

Minimizing Non-Holonomicity: Finding Sheets in Fibrous Structures

Babak Samari[§], Tabish Syed[§], and Kaleem Siddiqi

[§]equal contribution

School of Computer Science and Center For Intelligent Machines
McGill University, Canada
{babak,tabish,siddiqi}@cim.mcgill.ca

Abstract. Oriented elements compose fibrous structures in biological tissue, and their geometry plays an important role in organ function. In the heart, for example, myocytes are stacked end on end in a particular fashion to facilitate electrical conductivity and efficient mechanical contraction. In the brain, white matter fiber tracts are neuro-anatomically partitioned into specific bundles which connect distinct brain regions. In both cases, the local geometry has been qualitatively described as being sheet-like in particular regimes. Yet, to date, few if any quantitative methods exist for finding these sheets from imaging data. We here introduce a novel computational solution to this problem, motivated by the property that a holonomic vector field is locally normal to a family of smooth surfaces. We propose an algorithm which, given an input vector field, finds a second one with which it best spans a sheet-like structure locally, by minimizing non-holonomicity. We show that our algorithm converges in theory and in practice, under reasonable assumptions about the input data, and we provide high quality sheet reconstructions from both heart wall DTI data and labeled tracts in the human brain, along with a sheet likeliness measure. Whereas sheet-like geometries have been described qualitatively in past literature, ours is the first method to provide a reconstruction of them from a single direction field. Our algorithm also admits a parallel implementation that exploits GPUs and is hence very efficient. We thus anticipate that it will find use in the community for retrieving sheets on which oriented fibrous structures lie from imaging data.

Keywords: Fibrous Structures · Sheet Estimation · Non-Holonomicity.

1 Introduction

Fibrous biological tissues are reported to be organized in sheet-like geometries in certain regimes. Examples include laminar sheets and cleavage planes in the heart wall [6, 7], and sheets on which fiber tract systems in the mammalian brain are posited to lie [15]. Quantitative geometric characterizations of sheets in biological tissue are often based on orientation measurements from in-vivo

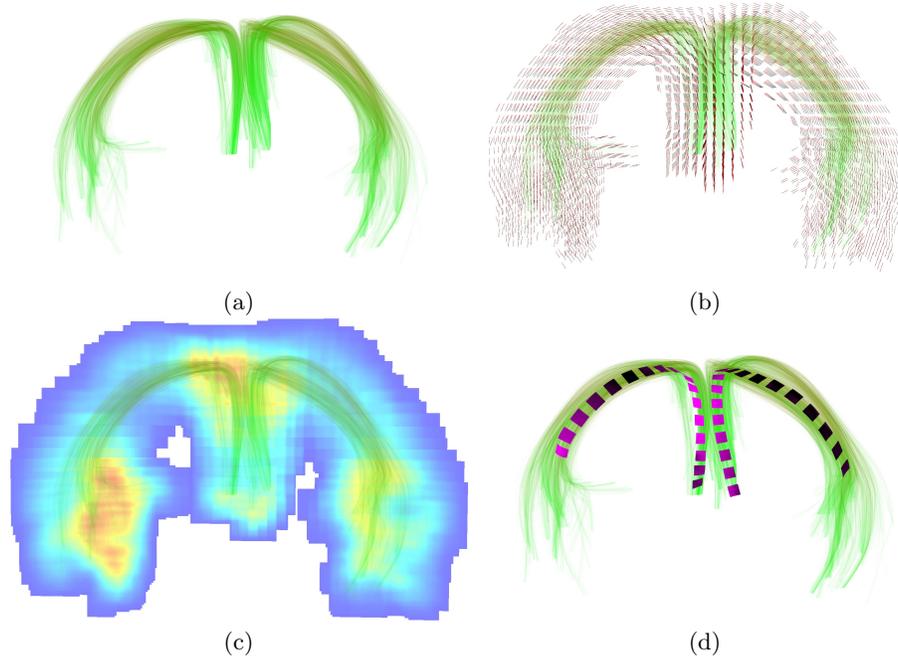


Fig. 1: (Best viewed by zooming-in on the PDF) (a) We assume as input a set of streamlines, in this example those corresponding to the left and right Fornix tracts from a Human Connectome Project (HCP) atlas. (b) We then extract a local direction field on a discrete grid, using the tangents to the streamlines. (c) The local non-holonomicity value, after convergence of our algorithm, gives a measure of the likelihood of a local sheet, with sheet probability decreasing from light blue to red using a jet colormap volumetric rendering. (d) Our estimated sheets are overlaid on the streamlines as magenta surfaces, but only in regimes found to support sheet-like geometries.

or ex-vivo specimens, using diffusion imaging. In analyses of the heart wall, the principal and second eigenvectors of a diffusion tensor are commonly assumed to span a local sheet, with the third eigenvector giving the direction of the sheet normal [4]. This approach has shown promise in the development of models of electrical conduction in the heart wall for electrophysiology [16], and in analyses of heart wall mechanics [3]. Statistical investigations of diffusion tensor data from populations of mammalian hearts have shown, however, that the second and third eigenvalues can be very close to one another in magnitude, and thus their directions can be locally ambiguous [11, 8]. As such, estimates of sheet geometries in myocardial tissue using all three eigenvectors of a diffusion tensor may not be reliable. The same is true of the mammalian brain; in regions where multiple tract systems co-exist and cross one another fiber geometry is not well characterized by a diffusion tensor.

