Published in final edited form as:

J Neurosci Methods. 2011 September 30; 201(1): 213–219. doi:10.1016/j.jneumeth.2011.07.025.

Spatial and Orientational Heterogeneity in the Statistical Sensitivity of Skeleton-Based Analyses of Diffusion Tensor MR Imaging Data

Richard A Edden^{1,2,3} and Derek K Jones^{3,4}

¹Russell H Morgan Department of Radiology and Radiological Science, The Johns Hopkins University School of Medicine, Baltimore, MD

²FM Kirby Center for Functional Brain Imaging, Kennedy Krieger Institute, Baltimore, MD

³Cardiff University Brain Research Imaging Centre (CUBRIC), School of Psychology, Cardiff University, Cardiff, United Kingdom

⁴Neurosciences and Mental Health Research Institute, Cardiff University, Cardiff, United Kingdom

Keywords

Diffusion tensor imaging; FA; Voxel-based analyses; TBSS; Bias; Nonstationarity; White Matter

1. INTRODUCTION

Diffusion MRI (Jones, 2010b; Le Bihan et al., 1986) is currently the only technique to provide non-invasive assessment of tissue microstructure at the micron-level scale. In the brain, the primary focus of diffusion MRI research has been on white matter since here the axons that form the connections between different brain regions, make coherent boundaries to the random motion of water molecules. Thus from the orientational differences in the mean-squared displacement of water molecules, it is possible to draw inferences about the microstructural make-up of the tissue contained within each voxel (Jones, 2008). By far the most prevalent approach to analyzing multi-direction diffusion-weighted data is to make the assumption of a uni-modal Gaussian displacement profile, and model this with a second-rank tensor (Basser et al., 1994). In this way, diffusion tensor MRI (Basser et al., 1994; Pierpaoli et al., 1996) yields a tensor matrix in each voxel of the image – which is most frequently summarized by computing the trace of the tensor (sum of the three eigenvalues) and the fractional anisotropy (*FA*) – which is the standard deviation of the eigenvalues, divided by their root-mean square (Pierpaoli and Basser, 1996).

Given the exquisite sensitivity of DT-MRI to changes in tissue microstructure, its popularity in clinical research and neuroscientific studies is understandable (Assaf and Pasternak, 2008; Horsfield and Jones, 2002). Since the white matter pathways form connections between cortical regions, the promise of DT-MRI is that it allows researchers to investigate the

© 2011 Elsevier B.V. All rights reserved.

Corresponding Author Details: Richard A. E. Edden, Division of Neuroradiology, The Johns Hopkins University School of Medicine, 600 N Wolfe St, Park 367C, Baltimore, MD 21287, rae2@jhu.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

physical connectivity of the brain (Jones, 2010a, 2008). The problem that the researcher faces, however, is that the metrics derived from DT-MRI are sufficiently variable both on an intra- and inter-subject basis, that making an inference on a single subject's data set is impossible. Thus, the lion's share of DT-MRI studies are based on group comparisons. When there is a clear *a priori* hypothesis about the location of any patient-control differences, or perhaps performance-microstructure correlations, then region-of-interest (ROI) –based approaches are appropriate (Jones and Cercignani, 2010), where the region of interest is either drawn by hand, with an intensity thresholding approach, or via the use of tractography. However, it is often the case that no such *a priori* hypotheses exists or that an investigator wishes to exclude the possibility that regions beyond the hypothesized ROI are implicated. Such concerns are most frequently addressed by conducting global-searches, using voxel-based analysis (VBA).

