

Iterative Subspace Screening for Rapid Sparse Estimation of Brain Tissue Microstructural Properties

Pew-Thian Yap², Yong Zhang¹, and Dinggang Shen²

¹ Department of Radiology and Biomedical Research Imaging Center,
The University of North Carolina at Chapel Hill, USA

² Department of Psychiatry and Behavioral Sciences,
Stanford University, USA
ptyap@med.unc.edu

Abstract. Diffusion magnetic resonance imaging (DMRI) is a powerful imaging modality due to its unique ability to extract microstructural information by utilizing restricted diffusion to probe compartments that are much smaller than the voxel size. Quite commonly, a mixture of models is fitted to the data to infer microstructural properties based on the estimated parameters. The fitting process is often non-linear and computationally very intensive. Recent work by Dadducci et al. has shown that speed improvement of several orders of magnitude can be achieved by linearizing and recasting the fitting problem as a linear system, involving the estimation of the volume fractions associated with a set of diffusion basis functions that span the signal space. However, to ensure coverage of the signal space, sufficiently dense sampling of the parameter space is needed. This can be problematic because the number of basis functions increases exponentially with the number of parameters, causing computational intractability. We propose in this paper a method called iterative subspace screening (ISS) for tackling this ultrahigh dimensional problem. ISS requires only solving the problem in a medium-size subspace with a dimension that is much smaller than the original space spanned by all diffusion basis functions but is larger than the expected cardinality of the support of the solution. The solution obtained for this subspace is used to screen the basis functions to identify a new subspace that is pertinent to the target problem. These steps are performed iteratively to seek both the solution subspace and the solution itself. We apply ISS to the estimation of the fiber orientation distribution function (ODF) and demonstrate that it improves estimation robustness and accuracy.

1 Introduction

Microstructural tissue properties can be inferred with the help of diffusion MRI (DMRI) thanks to its sensitivity to the restricted motion of water molecules owing to barriers such as cellular membranes. Microstructural information is typically obtained from diffusion parameters estimated via fitting to the acquired data some biophysical models. Fitting models such as the tensor model [1] is relatively simple and straightforward. But fitting models that are more sophisticated, such as the multi-compartmental models used in AxCaliber [2] and NODDI [3], is much more involved with significantly greater

computational load. Computational complexity is further increased when certain structure is imposed on the solution, such as sparsity [4,5].

In a recent work called AMICO [6], the authors show that it is possible to speed up AxCaliber and NODDI estimation several orders of magnitude by re-formulating the fitting problem as a linear system that can be efficiently solved using very fast algorithms. However, this work is limited to a sparse sampling of the parameter space (30 basis functions for AxCaliber and 145 basis functions for NODDI) and is also limited to voxels containing only a single principal diffusion direction. When extending AMICO for the more realistic case of multiple white matter (WM) directions per voxel [7], the number of the basis functions would need to be significantly increased, causing the optimization problem to be very high dimensional and computationally very expensive. A similar situation occurs in the estimation of the fiber orientation distribution function (ODF), as in [5], when one needs to increase the number of angular directions of the basis functions for improving accuracy.

In this paper, we propose a method called iterative subspace screening (ISS) to solve this kind of ultrahigh dimension problem more efficiently. ISS requires only solving the target problem in a medium-size subspace with a dimension that is much smaller than the original space spanned by all the basis functions, but is larger than the expected cardinality of the support of the solution. ISS is a subspace pursuit algorithm [8] involving the following iterative steps:

1. **Subspace Selection:** Select a potential solution subspace by screening out irrelevant basis functions. This is done by element-wise regression of the fitting residual similar to iterative sure independence screening (ISIS) [9] and subspace pursuit (SP) [8].
2. **Subspace Solution:** Solve for the solution, assumed sparse, in the selected subspace and compute the fitting residual. The fitting residual is then used in the subspace selection step above to refine the subspace so that irrelevant basis functions can be discarded and relevant ones can be included.

Such subspace selection approach allows us to remove from consideration many basis functions that might never be active and contribute to the solution. This is motivated by the fact that, in our case, the solution is sparse and hence resides in a small subspace. Since the dimension of the ISS subspace is typically much smaller, the solution can be obtained much faster with potentially lesser local minima that may result from the high dimensionality. We demonstrate the effectiveness of ISS on the problem of fiber ODF estimation and show that significant speed up can be achieved. Extension of ISS for AxCaliber and NODDI, similar to AMICO, should be straightforward.

