NONNEGATIVE ODF ESTIMATION VIA OPTIMAL CONSTRAINT SELECTION

Sören Wolfers, Evan Schwab, and René Vidal

Center for Imaging Science, Johns Hopkins University

ABSTRACT

We consider the problem of estimating a nonnegative orientation distribution function (ODF) from high angular resolution diffusion images. Since enforcing nonnegativity of the ODF for all directions on the sphere leads to an optimization problem with infinitely many constraints, prior work cannot guarantee the nonnegativity of the estimated ODF. The first contribution of this paper is to show that, under certain conditions, a single constraint is sufficient to guarantee the nonnegativity of the estimated ODF in all directions. Otherwise, when these conditions are violated, we propose an iterative algorithm that enforces one constraint at a time and is guaranteed to converge to the optimal nonnegative ODF. Experiments on synthetic and real data show that our methods produce more accurate solutions than prior work at a reduced runtime.

Index Terms— Diffusion MRI, HARDI, estimation of nonnegative ODFs, semi-infinite optimization.

1. INTRODUCTION

High-angular resolution diffusion imaging (HARDI) [1] is a non-invasive imaging technique that can be used to study the architecture of the brain. This can be done by estimating the local orientation of white matter fiber tracts as the peaks of the orientation distribution function (ODF), which measures the probability of water diffusion in each direction on the sphere.

Current methods for estimating the ODF from HARDI data represent the ODF in terms of a spherical harmonic (SH) basis [2] and solve for the SH coefficients using least squares [3, 4]. However, ODFs estimated by least squares from noisy HARDI data may have negative values, which is physically and theoretically incorrect for a probability density function.

Solving for an ODF subject to nonnegativity constraints requires solving an optimization problem with infinitely many constraints, one per each direction on the sphere. The work of [5] aims to address this issue by enforcing nonnegativity at a finite number of fixed directions. However, this does not guarantee the nonnegativity of the ODF in all directions. A number of methods have attempted to go further and enforce nonnegativity on the continuous domain. The work of [6] attempts to iteratively enforce nonnegativity of the fiber orientation distribution (FOD) function, but their model does not guarantee nonnegativity everywhere. The work of [7] esti-

mates the SH coefficients of the square root of the ODF and then simply squares the resulting function to obtain a nonnegative ODF. However, this result need not coincide with the optimal nonnegative ODF. The work of [8] proposes to enforce that a linear transformation of the homogeneous polynomial basis be positive definite. However, this requires a full rank assumption, which may not guarantee nonnegativity everywhere. The work of [9] shows that enforcing nonnegativity in all directions is equivalent to enforcing the positive semi-definiteness of an infinite matrix built from the SH coefficients. However, we will see that their solution does not coincide with the optimal solution of the original optimization problem with infinite constraints.

In this paper, we show that, under certain conditions, the nonnegativity of the ODF can be guaranteed by enforcing only *one* optimal constraint. Otherwise, nonnegativity of the ODF can be enforced iteratively by adding the most violated constraints one by one until convergence. Our experiments on synthetic and real HARDI data show that our methods not only reduce the amount of negativity of existing methods, but also produce ODFs that are much closer to the optimum of the original optimization problem with infinite constrains at a significantly reduced runtime.

2. ODF ESTIMATION PROBLEM

2.1. Spherical Harmonic Representation

We model the distribution of water diffusion at each voxel with an ODF, i.e., a probability density function on the sphere, $p(x)$ for $x \in \mathbb{S}^2$. The ODF is related to the HARDI signal $S(g)$ along a spatial gradient direction $g \in \mathbb{S}^2$ by [4]

$$p(x) = \frac{1}{4\pi} + \frac{1}{16\pi^2} \operatorname{FRT} \left\{ \nabla_b^2 \ln(-\ln(\frac{S(g)}{S_0})) \right\}, \quad (1)$$

where FRT is the Funk-Radon transform, ∇_b^2 is the Laplace-Beltrami operator on $\mathbb{S}^2$, and S_0 is the baseline MRI signal.

