Bilateral Filtering of Multiple Fiber Orientations in Diffusion MRI

Ryan P. Cabeen and David H. Laidlaw

Abstract We present and evaluate a bilateral filter for smoothing diffusion MRI fiber orientations with preservation of anatomical boundaries and support for multiple fibers per voxel. Two challenges in the process are the geometric structure of fiber orientations and the combinatorial problem of matching multiple fibers across voxels. To address these issues, we define distances and local estimators of weighted collections of multi-fiber models and show that these provide a basis for an efficient bilateral filtering algorithm for orientation data. We evaluate our approach with experiments testing the effect on tractography-based reconstruction of fiber bundles and response to synthetic noise in computational phantoms and clinical human brain data. We found this to significantly reduce the effects of noise and to avoid artifacts introduced by linear filtering. This approach has potential applications to diffusion MR tractography, brain connectivity mapping, and cardiac modeling.

1 Introduction

In this paper, we present and evaluate a method for smoothing orientation image data that preserving edges and supports multiple orientations per voxel. We apply this to diffusion MR imaging, a technique for measuring patterns of water molecule diffusion with clinical applications to the in-vivo characterization of tissue. While many of properties of tissue microstructure can be measured, we consider fiber orientations, which are a feature of most diffusion models. In brain white matter imag-

Brown University, Providence, RI, USA, {cabeen,dhl}@cs.brown.edu

The final publication will be available at http://link.springer.com/book/10.1007/978-3-319-11182-7. Computational Diffusion MRI. MICCAI Workshop, Boston, USA, September 18, 2014. L. O'Donnell, G.Nedjati-Gilani, Y. Rathi, M. Reisert, and T. Schneider. (Eds.),

ing, fiber orientations provide the basis for the reconstruction of fiber bundles and mapping of brain connectivity [2], and in cardiac imaging, fiber orientations aid the understanding of myocardial structure and the electrical and mechanical function of the heart [22]. A common issue with diffusion imaging is the presence of noise, which can lead to errors in fiber orientation estimation. We examine a model-based approach to regularization that extends the bilateral filter to single and multiple orientation data.

Standard approaches for diffusion MR filtering either operate the diffusion-weighted signal measured in each voxel and gradient encoding direction [27] or incorporate robust statistics when fitting diffusion models [3]. Alternatively, model-based image processing has the potential advantage of greater efficiency and the ability to incorporate anatomical knowledge [17]. However, one challenge is that the geometric structure of diffusion models must be incorporated. There has been much success in developing such model-based frameworks with differential manifolds, such as with the tensor model [17] and orientation distribution functions [11]. For multi-compartment models, an additional combinatorial problem arises, where correspondence must be made between fibers in different voxels. Local estimators that incorporate clustering are one solution and have been applied to multi-tensors [25] and orientations [6]. Other more global approaches for orientation regularization have also been studied for single [19][28][8] and multi-fiber models [10][23].

In this paper, we derive a bilateral filter that extends previous work on local linear filters for fiber orientations [6]. Such bilateral filters are well-studied for scalar and vector images [26] and are closely related to normalized convolution [14], anisotropic diffusion, and kernel regression [15]. Model-based bilateral filters have been developed for the regularization of single diffusion tensor [12] and functional [21] MR images. Related data-adaptive filters for 2D image orientations have also been proposed for 2D smoothing [18] and hair modeling [16] applications. Our work's distinguishing features are the handling of 3D orientations, support for multiple fibers per voxel, and a computationally efficient formulation.

In the rest of the paper, we first discuss computational analysis of single and multi-fiber orientations and derive the bilateral filter. We then evaluate our approach with computational phantoms and human brain data, measuring the effect on tractography-based reconstruction of fiber bundles and the voxelwise response to synthetic noise. We show that the proposed filter improves bundle reconstructions, significantly reduces the effects of noise, and offers an improvement over linear filtering at anatomical boundaries.

2 Methods

In this section, we first describe models of fiber orientations in diffusion MRI and discuss the computation of distances and averages of single and multiple orientations. We then apply these results to derive the proposed bilateral filter.

