissen et al., 2014). As such, the convolution reduces to a multiplication of the coefficients of corresponding order ℓ, i.e., $s_b(\ell, m) = \sum_t \sqrt{\frac{4\pi}{2\ell+1}} h_{t,b}(\ell) f_t(\ell, m)$ with $\ell \in \{0, 2, \ldots, \ell_{\max}\}$ and $m \in [-\ell, \ell]$. The response functions are axially-symmetric, and therefore constrained to the spherical harmonics of phase m=0, known as zonal spherical harmonics.

For this work, we structure the SH coefficients of the DWI signal in tensor $\bar{S}$, indexed by the voxel v and shell b, and rewrite (1) as

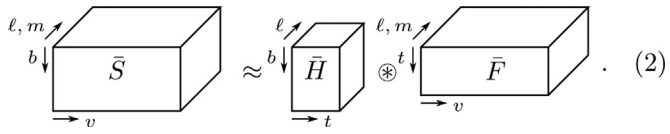

(2)

In this equation, $\bar{H}$ contains the zonal SH coefficients of the response functions, indexed by component t and shell b. $\bar{F}$ contains the SH coefficients of the ODFs, indexed by voxel v and component t. The operator $\circledast$ is introduced to denote spherical convolution in the SH basis, and corresponds to the matrix product of every slice $F_{:,:,(\ell,m)}$ with slice $H_{:,:,\ell}$ of corresponding order ℓ. Note that the $\ell = 0$ coefficients of $\bar{F}$ represent the isotropic volume fraction or density of each tissue.

2.2. Convexity- and nonnegativity-constrained spherical factorization

Considering both the response functions $\bar{H}$ and the ODFs $\bar{F}$ as unknown, expression (2) can be seen as an NMF or blind source separation problem, in which a data matrix is decomposed as the product of a source matrix and a nonnegative weight matrix (Paatero and Tapper, 1994; Lee and Seung, 1999; Wang and Zhang, 2013). In this case, the unknown sources are the response functions of separate components, the weights are the associated ODFs, and we aim to find

$$\bar{H}^\star, \bar{F}^\star = \arg\min_{(\bar{H},\bar{F})} \| \bar{S} - \bar{H} \circledast \bar{F} \|_F^2 \text{ s.t. } A\,\mathbf{f}_{v,t,:} \geq 0. \tag{3}$$

The matrix A evaluates the SH basis across a dense set of directions, to impose nonnegativity of the estimated ODFs denoted by vector slices $\mathbf{f}_{v,t,:}$. The vector $\mathbf{f}_{v,t,:}$ thus contains the SH coefficients $\bar{F}$ at index (v,t) for all (ℓ, m). The only parameters in this framework are the number of components n and the maximal harmonic order $\ell_{\max}$ of each component.

However, the solution to (3) is not unique. As illustrated in Fig. 1, the response functions $\bar{H}$ span a n-gonal simplicial cone in the high-dimensional data space, radiating outwards from the origin 0. Only voxels "within" this cone are represented exactly; data points "outside" this cone give rise to the residual under minimization in (3). As such, any combination of RFs that envelops all observed data points gives rise to a zero residual, but may not necessarily be physically meaningful. Therefore, we impose a convexity constraint (Ding et al., 2010), which ensures that all sources H_t are a convex combination of the measured signal $\bar{S}$ after reorientation. In other words, the convexity constraint ensures that all response functions are observed in the data, typically in voxels with low PVE in both spatial and angular domains.