The transformed signal $s(g) \triangleq \ln(-\ln(\frac{S(g)}{S_0}))$ is a real spherical function, which can be written as a linear combination of modified SH basis functions $Y_j : \mathbb{S}^2 \to \mathbb{R}$ (with notation as in [2]) as $s(g) = \sum_{j=1}^{R} c_j Y_j(g)$, where $\mathbf{c} \triangleq [c_j]_{j=1}^{R}$ are the SH coefficients with respect to a modified SH basis of order up to L, and $R = \frac{(L+1)(L+2)}{2}$. Recall that the modified SH basis functions Y_j are eigenfunctions of the Funk-

Radon transform and of the Laplace-Beltrami operator, i.e., $\mathrm{FRT}(Y_j) = \lambda_j Y_j$ and $\nabla_b^2(Y_j) = \ell_j Y_j$. Thus, we can express $p(x)$ as a linear combination of modified SH functions as

$$p(x; \mathbf{c}) = \frac{1}{4\pi} + \frac{1}{16\pi^2} \sum_{j=1}^{R} c_j \lambda_j \ell_j Y_j(x). \qquad (2)$$

Since Y_j integrates to 0 over $\mathbb{S}^2$, $p(x; \mathbf{c})$ integrates to 1.

2.2. Nonnegative ODF Estimation Problem

To estimate the ODF $p(x; \mathbf{c})$ from samples of the HARDI signal $s(g)$, we must recover the SH coefficient vector $\mathbf{c}$ subject to the constraint that $p(x; \mathbf{c})$ be nonnegative for all $x \in \mathbb{S}^2$. This leads to the *nonnegative ODF estimation problem*:

$$\min_{\mathbf{c} \in \mathbb{R}^R} \|\mathbf{Bc} - \mathbf{s}\|_2^2 \quad \text{s.t.} \quad p(x; \mathbf{c}) \geq 0 \quad \forall x \in \mathbb{S}^2, \qquad (\mathrm{P}_\infty)$$

where $\mathbf{B} \triangleq [Y_j(\mathbf{g})]_{j=1}^R$ is the $G \times R$ matrix of SH basis functions sampled at G fixed gradient directions $\mathbf{g} \triangleq [g_i]_{i=1}^G$ and $\mathbf{s} = s(\mathbf{g}) \in \mathbb{R}^R$ is the sampled HARDI signal. We assume that $\mathbf{B}$ is of full column rank. We refer to this problem as (P_∞) since we have a continuum of infinitely many constraints.

2.3. Previous ODF Estimation Methods

Previous ODF estimation methods [3, 4] recover the coefficient vector $\mathbf{c}$ by disregarding the nonnegativity constraints, which leads to the following least-squares (LS) solution

$$\mathbf{c}_0 = \arg\min_{\mathbf{c} \in \mathbb{R}^R} \|\mathbf{Bc} - \mathbf{s}\|^2 = (\mathbf{B}^\mathsf{T}\mathbf{B})^{-1}\mathbf{B}^\mathsf{T}\mathbf{s}. \qquad (\mathrm{LS})$$

The approximate method of [5] enforces nonnegativity at a finite number of fixed grid points $\mathbf{x} \triangleq [x_i]_{i=1}^M$ and solves a quadratic optimization problem with a finite set of constraints:

$$\min_{\mathbf{c} \in \mathbb{R}^R} \|\mathbf{Bc} - \mathbf{s}\|_2^2 \quad \text{s.t.} \quad p(x_i; \mathbf{c}) \geq 0, \ i = 1, \dots, M. \qquad (\mathrm{DC})$$

We refer to this method as Discretely Constrained (DC). However, notice that enforcing nonnegativity at a finite set of fixed x_i's does not guarantee nonnegativity for all $x \in \mathbb{S}^2$.

