Fast DSI Acquisition and Reconstruction Based on Sparse Diffusion Propagator Representations

Antonio Tristán-Vega[1], and Carl-Fredrik Westin[1]
[1]*Laboratory of Mathematics in Imaging, Brigham and Women's Hospital, Boston, Massachusetts, United States*

Motivation: Estimating the Ensemble Average Propagator (EAP) via Diffusion Spectrum Imaging (DSI) is unpractical due to the need to acquire large amounts of data. Compressed Sensing (CS) grants the possibility of considerably reducing the number of samples required to describe a signal far below its Nyquist rate [1], assuming it is highly sparse. Unfortunately, it results hard to represent the EAP in a basis providing sparse enough representations. We address this limitation by (1) introducing a novel acquisition model that explicitly isolates the non-sparse residual of the EAP, which is assumed to be a low energy, noise-like signal; (2) proposing a suitable wavelet basis to sparsely represent the EAP, as opposed to the sparse-in-nature approach considered in [2]. Our model, together with two novel reconstruction methods proposed here (based on either ℓ_0 or ℓ_1 optimization), allows improving the reconstruction accuracy of the EAP in realistic scenarios with respect to the CS approach.

Methods: Since the EAP and the acquired signal $E(\mathbf{q})$ are related as a pair of direct/inverse Fourier transforms, $P(\mathbf{R}) = \iiint_{\mathbb{R}^3} E(\mathbf{q}) \exp(-j2\pi\mathbf{q}\cdot\mathbf{R}) d\mathbf{q}$, we arrange the discretized EAP as an $N \times 1$ vector $\mathbf{x}$, and the sampled $E(\mathbf{q})$ as an $M \times 1$ vector $\mathbf{y}$, related by the 3-D DFT operator $\mathfrak{F}$. We assume $\mathbf{x}$ is nearly sparse in some frame represented by the $N \times P$ matrix Φ ($P \geq N$), with coefficients arranged in a $P \times 1$ vector $\mathbf{a}$ such that: $\mathbf{y} = \mathbf{S}\mathfrak{F}\mathbf{x} + \mathbf{w}; \Phi^{-1}\mathbf{x} = \mathbf{a} + \mathbf{r}$, where Φ^{-1} is the (pseudo) inverse of Φ and $\mathbf{S}$ is an $M \times N$ selection matrix whose rows are all zeros except for one that selects the wave vector sampled. The vector $\mathbf{w}$ accounts for the noise in the acquisition, meanwhile $\mathbf{r}$ is the non-sparse residual, i.e. the part of the signal that cannot be sparsely represented in the frame Φ. Hence, we do not need the columns of Φ to be faithful replicas of the expected $\mathbf{x}$, but only that most of the energy of $\mathbf{x}$ is concentrated in very few terms. Since wavelet bases usually fulfill this condition, we are provided with a wide pool of useful Φ. The model concretes in the following minimization framework [3]:

$$(\mathbf{x}^*, \mathbf{a}^*) = \arg\min_{\mathbf{x},\ \mathbf{a}} \{ \|\mathbf{a}\|_p + \overbrace{\lambda^{-1}\|\mathbf{y} - \mathbf{S}\mathfrak{F}\mathbf{x}\|_2^2}^{\text{noise}} + \overbrace{\mu^{-1}\left\|\Phi^{-1}\mathbf{x} - \mathbf{a}\right\|_2^2}^{\text{non-sparse residual}} \},$$

with $\lambda, \mu > 0$. The pseudo-norm $\|\mathbf{a}\|_p, 0 \leq p \leq 1$, is the term enforcing sparsity [1].

Numerics: The optimization problem is solved with a two-step, iterative algorithm, both for $p = 0$ and $p = 1$, assuming Φ is a Parseval frame. **Step 1:** For a fixed approximation of $\mathbf{a}$, say $\mathbf{a}_n$, we have a quadratic problem in $\mathbf{x}$, for which we obtain the next approximation $\mathbf{x}_{n+1} = \left(\mathbf{I}_N + \frac{\mu}{\lambda}\Re\left\{\mathfrak{F}^H\mathbf{S}^T\mathbf{S}\mathfrak{F}\right\}\right)^{-1}\left(\frac{\mu}{\lambda}\mathfrak{F}^H\mathbf{S}^T\mathbf{y} + \Phi\mathbf{a}_n\right)$. Whenever $\mu < \lambda$, and because $\mathfrak{F}$ is Hermitian, this expression may be proven to be equivalent to: $\mathbf{x}_{n+1} = \frac{\mu}{\mu+\lambda}\mathfrak{F}^H\mathbf{S}^T\mathbf{y} + \left(\mathbf{I}_N - \frac{\mu}{\mu+\lambda}\mathfrak{F}^H\mathbf{S}^T\mathbf{S}\mathfrak{F}\right)\Phi\mathbf{a}_n$, which reduces to 3-D DFT operations, wavelet reconstructions, and wave-vector selections. **Step 2:** For a fixed approximation $\mathbf{x}_{n+1}$, we have an ℓ_p problem that is solved with the standard hard (ℓ_0) or soft (ℓ_1) thresholding operations: $\mathbf{a}_{n+1} = \Theta_{\sqrt{\mu}}^{\text{hard}}\left(\Phi^T\mathbf{x}_{n+1}\right)$, $\mathbf{a}_{n+1} = \Theta_{\mu/2}^{\text{soft}}\left(\Phi^T\mathbf{x}_{n+1}\right)$. **Summary:** By alternating steps 1 and 2 until convergence, it may be formally proven that the ℓ_0 problem converges to a <u>local</u> minimum and the ℓ_1 to a <u>global</u> minimum of the solution whenever $\mu < \lambda$.

