<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *IEEE 13th International Symposium on Biomedical Imaging (ISBI)*.

Citation for the original published paper:

Jörgens, D., Smedby, Ö., Moreno, R. (2016)
Steering second-order tensor voting by vote clustering.
In: (pp. 1245-1248). IEEE Computer Society
<https://doi.org/10.1109/ISBI.2016.7493492>

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-190103>

STEERING SECOND-ORDER TENSOR VOTING BY VOTE CLUSTERING

Daniel Jörgens, Örjan Smedby, Rodrigo Moreno

School of Technology and Health, KTH Royal Institute of Technology
Alfred Nobels Allé 10, 141 52 Huddinge, Sweden
{daniel.jorgens, orjan.smedby, rodrigo.moreno}@sth.kth.se

ABSTRACT

Among the various diffusion MRI techniques, diffusion tensor imaging (DTI) is still most commonly used in clinical practice in order to investigate connectivity and fibre anatomy in the human brain. Besides its apparent advantages of a short acquisition time and noise robustness compared to other techniques, it suffers from its major weakness of assuming a single fibre model in each voxel. This constitutes a problem for DTI fibre tracking algorithms in regions with crossing fibres. Methods approaching this problem in a postprocessing step employ diffusion-like techniques to correct the directional information. We propose an extension of tensor voting in which information from voxels with a single fibre is used to infer orientation distributions in multi fibre voxels. The method is able to resolve multiple fibre orientations by clustering tensor votes instead of adding them up. Moreover, a new vote casting procedure is proposed which is appropriate even for small neighbourhoods. To account for the locality of DTI data, we use a small neighbourhood for distributing information at a time, but apply the algorithm iteratively to close larger gaps. The method shows promising results in both synthetic cases and for processing DTI-data of the human brain.

Index Terms— Tensor voting, spherical clustering, DTI

1. INTRODUCTION

Diffusion MRI techniques have been used to study connectivity between different regions in the human brain. Based on local diffusion profiles in the white matter, tractography methods can extract fibre paths that connect different regions.

The most mature diffusion method is diffusion tensor imaging (DTI) in which anisotropy and orientation of fibres are locally modelled through second-order tensors. The main drawback of DTI is that it is unable to resolve crossings of fibres, which are not unusual in the brain. There are two strategies for addressing this issue: either by using more advanced acquisition protocols such as high angular resolution diffusion imaging (HARDI) or by performing postprocessing of DTI data. Most of the previous work has followed the first strategy. However, it is noteworthy to mention that DTI has much shorter acquisition times and is less prone to motion artifacts than HARDI, which makes DTI more appropriate for clinical use. This means that methods following the latter

strategy have more potential to be used in clinics.

Two different approaches have been followed for resolving crossings from DTI data. On the one hand, tractography algorithms can apply global strategies for disentangling the fibres [1]. On the other hand, the DTI data can be preprocessed before tractography [2, 3]. These methods assume that crossings can be resolved by considering neighbouring data. The proposed method follows the second approach.

The authors of [2] apply diffusion for obtaining improved orientation distribution functions (ODFs) based on information from the neighbourhood. In a second step, different fibre orientations can be extracted from the estimated ODFs. In [3], the ODF representation is avoided by performing all computations with tensors. First, higher-order tensors were obtained from DTI data based on outer products. Second, regions with low anisotropy are being “*inpainted*” by using an extended version of tensor voting which is applicable to higher-order tensors. Finally, the fibre orientations at the crossings are extracted from those inpainted tensors.

Although the proposed method in this paper is related to the one described in [3], there are two main differences between the two approaches. First, instead of higher-order tensor voting, we apply the classical second-order technique, which is less computationally expensive. Second, instead of summing up the votes cast by neighbours, we perform a directional analysis of these votes in order to assess the most likely orientations of the fibres. Thus, it is not necessary to perform higher-order tensor decomposition, which is not a straightforward problem.

The paper is structured as follows. First, we elaborate how we model the information provided by DTI within the tensor voting framework. From analysing the distribution of received tensor votes we derive how to steer their aggregation properly. Finally, an iterative algorithm is presented that accounts for valid assumptions on DTI measurements. Experimental results on synthetic data including comparison with [3] as well as on human brain data are shown.