3. PROPOSED ODF ESTIMATION METHODS

In this paper, we propose (to the best of our knowledge) the first method that is able to guarantee nonnegativity of the optimal ODF in all directions on the sphere. The key idea behind our approach is that (P_∞) is equivalent to

$$\min_{\mathbf{c} \in \mathbb{R}^R} \|\mathbf{Bc} - \mathbf{s}\|_2^2 \quad \text{s.t.} \quad \min_{x \in \mathbb{S}^2} p(x; \mathbf{c}) \geq 0. \qquad (3)$$

Therefore, in theory, we only need to enforce *one* constraint, namely $p(x^*, \mathbf{c}) \geq 0$, where $x^* = \arg\min_x p(x; \mathbf{c})$. In practice, however, we face two challenges. First, x^* depends on $\mathbf{c}$, which is unknown. Thus, we cannot compute x^* without knowing $\mathbf{c}$ in the first place. Second, x^* may not be unique. This means that we may need to enforce more than one constraint, one per global minimizer. In what follows, we show

how to address each of these issues. In §3.2 we show that if there is a unique active constraint at the optimum, this constraint can be found by solving an optimization problem in $\mathbb{S}^2$. Then, in §3.3 we show that if the minimizer of (P_∞) is not unique, one can still find the set of active constraints through a provably convergent iterative procedure. Before proceeding further, in §3.1 we study the geometry of the set of constraints.

3.1. Geometry of the Set of Constraints

Let us consider the general optimization problem

$$\min_{\mathbf{c}} f(\mathbf{c}) \quad \text{s.t.} \quad \mathbf{n}^\mathsf{T}\mathbf{c} \leq b \quad \forall (\mathbf{n}, b) \in C, \qquad (\mathrm{P}0)$$

where $f\colon \mathbb{R}^R \to \mathbb{R}$ is strictly convex and C is compact. Since for each $x \in \mathbb{S}^2$, the constraint $p(x; \mathbf{c}) \geq 0$ is a half-space of the form $\mathbf{n}^\mathsf{T}\mathbf{c} \leq b$ with normal vector $\mathbf{n} = -\tilde{\mathbf{Y}}(x) \triangleq [\tilde{Y}_j(x)]_{j=1}^R \in \mathbb{R}^R$, where $\tilde{Y}_j(x) \triangleq \lambda_j \ell_j Y_j(x)$, and intercept $b = 4\pi \in \mathbb{R}$, we see that (P_∞) is a particular case of (P0) with $f(\mathbf{c}) = \|B\mathbf{c} - \mathbf{s}\|^2$ and $C = \{(-\tilde{\mathbf{Y}}(x), 4\pi)\colon x \in \mathbb{S}^2\}$. Now, since half-spaces are convex, the intersection of such half-spaces forms a convex set that we refer to as $\mathcal{C} = \{\mathbf{c} : \mathbf{n}^\mathsf{T}\mathbf{c} \leq b \ \forall (\mathbf{n}, b) \in C\}$. Thus, (P0) is a convex optimization problem with a unique global optimum $\mathbf{c}^*$. Moreover, the boundary of $\mathcal{C}$ is defined by a subset of all the constraints, as stated next.

Lemma 1. *For each point* $\mathbf{c} \in \partial\mathcal{C}$ *in the boundary of* $\mathcal{C}$*, there is a constraint* $(\mathbf{n}, b) \in C$ *passing through it, i.e.,* $\mathbf{n}^\mathsf{T}\mathbf{c} = b$*. Moreover, the boundary is differentiable at* $\mathbf{c}$ *if and only if there is exactly one constraint* $(\mathbf{n}, b) \in C$ *passing through it.*

Thanks to Lemma 1, the optimization problem (P0) can be illustrated as in Figure 1. Here, we assume that the solution $\mathbf{c}_0$ to the unconstrained problem (LS) does not belong to $\mathcal{C}$, otherwise, $p(x; \mathbf{c}_0)$ would be already nonnegative. Now, since $\mathbf{c}_0 \notin \mathcal{C}$, the global minimizer of (P0), $\mathbf{c}^*$, must lie on the boundary of $\mathcal{C}$.
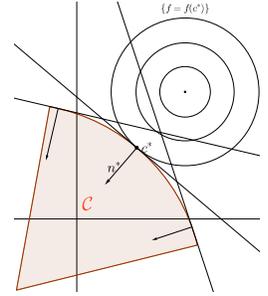
Fig. 1: Feasible set $\mathcal{C}$, level sets of f and normal $\mathbf{n}^*$ at $\mathbf{c}^*$.

