A Framework to Compare Tractography Algorithms Based on Their Performance in Predicting Functional Networks

Fani Deligianni, Chris A. Clark, and Jonathan D. Clayden

Imaging and Biophysics Unit, Institute of Child Health, University College London, London, United Kingdom
{f.deligianni,j.clayden,christopher.clark}@ucl.ac.uk
http://www.ucl.ac.uk/ich/research-ich/imaging-and-biophysics

Abstract. Understanding the link between brain function and structure is of paramount importance in neuroimaging and psychology. In practice, inaccuracies in recovering brain networks may confound neurophysiological factors and reduce the sensitivity in detecting statistically robust links. Hence, reproducibility and inter-subject variability of tractography approaches is currently under extensive investigation. However, a reproducible network is not necessarily more accurate. Here, we build a statistical framework to compare the performance of local and global tractograpy in predicting functional brain networks. We use a model selection framework based on sparse canonical correlation analysis and an appropriate metric to evaluate the similarity between the predicted and the observed functional networks. We demonstrate compelling evidence that global tractography outperforms local tractography in a cohort of healthy adults.

Keywords: structural connectivity, global tractography, prediction.

1 Introduction

Investigating the relationship between functional and structural brain connectivity is vital in understanding and interpreting neurophysiological findings. It is well established that during rest the brain shows spontaneous activity that is highly correlated between multiple brain regions. This activity can be captured with resting-state functional magnetic resonance imaging (rs-fMRI) and it results in reproducible functional networks across subjects. On the other hand, diffusion weighted images (DWI) measure the anisotropic diffusion of water molecules in the brain and carry valuable information regarding interregional brain connections. However, reconstructing the neuronal pathways relies on tractography algorithms that generate networks with poor reproducibility across subjects and studies. Although compelling evidence has emerged that there are strong structural connections between regions that are functionally linked [1], it is challenging to quantify the influence of fiber reconstruction errors.

L. Shen et al. (Eds.): MBIA 2013, LNCS 8159, pp. 211–221, 2013.

Tractography algorithms are tested for their reproducibility within and across subjects. However, reproducibility does not necessarily correlate with accuracy. For example, deterministic tractrography algorithms provide reproducible results but they are inaccurate when fibers cross and diverge. Phantom evaluation is a principled way of examining the performance of tractography algorithms [2]. Phantoms are made based on simplified assumptions that aim to test the performance of tractography under certain scenarios. Nevertheless, the increased fiber complexity of real tissues cannot be captured with the fiber cup phantom, which is an over-simplified scenario of a few crossing fibers.

In this work, we aim to compare the performance of two well known tractography algorithms [3, 4], which are categorised as local and global, respectively. Tractography techniques are categorised as global or local according to whether they account for information along the whole tract or not, respectively. The main advantages of local tracking are the low time complexity and that each tract is independent of the others. However, a major drawback that limits their robustness is that errors in propagating can accumulate along the local steps [3]. This can significantly affect the results and also contributes to distance-related biases. On the other hand, global tracking represents a new approach for identifying brain networks, which involves the simultaneous reconstruction of all the neuronal fibers by finding a solution that best fits the measured diffusion data. This approach has a better ability to resolve ambiguous fiber orientations, since it considers information along the whole neuronal pathway.

A phantom evaluation demonstrated that global tracking is superior to local tractography approaches [4, 5]. However, further experiments based on tracing results in macaques and a cohort of healthy subjects showed only a small improvement in the results [5]. Here, we suggest a systematic way to examine how different tractography approaches affect the observed relationship between structure and function. To our knowledge, this is the first systematic approach to examine whether a predictive framework based on resting-state fMRI is sensitive to differences in tractography.

To this end, we have developed a predictive framework based on sparse canonical correlation analysis (SCCA) [6]. SCCA is a special case of sparse reduced rank regression [7]. Firstly, we characterise functional connectivity as the inverse covariance matrix based on a shrinkage approach that guarantees well-defined, symmetric positive definite (SPD) matrices [8]. Subsequently, we introduce a model selection framework based on cross-validation to quantify the out-of-sample error related to each tractography approach. Finally, the distance between the predicted functional networks and the 'ground truth' rs-fMRI brain networks is estimated based on an intrinsic metric suitable to quantify differences in covariance matrices, independently of their parameterisation, ie. covariance versus the inverse of covariance [9–11].