The overall principle underpinning VBA is to align every participant's dataset to the same space so that a voxel at address (i, j, k) in each subject corresponds to exactly the same anatomical structure. In a DT-MRI based setting, this then allows for a group comparison of the patient *vs.* control group in each and every voxel. It is widely argued that by searching every voxel, and using a fully-automated pipeline, VBA methods are free of operator-induced bias. In practice, however, there are additional steps in the VBA pipeline where the operator may have impact on the output. Even with a high dimensional-warp, meeting the strict requirement that in the aligned space, voxel (i, j, k) contains exactly the same part of the same structure across individuals, is a tall order. Voxel-based morphometry (Ashburner and Friston, 2000) was developed for comparison of white matter and grey matter volume / densities – and operates on a T1-weighted scan. In such an image, each tissue type (white matter, grey matter, CSF) appears fairly homogenous – and thus misalignment of a white matter voxel within a sea of surrounding white matter voxels will have little impact on the outcome of the analysis. In contrast, DT-MRI-based indices like *FA* are incredibly *heterogenous* (Pierpaoli et al., 1996) – and thus, achieving point-to-point correspondence is of paramount importance in order for a voxel-based analysis to make sense, see also (Ashburner and Friston, 2001; Bookstein, 2001). Early adopters of the VBA pipeline addressed this issue by applying a smoothing kernel to each data set prior to statistical comparison – but there is no principled way of choosing the spatial extent of this kernel. Due to the Matched Filter Theorem (Rosenfeld and Kak, 1982), the choice of kernel size will tailor the search space – and therefore, even within the same data set, one can obtain very different results (and therefore draw very different inferences from the data) by changing the size of the kernel alone (Jones et al., 2005). Moreover, as the size of the smoothing kernel increases, the ability to localize the group difference to a particular anatomical location becomes highly questionable (Jones and Cercignani, 2010).

To address these issues of achieving point-to-point correspondence and smoothing-kernel-dependencies, skeleton-projection approaches have been proposed (Smith et al., 2006; Smith et al., 2007). In this approach, an initial alignment of the data sets is performed with a high-dimensional warp. Rather than continue to correct for residual misalignment by smoothing, the skeleton-projection approach first creates a skeletonized version of the average of the data sets. This is achieved using morphological operators to erode the averaged map down to a 'one voxel thick' skeleton. It is assumed that the skeleton at every voxel represents the 'centre' of the structure it passes through. Then in each data set and at every point on the skeleton, searches in the perpendicular plane for a neighbouring voxel that is deemed to be more likely to represent the centre of the object. In the tract-based spatial statistics (TBSS) approach (Smith et al., 2006), it is assumed that the voxel in the perpendicular plane with the highest anisotropy represents the location of the centre of the tract, and that groupwise comparisons should be made on this 'centre' voxel. Thus, this voxel's value is projected

onto the skeleton, and once the projection stage is completed for each subject, a voxel-wise statistical comparison for each voxel on the skeleton is performed.

Given the fully automated pipeline and the need to choose a smoothing kernel obviated, such approaches have been rapidly adopted in the literature (> 100 publications in PubMed with search term = 'TBSS') – and seem to have become the analysis method of choice for group-comparisons of *FA*. It is therefore important to ensure that such approaches are free of bias. Our overall aim, therefore, in this work is to ascertain whether the power to detect a group difference (e.g. patients *vs* controls) or voxel-wise correlation (e.g., performance on a task *vs* *FA*) is uniform across the search space, or whether certain areas of the brain are preferentially biased towards finding an effect. Our first concern in this regard is whether the intrinsic variability of points on the skeleton is uniform? If there is larger variability within a group (e.g., within the control group) at a particular anatomical location, then evidently the effect size will need to be much larger than in areas where the intra-group variability is lower. Failure to recognize this will obviously impact on interpretation of findings.

Our second, and main concern in this study is whether the skeletonization process itself introduces a bias in the analysis such that the power to detect a difference between two groups of subjects, or to detect a correlation between a DT-MRI index and a covariate of interest, will depend on the orientation of the structure with respect to the imaging matrix. This concern was based on a trivial geometrical observation that runs to the heart of the skeletonization procedure. Simply put, if the structure to be skeletonized runs parallel to any of the three principal axes of the image acquisition matrix, then the skeleton will be exactly one voxel-thick. In contrast, if the structure is at an angle θ to a given axis, then its projection onto that axis will be $\cos(\theta)$ and, since the data are discretized (into voxel-sized chunks), the skeleton thickness will increase. For a fibre oriented in the xy-plane at 45 degrees to the x-axis, the skeleton will have maximal thickness of $\sqrt{2}$ voxels. Thus, the more obliquely oriented the fibre to the principal axes of the imaging matrix, the thicker the skeleton becomes. This has two knock-on effects:

1. the number of voxels being compared for group differences at a single point in a tract effectively increases, thereby increasing the chances of finding a differences – and also increasing the risk of a type I error;
2. if an inferential statistic based on spatial extent (cluster-based statistics) is used to make an inference, the larger number of voxels in the vicinity could lead to a larger cluster – and therefore bias the findings.