3.2. Optimal Constraint Selection (OCS) Method

Assume that $\partial\mathcal{C}$ is differentiable at $\mathbf{c}^*$ so that there is only one active constraint $(\mathbf{n}^*, b^*) \in C$ passing through $\mathbf{c}^*$. Then, we can minimize f subject to this constraint alone (see Theorem 2). Moreover, by looking at Figure 1, we can see that the optimal constraint $(\mathbf{n}^*, b^*)$ is the one that maximizes the value of f among those that do not contain $\mathbf{c}_0$ (see Theorem 3).

Theorem 2. *If there is only one active constraint* $(\mathbf{n}^*, b^*)$ *passing through* $\mathbf{c}^*$*, then problem* (P0) *is equivalent to*

$$\min_{\mathbf{c}} f(\mathbf{c}) \quad \text{s.t.} \quad \mathbf{n}^{*\mathsf{T}}\mathbf{c} \leq b^*. \qquad (\mathrm{P}1)$$

Proof. Since $\mathbf{c}^*$ lies in the feasible set of (P1), it suffices to show that $f(\mathbf{c}^*) < f(\mathbf{c})$ for each $\mathbf{c} \neq \mathbf{c}^*$ such that $\mathbf{n}^{*\mathsf{T}}\mathbf{c} \leq b^*$. Assume for the sake of contradiction that there exists $\tilde{\mathbf{c}}$ such that $\mathbf{n}^{*\mathsf{T}}\tilde{\mathbf{c}} \leq b^*$ and $f(\tilde{\mathbf{c}}) \leq f(\mathbf{c}^*)$. For the case $\mathbf{n}^{*\mathsf{T}}\tilde{\mathbf{c}} < b^*$, the strong convexity of f implies that $f(\mathbf{c}^* + t(\tilde{\mathbf{c}} - \mathbf{c}^*)) < f(\mathbf{c}^*)$ for $0 < t \leq 1$. But by Lemma 1, ∂C is differentiable at $\mathbf{c}^*$, so $\mathbf{c}^* + t(\tilde{\mathbf{c}} - \mathbf{c}^*)$ must be in $\mathcal{C}$ for small t because $\mathbf{n}^{*\mathsf{T}}(\tilde{\mathbf{c}} - \mathbf{c}^*) < 0$. This contradicts the optimality of $\mathbf{c}^*$. The case $\mathbf{n}^{*\mathsf{T}}\tilde{\mathbf{c}} = b^*$ can be treated as above if we replace $\tilde{\mathbf{c}}$ by $\tilde{\mathbf{c}} - \epsilon \mathbf{n}^*$, where ϵ is chosen such that $f(\tilde{\mathbf{c}} - \epsilon \mathbf{n}^*) < f(\mathbf{c}^*)$. $\square$

Theorem 3. *If there is only one active constraint $(\mathbf{n}^*, b^*)$ passing through $\mathbf{c}^*$, this constraint can be found by solving*

$$\max_{(\mathbf{n},b) \in C : \mathbf{n}^{\mathsf{T}}\mathbf{c}_0 > b} \quad \min_{\mathbf{c}:\mathbf{n}^{\mathsf{T}}\mathbf{c}=b} f(\mathbf{c}). \tag{P2}$$