2.1 Fiber Orientation Modeling

Many methods exist for estimating fiber orientations. In this paper, we use the multi-direction ball-and-stick compartment model to obtain both fiber orientations and their volume fractions. This is a parametric mixture model consisting of an isotropic ball compartment and multiple tensor sticks that are constrained to be completely anisotropic [4]; our focus in this paper is on fiber orientations. These orientations have no preferred direction, so they are typically considered axial or line data and can be represented by unit vectors with no associated sign. For the case of multiple fibers, we consider a model M to be a weighted combination of N fiber volume-fraction and orientation pairs $M = \{(f_i, v_i)\}_{i=1}^{N}$ that lie in a single voxel. To perform analysis of these models, we need suitable distances, weighted averages, and related efficient computational routines. In the following sections, we'll discuss such ideas for both individual fiber orientations and their weighted combinations.

2.2 Single Fiber Analysis

We measure the distance $d_f(a,b)$ between single fiber orientations a and b by the sine of their angle. Although the angle between axes may seem more natural, the sine angle distance allows for a desirable representation and offers robustness to outliers [5]. This distance can be found by considering the representation $\phi(v) = vv^T$, known as the Veronese-Whitney embedding, the dyadic product, or Knutsson mapping [13][20]. This representation induces the sine angle fiber distance d_f^2 via the Euclidean distance d_e^2 in the embedding ϕ:

$$d_f^2(a,b) = d_e^2(\phi(a),\phi(b)) = \|\phi(a) - \phi(b)\|^2 \tag{1}$$

$$= \mathrm{Tr}((\phi(a) - \phi(b))^T(\phi(a) - \phi(b)) \tag{2}$$

$$= 2(1 - (a \cdot b)^2) = 2sin^2(\theta) \tag{3}$$

Weighted averages can then be computed with respect to this distance by $\mu = \sum_i w_i \phi(v_i)$. As this is an extrinsic mean, the result may no longer lie in the embedding, so it must be projected to the nearest point $\mathrm{argmin}_v d_e^2(\phi(v), \mu)$. A closed form expression for this is given by the principal eigenvector of the matrix μ [5]. This formulation also has a statistical interpretation, as the fiber distance is equivalent to the Bregman divergence between Watson distributions [7], and the weighted average of fibers is the maximum likelihood estimate of the direction of a Watson distribution [24]. We also note that both the fiber distance and the embedding provide a computationally efficient approach for optimization with orientations.

2.3 Multi-fiber Analysis

In addition to measuring distances between fibers, we also wish to measure a distance $d_m(M,\hat{M})$ between weighted combinations of fibers:

$$d_m^2(M,\hat{M}) = \min_{\pi} \sum_j f_j d_f^2(v_j, \hat{v}_{\pi(j)}) \tag{4}$$

which is selected across all possible mappings π from left to right fibers. Intuitively, this finds the weighted sum of squared fiber distances from each of the left fibers to its nearest right fiber. Similar combinatorial distances have been applied to multi-fiber analyses by [25] [23]. We note that d_m^2 is asymmetric with respect to its inputs and invariant to the specific order of fiber compartments in M. Of course, when one fiber per voxel is present, this distance reduces to the single fiber distance d_f^2. We also need to compute the weighted average $\hat{M}$ under this distance, which can be defined and simplified as follows:

$$\hat{M} = \operatorname*{argmin}_M \sum_i^C w_i d_m^2(M_i, M) = \operatorname*{argmax}_{M,\pi} \sum_i^C \sum_j^{N_i} w_i f_{ij} (v_{ij} \cdot v_{\pi(ij)})^2 \tag{5}$$

For a fixed number of fibers in $\hat{M}$, this objective can be minimized by an iterative Expectation Maximization procedure similar to k-means clustering. In fact, this is equivalent to the procedure for hard Mixture of Watsons clustering of Sra et al [24]. With this in mind, we now move to the task of defining bilateral filtering for multi-fiber models.