In the most widely-used version of skeleton-projection methods (Smith et al., 2006), the statistical inference is based on threshold-free cluster enhancement (TFCE, (Smith and Nichols, 2009)) which uses spatial extent of the group difference to infer significance. Hence, the orientational property of the skeletonization effects are of extra concern and we hypothesize that it will impact on the statistical sensitivity of the analysis pipeline.

We begin with a trivial geometric simulation which demonstrates that the skeletonization process is confounded by fibre orientation, and then proceed to demonstrate with real data that this issue extends beyond the theoretical and may therefore have confounded prior studies that have not accounted for the orientational bias.

METHODS

2.1 Simulations

To illustrate the fundamental and basic issue, a fibre was simulated in two dimensions such that it had a Gaussian profile of anisotropy across its diameter, and oriented at an angle θ to the x-axis of the imaging matrix, i.e.

$$FA(x, y) = \exp(-\alpha(x\sin\theta - y\cos\theta)^2)$$

Noise was added to the resulting *FA* map (max value = 0.1 % of maximum *FA*) to ensure that no two adjacent pixels had the same value of *FA*. A total of 72 maps were made (for $0 \leq \theta \leq \pi$, in equal increments of 5°). A circular aperture was then used to ensure that all simulated fibres had the same length. For each map, the skeletonization algorithm included as part of the TBSS processing pipeline was applied, and the number of voxels included in the skeleton was computed.

2.2 In Vivo Data

2.2.1 Data Acquisition—20 normal healthy participants were recruited to the study, comprising 15 males and 5 females (age = 33 $\pm$ 6.2 years). Diffusion weighted data were acquired on a GE 3T Signa HDx MR system, with an 8-channel receive coil, and using a peripherally-gated twice-refocused pulsed-gradient spin echo EPI sequence with 60 diffusion-weighted images (isotropically distributed over space) at $b = 1000 \text{ s/mm}^2$ and 6 images at $b = 0 \text{ s/mm}^2$ at each of 60 slice locations. The field of view was 230 cm, with an acquisition matrix of 96×96 and slice thickness of 2.4 mm. (TE = 87 ms; TR = 20 R-R intervals (triggered with the pulse-oximeter); acceleration factor = 2) (Jones et al., 1999; Jones and Leemans, 2011; Jones et al., 2002).

2.2.2 Data Pre-Processing—Data were corrected for eddy-current induced distortion and subject motion by registering each diffusion-weighted image to the first non-diffusion-weighted image, with subsequent rotation of the encoding vectors (Leemans and Jones, 2009) and correction of signal intensities for volumetric change of the voxels by the determinant of the Jacobian (Jones, 2010c) in *ExploreDTI* (Leemans et al., 2009). Subsequently, a single tensor was fitted in each voxel using non-linear least squares estimation and the fractional anisotropy computed in each voxel.

2.2.3 Assessment of Inter-Group Variance—The first stage in the analysis involved running the 20 *FA* maps through the TBSS pipeline and computing the mean and standard deviation of the vector of 20 *FA* values in each voxel of the skeleton.

2.2.4 Assessment of Orientational Bias—The assessment of the impact of orientation of the fibres with respect to the imaging matrix involved a multi-step procedure:

1. Each of the 20 participant data sets was randomly assigned into 2 equal sized ($n = 10$) groups, ('Group A' and 'Group B').
2. In addition to analyzing the data sets in their 'native' orientations, the effect of orientation on the statistical sensitivity was investigated. Immediately prior to the skeletonization process, the entire set of 20 data volumes was rotated in the axial plane about the z -axis (through-plane axis). First, the data sets were rotated by $+\pi/8$ rads using FLIRT and these rotated data stored as one data set ('clockwise-rotated'); next the original set of 20 data sets were rotated $-\pi/8$ rads – and the rotated data stored as a second data set ('anti-clockwise rotated'). The angle

between the two new data sets was therefore $\pi/4$ rads and, importantly, as both data sets were rotated by the same amount, the amount of interpolation was identical.

3. A uniform offset in FA was then applied to particular structures in the data sets belonging to 'Group B' (see Step 1). Four different magnitudes of offset: 0, +0.10, +0.15 and +0.25 were applied to each of two structures – the right internal capsule (IC) and the right optic radiation (OR), generating a total of 24 datasets, i.e. all permutations of native/clockwise/anticlockwise; IC/OR; 0/0.1/0.15/0.25.
4. Finally, for each of the two data sets (clockwise and anti-clockwise rotated), and for each of the FA offsets (0, +0.10, +0.15, +0.25) applied to the 2 structures, the data were passed through the standard TBSS processing pipeline to compare Group A with Group B.

