ℓ_0-deconvolution for compressive diffusion MRI

Alessandro Daducci*, Anna Auría*, Jean-Philippe Thiran*†, Yves Wiaux*†‡
* Signal Processing Lab (LTS5), École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
† University Hospital Center (CHUV) and University of Lausanne (UNIL), Switzerland
‡ Department of Radiology and Medical Informatics, University of Geneva (UniGE), Geneva, Switzerland

Diffusion MRI is a well established imaging modality providing a powerful way to non-invasively probe the structure of the white matter. Despite the potential of the technique, the intrinsic long scan times of these sequences have hampered their use in clinical practice. For this reason, a wide variety of methods have been proposed to shorten acquisition times. Among them, *spherical deconvolution approaches* have gained a lot of interest for their ability to reliably recover the intra-voxel distribution of neuronal fiber orientations (FOD), represented as a function on the sphere $\boldsymbol{x}$, with a relatively small number of samples. To overcome the ill-posedness of deconvolution, these methods make use of regularization schemes generally based on the assumption that the FOD is sparse due to the small number of fiber populations present in each voxel. On one hand, the well-known Constrained Spherical Deconvolution [1] approach (herein `CSD`) relies on an ℓ_2 prior which presents the drawback of not promoting sparsity explicitly. On the other hand, convex optimization methods have recently been advocated in a compressed sensing perspective. A recent approach [2] (herein `L2L1`) relies on some ℓ_1 minimization which unfortunately conflicts with the physical constraint that the fiber compartments must sum to unity: $||\boldsymbol{x}||_1 = 1$.

We here review a recent work [3] where we propose to further exploit the versatility of compressed sensing and convex optimization with the aim to characterize the FOD sparsity more optimally. We re-formulate the spherical deconvolution problem as a constrained ℓ_0 minimization:

$$\underset{\boldsymbol{x}\geq 0}{\operatorname{argmin}} \; ||\Phi\boldsymbol{x} - \boldsymbol{y}||_2^2 \;\; \text{subject to} \;\; ||\boldsymbol{x}||_0 \leq k, \qquad (1)$$

where $||\cdot||_0$ explicitly counts the number of nonzero coefficients and k represents the expected number of fiber populations in a voxel. We call this formulation *Reweighted Sparse Deconvolution* or `RSD`.
Surely, ℓ_0 problems as such are intractable. However, the reweighting scheme proposed in [4] preserves the tractability of the problem by sequentially solving weighted ℓ_1 problems of the form (1), where the ℓ_0 norm is substituted by a weighted ℓ_1 norm defined as $||\boldsymbol{w}\boldsymbol{\alpha}||_1 = \sum_i \boldsymbol{w}_i |\boldsymbol{\alpha}_i|$, for positive weights $\boldsymbol{w}_i$ and where i indexes vector components. At each iteration, the weights are essentially set as the inverse of the values of the solution of the previous problem, i.e. $\boldsymbol{w}_i^{(t)} \approx 1/\boldsymbol{x}_i^{(t-1)}$. At convergence, this set of weights makes the weighted ℓ_1 norm independent of the precise value of the nonzero components, thus mimicking the ℓ_0 norm behavior.

To test our reconstruction method, we compared it against `CSD` and `L2L1` on two human brain datasets acquired using standard clinical protocols, respectively using 30 and 20 diffusion directions, with $b = 1000$ s/mm^2. The results are reported in the Figure. Subplots A, B and C correspond to the dataset acquired using 30 samples. Even though the acquisition scheme used for this dataset is not the setting where our numerical simulations (not reported here for brevity) highlighted the most substantial differences between the three approaches, important conclusions can be drawn. First, the ability of both `L2L1` (center plots) and `RSD` (righthand plots) to properly model the isotropic compartment in voxels with full or partial contamination with CSF is clearly visible. Moreover, comparing B and C we can observe that `RSD` successfully differentiates gray matter (light gray regions) from CSF voxels with pure isotropic and fast diffusion (very bright areas), while `L2L1` appears unable to distinguish them. The yellow frames in the figures highlight the corona radiata, a well-known region in the white matter containing challenging crossing fibers. We observe that `RSD` clearly results in sharper and more defined profiles than `L2L1`, whereas the improvements with respect to `CSD` (lefthand plots) are confined only to few voxels.

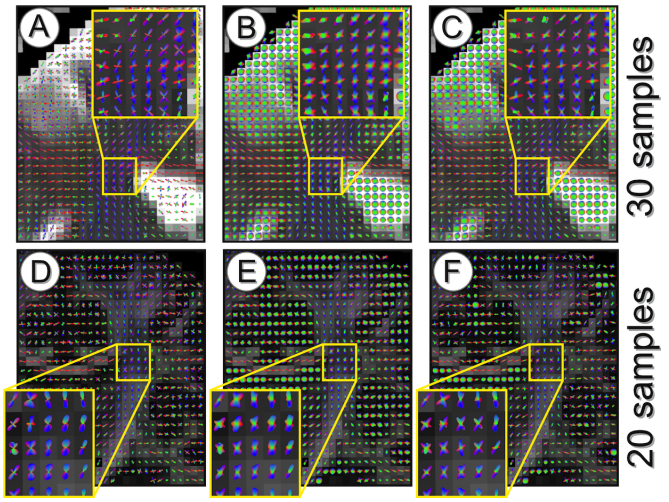

The performances of the three methods sensibly change when decreasing the acquisition samples to 20 (subplots D, E and F). Reconstructions with `RSD` are definitely much better resolved than both `CSD` and `L2L1`. In fact `CSD` clearly breaks, missing many fiber compartments probably due to limitations in the internal Spherical Harmonics representation. The same deterioration happens to `L2L1`, whose reconstructions appear very blurred and noisy. These results show that our proposed regularization scheme is indeed very effective and that the improvements are most remarkable in a *high q-space under-sampling regime*, thus opening the way for a further scan time reduction of high angular resolution acquisitions closer to DTI.

References

[1] J. Tournier, F. Calamante, D. Gadian, and A. Connelly, "Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution." *Neuroimage*, vol. 23, pp. 1176–85, 2004.
[2] B. Landman, J. Bogovic, H. Wan, F. Elshahaby, P. Bazin, and J. Prince, "Resolution of crossing fibers with constrained compressed sensing using diffusion tensor MRI," *Neuroimage*, vol. 59, pp. 2175–86, 2012.
[3] A. Daducci, D. Van De Ville, J.-P. Thiran, and Y. Wiaux, "Reweighted sparse deconvolution for high angular resolution diffusion MRI," 2012, pre-print in arXiv:1208.2247, submitted to Neuroimage.
[4] E. Candès, M. Wakin, and S. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *J Fourier Anal Appl*, vol. 14(5), pp. 877–905, 2008.