In the human brain, the hypothesis that distinct tract systems lie on 2D sheets which intersect one another at approximately 90 degree angles to generate a local grid-like pattern, has been supported by qualitative considerations and visualization in [15]. Others have pointed out that distinct tract systems might span a sheet without being locally orthogonal to one another [2] and that quantitative measures of sheet geometries derived from diffusion data are desirable [13, 14, 1]. Yushkevich et al. have fit deformable medial models to segmented tract systems to obtain sheet-like representations of fasciculi where appropriate [18]. Motivated by the property that two vector fields span an integrable surface when the normal component of the associated Lie bracket goes to zero, Tax et al. have used this quantity to define a local sheet probability index along with robust algorithms to estimate it [13, 14]. In their work the local vector field directions are chosen from the peaks of a fiber orientation distribution function. The normal component of the Lie bracket has also been shown to be effective for sheet structure visualization, via the construction of a sheet tensor in [5], while Ankele and Shultz have applied this measure directly to diffusion tensor data in [1]. Motivated by these formal approaches to sheet structure estimation based on the integrability of vector fields, we study a more general problem here, which is the reconstruction of local sheets from a *single* direction field. Such a field might arise from the principal eigenvector direction of a diffusion tensor (e.g. in the case of the heart wall) or from the tangent vectors to precomputed streamlines, e.g., in the case of labeled tract systems in the mammalian brain. Fig. 1 presents an overview of our approach. We depart from past approaches by treating sheet structure estimation as an energy minimization problem, where, given the single direction field as input, we find a second vector field that is optimal with respect to spanning a local sheet with it. The optimality comes from the formal notion of holonomicity. As it turns out, this is equivalent to finding the second field as a local minimizer of the normal component of a Lie bracket. Based on these ideas, we present a gradient descent based algorithm to recover sheets, along with a proof of convergence. We also report a very efficient implementation of the algorithm using GPUs. Our application of this algorithm to mammalian heart wall orientation data from two species, rat and canine, reveals sheet geometries consistent with what has been reported only in qualitative descriptions thus far from DTI, or in observations from histological slices of heart wall tissue. Of more significance is our results on labeled tract systems from the Human Connectome Project, where we test the sheet hypothesis in a formal way. Given a tract system described as a collection of streamlines in 3D as input, we recover sheet geometries where there is support for them. As a by-product of our approach we provide high quality visualizations of the sheets in biological structures using software we have written. We anticipate that such visualizations will be useful to anatomists when examining fibrous structures in the context of their local surroundings.

2 Minimizing Non-Holonomicity: Theory and Analysis

2.1 Background : Sheets and Holonomicity

Given a unit vector field \mathbf{n} which is orthonormal to a family of surfaces, it can be shown that $\langle \mathbf{n}, \text{curl } \mathbf{n} \rangle = 0$ [17], where we use $\langle \cdot, \cdot \rangle$ to denote the inner product. Conversely, any vector field \mathbf{n} such that $\langle \mathbf{n}, \text{curl } \mathbf{n} \rangle = 0$ is orthonormal to a family of surfaces and is said to be *holonomic*. For a general vector field \mathbf{n} , its degree of non-holonomicity ρ is defined as follows [17]:

$$\rho = \langle \mathbf{n}, \text{curl } \mathbf{n} \rangle.$$

Consider two orthonormal vector fields \mathbf{u} and \mathbf{v} such that $\mathbf{n} = \mathbf{u} \times \mathbf{v}$. It is easy to show then, that *non-holonomicity* ρ expressed in terms of \mathbf{u} and \mathbf{v} reduces to

$$\rho^{\mathbf{uv}} = \langle \mathbf{u} \times \mathbf{v}, [\mathbf{u}, \mathbf{v}] \rangle,$$

where $[\mathbf{u}, \mathbf{v}]^i = u^j \frac{\partial v^i}{\partial x^j} - v^j \frac{\partial u^i}{\partial x^j}$ is the Lie Bracket of \mathbf{u} and \mathbf{v} . For a pair of vector fields which span a surface, the non-holonomicity ρ vanishes identically. While this is true in the continuous case, we expect the non-holonomicity value to be small for discrete fields spanning a surface. This is the key idea behind our approach.