Proof. From Theorem 2 we have $\min_{\mathbf{n}^{*\mathsf{T}}\mathbf{c}=b^*} f(\mathbf{c}) = f(\mathbf{c}^*)$. Hence, we need to show that $\min_{\mathbf{n}^{\mathsf{T}}\mathbf{c}=b} f(\mathbf{c}) < f(\mathbf{c}^*)$ for all $(\mathbf{n}, b) \in C$ such that $(\mathbf{n}, b) \neq (\mathbf{n}^*, b^*)$ and $\mathbf{n}^{\mathsf{T}}\mathbf{c}_0 > b$. For, let $\tilde{\mathbf{c}} = \mathbf{c}^* + \frac{b - \mathbf{n}^{\mathsf{T}}\mathbf{c}^*}{\mathbf{n}^{\mathsf{T}}\mathbf{c}_0 - \mathbf{n}^{\mathsf{T}}\mathbf{c}^*}(\mathbf{c}_0 - \mathbf{c}^*)$ and note that $\mathbf{n}^{\mathsf{T}}\tilde{\mathbf{c}} = b$. It follows from the strong convexity of f that

$$\min_{\mathbf{c}:\mathbf{n}^{\mathsf{T}}\mathbf{c}=b} f(\mathbf{c}) \leq f(\tilde{\mathbf{c}}) < f(\mathbf{c}^*), \tag{4}$$

because $\tilde{\mathbf{c}}$ is on the line segment connecting $\mathbf{c}^*$ and $\mathbf{c}_0$. $\square$

The importance of these results is that they allow us to solve (P_∞) by solving a much simpler optimization problem over the sphere. To see this, notice that the minimization of f subject to a single constraint can be done in closed form as

$$\min_{\mathbf{c}:\mathbf{n}^{\mathsf{T}}\mathbf{c}=b} f(\mathbf{c}) = f(\mathbf{c}_0) + \frac{(b - \mathbf{n}^{\mathsf{T}}\mathbf{c}_0)^2}{\mathbf{n}^{\mathsf{T}}(\mathbf{B}^{\mathsf{T}}\mathbf{B})^{-1}\mathbf{n}}. \tag{5}$$

Recall now that $b = 4\pi$ and $\mathbf{n} = -\tilde{\mathbf{Y}}(x) \triangleq [\tilde{Y}_j(x)]_{j=1}^R$ with $\tilde{Y}_j(x) \triangleq \lambda_j \ell_j Y_j(x)$ as per (2). Therefore, the optimization problem in (P2) reduces to

$$\max_{x \in \mathbb{S}^2} \frac{(4\pi + \tilde{\mathbf{Y}}(x)^{\mathsf{T}}\mathbf{c}_0)^2}{\tilde{\mathbf{Y}}(x)^{\mathsf{T}}(\mathbf{B}^{\mathsf{T}}\mathbf{B})^{-1}\tilde{\mathbf{Y}}(x)} \quad \text{s.t.} \quad 4\pi + \tilde{\mathbf{Y}}(x)^{\mathsf{T}}\mathbf{c}_0 < 0, \tag{6}$$

which is in turn equivalent to the unconstrained problem

$$\min_{x \in \mathbb{S}^2} \frac{4\pi + \tilde{\mathbf{Y}}(x)^{\mathsf{T}}\mathbf{c}_0}{\sqrt{\tilde{\mathbf{Y}}(x)^{\mathsf{T}}(\mathbf{B}^{\mathsf{T}}\mathbf{B})^{-1}\tilde{\mathbf{Y}}(x)}}. \tag{7}$$

We can solve this problem by minimizing over a fixed grid of 10,000 spherical coordinates (θ, ϕ), where $x = x(\theta, \phi)$.

3.3. Iterative Constraint Selection (ICS) Method

It follows from Theorem 2 that, under certain conditions, the ODF associated with the solution to (P1) is guaranteed to be nonnegative. While we cannot verify the conditions of Theorem 2 *a priori*, it follows from Proposition 4 that if the ODF associated with the solution to (P1) is nonnegative, then it must be optimal. In other words, we can verify whether (P1) gives a solution to (P0) *a posteriori*.

