

Convex Non-negative Spherical Factorization of Multi-Shell Diffusion-Weighted Images

Daan Christiaens^{1,4}, Frederik Maes^{1,4}, Stefan Sunaert^{2,4}, and Paul Suetens^{1,3,4}

¹ KU Leuven, Dept. of Electrical Engineering, ESAT/PSI, Medical Image Computing

² KU Leuven, Dept. of Imaging and Pathology, Translational MRI

³ iMinds, Medical IT Department

⁴ UZ Leuven, Medical Imaging Research Center
Herestraat 49 – 7003, 3000 Leuven, Belgium

Abstract. Diffusion-weighted imaging (DWI) allows to probe tissue microstructure non-invasively and study healthy and diseased white matter (WM) in vivo. Yet, less research has focussed on modelling grey matter (GM), cerebrospinal fluid (CSF) and other tissues. Here, we introduce a fully data-driven approach to spherical deconvolution, based on convex non-negative matrix factorization. Our approach decomposes multi-shell DWI data, represented in the basis of spherical harmonics, into tissue-specific orientation distribution functions and corresponding response functions. We evaluate the proposed method in phantom simulations and in vivo brain images, and demonstrate its ability to reconstruct WM, GM and CSF, unsupervised and solely relying on DWI.

1 Introduction

Diffusion-weighted imaging (DWI) is a non-invasive MRI modality that allows to probe tissue microstructure in vivo by measuring its hinderance to the diffusion process of water [1]. DWI has found applications in neuroscientific research into brain organization and disease processes, as well as clinical applications such as diagnosing acute stroke and neurosurgical planning based on tractography.

Since it was recognised that diffusion anisotropy is indicative of axonal structure in the nervous system [2], research has largely focussed on reconstructing fibre orientation and connectivity in white matter (WM). On the one hand, a myriad of biophysical models have been introduced [3], which aim to estimate microstructure parameters such as axon diameter or dispersion. On the other hand, spherical deconvolution (SD) approaches estimate the fibre orientation distribution function (ODF), given a fibre response function (RF) that is typically estimated from and hence adapted to the data at hand [4]. However, this calibration of the RF may severely impact the reconstructed fibre ODF [5, 6]. Additionally, the presence of partial volume effects (PVE), originating from adjacent grey matter (GM) and cerebrospinal fluid (CSF), has been shown to affect SD [7]. Hence, the question arises how to accurately represent both WM and other tissues, as well as pathology.

Recent work has extended SD to multiple tissue types [8]. Yet, for calibrating the tissue response functions they rely on a prior tissue segmentation obtained from a T1-weighted MR image, which may not always be available or may not be well aligned to the DWI data, e.g., due to EPI distortion [9]. In this paper, we present a more general take on spherical deconvolution, formulated as convex non-negative matrix factorization (NMF) [10] in the basis of spherical harmonics, in which tissue-specific RFs and ODFs are estimated simultaneously from the DWI data itself, without the need for external input.

NMF has previously been applied to diffusion tensor imaging (DTI) data [11], which we now generalize to (multi-shell) high angular resolution imaging (HARDI). The dictionary learning approach introduced in [12] is also based on NMF but assumes sparsity of the reconstructed ODFs. Here, we do not impose any constraints on the ODFs except for non-negativity. 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.

We evaluate accuracy of the estimated tissue RFs and ODFs in simulations on *in silico* data, and demonstrate the method on real brain data. In addition, we investigate the required number of NMF sources in healthy subjects and in the presence of pathology, and show that they can be associated to known anatomy.

2 Method

Assuming linear partial volume effect (PVE), the diffusion signal in each voxel, for gradient direction \mathbf{g} and given b -value, $S_b(\mathbf{g})$, can be decomposed into N tissues, each of which are characterized by an axially symmetric response function $H_{t,b}(\theta)$. Each tissue t then contributes to the signal in a voxel as the spherical convolution of its RF and an orientation distribution $F_t(\theta, \phi)$, i.e.,

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

When S_b and F_t are represented in the basis of real, symmetric spherical harmonics (SH) of maximum order ℓ_{\max} , and the response functions $H_{t,b}$ are represented as zonal harmonics (phase $m = 0$), 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 $\ell \in \{0, 2, \dots, \ell_{\max}\}$ and $m \in [-\ell, \ell]$ [4, 8]. We then structure the SH coefficients of all voxels v and shells b in the tensor equation

$$\bar{S} \approx \bar{H} \circledast \bar{F} \quad , \quad (1)$$

where \bar{H} contains the RF coefficients and \bar{F} the coefficients of the tissue ODFs. The operator \circledast is used to denote spherical convolution in the SH basis, and corresponds to the matrix product of every slice $F_{\cdot,\cdot,(\ell,m)}$ with slice $H_{\cdot,\cdot,\ell}$ of corresponding order ℓ . For all voxels and tissues, the coefficients $\mathbf{f}_{v,t,\cdot}$ are constrained to represent non-negative functions. Note that the $\ell = 0$ coefficients of \bar{F} represent the isotropic volume fraction or density of each tissue.