2.4 Bilateral Filtering

Perhaps one of the most basic smoothing filters is the Gaussian blur, where a weighted average of pixel intensities is found based on spatial proximity to a given voxel. While this approach can remove noise, it also tends to smooth features that we'd rather preserve. In contrast, bilateral filtering is a non-linear technique that has been found to smooth images while preserving edges [26]. This is achieved by computing weights based on both spatial proximity and intensity similarity. For multi-fiber models, we can make a similar extension to the linear multi-fiber filter proposed in [6] by including weights for directional similarity of fiber models. We define such a filter on a per-voxel basis with a local estimation framework [15] as follows. Given an input voxel position and model (p_0, M_0) and local neighborhood $\{(p_i, M_i)\}_{i=1}^C$, with $M_i = \{(f_{ij}, v_{ij})\}_{j=1}^{N_i}$, the filtered model $\hat{M}$ is:

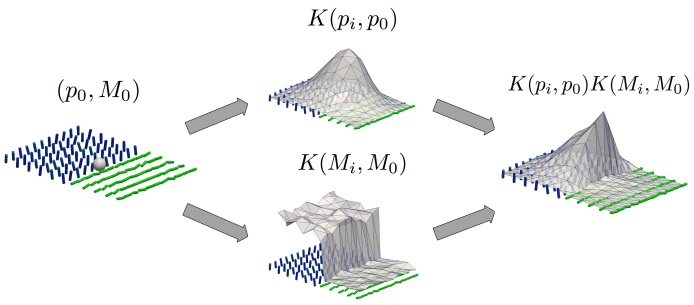

Fig. 1: An example of bilateral filter weights computed for a single voxel p_0 with associated fiber model M_0 (left), located near a boundary in a noisy phantom. For each voxel p_i and associated model M_i in a neighborhood of p_0, the bilateral weights (right) are given by the product of the linear Gaussian weights $K(p_i, p_0)$ (top) and data-adaptive weights $K(M_i, M_0)$ (bottom).

$$\hat{M} = \underset{M}{\operatorname{argmin}} \sum_{i=0}^{C} K\left(\frac{d_e^2(p_i, p_0)}{h_p^2}\right) K\left(\frac{d_m^2(M_i, M_0)}{h_m^2}\right) d_m^2(M_i, M) \tag{6}$$

given bandwidth parameters h_p and h_m and a kernel functions K, which we take to be the exponential $K(x) = \exp(-x)$. This defines non-linear filter weights that depend on M_0, as illustrated in Fig. 1. Each voxel may be processed separately by recomputing weights and solving Eq. 5 with the related Expectation Maximization procedure [24]. Two additional concerns are the number of fibers and the resulting volume fractions, which we estimate with standard bilateral filtering. The number of fibers is then a weighted average, which is rounded to the nearest integer, and the volume fractions are also weighted averages, but within groups defined by the optimal fiber correspondences π. Multiple passes through the volume may have some benefit, but we only consider a single pass.

3 Experiments

We performed evaluation with two experiments: the first applies tractography-based fiber bundle reconstructions in the human brain data, and the second measures the response to noise in computational phantoms and human brain data. We compared to the linear filtering approach [6] in both experiments.

3.1 Datasets

3.1.1 Computational Phantom Data

Two computational phantoms were constructed with single and double fiber models. The first represents the interface between two bundles, such as the corpus callosum/cingulum boundary. The second represents a similar interface with an additional bundle crossing both, such as the corpus callosum/corona radiata/superior longitudinal fasciculus juncture. These are represented by single and double fiber models, respectively, and both include fanning and curving to represent features of real data.

3.1.2 Human Brain Data

Diffusion MRIs were acquired from a healthy volunteer with a GE 1.5T scanner with a voxel size of 2mm^3, dimensions 128x128x72, seven T_2-weighted volumes, and 64 gradient encoding directions with b-value 1000 s/mm^2. Three repeated acquisitions of a single subject were concatenated to produce a high signal-to-noise volume. The repeated scans (high SNR) and a single acquisition (low SNR) were each processed with FSL to correct for motion, extract the brain, and fit single and multi-fiber models with Xfibres [4].