2 Approach

2.1 Problem Description

The problem we are interested in solving has the following form:

$$\min_{\mathbf{f} \geq 0} \{\phi(\mathbf{f}) = l(\mathbf{f}) + r(\mathbf{f})\}, \quad (1)$$

Algorithm 1. Iterative Subspace Screening

Input : $\mathbf{A}, \mathbf{s}, D$
Initialization:

1. $\mathcal{J}^{(0)} = \{D \text{ indices of the entries with the largest magnitudes in vector } \mathbf{A}^T \mathbf{s}\}$
2. $\mathbf{r}^{(0)} = \text{residual}(\mathbf{s}, \mathbf{A}_{\mathcal{J}^{(0)}})$
3. $\mathbf{f}^{(0)} = \mathbf{0}$

Iteration : At iterations $k = 1, 2, \dots$, go through the following steps:

1. $\mathcal{J}^{(k)} = \{D \text{ indices of the non-zeros entries of } \mathbf{f}^{(k-1)} \text{ and the entries with the largest magnitudes in vector } \mathbf{A}^T \mathbf{r}^{(k-1)}\}$
2. Set $\mathbf{f}^{(k)} = \text{solve}(\mathbf{s}, \mathbf{A}_{\mathcal{J}^{(k)}})$
3. $\mathbf{r}^{(k)} = \text{residual}(\mathbf{s}, \mathbf{A}_{\mathcal{J}^{(k)}})$
4. If $\|\mathbf{r}^{(k)}\| > \|\mathbf{r}^{(k-1)}\|$, quit the iteration.

Output : The solution $\mathbf{f} = \mathbf{f}^{(k-1)}$.

where in our case $l(\mathbf{f}) = \|\mathbf{A}\mathbf{f} - \mathbf{s}\|_2^2$ is a smooth and convex data fidelity term with respect to observation vector $\mathbf{s} \in \mathbb{R}^N$ and $r(\mathbf{f})$ is a sparsity-inducing regularization term that is not necessarily smooth or convex. In diffusion MRI, one would typically fill the columns of basis matrix $\mathbf{A} \in \mathbb{R}^{N \times P}$ with the set of basis functions derived from diffusion models with varying parameters, such as orientation and diffusivity. The vector $\mathbf{f} \in \mathbb{R}^P$ consists of the corresponding volume fractions associated with the basis functions. For example, for estimating the fiber ODF as in [5], one would fill the columns of \mathbf{A} with the signal vectors sampled from tensor models oriented uniformly in all directions. For fast AxCaliber [2] and NODDI [3] estimation using a linear system as described in AMICO [6], one would fill the columns of \mathbf{A} with signal vectors from compartment models orientated in a direction estimated via fitting a diffusion tensor. If the parameter space is sampled densely, the above estimation would be computationally very expensive and can be susceptible to local minima due to the high dimensionality and high correlation between the basis functions.

2.2 Iterative Subspace Screening (ISS)

We propose a method called iterative subspace screening (ISS) to tackle this kind of ultrahigh dimension problem more efficiently. ISS requires only solving the target problem (1) in a medium-size subspace with a dimension that is much smaller than the original space spanned by all the basis functions, but is larger than the expected cardinality of the support of the solution. This is based on the observation that the problem (1) can be rewritten as

$$\min_{\mathbf{f}_{\mathcal{J}} \geq \mathbf{0}, \mathbf{f}_{\mathcal{J}^c} = \mathbf{0}, |\mathcal{J}| = D} \{\phi(\mathbf{f}) = l_{\mathcal{J}}(\mathbf{f}) + r_{\mathcal{J}}(\mathbf{f})\}, \quad (2)$$

subject to $D \ll P$ is larger than the expected cardinality of the support of the solution. Here, $l_{\mathcal{J}}(\mathbf{f}) = \|\mathbf{A}_{\mathcal{J}} \mathbf{f}_{\mathcal{J}} - \mathbf{s}\|_2^2$ with $\mathbf{f}_{\mathcal{J}}$ being the subvector formed by the elements of

\mathbf{f} indexed by set $\mathcal{J} \subseteq \{1, 2, \dots, P\}$, $\bar{\mathcal{J}} = \{1, 2, \dots, P\} \setminus \mathcal{J}$, and $\mathbf{A}_{\mathcal{J}}$ being the submatrix formed by the columns of \mathbf{A} indexed by \mathcal{J} . The regularization term $r_{\mathcal{J}}(\mathbf{f})$ now penalizes only the elements in \mathbf{f} that are indexed by $\bar{\mathcal{J}}$. We can think of problem (2) as concurrently solving for the solution subspace and the solution itself. This problem is clearly non-convex, non-smooth, and discontinuous; but decoupling subspace identification from the problem allows us to devise an algorithm that focuses on determining the solution in a subspace that actually contains the solution and not in the original subspace, especially when the original problem (1) is very high dimensional (i.e., P is large). As we shall see later, this subspace can be progressively refined by including basis functions that will contribute to the solution and by discarding those that will not. Since the solution is sparse, most basis functions are irrelevant to the solution and can be removed from consideration, reducing significantly the computation cost. Note that if the solution to (1) resides in the identified subspace, solving for (2) will give the same solution as (1).