In this work, we do not assume that the response functions \bar{H} are known a priori. As such, expression (1) becomes a blind source separation problem in the form of non-negative matrix factorization (NMF) [10], in which a data matrix is decomposed as the product of a source matrix and a non-negative weight matrix. In this case, the unknown sources are the tissue-specific response functions, the weights are the tissue ODFs, and we aim to find

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

where A is a matrix that evaluates the SH basis across a dense set of directions and ensures non-negativity of the tissue ODFs. In addition, we impose that all sources H_t are a convex combination of the measured signal after reorientation. To achieve this, we compute an auxiliary tensor \bar{Z} that contains, for each voxel v , the coefficients of the best fitting zonal harmonics to the data S_v across all possible orientations of the symmetry axis. The tissue RFs are then represented as $h_{t,b,\ell} = \mathbf{z}_{\cdot,b,\ell} \cdot \mathbf{w}_{t,\cdot}$, with voxel weights $W \geq 0$ and $\|\mathbf{w}_{t,\cdot}\|_1 = 1$.

The resulting convex non-negative factorization problem is solved by alternately solving for \bar{F} and W , starting from an initialization for \bar{H} which we obtain from k -means clustering in the space of best-fitting zonal harmonics \bar{Z} . We solve (2) for the ODFs \bar{F} in each voxel by unfolding all tensors along the dimensions of shells and SH coefficients. This results in the constrained least-squares expression identical to multi-shell multi-tissue spherical deconvolution [8], which is solved with quadratic programming (QP) given the non-negativity constraints on \bar{F} . Subsequently, (2) can be solved for \bar{H} by unfolding across all voxels and SH coefficients. Imposing the convexity constraint furthermore requires unfolding across shells to solve directly for the weights W in one global QP problem. This procedure is repeated for 10 iterations, which we found to be sufficient for convergence (the relative decrease in residual was consistently $< 1\%$).

We implemented the factorization method in Python, using a custom SH basis implementation and CVXOPT (cvxopt.org) for optimizing the QP problems.

3 Experiments and Results

3.1 Monte-Carlo Simulations

We simulated a phantom dataset with the geometry shown in Fig. 1. The phantom contains 3 tissue classes, mimicking WM, GM and CSF regions, each simulated by a ground-truth RF that was originally estimated from in vivo data according to the procedure advised in [8]. The WM ODF is represented at $\ell_{\max} = 10$. GM and CSF are modelled isotropically ($\ell_{\max} = 0$). The original




Fig. 1. (A) Ground truth geometry of simulated WM (*blue*), GM (*green*), CSF (*red*), and WM-GM PVE (*cyan*). (B-D) Estimated ODFs of 3 components at SNR 20.




Fig. 2. Measures of accuracy and precision of the proposed method on the simulated phantom. *Top row:* Response function accuracy as a function of signal-to-noise ratio (SNR). *Bottom row:* ODF accuracy as a function of SNR, restricted to the isotropic volume fraction ($\ell_{\max} = 0$) in case of components 2 and 3. The dashed lines represent the ODF accuracy in case of direct deconvolution with the ground truth kernels.

phantom was generated at 100×100 voxels and subsampled to 20×20 voxels to simulate PVE. Finally, we sampled the signal with the gradient scheme of the Human Connectome Project [13], and added Rician noise for signal-to-noise ratios (SNR) ranging from 5 to ∞ . At each SNR, 100 noise realizations of the data were generated.