Proposition 4. (P2) *attains a finite maximum even if ∂C is not differentiable at $\mathbf{c}^*$. Let $(\bar{\mathbf{n}}, \bar{b})$ be one of the maximizing constraints. Then, the minimizer $\bar{\mathbf{c}}$ of* (P1) *with constraint $(\bar{\mathbf{n}}, \bar{b})$ is the minimizer of* (P0) *if and only if it is in $\mathcal{C}$.*

When the condition of Proposition 4 fails, i.e., if there exists $(\mathbf{n}, b) \in C$ such that $\mathbf{n}^{\mathsf{T}}\bar{\mathbf{c}} > b$, we can find a new solution for $\mathbf{c}$ subject to $\mathbf{n}^{\mathsf{T}}\mathbf{c} \leq b$. We can then repeat this process of adding one constraint at a time and minimizing f subject to these constraints until the solution for $\mathbf{c}$ belongs to $\mathcal{C}$. This generates a sequence of solutions $\mathbf{c}$ as described in Alg. 1. Theorem 5 shows that this iterative procedure converges to $\mathbf{c}^*$.

Algorithm 1 (Iterative Constraint Selection)

1: Initialization: Choose a (finite) set of constraints $I \subset C$.
2: Set $\mathbf{c} \leftarrow \arg\min f(\mathbf{c})$ s.t. $\mathbf{n}^{\mathsf{T}}\mathbf{c} \leq b \quad \forall (\mathbf{n}, b) \in I$.
3: **if** $\max_{(\mathbf{n},b) \in C} \mathbf{n}^{\mathsf{T}}\mathbf{c} - b > 0$, **then**
4:    $I \leftarrow I \cup \{\arg\max_{(\mathbf{n},b) \in C} \mathbf{n}^{\mathsf{T}}\mathbf{c} - b\}$,
5:    go to step 2;
6: **else**
7:    $\mathbf{c}^* \leftarrow \mathbf{c}$
8: **end if**

Theorem 5. *Algorithm 1 converges to $\mathbf{c}^*$.*

Proof. Let $\mathbf{c}_i$ and I_i denote, respectively, the values of $\mathbf{c}$ and I at the ith iteration of Algorithm 1. Let $\mathcal{C}_i = \{\mathbf{c} \in \mathbb{R}^R | \mathbf{n}^{\mathsf{T}}\mathbf{c} \leq b \, \forall (\mathbf{n}, b) \in I_i\}$. To show that $\mathbf{c}_i$ converges to $\mathbf{c}^*$ it suffices to show that every subsequence $\mathbf{c}_{i_k}$ has a subsubsequence that converges to $\mathbf{c}^*$. By definition, we have $\mathcal{C}_i \supset \mathcal{C}$ for every $i \in \mathbb{N}$. Thus, $f(\mathbf{c}_i) \leq f(\mathbf{c}^*)$ and since $\{\mathbf{c} | f(\mathbf{c}) \leq f(\mathbf{c}^*)\}$ is bounded, it follows that there is a converging subsubsequence $\mathbf{c}_{i_{k_l}} \to \bar{\mathbf{c}}$. Without loss of generality we assume that $\mathbf{c}_i \to \bar{\mathbf{c}}$. First, we show that $\bar{\mathbf{c}} \in \mathcal{C}$. Assuming that $\bar{\mathbf{c}} \notin \mathcal{C}$, we can find $(\bar{\mathbf{n}}, \bar{b}) \in C$ such that $\epsilon \triangleq \bar{\mathbf{n}}^{\mathsf{T}}\bar{\mathbf{c}} - \bar{b} > 0$. Choose $K \triangleq \max_{(\mathbf{n},b) \in C} \|\mathbf{n}\|$ and N such that $\|\mathbf{c}_i - \bar{\mathbf{c}}\| < \frac{\epsilon}{4K}$ for all $i > N$. With $(\mathbf{n}, b) = I_{N+1} \backslash I_N$, it follows that