ϵ (%)		ℓ_0 algo.			ℓ_1 algo.			[2]		
$1/\sigma$	M'	can	mey	sym	can	mey	sym	can	mey	sym
	64	28	58	36	**22**	60	37	32	69	43
5	128	32	41	24	29	38	**20**	32	50	32
	256	43	43	23	39	31	**17**	42	46	40
	64	15	50	31	**13**	30	17	18	17	14
10	128	**13**	30	17	15	30	**13**	15	32	16
	256	18	17	12	13	17	**9**	14	18	10
	64	10	47	28	**4**	51	27	5	53	28
30	128	6	27	15	**3**	22	11	5	23	12
	256	5	8	7	**3**	7	5	4	8	5

Table 1: Relative errors in the reconstructed EAP.

Experiments: We simulate a mixture of $K = 2, 3$ Gaussians with equal partial volumes and eigenvalues $[1.7, 0.3, 0.3] \cdot 10^{-3}$ s/mm^2 (results are shown only for $K = 2$; for $K = 3$), rotated to a random 3-D orientation. The sampling frequency is chosen such that the maximum *b*-value is $10,000$ s/mm^2. $E(\mathbf{q})$ and $P(\mathbf{R})$ are discretized into 16×16×16 matrices. $M' \ll 16 \cdot 16 \cdot 16$ samples are randomly kept for $E(\mathbf{q})$ using a binomial law [2], and then contaminated with Rician noise for $\text{SNR} = 1/\sigma = 5, 10, 30$. The relative error between the reconstructed EAP and the ground-truth is computed. We have restricted Φ to discrete, orthogonal wavelets: discrete Meyer wavelets (<u>mey</u>) are infinitely smooth, while symlets (<u>sym</u>) are compact-supported. The canonical basis (<u>can</u>) is included for the sake of comparison with [2]. In each scenario, the parameters used are those producing the minimum relative error in the reconstruction.

ℓ_0 algo.			ℓ_1 algo.			[2]		
can	mey	sym	can	mey	sym	can	mey	sym
1.2	0.8	0.5	9	1.6	1.3	13	23	2.2
0.9	1.2	0.8	18	2.6	2.0	36	22	4.3
8.4	6.6	0.8	66	2.1	1.4	54	31.5	8.5
1.5	1.3	0.9	11	12	2.0	10	2.5	2.0
1.7	1.8	1.0	18	14	3.4	29	14	3.5
9.2	4.3	2.2	44	39	6.1	45	39	6.2
2.3	1.3	0.9	5.4	8.1	1.7	5.4	2.1	2.2
5.0	1.8	1.0	9.4	9.0	3.9	9.7	8.0	4.0
4.7	7.4	5.2	17	17	9	26	17	9

Table 2: Percentages of non-zero coefficients of $\mathbf{a}$.

Results and discussion: Table 1 shows the relative reconstruction errors. Overall, <u>sym</u> seems to be the more suitable basis to represent the EAP, unless very high SNR is considered. At the same time, <u>can</u> behavior is not consistent because the EAP is not truly sparse in this basis (see Table 3), hence the CS theory does not hold (the ℓ_1 solution is not guaranteed to approach the sparsest solution) [1]. The situation is also similar for <u>mey</u>. Table 2 shows the sparsity of the solution: consistently with our previous considerations, <u>sym</u> is the basis producing the sparsest solution. Though the ℓ_0 problem provides the sparsest solution, it is worth noticing this is a suboptimal solution (local minimum), and consequently the final error is larger than that of ℓ_1. Note the sparsity of the ℓ_1 problem is similar to that of the ℓ_0 problem only for <u>mey</u> and <u>sym</u> (for the most representative SNR 5 and 10): as expected after [1], both solutions will be similar only in case the global ℓ_0 solution is highly sparse. In any case, note the inclusion of the non-sparse residual greatly improves the actual sparsity of the solution over the pure CS approach in [2], and accordingly we are able to improve the reconstruction errors in [2] for realistic SNR. Contrary to the claim in [2], we may conclude that the EAP cannot be considered strictly sparse in the canonical basis (though our model is able to cope with this situation to some extent), and that <u>sym</u> provide a suitable frame to represent the EAP. Fig. 1 graphically illustrates these considerations in a representative scenario: symlets are able to properly detect the diffusion directions with less outliers than the canonical basis.

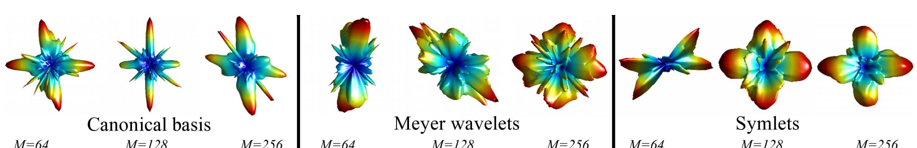

Figure 1: Examples of ODFs computed from the reconstructed EAP in a noisy scenario with our ℓ_1 method. Two fibers crossing in an angle of 90º with equal partial volume fractions have been simulated.

[1] Donoho, D.L.: Compressed sensing. IEEE Trans. Info. Th. 52(4) (2006) 1289–1306.

[2] Merlet, S., Deriche, R.: Compressed sensing for accelerated EAP recovery in diffusion MRI. In: MICCAI Workshop Comp. Diffusion MRI, Beijing (China) (2010).

[3] Portilla, J.: Image restoration through l0 analysis-based spare optimization in tight frames. In: Procs. IEEE Intl. Conf. Im. Proc., Cairo (2009) 3909–3912.