2. RESULTS

3.1 Simulations

Figure 1 shows the skeletonization of two simulated fibre tracts, one close to vertical (5°) and one close to diagonal (45°). Overlaid on the FA map are the quantized search vectors along which the skeletonization search occurs. Searching along these vectors for voxels with the highest FA produces the skeletons shown in Figure 1 c and d, respectively. As anticipated above from simple geometry, the 'one-voxel thick' skeleton is thicker on the diagonal than on the vertical, leading to the variation in the total number of voxels in the skeleton as a function of fibre orientation, as shown in Figure 1e. The maxima occur at the predicted angles ($\pi/4$ and $3\pi/4$ rads), and the number of voxels is close to $\sqrt{2}$ times more than when the tract is aligned with the principal axes of the image matrix.

3.2 In Vivo Data - Inter-Group Variance

Figure 2 shows the mean and standard deviation of the FA values on the skeleton (thresholded at $FA > 0.2$) derived from the 20 participants, together with histograms of the standard deviation and standard deviation. Immediately striking is the width of the histograms – and of the spatial heterogeneity of the variance in the FA. Centrally located 'core' structures, such as the internal capsule have much lower variance than peripheral structures.

3.3 In Vivo Data – Group Comparison

When Group A and Group B were compared prior to the addition of any FA bias, TBSS did not find any voxels where there was a significant difference between the groups (at a significance threshold of $p = 0.05$).

Figure 3 shows the results of the TBSS analysis of Group A vs Group B after the addition of varying degrees of FA bias to Group B, in the native orientation. Each voxel in the two regions of interest (internal capsule / optic radiation) is colour-coded according to minimum FA-offset at which the two groups were deemed to be significantly different (at $p = 0.05$, using TFCE). This result clearly demonstrates the spatial heterogeneity of statistical sensitivity and shows that the internal capsule is far more sensitive to artificially added bias in the data than the optic radiation. There are two observations worth stressing. First, an FA offset of 0.25 is substantial (given values reported in the literature for group differences), and yet there are several areas in the regions of interest where the processing pipeline deemed the groups not to be significantly different. Second, although the standard deviation of the FA values in the optic radiation is undoubtedly larger than in the internal capsule (Figure 3), the difference in statistical sensitivity between the two regions is larger than one might initially expect.

3.4 In Vivo Data – Effect of Orientation

Figure 4 compares the results of skeletonization and statistical testing between two datasets rotated 45° relative to each other. Consistent with the results of simulated fibre tracts, the skeleton at each location contains more voxels when is oblique to the imaging matrix (shown for example by expansions of the optic radiations and internal capsule. The minimum group difference that is revealed as statistically significant clearly differs with orientation.

3. DISCUSSION

This study makes four main observations:

1. That skeleton thickness is a function of fibre orientation;
2. That the variance in *FA* in the skeleton is non-uniform;
3. That statistical power to reveal group differences is non-uniform; and
4. That statistical power to reveal group differences is explicitly dependent on fibre orientation with respect to the imaging matrix.

The spatial heterogeneity in *FA* variance (point 2) impacts statistical power to detect a group difference (point 3). From the variance, it is possible to compute the minimum detectable difference, δ , between two groups (Zar, 1999) that will be deemed statistically significant, i.e.

$$\delta = \sqrt{\frac{\sigma^2}{n}} (t_{\alpha,v} + t_{\beta,v})$$

where n is the sample number, σ is the standard deviation, and $t_{\alpha,v}$ and $t_{\beta,v}$ are t statistics for power and significance thresholds α and β , and v degrees of freedom. This formula was used to generate the alternative scale for the minimum detectable difference shown in Figure 2. Note that this is for a single-voxel group comparison, and after multiple-comparisons correction this value would be upscaled uniformly across the image – but nevertheless, the spatial distribution of differences in sensitivity would be preserved. This interpretation of Figure 2 makes clear the implication of non-uniform variance for statistical power. .