This work builds upon previous inference approaches to investigate influences between structural connections and functional links [11, 12]. Deligianni et al. presented a framework based on principal component analysis (PCA) and canonical correlation analysis (CCA). Whereas this approach achieves dimensionality

reduction, SCCA is based on a lasso penalty that aims to select the most relevant connections resulting in simultaneous dimensionality reduction and model selection. Model selection based on an $l1$ penalty is also presented in [11]. However, [11] relies on structural networks to regularise the estimation of functional connectivity based on the assumption that a functional connection between two regions only exists when there is a direct structural connection. An advantage of our approach is that it combines a multivariate technique, CCA, with model selection. Therefore, the predicted functional connectivity is well conditioned and it does not necessarily requires further regularisation to be constrained into symmetric positive definite space of matrices.

We apply the proposed approach in a cohort of healthy subjects with multiple structural and functional scans. We demonstrate that a prediction framework based on resting-state fMRI data can capture systematic differences between global and local tractogaphy. We present detailed quantitative results suggesting that global connectivity outperforms local connectivity in predicting functional networks from structural brain networks. Our results suggest that a prediction framework based on functional data could potentially be applied to compare the performance of tractography algorithms in *in-vivo* human imaging data and highlight specific connections that influence prediction.

2 Methods

2.1 Brain Network Construction

Defining Regions of Interest (ROIs). To define corresponding nodes in both functional and structural networks we used atlas-based regions of interest (ROIs) derived from Freesurfer cortical parcellation of the T1-weighted images [13]. BOLD fluctuations are profound in gray matter (GM), whereas tractography is more reliable in delineating white matter (WM) fibers. Hence, we focus on brain networks defined by cortical GM ROIs that are connected via WM fibers. We propagate the anatomical labels from T1 space to native space for both fMRI and DWI using non-rigid registration [14].

Preprocessing. The first five volumes of rs-fMRI data are removed to avoid T1 effects and preprocessing of the functional data involves motion correction, low pass filtering and spatial smoothing with FSL [15]. To construct corresponding functional networks the fMRI signal is averaged across voxels within each GM ROI. The signal in WM and cerebrospinal fluid (CSF) is also averaged and along with the six motion parameters provided from FSL is linearly regressed out.

We used TractoR for preprocessing of the diffusion weighted images (DWI) [16]. This involves converting DICOM files into a 4D NIfTI, identifying the volume with no diffusion weighting to use as an anatomical reference, creating a mask to identify voxels which are within the brain and correcting the data set for eddy current induced distortions. The last two stages are performed using FSL. Furthermore, the gradient vectors are corrected retrospectively to account for the eddy current correction.

Functional Brain Networks. Assuming that the brain activity patterns are described by a Gaussian multidimensional stationary process, the covariance matrix characterizes fully the statistical dependencies among the underlying signals [17]. Hence, we use the inverse covariance, normalised to unit diagonal to characterise functional connectivity. The inverse covariance is directly related to partial correlation, which provides a measure of connectivity strength between two regions once the influence of the others has been regressed out. This is particularly suited to relate functional connectivity with structural brain connectivity, since the latter describes direct pathways between cortical regions. To produce a well-conditioned, symmetric positive definite, ($\mathcal{S}ym_p^+$), sample covariance matrix we use the shrinkage estimator [8]:

$$\hat{\Sigma_\lambda} = \lambda\hat{\mathbf{T}} + (1-\lambda)\hat{\Sigma} \tag{1}$$

where the sample covariance matrix $\hat{\Sigma_\lambda}$ is estimated as a convex linear combination of the unrestricted sample covariance matrix $\hat{\Sigma}$ and the estimator $\hat{\mathbf{T}}$, which is the identity matrix $\mathbf{I}$. In this case, the optimal regularization parameter $\lambda \in [0, 1]$ is determined analytically based on the Ledoit-Wolf theorem [18]. This approach provides a systematic way to regularise the sample covariance matrix and it has been shown to greatly enhance inference of gene association networks [19], where the number of variables n is much greater than the number of observations p. Therefore, it is suited to examine brain networks where the number of connections grow quadratically with the number of ROIs.