We applied the proposed method to estimate $N = 3$ components in the phantom. Accuracy and precision of the RFs and ODFs are measured by the relative root-mean-square (RMS) error of the amplitude. Results for one noise realization at SNR 20 are shown in Fig. 1. Component 1 corresponds well with WM and the fibre directions are nicely recovered. Components 2 and 3 correspond to GM and CSF. Although the estimated ODFs contain spurious peaks, the isotropic density ($\ell = 0$ coefficient) closely matches the isotropic ground truth. The RF and ODF accuracy and precision are plotted in Fig. 2 and confirm these observations. The accuracy of the estimated RFs is good (error $< 5\%$ for $\text{SNR} \geq 10$)




Fig. 3. Model selection for dataset 1. Plot of the Bayesian information criterion for varying no. anisotropic components (blue) and one anisotropic and a range of isotropic components (green), annotated with the volume fractions of the estimated components.

and improves with increasing SNR. The accuracy of the WM ODF and of the density of GM and CSF is in the same range. Notably, our approach outperforms direct SD with the ground-truth RFs in the low SNR range (dashed line). We attribute this to the Rician noise floor, particularly at high b -values, which the joint RF estimation may be able to model, whereas direct SD can not.

3.2 Results on Human Brain Data

We evaluate our method in data of two neurologically healthy subjects and in one post-surgery image of a patient who suffered a grade IV glioma in the right temporal lobe. Dataset 1 was provided by the Human Connectome Project [13]. Dataset 2 and 3 were acquired in house on a Philips Achieva 3T system, using an isotropic voxel size of 2.5 mm, and 10, 25, 40, 75 gradients at b -values 0, 700, 1000, and 2800 s/mm² respectively. Dataset 2 was corrected for motion, eddy current, and EPI distortions using reverse-phase encoding in the b_0 [9].

In order to keep the computational requirements low, we apply our method in a random subset of 1000 voxels, which takes about 8 min. Afterwards, the ODFs of the entire dataset are computed as the deconvolution with the estimated RFs. The estimated RFs were verified to be robust to this subsampling (e.g., the precision of the WM response on 50 runs is $< 5\%$).

First, we investigate the required number of sources and their maximal SH order for the healthy human brain. We vary the number of sources from 1 to 5 and compare two cases, one in which all sources are modelled up to $\ell_{\max} = 10$, and one in which only the first is modelled in this way and all others are modelled isotropically. We use the Bayesian information criterion (BIC) [14] to find the optimal balance between model fit and model complexity. The result is plotted in Fig. 3 and shows that minimal BIC is obtained with one anisotropic and at least two isotropic components. Slices of the corresponding tissue volume fractions are shown as well and allow for a better interpretation of this result. The first 3 components can be associated to WM, GM and CSF. Additional components can not be related to known anatomy. Similar results were obtained in dataset 2.




Fig. 4. Estimated response functions for dataset 2. From left to right: (A) anisotropic response function, and (B) radial component of all response functions (full lines), compared to the WM, GM and CSF response estimated as in [8] (dashed lines); (C) residual compared to the supervised method; (D) weights of the response estimation.




Fig. 5. Estimated tissue ODFs. (A-D) for dataset 1. (E-H) for dataset 2.

Based on the optimal model complexity, we apply the proposed method to decompose both DWI datasets into 3 tissue classes, two of which are modelled isotropically. The response functions of dataset 2 are depicted on the left in Fig. 4, and compared to the response functions of WM, GM, and CSF, estimated as described in [8] based on a prior segmentation of the T1 image (dashed lines). The RFs estimated in our (unsupervised) method qualitatively correspond to those obtained from the supervised method, up to a scaling factor. Moreover, we consistently obtained a smaller residual fitting error given the same model complexity, as shown in Fig. 4. Notice how the algorithm converges after about 5 iterations. The graphs on the right of Fig. 4 show the weights W by which each voxel contributes to the estimated kernels. Interestingly, these are naturally sparse in convex NMF, as pointed out in [10]. In Fig. 5, we show the tissue ODFs for both datasets. The ODF of component 1 closely resembles the fibre ODF in [8] and nicely captures partial voluming at the WM-GM interface.




Fig. 6. Results of dataset 3. (A-B) T1 and T2. (C-F) Tissue ODFs of 4 components. (G-H) Close-up of the WM ODF, overlaid onto the estimated edema density, and T2.

For dataset 3, the model selection as described above suggested one anisotropic and 3 isotropic kernels, and the resulting tissue ODFs are shown in Fig. 6. Results indicate that, in addition to WM, GM, and CSF, edema is detected as a 4th component. Notice that CSF is detected in the surgical cavity.