2. METHOD

In the following we give a brief introduction into the framework of tensor voting. The central idea of the technique is to encode orientation features as second-order tensors. Based

on its eigendecomposition each symmetric positive definite tensor T can be decomposed into a weighted sum of corresponding rank-1, rank-2 and rank-3 tensors amounting in

$$T = s_1 S + s_2 P + s_3 B. \quad (1)$$

The tensors S , P and B describe the linear, planar and isotropic part of T and their weights $s_1 = \lambda_1 - \lambda_2$, $s_2 = \lambda_2 - \lambda_3$ and $s_3 = \lambda_3$ defined from the eigenvalues of T represent their saliencies. As it is most convenient for our purpose, we choose S and P to encode the tangent space of a particular feature like a curve or a surface, respectively. By modelling three perceptual principles, namely *good continuation*, *proximity* and *similarity*, the particular features are distributed in a neighbourhood of each tensor T , denoted the *voting* step. On the final stage, all votes received at a particular site are summed up and the resulting tensor is being analysed in terms of the basic features from above.

Figure 1 visualises schematically the distribution of the one-dimensional feature S , also referred to as *stick voting*. The computation of such a vote from a position $\mathbf{q}$ to $\mathbf{p}$ is given by $SV(\mathbf{p} - \mathbf{q}, S_{\mathbf{q}}) = w_s R_{2\theta_{pq}}(S_{\mathbf{q}})$, where $R_{2\theta_{pq}}$ performs a rotation of $S_{\mathbf{q}}$ so that the result agrees with the tangent in $\mathbf{p}$ as depicted in Fig. 1. The weighting factor w_s decays exponentially with larger distance $\|\mathbf{q} - \mathbf{p}\|$ and larger curvature of the osculating circle defined by $\mathbf{p}$ and $\mathbf{q}$. Basically two parameters, σ and c , steer the influence of distance, i.e. scale, and curvature. Voting is restricted to a cone of 45° deviation from the orientation of $S_{\mathbf{q}}$. As explained further in Subsect. 2.1, we only employ stick voting in our method.

Fig. 1. Visualisation of stick voting. The tensor $S_{\mathbf{q}}$ is voted to $\mathbf{p}_1$, $\mathbf{p}_2$ and $\mathbf{p}_3$. The rotation axis lies perpendicular to the image plane and the osculating circle is indicated by the particular dashed line. Weighting promotes straight continuation ($\mathbf{p}_2$) and punishes high curvature ($\mathbf{p}_1$) and large distance ($\mathbf{p}_3$).

2.1. Modelling diffusion information in tensor voting

As the data produced by DTI is already tensor valued, it is natural to analyse them in terms of Eq. 1. Clearly, S coincides with the principle diffusion direction of the corresponding diffusion tensor and its weight s_1 captures the confidence about that orientation. Following the argumentation of [3] we employ the linearity measure $c_l = \frac{\lambda_1 - \lambda_2}{\lambda_1}$ introduced in [4] to define whether S is likely to be well defined. If c_l is sufficiently large we assume that the DTI model assumption of a single fibre configuration is valid and $s_1 S$ is taken for initialisation of tensor voting in that point. This group of voxels will be referred to as the *voters* as their information

will be spread in the stick voting step.

As opposed to that, from a diffusion tensor with a low linearity measure we can not deduct a single direction. Therefore, this group of voxels, referred to as *receivers*, collects votes from their local neighbourhood in order to infer orientation information. However, if the weight of its planar contribution, that is s_2 , is large enough we can still conclude that diffusion is mainly taking place in the plane defined by P . For that reason the collected votes are weighted with the size of their projection onto that plane. By that, precedence is given to those orientations that coincide with the local planar information.

Finally, we cannot extract any orientation information from the three-dimensional feature B and therefore exclude it from further considerations.

As mentioned above the scale parameter σ controls how fast the weighting function w_s decays with distance and thereby implicitly determines the actual size of the voting neighbourhood. Usually, its choice depends on the distance between data points that should be grouped into the same structure. Regarding that diffusion tensors describe the *local* distribution of diffusion in the region of a single voxel, we expect the prediction of fibre orientation through tensor voting only to be valid in a reasonably small neighbourhood. For that reason we choose $\sigma = 1$ resulting in a maximal weight of 0.368 at a neighbouring voxel center. Additionally choosing $c = 60$ amounts in a weight of 0.011 for an angular deviation of 15° at unit distance.

2.2. Analysing the local distribution of votes

Due to the fact that tensor voting is—just like DTI—employing second-order tensors for encoding orientation information it suffers from the inability to resolve more than one principle orientation. In case all received votes at a particular site agree in orientation (up to a certain deviation), the linear part, S , of their sum is a good representation for their mean orientation and its corresponding weight s_1 measures the mentioned deviation. However, if the distribution of the votes exhibits multiple peaks it does not make sense to describe it by one major orientation. Based on this observation, we propose to group the received votes into clusters of similar orientation. This makes it possible to analyse each of those clusters separately using the same mechanism provided by classical tensor voting.