Structural Brain Networks. Two sets of structural networks are derived based on local and global tractography, respectively. The local tractography is implemented in TractoR based on the classic ball and stick model [3].

Global tractography is used to produce a second set of brain networks from the same DWI. Global tractography approaches try to reconstruct neuronal fibers simultaneously by finding line configurations that describe best the measured data. This makes them robust to noise, crossing fibers and imaging artefacts. However, their practical application was hindered by the computational complexity and time requirements. Reisert et al. proposed a time efficient approach that minimises a sum of two energies with respect to the fiber model, M, and the observed data, D. These energies control for the length of fibers (internal energy, E_{int}) and the difference between the model and the observed data (external energy, E_{ext}), respectively [4].

$$E(M) = E_{int}(M) + E_{ext}(M, D) \tag{2}$$

The model consists of small line segments. Each segment is described as a spatial location and orientation. From the arrangement of all segments a predicted MR-signal is computed based on the ball and stick model. The internal energy is based on attraction forces between connected segments that encourages them to stay together and have similar orientation. The external energy is the square difference between the actual signal and the predicted signal after the local mean is subtracted from both model and measured data. The optimisation is done based on the Metropolis-Hasting sampler.

To produce the corresponding structural connectivity matrices we consider only fibers that connect cortical regions via white matter and disregard the rest. We estimate the 'strength' of connection between two cortical regions as the number of fibers that connect the regions divided by the average number of white matter voxels that surround them.

2.2 A Predictive Framework Based on Sparse Canonical Correlation

Sparse Canonical Correlation. Canonical correlation analysis (CCA) is generally applied when one set of predictor variables $\mathbf{X}$ is to be related to another set of predicted variables $\mathbf{Y}$ and observations are available for both groups. Note that CCA is designed to deal with situations where the underlying variables are not statistically independent and, hence, they are inherently inter-correlated. The ultimate goal of CCA is to find two basis vectors (canonical vectors) u, v, one for each variable, so that the projections of $\mathbf{X}$, $\mathbf{Y}$ onto these vectors, respectively are maximally linearly correlated.

In CCA all variables from both sets are included in the fitted canonical vectors. However, for the purpose of studying brain connectivity, we are interested in sparse sets of associated variables, that would allow interpretable links between function and structure to emerge. Hence, we adapt sparse canonical correlation analysis (SCCA) to optimise the CCA criterion, subject to certain constrains [6]:

$$\begin{gathered} \text{maximise}_{\mathrm{u,v}} u^T \mathbf{X}^T \mathbf{Y} v \\ \text{subject to :} \|u\|^2 \leq 1,\ \|v\|^2 \leq 1,\ \|u\|_1 \leq c_1,\ \|v\|_1 \leq c_2 \end{gathered} \tag{3}$$

$\|u\|_1 \leq c_1$ and $\|v\|_1 \leq c_2$ represent the L_1 (or *lasso*) penalty and they result in sparse canonical vectors u, v when the penalties c_1 and c_2, respectively, are chosen appropriately. Note that with u fixed, the criterion in eq. 3 is convex in v , and with v fixed, it is convex in u . Therefore, the objective function of this biconvex criterion increases in each step of an iterative algorithm [6]:

$$\begin{gathered} u \leftarrow \text{argmax}_{\mathrm{u}} u^T \mathbf{X}^T \mathbf{Y} v \ \text{ subject to : } \|u\|^2 \leq 1,\ \|u\|_1 \leq c_1 \\ v \leftarrow \text{argmax}_{\mathrm{v}} u^T \mathbf{X}^T \mathbf{Y} v \ \text{ subject to : } \|v\|^2 \leq 1,\ \|v\|_1 \leq c_1 \end{gathered} \tag{4}$$