3.2 Design

The first experiment tested the effect of filtering on streamline tractography of the superior longitudinal fasciulus I [9], as shown in Fig. 2. We compared bundles with the Dice coefficient, fiber count, and volume, taking the high SNR acquisition as a reference and applying filtering to the low SNR scan. The second experiment tested the response to noise by randomly perturbing fiber orientations in the phantom and real data, as shown in Fig. 3. We measured error by the volume-fraction weighted minimum angular difference in degrees between models across all one-to-one pairings (the same metric used in [6]) and estimated the error rates by a Monte Carlo simulation with 1000 noise iterations with $h_p = 3.0$ and $h_m = 0.75$. We tested for a reduction in error by a one-sided paired t-test at each voxel, where samples were paired by noise iteration. We measured the per-voxel effect size by a paired Cohen's d-score. In both experiments, we compared linear and bilateral filtering.

3.3 Results

In the first experiment, we found bilateral filtering to produce more similar bundles to the reference than either the source or linear filtered volumes, as shown by an increased Dice score and similar fiber counts and volumes. In the second experiment, we found bilateral filtering to significantly reduce noise-induced error in all voxels ($d > 1.0$, $p < 0.05$) and found linear filtering to reduce noise in most areas, though not near some boundaries. Near these boundaries, the bilateral error was significantly lower with a large effect size ($d > 1.0$, $p < 0.05$). We also measured error as a function of adaptive bandwidth h_m and found a nonlinear trend that varied between high error as $h_m \to 0$ and the linear filtering error as $h_m \to \infty$ with a single global minimum between. On a 1.3 GHz Intel Core i5, our implementation ran in two minutes for a full brain volume.

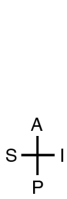

	Reference 3x64 directions	Unfiltered 64 directions	Linear Filtered 64 directions	Bilateral Filtered 64 directions
Dice	1.0	0.32	0.47	0.53
Count	130	36	312	125
Volume	2320 mm^3	608 mm^3	5288 mm^3	2432 mm^3

Fig. 2: Filtering effects on tractography of the left superior longitudinal fasciculus I, which runs anterior-posterior through dorsal frontal and parietal white matter. The high SNR (3x64 dir) taken as a reference, the low SNR scan (64 dir) was taken as a test case, and linear and bilateral filtering were applied to the test case. Using TrackVis, two spherical regions were manually chosen to delineate the bundle consistently across cases, and agreement of bundles with the reference was computed with the Dice coefficient, fiber count, and volume. We found bilateral filtering to be most similar to the reference, while linear filtering inflated both volume and count.