Given a *single* vector field \mathbf{u} , we consider the problem of estimating an orthonormal vector field \mathbf{v} , such that \mathbf{u} and \mathbf{v} span a sheet locally. We propose a strategy which starts with a *single* input vector field \mathbf{u} , and estimates an orthonormal vector field \mathbf{v} by iterative minimization of an appropriate energy function.

2.2 Problem Formulation and Setup

Let \mathbf{v} be the current estimate of the vector field such that $\langle \mathbf{u}, \mathbf{v} \rangle = 0$. Further, let $\rho^{\mathbf{uv}}$ be the non-holonomicity function of \mathbf{u} and \mathbf{v} . Consider the non-holonomicity $\rho(\theta)$ corresponding to a perturbed vector field $\hat{\mathbf{v}} = \mathbf{v}(\theta)$, where θ is a scalar function which parametrizes the field $\hat{\mathbf{v}}$ with respect to field \mathbf{v} in the plane perpendicular to \mathbf{u} . Using the definition of non-holonomicity $\rho^{\mathbf{uv}}$ given above, a straightforward computation shows that $\rho(\theta)$ is given by

$$\rho(\theta) = \rho^{\mathbf{uv}} \cos^2 \theta + \rho^{\mathbf{un}} \sin^2 \theta + \alpha^{\mathbf{uv}} \sin 2\theta + \nabla_{\mathbf{u}} \theta = \rho_s(\theta) + \nabla_{\mathbf{u}} \theta \quad (1)$$

where, $\alpha^{\mathbf{uv}} = \frac{\langle \mathbf{n}, [\mathbf{u}, \mathbf{n}] \rangle - \langle \mathbf{v}, [\mathbf{u}, \mathbf{v}] \rangle}{2}$ and, $\rho^{\mathbf{uv}}$ and $\rho^{\mathbf{un}}$ are the non-holonomicity functions of the unperturbed fields \mathbf{u} and \mathbf{v} , and \mathbf{u} and \mathbf{n} .

Notice that for a constant perturbation function θ , $\nabla_{\mathbf{u}} \theta = 0$. Therefore $\rho(\theta) = \rho_s(\theta)$, the sinusoidal part of the $\rho(\theta)$ function which has a period of π . Further, a 180° turn of \mathbf{v} leaves $\rho(\theta)$ unchanged. We know that for regions with sheet-like geometry, the spanning vector fields \mathbf{u} and \mathbf{v} are such that $\rho^{\mathbf{uv}} = 0$. We therefore define an energy function as follows:

$$E(\mathbf{u}, \mathbf{v}, \theta) = \rho^2(\theta). \quad (2)$$

This energy is zero in regions where \mathbf{u} and \mathbf{v} span a sheet, and high in regions which are less sheet like. We can therefore pose the estimation of the vector field \mathbf{v} as the following energy minimization problem:

$$\begin{aligned} \mathbf{v}^* &= \arg \min_{\mathbf{v}} E(\mathbf{u}, \mathbf{v}, \theta) \\ &\text{subject to } \langle \mathbf{u}, \mathbf{v} \rangle = 0. \end{aligned} \quad (3)$$

2.3 Minimization Algorithm and Analysis

We solve the minimization problem (3) using an iterative gradient descent approach. We initialize with a \mathbf{v} , such that the orthogonality constraint $\langle \mathbf{u}, \mathbf{v} \rangle$ is satisfied. We maintain this constraint in the subsequent iterations by forcing each update to lie in the plane orthonormal to \mathbf{u} . At each step the varying vector field $\mathbf{v}(\theta)$ is updated using the gradient of the energy function $E(\mathbf{u}, \mathbf{v}, \theta)$. We update θ using a general discrete gradient descent update as follows:

$$\begin{aligned} \frac{\theta^{t+1} - \theta^t}{\eta} &= -\frac{\partial E(\theta)}{\partial \theta} \\ \implies \mathbf{v}^{t+1} &= \mathbf{v}(\theta^{t+1}), \end{aligned}$$

where η is the size of the time step. The update for \mathbf{v} is carried out implicitly by rotating \mathbf{v} about vector \mathbf{u} without explicitly using θ^{t+1} . To rotate the vector \mathbf{v} by an angle β we use the rotation matrix $\mathbf{R}_\beta^{\mathbf{u}} = \cos \beta \mathbf{I} + \sin \beta [\mathbf{u}]_\times + (1 - \cos \beta)(\mathbf{u}\mathbf{u}^T)$, where $[\mathbf{u}]_\times$ is the cross-product matrix corresponding to \mathbf{u} . This ensures that the orthonormality constraint is satisfied at every iteration, without having to explicitly express \mathbf{v} as a function of θ in the local coordinates.