Proposed Solution. We propose to solve (2) by alternating between solving for \mathcal{J} and solving for \mathbf{f} . At each iteration, one only needs to solve for \mathbf{f} in a medium-size subspace with dimension $D \ll P$. To select the subspace, we perform an element-wise regression of the fitting residual with the basis functions. That is, at the k -th iteration, we first compute the fitting residual $\mathbf{r}^{(k-1)} = \text{residual}(\mathbf{s}, \mathbf{A}_{\mathcal{J}^{(k-1)}}) = \mathbf{s} - \mathbf{A}_{\mathcal{J}^{(k-1)}} \mathbf{f}_{\mathcal{J}^{(k-1)}}$ and then use it to determine a subspace of dimension D . This subspace is spanned by the basis functions corresponding to the non-zero elements of $\mathbf{r}^{(k-1)}$ and, in addition, the basis functions corresponding to the largest entries of $\mathbf{A}^T \mathbf{r}^{(k-1)}$. Based on this new subspace, with the constituent basis functions indexed by $\mathcal{J}^{(k)}$, we solve for $\mathbf{f}^{(k)}$, i.e., $\mathbf{f}^{(k)} = \text{solve}(\mathbf{s}, \mathbf{A}_{\mathcal{J}^{(k)}})$, via (2) with $\mathcal{J}^{(k)}$ fixed. This solution can then be used to recompute the residual and refine the subspace. Since at each iteration, the residual vector will be used to screen out a significant number of basis functions, thus removing a big portion of the original subspace, we call our method iterative subspace screening (ISS). See Algorithm 1 for a step-by-step summary of ISS. In addition to reducing significantly the computation time by not having to solve the problem in the original high-dimensional space, ISS can also deal with collinearity between basis functions. By performing subspace screening using the residual vector, some unimportant basis functions that are highly correlated with the important basis functions can be discarded, as observed in [9].

Grouped Variant. We can take advantage of the relationships between basis functions by grouping them [9]. We can divide the pool of P basis functions into disjoint groups, each with a number of basis functions. ISS can then be applied to the selection of basis function groups instead of the individual basis functions. This will reduce the chance of missing important basis functions by taking advantage of the joint information among them, making the estimation more reliable. This can be achieved by modifying Step 1 in the k -loop of Algorithm 1 to select the top D groups of basis functions with the greatest ℓ_2 -norm of inner products with the residual vector. We shall show how grouping can be used in ISS to help solve sparse-group approximation problem similar to the one described in [10].




Fig. 1. Computation times.

A Specific Case. To demonstrate the utility of ISS, we select to use the general sparse-group regularization $r(\mathbf{f}) = \gamma \left[\alpha \|\mathbf{f}\|_0 + (1 - \alpha) \sum_{g \in G} \mathcal{I}(\|\mathbf{f}_g\|_2) \right]$, where $\mathcal{I}(z)$ is an indicator function returning 1 if $z \neq 0$ or 0 if otherwise. The ℓ_0 -“norm” gives the cardinality of the support, i.e., $\|\mathbf{f}\|_0 = |\text{supp}(\mathbf{f})| = |\{k : f_k \neq 0\}|$. Parameters $\alpha \in [0, 1]$ and $\gamma > 0$ are for penalty tuning, analogous to those used in the sparse-group LASSO [10]. Note that $\alpha = 1$ gives the ℓ_0 fit, whereas $\alpha = 0$ gives the group ℓ_0 fit. Note that in contrast to the more commonly used ℓ_1 -norm penalization, we have chosen here to use ℓ_0 -“norm” cardinality-based penalization. As reported in [11], ℓ_1 -norm penalization [5] conflicts with the unit sum requirement of the volume fractions and hence results in suboptimal solutions.

