Regularization

Finding an appropriate value for hyper-parameters is always a challenge in machine learning problems. To select an appropriate regularization parameter, we want to (a) satisfy biological plausibility conditions of fascicles as well as (b) reconstruct the diffusion signal Y. We started with a larger regularization parameter and decreased it until we found the reconstruction error did not change much, to find a suitable balance. We found that picking this parameter was not too difficult, though tuning it more would likely result in better performance.

Prior Tractography Results

We do not yet compare to prior tractography algorithms. Rather, we first ask: how effective is our new formulation and algorithms, given ground truth? We use LiFE, which is a tractography evaluation method, with one tractography solution to give us ground truth. The availability of a ground truth Phi let's us focus our investigation into the soundness of the formulation and the algorithms. The computational complexity of solving the formulation is non-trivial, and understanding the soundness of the proposed algorithms is critical before deploying at a larger scale. Introducing this new formulation and developing and understanding sound approximation algorithms for this formulation is the right size for a NeurIPS paper; also demonstrating performance for full tractography solutions, across brains and against many other tractography methods, would be large journal paper let alone NeurIPS paper. The next step, though, is absolutely to do this comparison.

Problem Formulation

Reviewers 1 and 2 both asked how the dictionary D was obtained and Reviewer 1 asked about N_θ. dMRI measurements are collected with and without a diffusion sensitization magnetic gradient and along several (i.e. N_θ number of) **gradient directions** ($\theta \in R^3$). The data was collected for $N_\theta = 96$ different angles of gradient direction. Then matrix $D \in R^{N_\theta \times N_a}$ is a dictionary of diffusion predictions with the same type as diffusion signal, whose columns (i.e. N_a atoms) correspond to precomputed fascicle orientations, and intuitively shows that what diffusion signal we expect to receive from any axon in the direction of any possible fascicles orientations (a) in space by sensitizing magnetic gradient in each direction of (θ). More specifically, the entries of the dictionary were computed as follows: $D(\theta, a) = e^{-b\theta^T Q_a \theta} - \frac{1}{N_\theta} \sum_\theta e^{-b\theta^T Q_a \theta}$, in which Q_a is an approximation of diffusion tensor per fascicle-voxel. $\theta^T Q_a \theta$ gives us the diffusion at direction θ generated by fascicle f. We will include this information in section 2 of the paper for the final draft.

Reviewer 1

With the additional space in the camera-ready, we can include more background and discussion on ENCODE.

Reviewer 2

Our solution exploits the inherent sparseness in the optimization. Each voxel has a very small subset of potential fascicles and orientations. The algorithm scales well with an increase in the size of the tensor. The greedy algorithm to select orientations does have to iterate over all orientations once, for each voxel, before the optimization is run. However, this does not need to be run for each step of the optimization and is not expensive to do once upfront. The optimization then scales effectively linearly with the number of voxels considered, because a relatively small maximum number of orientations is considered for each. Additionally, the optimization is very easily parallelizable, allowing our approach to scale with the amount of available computational resources.

We did not include results across different subjects due to space constraints. We can include a few more subjects in the appendix of the final draft.

The greedy algorithm used for visualization has been carefully designed to reflect biological attributes of fascicles such as smoothness and continuity. It finds the best possible permutation of the segments for a fascicle in each voxel-based (a) on the two-by-two angular distance of each segment's orientations and (b) neighboring voxels. We believe the algorithm is reliable, both because of the principle in its design and because we extensively visualized several ground-truth and learned solutions to find any issues. You are right that we cannot guarantee that it provides a perfect visualization, and there is more to be understood about how to faithfully visualize our solutions.

We will include details about how fascicles are screened in the main text.

Reviewer 3

The main concern here is that the regularizer + OMP does poorly, suggesting the regularizer on its own is inadequate. However, this is a slight misunderstanding. OMP and GreedyDirection both first screen the set of considered orientations. This introduces an approximation, to make the optimization much faster (so that it is a feasible optimization). The best for the regularizer would be to do no screening. The ordering would be "Regularizer Alone" is better than "Regularizer + GreedyDirection" better than "Regularizer + OMP". The fact that "Regularizer + GreedyDirection" does well suggests the regularizer is actually reasonable and that we might even be able to better if we can improve the orientation screening. We only include OMP to demonstrate that we needed a better screening method than existing methods (i.e., OMP).