One can observe that the energy E is a function not only of the values of the vector fields \mathbf{u} and \mathbf{v} at a point, but also of their derivatives. Therefore, the energy at a point depends on the value of the field at that point and its neighbourhood. It is quite possible, therefore, that a local point wise gradient descent update at a point, as described above, may increase the energy in the neighbourhood. In fact, a key contribution of our method is that it converges due to the following property

Proposition 1. *For an incompressible smooth vector field \mathbf{u} , there exists a positive η such that a gradient descent update of θ reduces the energy everywhere.*

Proof. Let \mathbf{u} be the fixed input vector field and \mathbf{v} be the current estimate of the second vector field. Consider the gradient of the energy function (2) at time t given by

$$\frac{\partial E(\theta)}{\partial \theta} = 2\left(\rho_s(\theta) + \nabla_{\mathbf{u}}\theta\right)\left((\rho^{\mathbf{u}\mathbf{u}} - \rho^{\mathbf{u}\mathbf{v}})\sin 2\theta + 2\alpha^{\mathbf{u}\mathbf{v}}\cos 2\theta + \text{div } \mathbf{u}\right),$$

where θ represents the perturbation with respect to the current \mathbf{u} and \mathbf{v} .

Before the update, $\theta \equiv 0$, therefore we have

$$\left.\frac{\partial E(\theta)}{\partial \theta}\right|_{\theta=0} = 2\rho^{\mathbf{u}\mathbf{v}}(2\alpha^{\mathbf{u}\mathbf{v}} + \text{div } \mathbf{u}) = \mathcal{E}_0^p.$$

Then, θ^{t+1} is given by

$$\theta^{t+1} = -\eta \mathcal{E}_0^\rho.$$

To prove our claim, it is sufficient to show that this update reduces the energy, so that

$$E(0) > E(-\eta \mathcal{E}_0^\rho). \quad (4)$$

For small θ we let $\sin \theta \approx \theta$, $\cos \theta \approx 1$. The condition for convergence in (4) for small positive η then reduces to

$$4\alpha^{\mathbf{uv}}(2\alpha^{\mathbf{uv}} + \operatorname{div} \mathbf{u}) + \nabla_{\mathbf{u}} \mathcal{E}_0^\rho > 0.$$

Since we have assumed that our initial fixed vector field \mathbf{u} represents smoothly varying local orientation in fibrous tissue, we can assume that $\operatorname{div} \mathbf{u}$ is small, and then the convergence condition reduces to $\nabla_{\mathbf{u}} \mathcal{E}_0^\rho \gtrsim -8(\alpha^{\mathbf{uv}})^2$. For a smooth enough vector field with small divergence, we can choose a positive time step η such that the energy is reduced.

It is possible to consider a stronger smoothness condition: $|\nabla_{\mathbf{u}} \mathcal{E}_0^\rho| < 8(\alpha^{\mathbf{uv}})^2$, which helps us improve our error function and select the additional parameters, as discussed in the following section.

3 Method Evaluation and Parameter Selection

Data In order to validate our method we carried out a few experiments with data obtained from two sources. The first is a dataset of rat and canine hearts used in [12], which is available for public download. It consists of ex-vivo diffusion tensor (DT) fits to myocardial tissue, with resolutions of $64 \times 64 \times 128$ (rat) and $300 \times 300 \times 333$ (canine) voxels, respectively. We use the principal eigenvector of the tensor, which gives an acceptable proxy for myofiber orientation [8], as our fixed vector field. The second source is a dataset of labeled fiber tract systems in the human brain, constructed from a fiber bundle atlas generated from data from the Human Connectome Project, used in the ISMRM 2015 Tractography challenge [9]. This dataset is available for public download as well. The tangents at a list of points along fiber tracts from the data were used to generate our fixed direction vector field.

Modifying the error function As a corollary of our convergence analysis above, we infer from the bound on the energy gradient that a smoother energy function will have better convergence properties. This is in fact the case in practice, and we therefore used a modified error function \hat{E} defined as

$$\hat{E}(\mathbf{u}, \mathbf{v}, \theta(\mathbf{x})) = \sum_{\hat{\mathbf{x}} \in Nbd(\mathbf{x})} \rho^2(\hat{\mathbf{x}}) \quad (5)$$

where $Nbd(\mathbf{x})$ is the neighbourhood of \mathbf{x} . We used a neighbourhood size of 3×3 in our experiments. This energy converges for a wider range of choices of the step

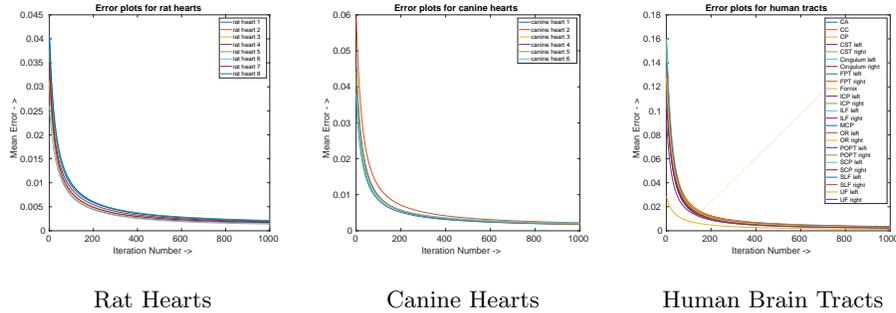


Fig. 2: (Best viewed by zooming-in on the PDF.) Mean error as a function of iteration number, for all examples reported in this paper.

size η . Intuitively one may understand this modification as increasing the size of the neighbourhood while calculating the deviation from sheet-like geometry at a point.