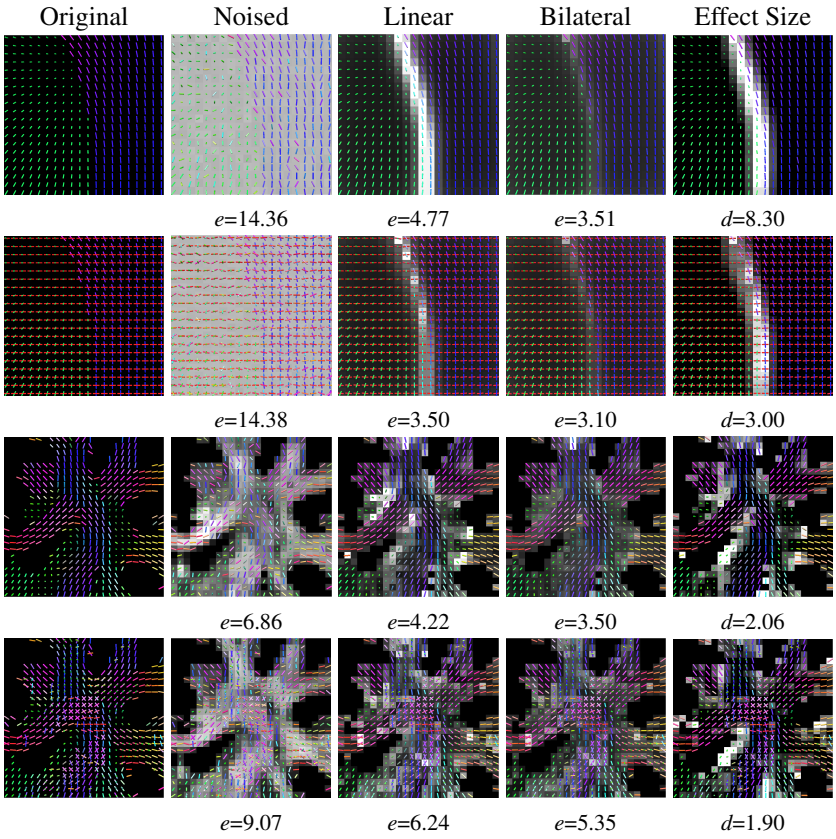

Fig. 3: Filter response to synthetic noise in two phantoms and two human brain cases. Orientation color-coded fiber models are superimposed on grayscale renderings of estimated voxel-wise error (e) and effect size (d) of the error reduction of bilateral vs. linear filtering (described in Section 2). Rows show (top to bottom): single and double fiber phantom, and coronal slices of single and double human brain data. Columns show (left to right): original fibers, noisy fibers, linear filtering error, bilateral filtering error, and significant improvements over linear filtering. We found both filters to significantly reduce noise-induced error, with linear filtering having higher error near junctions and boundaries but bilateral filtering showing reduced errors in these regions.

4 Conclusion

In this paper, we presented a bilateral filter for fiber orientation data. Our experiments suggest this approach is valuable for regularization and improving fiber bundle reconstruction. Other practical applications include brain connectivity mapping [2], heart modeling [22], and texture analysis [1]. Open problems include incorporation of microstructure properities, testing performance with partial volume effects, comparison with diffusion-weighted image smoothing techniques, and evaluation for clinical study. In conclusion, we found this approach to offer an efficient way to improve fiber-based modeling in diffusion MR images, as demonstrated with our experiments on synthetic phantom and real human brain datasets.

References

1. Aach, T., Mota, C., Stuke, I., Muhlich, M., Barth, E.: Analysis of superimposed oriented patterns. Image Processing, IEEE Transactions on **15**(12), 3690–3700 (2006)
2. Basser, P.J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, a.: In vivo fiber tractography using DT-MRI data. Mag. Res. in Med. **44**(4), 625–32 (2000)
3. Basu, S., Fletcher, T., Whitaker, R.: Rician noise removal in diffusion tensor mri. In: Medical Image Computing and Computer-Assisted Intervention MICCAI 2006, *Lecture Notes in Computer Science*, vol. 4190, pp. 117–125 (2006)
4. Behrens, T.E.J., Berg, H.J., Jbabdi, S., Rushworth, M.F.S., Woolrich, M.W.: Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? NeuroImage **34**(1), 144–155 (2007)
5. Bhattacharya, R., Patrangenaru, V.: Large Sample Theory of Intrisic Sample Means on Manifolds I. Annals of Statistics **31**(1), 1–29 (2013)
6. Cabeen, R.P., Bastin, M.E., Laidlaw, D.H.: Estimating constrained multi-fiber diffusion mr volumes by orientation clustering. In: Medical Image Computing and Computer-Assisted Intervention MICCAI 2013, *Lecture Notes in Computer Science*, vol. 8149, pp. 82–89 (2013)
7. Cabeen, R.P., Laidlaw, D.H.: White matter supervoxel segmentation by axial dp-means clustering. In: Medical Computer Vision. Large Data in Medical Imaging, Lecture Notes in Computer Science, pp. 95–104 (2014)
8. Coulon, O., Alexander, D., Arridge, S.: Diffusion tensor magnetic resonance image regularization. Medical Image Analysis **8**(1), 47–67 (2004)
9. De Schotten, M.T., Dell'Acqua, F., Forkel, S.J., Simmons, A., Vergani, F., Murphy, D.G., Catani, M.: A lateralized brain network for visuospatial attention. Nature neuroscience **14**(10), 1245–1246 (2011)
10. Demiralp, C., Laidlaw, D.H.: Generalizing diffusion tensor model using probabilistic inference in Markov random fields. In: Proceedings of MICCAI CDMRI Workshop (2011)
11. Goh, A., Lenglet, C., Thompson, P., Vidal, R.: A nonparametric riemannian framework for processing high angular resolution diffusion images (hardi). In: Computer Vision and Pattern Recognition, CVPR. IEEE Conference on, pp. 2496–2503 (2009)
12. Hamarneh, G., Hradsky, J.: Bilateral filtering of diffusion tensor magnetic resonance images. IEEE Trans. on Image Proc. **16**(10), 2463–2475 (2007)
13. Knutsson, H.: Representing local structure using tensors. In: Proceedings of the 6th Scandinavian Conference on Image Analysis (1989)
14. Knutsson, H., Westin, C.F.: Normalized and differential convolution. In: Computer Vision and Pattern Recognition, IEEE Computer Society Conference on, pp. 515–523 (1993)