For grouping the received votes we employ *mean shift* with the von Mises-Fisher (vMF) density kernel for clustering data on S^2 (cf. [5]). The input is given by the first eigenvector $\mathbf{e}_1$ of each tensor vote. Since we are interested in finding clusters of orientations (as opposed to directions), we use just the antipodally symmetric part of the von-Mises-Fisher kernel (cf. [6]) given by $f(\mathbf{e}_1, \boldsymbol{\mu}; \kappa) = C_f(\kappa) \cosh(\kappa \mathbf{e}_1 \cdot \boldsymbol{\mu})$ with mean $\boldsymbol{\mu}$ and concentration parameter κ . To further account for the different strengths of the votes we employ weighting of the data points by w_s in the computation of the mean shift,

which is in accordance with [7]. κ was chosen as 15 in the experiments of Sect. 3.

Fig. 2. Impact of the voxel grid on classical tensor voting when employed for tangential curve inference in differently oriented line structures. Dotted connections indicate voting to the central position, red ellipses qualitatively indicate the resulting tensors.

2.3. Sampling of the voting field

In order to analyse the distribution of the received votes, we have to ensure that they carry the expected information. For that purpose we consider the test case of a line segment with a width of three pixels (Fig. 2). In classical tensor voting only the input tokens—i.e. in our case the voxel positions—do vote to each other. Together with the choice of a small voting neighbourhood (cf. Subsect. 2.1) this results in a relatively small number of received votes. In this setting we observe that the voxel grid makes the classical method not rotational invariant. As depicted in Fig. 2, the resulting tensor obtained by summation of the votes coincides with the line orientation in case of proper alignment with the grid (45°). However, the method yields an erroneous result if the line is rotated by 30° .

The central problem with that behaviour is that tensor votes are distributed only into a few directions given by the alignment of the voting tensor with the grid which introduces a bias to the distribution of the received votes. To account for that problem, we perform voting in a continuous coordinate space. Each tensor is voting to randomly chosen positions within a spherical neighbourhood. By this, we make sure that the evolution of a tensor is, given a sufficient number of sampling points, equally likely in all directions and only influenced by the weighting factor from stick voting.

In a second step, each receiver collects all votes within a certain neighbourhood which we choose to be round shaped in order to again give all directions the same precedence. In that sense, the resulting clusters represent a local average of the previously created pseudo-dense field of tensor votes. As a side effect, the size of the averaging neighbourhood (radius of 1.2 voxels in our experiments) also has a smoothing effect in the result.

2.4. Iterative algorithm

Due to the way we define the voter masks in Subsect. 2.1, we do not have direct control over the size of the gaps created in it. But since the size of the voting neighbourhood is limited, it might not be possible to close the gaps within one pass of tensor voting. For this reason we employ the method in an iterative manner. Cluster orientations extracted in one step become voters in the next iteration. Nevertheless, the receiver

mask remains unchanged to successively refine the orientation information.

In order to maintain stability of the orientation distributions over iterations we keep the votes from previous steps, but reduce their weights by a constant factor in each iteration. Votes from the initial voter group are distributed only once and remain unchanged. To have a comparable range of weights, both groups, initial votes and those induced by extracted clusters, are normalised prior to each voting step. An additional factor adjusts their strengths relative to each other. Those factors were respectively empirically determined as 1.2 and 0.78 in our experiments.

3. EXPERIMENTS

3.1. Synthetic data

To quantitatively assess the accuracy of our method in the crossing region, we measure the mean and standard deviation, denoted $\Delta\theta$ and $\sigma_{\Delta\theta}$, of the angular error of the extracted fibre orientations with respect to the ground truth available in synthetic data experiments. Further, the fraction m of voxels in which each fibre from the ground truth can be matched to a cluster, is counted. Only these voxels are regarded for computation of $\Delta\theta$.

Structures exhibiting different characteristics are used for experiments. Besides the 90° -crossing already introduced in [3] we further employ a crossing of 70° (cf. Fig. 4 (a)) as well as curved fibres crossing a straight line structure (cf. Fig. 4 (b)). To prove rotation invariance, the crossing examples are rotated by 15° . The data was simulated with a b-value of 1000 s/mm^2 and noise was added in order to get a signal-to-noise ratio of 15 (w.r.t. the baseline signal S_0).

Table 1 shows a comparison of our clustering approach with the higher order method from [3]. As an implementation of an iterative version of [3] was not available, results after one iteration with a sufficiently large scale $\sigma = 2$ are used. To allow for a fair comparison both methods use the randomised sampling technique introduced in Subsect. 2.3. Both approaches lead to accurately reconstructed orientations in noiseless scenarios. Employing the same respective parameter settings for the methods in the noisy case, they provide a similar number of reconstructed voxels, but our method shows a better accuracy in our specific experiments.