Choosing η : The convergence bound can also act as a guide to choosing the time step η , which is the only parameter other than the neighborhood size, in our approach. In accordance with the final convergence condition, we observed empirically that for brain tracts, the range of feasible η choices is smaller. This can be attributed to the more complex tract geometry in the brain compared to the heart. In fact, while larger η 's worked well for simple tracts like the Inferior Longitudinal Fasciculus (ILF), more complex tract systems such as the corpus callosum (CC) required the use of a smaller η for convergence. Nonetheless, an η of 0.1, as shown in Fig. 2, resulted in smooth convergence for every tract system, as predicted by our analysis.

Convergence rate: As shown by the plots in Fig. 2, for our 3 datasets, with $\eta = 0.1$, the mean error over all voxels in a dataset starts to flatten out after about 500 iterations. A quantitative analysis of the convergence rate is however beyond the scope of the present article, but we aim to take this up in the future.

Implementation and runtime We implement our algorithm in the PyTorch framework [10]. To achieve this the derivatives and the spatial average calculations for the modified energy function are formulated as convolutions. This allows for GPU-based computation which significantly reduces run-time and makes it feasible to analyze large datasets, in practical exploratory settings. We provide a comparison of run times on a CPU (Intel Core i9-7900x) and a GPU (Titan Xp), for volumes of different sizes in Table 1.

4 Sheet Reconstruction and Visualization

The output of our algorithm is a vector field \mathbf{v} which is locally best in the sense of spanning a sheet with \mathbf{u} , together with an energy value at each voxel that is proportional to non-holonomicity. We use this final energy value as a

Sample	Size (Voxels)	CPU (s)	GPU (s)
Rat Heart	$64 \times 64 \times 128$	7.9	0.1
Canine Heart	$300 \times 300 \times 333$	279.3	5.6
Human Brain Tracts	$90 \times 108 \times 90$	11.0	0.1

Table 1: Time taken in seconds, per iteration, for a CPU based implementation versus that taken by a GPU based implementation, for different volumes.

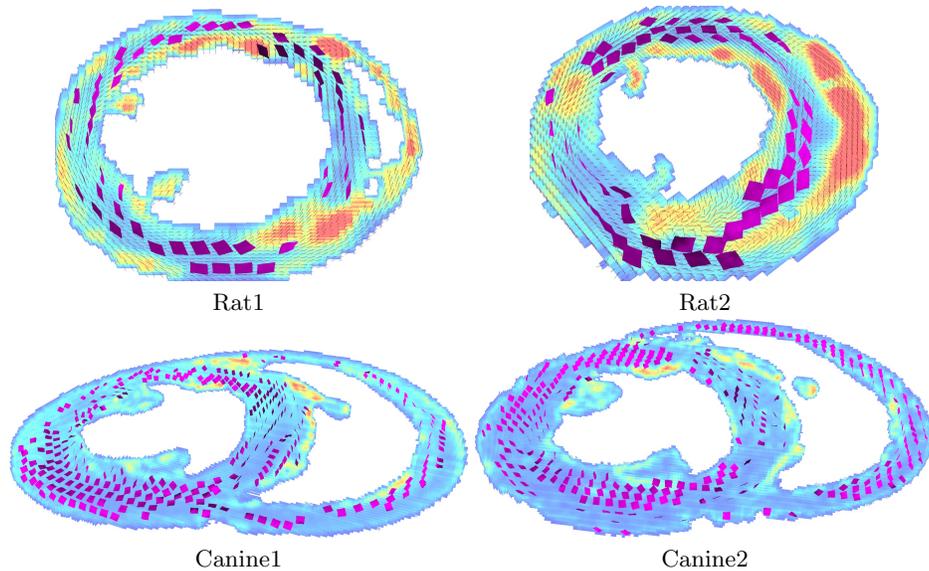


Fig. 3: (Best viewed by zooming-in on the PDF.) Sheet likelihood estimates in short axis slices of two canine and two rat hearts, illustrated with volumetric jet colormap rendering, along with our estimated sheets shown as magenta surfaces.