Rotation of the data sets with respect to the imaging matrix provides direct confirmation of the concerns outlined above, i.e., that the relative orientation of the structure with respect to the imaging matrix affects the thickness of the skeleton and, in turn, this modulates the statistical sensitivity of the group comparisons. For example, the optic radiation shows more yellow voxels (minimum *FA* bias required to produce significant difference = 0.25), when running parallel to the image matrix, and more red voxels (minimum *FA* bias to produce significant difference = 0.15). Similarly, the posterior limb of the internal capsule shows more blue voxels (the minimum *FA* bias for a significant difference = 0.10) when it is obliquely oriented to the image matrix compared with a parallel orientation.

In combination, these results strongly suggest that the differences in statistical power between the internal capsule and the optic radiations seen in Figure 3 are driven by two main effects: differences in variance (as seen in Figure 2); and differences in orientation of fibres relative to the imaging matrix (as demonstrated in Figure 3).

In the absence of *a priori* hypotheses about the spatial location of patient-group differences in tissue microstructure or of a performance-microstructure correlation, there is a definite need for global-search approaches to analyzing diffusion MRI data. Moreover, the use of

such methods is useful to rule out un-hypothesized group differences / performance-microstructure correlations, to boost the specificity of the findings. Thus, the widespread adoption of voxel-based analysis methods for DT-MRI data is completely understandable. However, it is clear that applying methods previously developed for structural T1-weighted data (Ashburner and Friston, 2000) to DT-MRI may not be the optimal approach to the problem (Jones, 2010a; Jones and Cercignani, 2010; Jones et al., 2007; Jones et al., 2005). The advent of skeleton-projection based methods, such as TBSS (Smith et al., 2006), is without question an improvement. However, it is clear from informal discussions with peers that there are some implicit assumptions adopted when interpreting the results from such an analysis. First, there is the assumption that there is equal power to detect a group difference everywhere within the brain. We have clearly shown this not to be the case. Moreover, here we report an additional contribution to the spatial heterogeneity in statistical sensitivity that arises from simple geometry- that is, the orientation of the skeletonized white matter in relation to the imaging matrix alters the thickness of the skeleton. In turn, different skeleton thickness changes the number of voxels that are in a local neighbourhood to lend support to the existence of a cluster of significant voxels.

In this study, we report on two main effects: spatial heterogeneity in the variance of data on the skeleton; and orientational dependence of the statistical sensitivity. An interesting question is whether these observations represent vagaries of biology, or are simply a result of the skeleton-based processing method. With regard to the former observation (heterogeneity of variance), it is likely that there is a strong interaction between biology and processing. Firstly, after the FA volumes have been warped to the template image with a high dimensional warp (Rueckert et al., 1999), it is seen that 'core' structures (e.g. genu, splenium, internal and external capsule, and cerebral peduncles) are well normalized, while more peripheral white matter structures show much more anatomical variability. (This can be verified by looping through warped FA data sets where more 'wiggle' is seen in peripheral structures than central structures, or by computing the standard deviation across individuals). Thus, the 'projection' step has to work harder to find the correct voxel to project onto the skeleton in peripheral regions than in more central regions. Indeed in other investigations (data not shown), we found a positive correlation between variance on the skeleton and the mean projection distance. It is increasingly less likely that the 'correct' voxel is projected onto the skeleton in such scenarios. Thus, the effect of increased biological variability is amplified in these peripheral regions. It is possible that more sophisticated search techniques might provide a more robust result and this is the subject of future research.

With regard to the orientational dependence of the statistical sensitivity, given that in our simulations we did not change biological variability – but just the orientation, we are confident that this variability in statistical sensitivity is attributable purely to the analysis. Although the general problem of non-stationarity is beginning to be addressed by researchers such as Salimi-Khorshidi et al (Salimi-Khorshidi et al., 2011, 2009), it is possible that accounting for the number of contiguous skeleton voxels in the neighbourhood that are available to lend support to a cluster in the TFCE approach might make the approach less orientationally biased. It is worth stressing three important observations: (1) There have already been multiple studies published in the literature where this orientational bias has not been taken into account; (2) adoption of a general correction for non-stationarity does not appear to be widespread as yet; (3) as Salimi-Khorshidi et al. state themselves, there is no such thing as a free lunch, in that correcting for non-stationarity will not simply improve sensitivity everywhere. On the contrary, it may be that applying such a correction to account for these biases will actually *reduce* the significance of the findings to the point that they no longer survive statistical thresholding, while in other places – the sensitivity will

be increased. This has clear implications for the veracity of conclusions drawn from the data.