3 Experiments

3.1 Data

Synthetic Data: For quantitative evaluation, we generated a synthetic dataset for evaluation of ISS. The dataset was generated using a mixture of four tensor models. Two of which are anisotropic and represent two white matter (WM) compartments that are at an angle of 60° with each other. The other two are isotropic and represent the gray matter (GM) and cerebrospinal fluid (CSF) compartments. The generated diffusion-attenuated signals therefore simulate the partial volume effects resulting from these compartments. The volume fractions and the diffusivities of the compartments were allowed to vary in ranges that mimic closely the real data discussed in the next section. Various levels of noise (SNR=10, 20, 30, with respect to the signal value at $b = 0 \text{ s/mm}^2$) was added.

Real Data: Diffusion weighted (DW) data from the Human Connectome Project (HCP) [12] were used. The $1.25 \times 1.25 \times 1.25 \text{ mm}^3$ data were acquired with diffusion weightings $b = 1000, 2000, \text{ and } 3000 \text{ s/mm}^2$ each applied in 90 directions. 18 baseline images with low diffusion weighting $b = 5 \text{ s/mm}^2$ were also acquired. All images were acquired with reversed phase encoding for correction of EPI distortion.

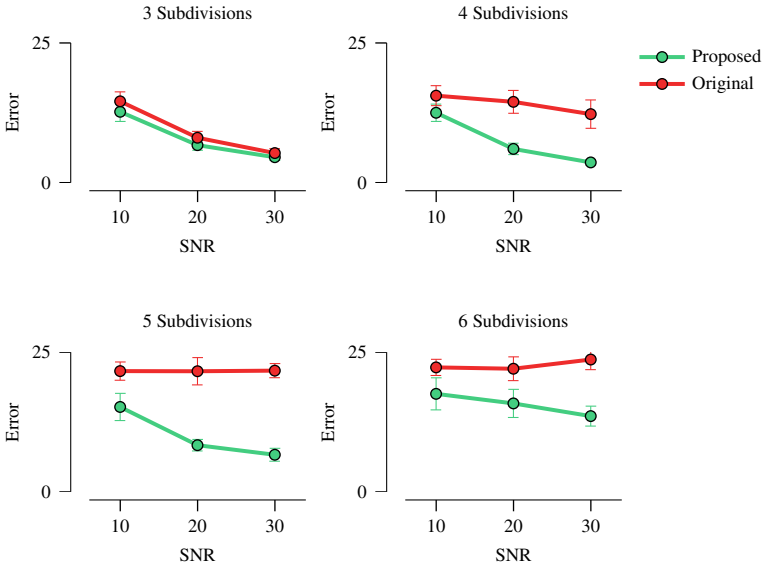


Fig. 2. ODF estimation error measured via orientational discrepancy (OD). The error bars indicate the standard deviations.

3.2 Methods of Evaluation

For the synthetic data, our aim is to estimate the fiber ODFs [13] and evaluate their accuracy by comparing their peaks (local maxima) with respect to the ground truth. The orientational discrepancy (OD) measure defined in [14] was used as a metric for evaluating the accuracy of peak estimation. For the real data, we want to evaluate whether consistent results are given by ISS compared with the original problem (1).

The fiber ODF was estimated by fitting to the data the compartment models for WM, GM, and CSF. Similar to [15], the WM is represented by a large number of anisotropic single-tensor models with orientations distributed evenly on a unit sphere. Dissimilar to [15], however, we allow the axial and radial diffusivities to vary. The GM and CSF are represented using isotropic tensor models with GM having a lower range of diffusivities compared with CSF. Mathematically, this is realized by filling the columns of \mathbf{A} with the above models with different parameters, i.e., orientations and diffusivities, and then solving problem (1) (or (2)) for the corresponding volume fractions \mathbf{f} . Structure is imposed on the problem by grouping the WM models for each direction as well as the GM models and the CSF models. The tuning parameters were set as follows: $\gamma = 1 \times 10^{-4}$ and $\alpha = 0.05$. D was set to about 15% of P .

3.3 Results

Synthetic Data: To show that ISS improves the speed of estimation, we evaluated the computation time of ISS in comparison with solving the original problem (1). We perform this for different numbers of subdivisions of the icosahedron, giving different




Fig. 3. WM ODF glyphs at the centrum semiovale. White boxes mark examples of improvements given by ISS.

numbers of directions for the WM models and hence different numbers of columns P for A. 3, 4, 5, and 6 subdivisions of the icosahedron give respectively 321, 1281, 5121, and 20481 directions on a hemisphere. Figure 1 shows that ISS improves the speed of convergence remarkably.