guide for exploring and reconstructing sheets in the brain and the heart, in an iterative breadth first approach. At each step we extend the sheet by a small quadrilateral sheetlet, composed of two triangular faces. The two triangular faces are generated by moving in the direction of \mathbf{u} followed by \mathbf{v} with a step size of around $ds = 0.2$ voxels, and then in direction \mathbf{v} followed by \mathbf{u} . For sheet-like regions the two triangular faces are expected to share an edge. In fact, the gap between the two triangular faces is proportional to the Lie bracket $[\mathbf{u}, \mathbf{v}]$ at a point. In our visualizations we only show sheets in locations with low energy. This reconstruction process is repeated n times (with $n \approx 7$) in all four $\pm\mathbf{u}, \pm\mathbf{v}$ directions, in a breadth-first traversal manner. Since the reconstructed points are not limited to voxel locations, we use distance from already drawn vertices, r , as a criterion for adding new points. The parameter r was fixed at half the step size. We developed an OpenGL based visualization tool for rendering all the results presented in this paper. For the heart data, we draw sheets with uniform

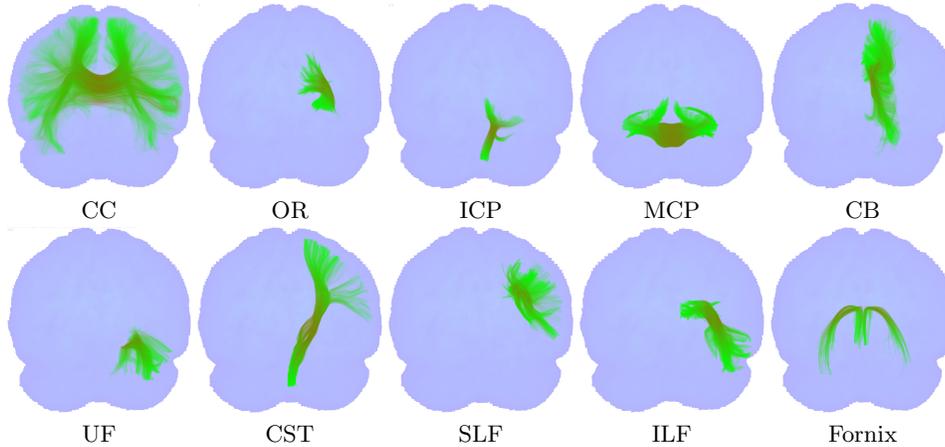


Fig. 4: The HCP tracts analysed in Fig. 5, shown in the context of their actual locations in the brain.

sampling in a given box (Fig. 3), while for the brain data, we draw sheets, when supported, at locations uniformly sampled along the tracts (Fig. 5).

5 Results

We encourage the reader to zoom-in on the electronic (PDF) version of each Figure. Fig. 3 presents sheets reconstructed from diffusion tensor data for two rat hearts (top row) and two canine hearts (bottom row). In all sub-figures we show a short axis slice, with a jet colormap volume rendering of the final sheet fitting energy (energy increases from cyan to red), and with our reconstructed sheets shown as magenta surfaces. For the rat hearts we also show the direction of the principal eigenvector as an orientation field. In both species, there is clear support for both axial sheets in the wall of the left ventricle (LV), consistent with the laminar organization reported via histology in early work [7], as well as regimes of more circumferential sheets. The geometry of sheets in the septum is more complex, being predominantly circumferential, and exhibiting a degree of fanning. The right ventricle (RV) also shows sheet-like geometries for the canine hearts, while for the rats the RV is squashed due to imaging conditions. Finally, the LV papillary muscles are associated with a higher energy, and thus a lower likelihood of sheets, consistent with the property that along them muscle cells are oriented in the long axis direction. The sheet fitting energy is also high at the junctions of the LV and RV.

Moving to the human brain, Fig. 1 (bottom left) shows sheet fitting energy rendered as a jet colormap for the Fornix tract, with our recovered sheets shown as magenta surfaces (bottom right). We then consider many other tract systems, which are shown in the context of their actual positions in the full brain in Fig. 4. For several of these we have chosen the left hemisphere tract to analyze. A