4 Discussion and Conclusion

In this work, we introduced a generalization of multi-tissue SD as a blind source separation problem, formulated as convex NMF in the SH basis. Like SD, our approach assumes non-negativity of the tissue ODFs and spatial invariance of their RFs, but jointly optimizes the RFs instead of assuming them as known.

As such, our method is an unsupervised RF estimation method that, in contrast to [8], does not require (alignment to) a T1-segmentation. Instead, we impose the convexity constraint to ensure that the RF is actually observed in the data, typically in a sparse set of voxels. Not only does this broaden the application of multi-tissue SD to DWI data that could not be corrected for distortions or for which no T1 is available, ultimately a better fit to the data was obtained.

For the normal brain, we found optimal model complexity with one anisotropic and two isotropic RFs, which can be associated to WM, GM, and CSF, and agrees with the model adopted in [8]. Instead of imposing such model explicitly, our method is more flexible in that it allows to estimate the most suitable number of sources and their SH order from the data itself. For instance in brain tumour data, our method detected the need for a 4th component that could be associated to edema. We expect that, by explicitly accounting for edema, the reconstructed WM ODF is less contaminated by the increased diffusion, which may facilitate tracing WM fibre bundles through edema. Possibly, this approach can be extended to characterize tissue heterogeneity within the tumour itself.

In future work, the presented method may be applied to explore tissue microstructure and growth patterns in brain tumours, as well as in other organs. Furthermore, a straightforward extension to multiple subjects can be used for group analyses of patients and controls.

Acknowledgements. D. Christiaens is supported by a Ph.D. grant of the Agency for Innovation by Science and Technology (IWT). Data were provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

1. Le Bihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., Laval-Jeantet, M.: MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology* 161(2), 401–407 (1986)
2. Beaulieu, C.: The basis of anisotropic water diffusion in the nervous system – a technical review. *NMR in Biomedicine* 15(7-8), 435–455 (2002)
3. Panagiotaki, E., Schneider, T., Siow, B., Hall, M.G., Lythgoe, M.F., Alexander, D.C.: Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison. *NeuroImage* 59(3), 2241–2254 (2012)
4. Tournier, J.-D., Calamante, F., Connelly, A.: Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* 35(4), 1459–1472 (2007)
5. Parker, G.D., Marshall, D., Rosin, P.L., Drage, N., Richmond, S., Jones, D.K.: A pitfall in the reconstruction of fibre ODFs using spherical deconvolution of diffusion MRI data. *NeuroImage* 65, 433–448 (2013)
6. Schultz, T., Groeschel, S.: Auto-calibrating spherical deconvolution based on ODF sparsity. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013, Part I. LNCS, vol. 8149, pp. 663–670. Springer, Heidelberg (2013)
7. Roine, T., Jeurissen, B., Perrone, D., Aelterman, J., Leemans, A., Philips, W., Sijbers, J.: Isotropic non-white matter partial volume effects in constrained spherical deconvolution. *Frontiers in Neuroinformatics* 8 (2014)
8. Jeurissen, B., Tournier, J.-D., Dhollander, T., Connelly, A., Sijbers, J.: Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data. *NeuroImage* 103, 411–426 (2014)
9. Andersson, J.L.R., Skare, S., Ashburner, J.: How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *NeuroImage* 20(2), 870–888 (2003)
10. Ding, C.H.Q., Li, T., Jordan, M.I.: Convex and semi-nonnegative matrix factorizations. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 32(1), 45–55 (2010)
11. Xie, Y., Ho, J., Vemuri, B.: Nonnegative factorization of diffusion tensor images and its applications. In: Székely, G., Hahn, H. (eds.) *Information Processing in Medical Imaging*. LNCS, vol. 6801, pp. 550–561. Springer, Heidelberg (2011)
12. Reisert, M., Skibbe, H., Kiselev, V.G.: The diffusion dictionary in the human brain is short: Rotation invariant learning of basis functions. In: Schultz, T., Nedjati-Gilani, G., Venkataraman, A., O'Donnell, L., Panagiotaki, E. (eds.) *Computational Diffusion MRI and Brain Connectivity*. Mathematics and Visualization, pp. 47–55. Springer, Heidelberg (2014)
13. Van Essen, D., Smith, S., Barch, D., Behrens, T., Yacoub, E., Ugurbil, K.: The WU-Minn human connectome project: An overview. *NeuroImage* 80, 62–79 (2013)
14. Schwarz, G.: Estimating the dimension of a model. *The Annals of Statistics* 6(2), 461–464 (1978)