$$\mathbf{n}^{\mathsf{T}}\mathbf{c}_{N+1} - b = \epsilon - (\bar{\mathbf{n}}^{\mathsf{T}}\bar{\mathbf{c}} - \bar{b}) + \mathbf{n}^{\mathsf{T}}\mathbf{c}_N - b + \mathbf{n}^{\mathsf{T}}(\mathbf{c}_{N+1} - \mathbf{c}_N)$$

$$\geq \epsilon - (\bar{\mathbf{n}}^{\mathsf{T}}\bar{\mathbf{c}} - \bar{b}) + \bar{\mathbf{n}}^{\mathsf{T}}\mathbf{c}_N - \bar{b} + \mathbf{n}^{\mathsf{T}}(\mathbf{c}_{N+1} - \mathbf{c}_N)$$

$$= \epsilon + \bar{\mathbf{n}}^{\mathsf{T}}(\mathbf{c}_N - \bar{\mathbf{c}}) + \mathbf{n}^{\mathsf{T}}(\mathbf{c}_{N+1} - \mathbf{c}_N)$$

$$\geq \epsilon - K\|\mathbf{c}_N - \bar{\mathbf{c}}\| - K\|\mathbf{c}_{N+1} - \mathbf{c}_N\|$$

$$> \epsilon - K(\frac{\epsilon}{4K} + 2\frac{\epsilon}{4K}) = \frac{\epsilon}{4},$$

which contradicts $\mathbf{c}_{N+1} \in C_{N+1}$. Now, $\bar{\mathbf{c}} = \mathbf{c}^*$ follows from $f(\bar{\mathbf{c}}) = \lim_{i \to \infty} f(\mathbf{c}_i) \leq f(\mathbf{c}^*)$. $\square$

While Alg. 1 is guaranteed to converge, the number of iterations could be infinite. In this case, the number of constraints could be infinite as well. In our implementation, we enforce nonnegativity up to an arbitrary threshold $\tau < 0$ in step 3 of Alg. 1, i.e., we check if $\mathbf{n}^{\mathsf{T}}\mathbf{c} - b > \tau$. This guarantees convergence in a finite number of iterations (typically 3), but does not guarantee nonnegativity everywhere. We find this to give a good compromise between accuracy and speed.
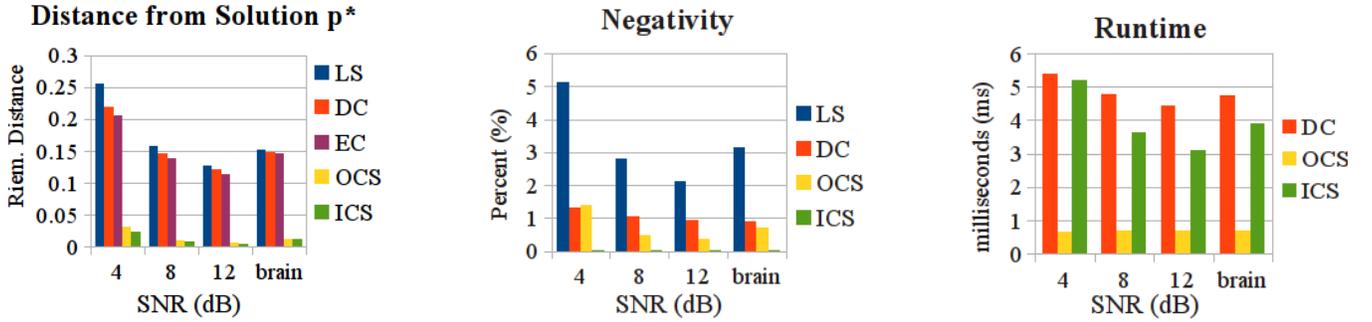
Fig. 2: Performance of the Least-Squares (LS), Discretely Constrained (DC) and Eigenvalue Constraint (EC) methods compared to that of our Optimal Constraint Selection (OCS), and Iterative Constraint Selection (ICS) methods on synthetic and brain data.