Next, we proceeded to evaluate whether such increase in speed implies a decrease in ODF estimation accuracy. Figure 2, perhaps surprisingly, indicate that ISS actually improves ODF estimation accuracy. This can be explained from the fact that solving the problem in a lower-dimensional subspace can help alleviate the problem of local minima, especially when the problem is non-convex as in our case.

Real Data: Figure 3 shows the WM ODF glyphs at a portion of the centrum semiovale. The images show that the centrum semiovale contains elements of the corticospinal tract (blue), corpus callosum (red), and superior longitudinal fasciculus (green), including voxels with two- and three-way intersections of these elements. The images indicate that ISS gives consistent results comparable to the original problem. However, we could observe that ISS in fact improves ODF estimation in some voxels (marked by white boxes) and gives ODF estimates that are more coherent and match better with the underlying anatomy.

4 Conclusion

We have proposed a method for improving the speed and robustness of solving least-squares problems that are regularized by sparse inducing norms. We applied our method called iterative subspace screening (ISS) to ODF estimation and showed that better and faster estimates of the fiber orientations can be obtained. In the future, ISS will be

applied to improve the speed and accuracy of estimation techniques for microstructure, such as AxCaliber and NODDI.

Acknowledgment. This work was supported in part by a UNC BRIC-Radiology start-up fund and NIH grants (EB006733, EB009634, AG041721, MH100217, and 1UL1TR001111).

References

1. Basser, P.J., Pierpaoli, C.: Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance Series B* 111(3), 209–219 (1996)
2. Alexander, D.C., Hubbard, P.L., Hall, M.G., Moore, E.A., Ptito, M., Parker, G.J., Dyrby, T.B.: Orientationally invariant indices of axon diameter and density from diffusion MRI. *NeuroImage* 52, 1374–1389 (2010)
3. Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C.: NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage* 61, 1000–1016 (2012)
4. Landman, B.A., Bogovic, J.A., Wan, H., ElShahaby, F.E.Z., Bazin, P.L., Prince, J.L.: Resolution of crossing fibers with constrained compressed sensing using diffusion tensor MRI. *NeuroImage* 59, 2175–2186 (2012)
5. Ramirez-Manzanares, A., Rivera, M., Vemuri, B.C., Carney, P., Mareci, T.: Diffusion basis functions decomposition for estimating white matter intra-voxel fiber geometry. *IEEE Transactions on Medical Imaging* 26(8), 1091–1102 (2007)
6. Daducci, A., Canales-Rodríguez, E.J., Zhang, H., Dyrby, T.B., Alexander, D.C., Thiran, J.P.: Accelerated microstructure imaging via convex optimization (AMICO) from diffusion MRI data. *NeuroImage* 105, 32–44 (2015)
7. Zhang, H., Dyrby, T.B., Alexander, D.C.: Axon diameter mapping in crossing fibers with diffusion MRI. In: Fichtinger, G., Martel, A., Peters, T. (eds.) *MICCAI 2011, Part II. LNCS*, vol. 6892, pp. 82–89. Springer, Heidelberg (2011)
8. Dai, W., Milenkovic, O.: Subspace pursuit for compressive sensing signal reconstruction. *IEEE Transactions on Information Theory* 55(5), 2230–2249 (2009)
9. Fan, J., Lv, J.: Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society, Series B* 70(Part 5), 849–911 (2008)
10. Simon, N., Friedman, J., Hastie, T., Tibshirani, R.: A sparse-group lasso. *Journal of Computational and Graphical Statistics* 22(2), 231–245 (2013)
11. Daducci, A., Ville, D.V.D., Thiran, J.P., Wiaux, Y.: Sparse regularization for fiber ODF reconstruction: From the suboptimality of ℓ_2 and ℓ_1 priors to ℓ_0 . *Medical Image Analysis* 18, 820–833 (2014)
12. Essen, D.C.V., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K.: The WU-Minn human connectome project: An overview. *NeuroImage* 80, 62–79 (2013)
13. Tournier, J.D., Calamante, F., Gadian, D.G., Connelly, A.: Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage* 23(3), 1176–1185 (2004)
14. Yap, P.T., Chen, Y., An, H., Yang, Y., Gilmore, J.H., Lin, W., Shen, D.: SPHERE: SPHERical Harmonic Elastic REgistration of HARDI data. *NeuroImage* 55(2), 545–556 (2011)
15. Jian, B., Vemuri, B.C.: A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. *IEEE Transactions on Medical Imaging* 26(11), 1464–1471 (2007)