5. CONCLUSION

In conclusion, this study has highlighted some important contributions to non-stationarity in skeleton-projection based analyses of DT-MRI data, including one that appears to have been overlooked previously - that of the effect of tract shape/orientation with respect to the imaging matrix. Such topological biases in statistical sensitivity need to be corrected for, or at least acknowledged, before drawing inferences from the results of skeleton-projection based analyses. Continued efforts in developing approaches for robust corrections for such biases are encouraged in order to make global voxel-based methods a reliable means for analyzing DT-MRI data.

REFERENCES

- Ashburner J, Friston KJ. Voxel-based morphometry--the methods. *Neuroimage*. 2000; 11:805–21. [PubMed: 10860804]
- Ashburner J, Friston KJ. Why voxel-based morphometry should be used. *Neuroimage*. 2001; 14:1238–43. [PubMed: 11707080]
- Assaf Y, Pasternak O. Diffusion tensor imaging (DTI)-based white matter mapping in brain research: a review. *J Mol Neurosci*. 2008; 34:51–61. [PubMed: 18157658]
- Basser PJ, Mattiello J, LeBihan D. MR diffusion tensor spectroscopy and imaging. *Biophys J*. 1994; 66:259–67. [PubMed: 8130344]
- Bookstein FL. “Voxel-based morphometry” should not be used with imperfectly registered images. *Neuroimage*. 2001; 14:1454–62. [PubMed: 11707101]
- Horsfield MA, Jones DK. Applications of diffusion-weighted and diffusion tensor MRI to white matter diseases - a review. *NMR Biomed*. 2002; 15:570–7. [PubMed: 12489103]
- Jones DK. Challenges and Limitations of quantifying connectivity in the human brain in vivo with diffusion MRI. *Imaging in Medicine*. 2010a; 2:341–55.
- Jones, DK. *Diffusion MRI: Theory, Methods and Applications*. Oxford University Press; 2010b.
- Jones, DK. The signal intensity MUST be modulated by the determinant of the Jacobian when correcting for eddy currents in diffusion MRI; *Proc. ISMRM 18th Annual Meeting*; Stockholm, Sweden. 2010c; p. 1644
- Jones DK. Studying connections in the living human brain with diffusion MRI. *Cortex*. 2008; 44:936–52. [PubMed: 18635164]
- Jones DK, Cercignani M. Twenty-five pitfalls in the analysis of diffusion MRI data. *NMR Biomed*. 2010; 23:803–20. [PubMed: 20886566]
- Jones, DK.; Chitnis, XA.; Job, D.; Khong, PL.; Leung, LT.; Marengo, S.; Smith, SM. What happens when nine different groups analyze the same DT-MRI data set using voxel-based methods?; *Proc. ISMRM 15th Annual Meeting*; Berlin. 2007; p. 74
- Jones DK, Horsfield MA, Simmons A. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magn Reson Med*. 1999; 42:515–25. [PubMed: 10467296]
- Jones DK, Leemans A. Diffusion tensor imaging. *Methods Mol Biol*. 2011; 711:127–44. [PubMed: 21279600]
- Jones DK, Symms MR, Cercignani M, Howard RJ. The effect of filter size on VBM analyses of DT-MRI data. *Neuroimage*. 2005; 26:546–54. [PubMed: 15907311]
- Jones DK, Williams SC, Gasston D, Horsfield MA, Simmons A, Howard R. Isotropic resolution diffusion tensor imaging with whole brain acquisition in a clinically acceptable time. *Hum Brain Mapp*. 2002; 15:216–30. [PubMed: 11835610]