4. EXPERIMENTS

We compare our methods OCS and ICS with LS, DC, and [9], which we call Eigenvalue Constraint (EC), on a synthetic field of 375 1-, 2-, and 3-fiber ODFs distorted with SNR of 4, 8, and 12 dB as well as a real HARDI brain data set. We evaluate the overall performance of each method using three metrics: 1) the average Riemannian distance to the optimal solution p^* of (P_∞), i.e., the optimal nonnegative ODF computed by the DC method with 1 million constraints obtained from 1 million discrete grid points on the sphere, 2) the average percentage of negative values remaining in the estimated ODF, evaluated over 1 million discrete grid points, and 3) the average runtime per ODF. In the left plot of Figure 2 we see that our methods produce ODFs that are very close to p^*, while the state-of-the-art EC method does no better than DC and LS. From the center plot we see that OCS exhibits some negative values whenever the conditions of Theorem 2 are violated. Here ICS also shows a slight percentage of negative values due to the use of a threshold $\tau = -5 \times 10^{-4}$ in step 3 of Alg. 1. The percentage of negativity of EC was always zero, but EC had a runtime on the order of 10s per ODF. In sharp contrast, OCS is the fastest method, with a runtime of 0.5ms per ODF. This is expected because OCS enforces only one constraint. On the other hand, the runtime of ICS is comparable to that of DC. Overall, ICS outperforms the state-of-the-art EC by producing ODFs that are closer to the optimal nonnegative ODF, reducing the amount of negativity, and improving runtime.

5. CONCLUSION

We proposed two methods for estimating a nonnegative ODF from HARDI data. The first (OCS) solves a quadratic problem subject to one constraint and is guaranteed to produce a nonnegative ODF under some conditions. The second (ICS) iteratively solves a quadratic problem subject to multiple linear constraints and is guaranteed to converge to the correct solution, which would have required infinitely many constraints. Our experiments showed that our methods produce more accurate solutions than prior work at a reduced runtime.

6. REFERENCES

[1] D.S. Tuch, T.G. Reese, M.R. Wiegell, N. Makris, J.W. Belliveau, and V.J. Wedeen, "High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity," *Magnetic Resonance in Medicine*, vol. 48, no. 4, pp. 577–582, 2002.

[2] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche, "Regularized, fast and robust analytical Q-ball imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 3, pp. 497–510, 2007.

[3] B. Jian and B. Vemuri, "A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI," *IEEE Transactions on Medical Imaging*, vol. 26, no. 11, pp. 1464–1471, 2007.

[4] I. Aganj, C. Lenglet, G. Sapiro, E. Yacoub, K. Ugurbil, and N. Harel, "Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle," *Magnetic Resonance in Medicine*, vol. 64, no. 2, pp. 554–566, 2010.

[5] A. Goh, C. Lenglet, P.M. Thompson, and R. Vidal, "Estimating orientation distribution functions with probability density constraints and spatial regularity," in *Medical Image Computing and Computer Assisted Intervention*, 2009, pp. 877–885.

[6] J-D. Tournier, F. Calamante, and A. Connelly, "Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution," *NeuroImage*, vol. 35, no. 4, pp. 1459–1472, 2007.

[7] J. Cheng, T. Jiang, and R. Deriche, "Nonnegative definite EAP and ODF estimation via a unified multi-shell HARDI reconstruction," in *Medical Image Computing and Computer-Assisted Intervention*, 2012, vol. 7511, pp. 313–321.

[8] L. Qi, G. Yu, and Y. Xu, "Nonnegative diffusion orientation distribution function," *Journal of Mathematical Imaging and Vision*, vol. 45, no. 2, pp. 103–113, 2013.

[9] E. Schwab, B. Afsari, and R. Vidal, "Estimation of non-negative ODFs using eigenvalue distribution of spherical functions," in *Medical Image Computing and Computer Assisted Intervention*, 2012, vol. 7511, pp. 322–330.