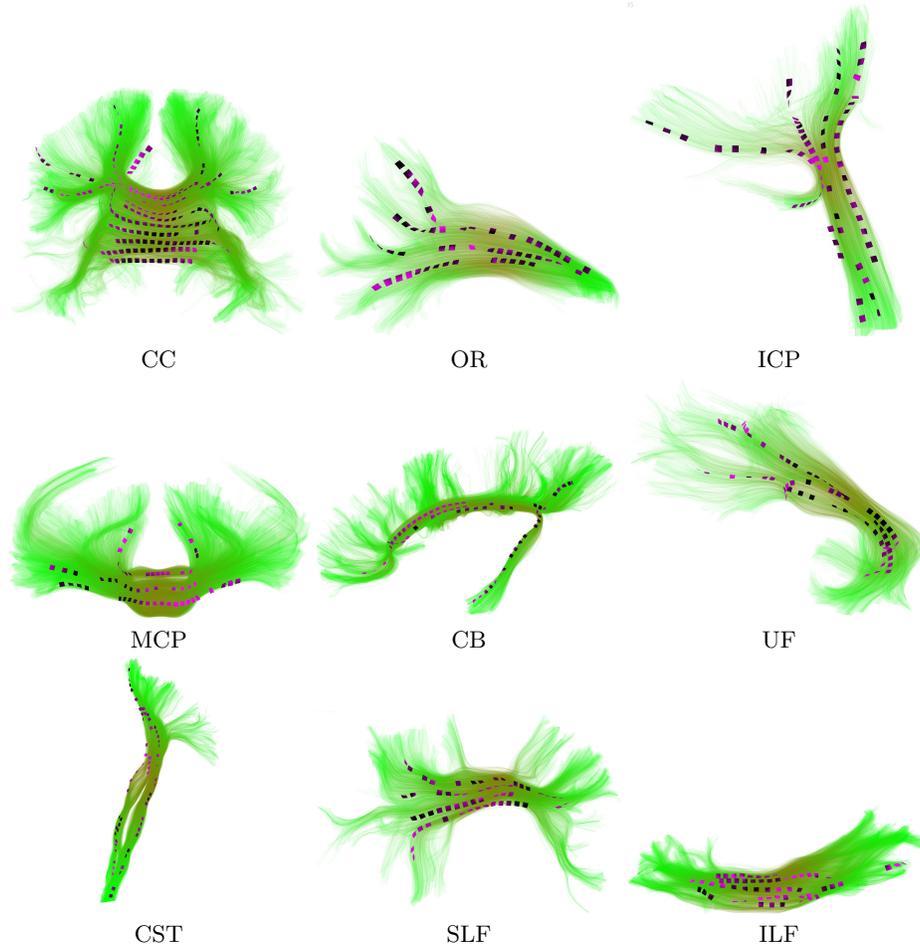


Fig. 5: (Best viewed by zooming-in on the PDF.) Tract systems from an HCP atlas are shown with green streamlines, with the estimated sheets (where present) shown as magenta surfaces: corpus callosum (CC), optic radiation (OR), inferior cerebellar peduncle (ICP), middle cerebellar peduncle (MCP), cingulum bundle (CB), uncinate fasciculus (UF), corticospinal tract (CST), superior longitudinal fasciculus (SLF), and inferior longitudinal fasciculus (ILF).

number of these tract systems (CC, CB, SLF, Fornix) have been considered in cross species qualitative investigations of sheet geometry in [15], but with no explicit sheet reconstruction. Our recovered sheets for these tracts, where we find support, are depicted as magenta surfaces in Fig. 5. The corresponding sheet fitting energies are rendered as a volumetric jet colormap in Fig. 6. It is clear from Figs. 1, 5 and 6 that whereas certain regimes of these tracts are indeed sheet-like (e.g. the two main arms of the Fornix, and the middle section of the CC),

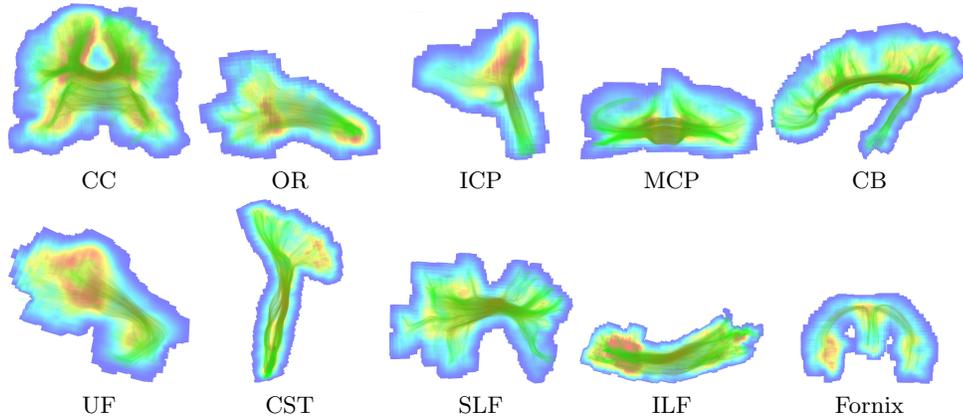


Fig. 6: (Best viewed by zooming-in on the PDF.) The HCP tracts analysed in Fig. 5, shown superimposed on the final sheet fitting energy, visualized using volume rendering with a jet colormap (energy increases from cyan to red).

others, such as the fanning regions, are not. Our method allows for navigation and labeling of a tract by sheet confidence, and for the quantitative recovery of subtle shape properties that thus far have been described only qualitatively. For example, the CST is more tube like, providing little local evidence of sheet geometries in most parts, while the MCP, SLF and ILF have large sheet-like regions.