Figure 3 visualises the effect of iterations with a scale of $\sigma = 1$ (cf. Subsect. 2.1). After the first pass of tensor voting, some voxels show incorrect orientations and not all clusters have a high saliency. However, after four iterations the crossing is well reconstructed. The result depicted in Fig. 4 (a) further shows that the method can also handle smaller angles (70°). Even the gradually changing circular structure in Fig. 4 (b) can be reconstructed. The lower saliencies within the circle structure compared to the straight line are due to the large gap to fill, the preference of tensor voting for straight feature continuation and the broader vote clusters in curved structures.

Table 1. Comparison of [3] (HOTV) with our method (CluTV) using scale $\sigma = 2$ and *one* iteration.

data set	HOTV			CluTV		
	m	$\Delta\theta$	$\sigma_{\Delta\theta}$	m	$\Delta\theta$	$\sigma_{\Delta\theta}$
Cross ₉₀						
SNR: ∞	1.00	1.8°	1.0°	0.99	1.7°	0.9°
SNR: 15	0.96	5.9°	6.6°	0.96	4.1°	4.1°
Cross ₇₀						
SNR: ∞	1.00	2.2°	1.5°	0.98	1.7°	1.0°
SNR: 15	0.83	7.6°	7.7°	0.83	5.9°	5.6°
CircLine						
SNR: ∞	0.97	10.5°	12.4°	0.98	6.3°	5.5°
SNR: 15	0.91	14.9°	17.1°	0.84	8.0°	6.1°

Fig. 3. Slices of vMF-ODFs obtained with our method ($\sigma = 1$) from Cross₉₀ with $b = 1000 \text{ s/mm}^2$ and SNR = 15.

3.2. Human brain

Following the examples from [2] we visualise estimated fibre orientations in the *centrum semiovale*, which is known to exhibit crossing fibres of three different orientations. Figure 5 shows how the region with isotropic diffusion tensors in the original data was replaced by von Mises-Fisher ODFs that feature multiple fibre orientations which fit smoothly into the crossing’s neighbourhood.

4. CONCLUSION

We have introduced an alternative method for inferring fibre orientations in crossing regions of diffusion MRI data that shows promising results in both synthetic and real data exper-

Fig. 4. vMF-ODFs obtained with our method ($\sigma = 1$) from synthetic examples with $b = 1000 \text{ s/mm}^2$ and SNR = 15.

Fig. 5. Coronal slice of vMF-ODFs obtained with our method ($\sigma = 1$, *three* iterations) from a crossing region (*centrum semiovale*) in diffusion MRI data ($b = 1000 \text{ s/mm}^2$, 64 gradient directions, 2 mm isotropic voxel size).

iments. Finding a suitable criterion to stop the iterations will be addressed in future work. Further extensions like including existing curvature augmented approaches of tensor voting as well as taking additional information (e.g. from DTI model estimation) into account during initialisation, can lead to even better performance.

5. REFERENCES

- [1] X Wu, M Xie, J Zhou, A Anderson, J Gore, and Z Ding, “Globally optimized fiber tracking and hierarchical clustering - a unified framework,” *Magn. reson. imaging*, vol. 30, no. 4, pp. 485–495, 2012.
- [2] V Prckovska, P Rodrigues, R Duits, BM ter Haar Romeny, and A Vilanova, “Extrapolating fiber crossings from DTI data: can we gain the same information as HARDI?,” 2010.
- [3] T Schultz, “Towards resolving fiber crossings with higher order tensor inpainting,” in *New Developments in the Visualization and Processing of Tensor Fields*, DH Laidlaw and A Vilanova, Eds., Math. and Vis., pp. 253–266. Springer, 2012.
- [4] C-F Westin, S Maier, B Khidhir, P Everett, F Jolesz, and R Kikinis, “Image processing for diffusion tensor magnetic resonance imaging,” in *Med. Image Comput. and Computer-Assist. Interv.-MICCAI’99*. Springer, 1999, pp. 441–452.
- [5] T Kobayashi and N Otsu, “Von Mises-Fisher mean shift for clustering on a hypersphere,” in *Pattern Recognit. (ICPR), 2010 20th Int. Conf. on*. IEEE, 2010, pp. 2130–2133.
- [6] M Moakher and PJ Basser, “Fiber orientation distribution functions and orientation tensors for different material symmetries,” in *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*, pp. 37–71. Springer, 2015.
- [7] Y Cheng, “Mean shift, mode seeking, and clustering,” *Pattern Anal. and Mach. Intell., IEEE Trans on*, vol. 17, no. 8, pp. 790–799, 1995.