- Le Bihan D, Breton E, Lallemand D, Grenier P, Cabanis E, Laval-Jeantet M. MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*. 1986; 161:401–7. [PubMed: 3763909]
- Leemans, A.; Jeurissen, B.; Sijbers, J.; Jones, DK. ExploreDTI: A Graphical Toolbox for Processing, analysing and visualising diffusion MR data. *Proc. ISMRM 17th International Meeting; Honolulu, Hawaii*. 2009; p. 3536
- Leemans A, Jones DK. The B-matrix must be rotated when correcting for subject motion in DTI data. *Magn Reson Med*. 2009; 61:1336–49. [PubMed: 19319973]
- Pierpaoli C, Basser PJ. Toward a quantitative assessment of diffusion anisotropy. *Magn Reson Med*. 1996; 36:893–906. [PubMed: 8946355]
- Pierpaoli C, Jezzard P, Basser PJ, Barnett A, Di Chiro G. Diffusion tensor MR imaging of the human brain. *Radiology*. 1996; 201:637–48. [PubMed: 8939209]
- Rosenfeld, A.; Kak, AC.; Digital, Picture. Processing 2. Academic Press; Orlando, FL: 1982.
- Rueckert D, Sonoda LI, Hayes C, Hill DL, Leach MO, Hawkes DJ. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Trans Med Imaging*. 1999; 18:712–21. [PubMed: 10534053]
- Salimi-Khorshidi G, Smith SM, Nichols TE. Adjusting the effect of nonstationarity in cluster-based and TFCE inference. *Neuroimage*. 2011; 54:2006–19. [PubMed: 20955803]
- Salimi-Khorshidi G, Smith SM, Nichols TE. Adjusting the neuroimaging statistical inferences for nonstationarity. *Med Image Comput Comput Assist Interv*. 2009; 12:992–9. [PubMed: 20426085]
- Smith SM, Jenkinson M, Johansen-Berg H, Rueckert D, Nichols TE, Mackay CE, Watkins KE, Ciccarelli O, Cader MZ, Matthews PM, Behrens TE. Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *Neuroimage*. 2006; 31:1487–505. [PubMed: 16624579]
- Smith SM, Johansen-Berg H, Jenkinson M, Rueckert D, Nichols TE, Miller KL, Robson MD, Jones DK, Klein JC, Bartsch AJ, Behrens TE. Acquisition and voxelwise analysis of multi-subject diffusion data with tract-based spatial statistics. *Nat Protoc*. 2007; 2:499–503. [PubMed: 17406613]
- Smith SM, Nichols TE. Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. *Neuroimage*. 2009; 44:83–98. [PubMed: 18501637]
- Zar, JH. *Biostatistical Analysis*. Fourth Edition. Prentice Hall; Upper Saddle River, NJ: 1999.

- Statistical sensitivity of skeleton-based diffusion MRI analyses is show to be spatially varying
- Variance of diffusion metrics (e.g. FA) is non-uniform across the skeleton
- Skeletonization is rotationally variant due to interaction between the imaging matrix and tract orientation
- Statistical methods using neighbour information can lead to orientational bias of statistical power

Figure 1. Skeletonization of simulated fibre tracts. Simulated 2D tracts oriented close to vertical (a) and diagonal (b) are shown with quantized tract-perpendicular direction overlaid. When these tracts are skeletonized (c and d), the resulting skeletons have different numbers of pixels. The number of pixels on the skeleton is plotted as a function of fibre orientation in (e).

Figure 2.

Statistics of data on the Skeleton. (a) Skeleton maps of FA mean (left) and standard deviation (right) across the group skeletonized data. (b) Histograms of voxelwise mean (left), standard deviation (right) for the whole skeleton.

Figure 3.

Results of TBSS comparison of Group A and Group B. (a) Skeletonized mean FA data. (b) Expanded regions of (a) showing the right internal capsule (IC) and right optic radiation (OR) and the orientation of the imaging matrix. (c) Within the regions of interest (right internal capsule and optic radiation), each voxel is coloured according to the minimum FA bias (of the 3 tested) at which the two groups were deemed to be significantly different ($p = 0.05$). Blue: minimum FA bias = 0.10; Red: minimum FA bias = 0.15; Yellow: minimum FA bias = 0.25; White: in this voxel, the two groups were not deemed to be significantly different at any of the three FA biases.

Figure 4.

Results of TBSS comparison of Group A and Group B after clock-wise / counter-clockwise rotation of the data sets (see Methods). (a) Mean FA data on the rotated skeletons. (b) Expanded regions of (a) showing the right internal capsule (IC) and right optic radiation (OR) and the orientation of the imaging matrix. These images are rotated back to the native orientation to facilitate comparison of the skeleton thickness. (c) and (d) Results of the statistical comparisons for OR and IC, respectively. Blue: Minimum FA bias = 0.10; Red: Minimum FA bias = 0.15; Yellow: Minimum FA bias = 0.25; White: In this voxel, the two groups were not deemed to be significantly different at any of the three FA biases.