6 Conclusion

Past analyses of tract geometry in the human brain have asked the question whether *two* local directions, often taken to be the peak directions of a fiber orientation distribution function [13, 14], or the first and second principal directions of a diffusion tensor [1], provide support for sheets. These attempts have had success in designing sheet probability measures by considering the normal component of the Lie bracket. The problem we have studied here is different in a subtle way, namely, we have asked whether a *single* direction field derived from fibrous tissue supports sheet-like geometries. We have designed an efficient algorithm to minimize non-holonomicity that converges, provides actual reconstructions and high quality visualizations of sheets on which myofibers in the heart are organized, or sheets on which fiber tract systems in the brain might lie, and gives a local measure of sheet likeliness. Such an algorithm could now be used to settle questions concerning the geometric organization of tract systems as suggested in [15], in a quantitative and principled way.

Acknowledgments We are grateful to NSERC (Canada) and FRQNT (Québec) for funding that has supported this research.

References

1. Ankele, M., Schultz, T.: A sheet probability index from diffusion tensor imaging. In: Computational Diffusion MRI, pp. 141–154. Springer (2018)
2. Catani, M., Bodi, I., Dell’Acqua: Comment on “the geometric structure of the brain fiber pathways”. *Science* **337**(6102), 1605 (2012)
3. Dou, J., Tseng, W.Y.I., Reese, T.G., Wedeen, V.J.: Combined diffusion and strain mri reveals structure and function of human myocardial laminar sheets in vivo. *Magnetic Resonance in Medicine* **50**(1), 107–113 (2003)
4. Gilbert, S.H., Benson, A.P., Li, P., Holden, A.V.: Regional localisation of left ventricular sheet structure: integration with current models of cardiac fibre, sheet and band structure. *Eur. J. Cardio-Thorac. Surg* **32**(2), 231–249 (2007)
5. Haije, T.C.J.D.: Finsler Geometry and Diffusion MRI. Ph.D. thesis, Eindhoven University of Technology (2017)
6. Helm, P., Beg, M.F., Miller, M.I., Winslow, R.L.: Measuring and mapping cardiac fiber and laminar architecture using diffusion tensor mr imaging. *Annals of the New York Academy of Sciences* **1047**(1), 296–307 (2005)
7. LeGrice, I.J., Smaill, B.H., Chai, L.Z., Edgar, S.G., Gavin, J.B., Hunter, P.J.: Laminar structure of the heart: ventricular myocyte arrangement and connective tissue architecture in the dog. *Am. J. Physiol.* **269**(2), H571–82 (1995)
8. Lombaert, H., Peyrat, J.M., Croisille, P., Rapacchi, S., Fanton, L., Cheriet, F., Clarysse, P., Magnin, I., et al.: Human atlas of the cardiac fiber architecture: study on a healthy population. *IEEE T-MI* **31**(7), 1436–1447 (2012)
9. Maier-Hein, K.H., Neher, P.F., Houde, J.C., Côté, M.A., Garyfallidis, E., Zhong, J., Chamberland, M., Yeh, F.C., et al.: The challenge of mapping the human connectome based on diffusion tractography. *Nature communications* **8**(1), 1349 (2017)
10. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
11. Peyrat, J.M., Sermesant, M., Pennec, X., Delingette, H., Xu, C., et al.: A computational framework for the statistical analysis of cardiac diffusion tensors: Application to a small database of canine hearts. *IEEE T-MI* **26**(11), 1500–1514 (Nov 2007)
12. Savadjiev, P., Strijkers, G.J., Bakermans, A.J., Piuze, E., Zucker, S.W., Siddiqi, K.: Heart wall myofibers are arranged in minimal surfaces to optimize organ function. *Proceedings of the National Academy of Sciences* **109**(24), 9248–9253 (2012)
13. Tax, C.M., Haije, T.D., Fuster, A., Westin, C.F., Viergever, M.A., Florack, L., Leemans, A.: Sheet probability index (spi): Characterizing the geometrical organization of the white matter with diffusion mri. *NeuroImage* **142**, 260–279 (2016)
14. Tax, C.M., Westin, C.F., Haije, T.D., Fuster, A., Viergever, M.A., Calabrese, E., Florack, L., Leemans, A.: Quantifying the brain’s sheet structure with normalized convolution. *Medical image analysis* **39**, 162–177 (2017)
15. Wedeen, V.J., Rosene, D.L., Wang, R., Dai, G., et al.: The geometric structure of the brain fiber pathways. *Science* **335**(6076), 1628–1634 (2012)
16. Young, R.J., Panfilov, A.V.: Anisotropy of wave propagation in the heart can be modeled by a riemannian electrophysiological metric. *Proceedings of the National Academy of Sciences* **107**(34), 15063–15068 (2010)
17. Yu, A.: The geometry of vector fields. Gordon & Breach Publ (2000)
18. Yushkevich, P.A., Zhang, H., Simon, T.J., Gee, J.C.: Structure-specific statistical mapping of white matter tracts. *NeuroImage* **41**(2), 448 – 461 (2008)