15. Milanfar, P., Takeda, H., Farsiu, S.: Kernel regression for image processing and reconstruction. IEEE Trans. on Image Proc. **16**(2), 349–66 (2007)
16. Paris, S., Briceño, H.M., Sillion, F.X.: Capture of hair geometry from multiple images. ACM Trans. Graph. **23**(3), 712–719 (2004)
17. Pennec, X., Fillard, P., Ayache, N.: A Riemannian framework for tensor computing. International Journal of Computer Vision **66**(January), 41–66 (2006)
18. Perona, P.: Orientation diffusions. IEEE Trans. on Image Proc. **7**(3), 457–67 (1998)
19. Poupon, C., Clark, C., Frouin, V., Régis, J., Bloch, I., Le Bihan, D., Mangin, J.F.: Regularization of Diffusion-Based Direction Maps for the Tracking of Brain White Matter Fascicles. NeuroImage **12**(2), 184–195 (2000)
20. Rieger, B., van Vliet, L.: Representing orientation in n-dimensional spaces. Computer Analysis of Images and Patterns pp. 1–8 (2003)
21. Rydell, J., Knutsson, H., Borga, M.: Bilateral Filtering of fMRI Data. IEEE Journal of Selected Topics in Signal Processing **2**(6), 891–896 (2008)
22. Scollan, D., Holmes, A., Zhang, J., Winslow, R.: Reconstruction of cardiac ventricular geometry and fiber orientation using magnetic resonance imaging. Ann. of Biomed. Eng. **28**(8), 934–944 (2000)
23. Sigurdsson, G., Prince, J.: Smoothing fields of weighted collections with applications to diffusion MRI processing. SPIE Medical Imaging (2014)
24. Sra, S., Jain, P., Dhillon, I.: Modeling data using directional distributions: Part II. Technical Report TR-07-05, Dept. of CS, Univ. of Texas at Austin (2007)
25. Taquet, M., Scherrer, B., Commowick, O.: A Mathematical Framework for the Registration and Analysis of Multi-Fascicle Models for Population Studies of the Brain Microstructure. IEEE Trans. on Med. Imaging pp. 1–14 (2013)
26. Tomasi, C., Manduchi, R.: Bilateral filtering for gray and color images. In: Computer Vision, pp. 839–846 (1998)
27. Tristán-Vega, A., Aja-Fernández, S.: Dwi filtering using joint information for dti and hardi. Medical Image Analysis **14**(2), 205–218 (2010)
28. Tschumperlé, D., Deriche, R.: Orthonormal vector sets regularization with PDE's and applications. International Journal of Computer Vision **50**(3), 237–252 (2002)