A Metric to Compare Covariance Matrices. Correlation and covariance matrices lie on the space of symmetric definite positive matrices $\mathcal{F} = \mathcal{S}ym_p^+$. The standard Euclidean distance on matrices, the Frobenius norm, does not account for the geometry of this space. Thus, this norm is ill-suited to quantify prediction errors. However, $\mathcal{S}ym_p^+$ can be parameterized as a Riemannian manifold using an intrinsic metric [9]:

$$d_{AI}(\mathbf{P}, \mathbf{G})^2 = tr\big(\log \mathbf{G}^{-\frac{1}{2}} \mathbf{P} \mathbf{G}^{-\frac{1}{2}}\big)^2 \tag{5}$$

This metric has beed used successfully to build statistical frameworks of precision matrices $\mathcal{S}ym_p^+$ [11]. d_{AI} is a distance metric, invariant to affine transformations and inversion, appropriate to quantify the distance between covariance matrices from biological data successfully [10].

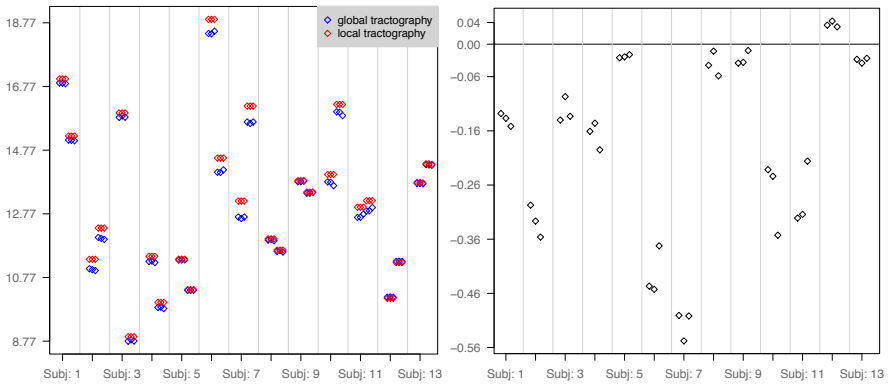

(a) d_{AI} (b) $d_{AI}\{global\} - d_{AI}\{local\}$

Fig. 1. This plot shows how global tractography compares with local tractography for each subject and structural scan. a) Demonstrates that the distance d_{AI} between the predicted and the observed functional connectivity is lower for global tractography than for local tractography. b) Shows the d_{AI} metric for each difference in performance between global tractography and local tractography, per subject and structural scan, according to equation 7.

Model Selection. To evaluate the performance of local and global tractography, we use model selection, which is based on cross-validation. For each subject $s = 1, \ldots, S$, the SCCA model is trained based on the remaining $S - 1$ datasets. The number of components is estimated as the min of the ranks of the predictor and predicted variables in CCA. The penalty values c_1, c_2 are optimised in each cross-validation loop using an approach, which permutes the rows of both the predictor and predicted variables of the SCCA [6]. Subsequently, we use the left-out structural connectivity matrix $\mathbf{A}$ to predict the functional connectivity $\hat{\mathbf{F}}$:

$$\hat{\mathbf{F}} = (u\mathbf{A})^{-1}\mathbf{D}v^{-1} \tag{6}$$

$\mathbf{D}$ is a diagonal matrix of the canonical correlation scores. Finally, the difference between the predicted, $\hat{\mathbf{F}}$, and 'ground truth' functional connectivity matrix $\mathbf{F}$ is estimated as: $d_{AI}(\hat{\mathbf{F}}, \mathbf{F})$.

3 Results

Imaging data was acquired from 13 healthy adults using a Siemens Avanto 1.5 T clinical scanner using a self-shielded gradient set with maximum gradient amplitude of 40 mT m^{-1} and standard 12 channel head coil. Echo-planar DWI were acquired along 60 non-collinear gradient directions at b=1000 s mm^{-2}, with three b=0 images for normalization. A voxel matrix of 96×96 was used and 45 contiguous axial slices acquired, each 2.5mm thick, with a 240 mm FOV,

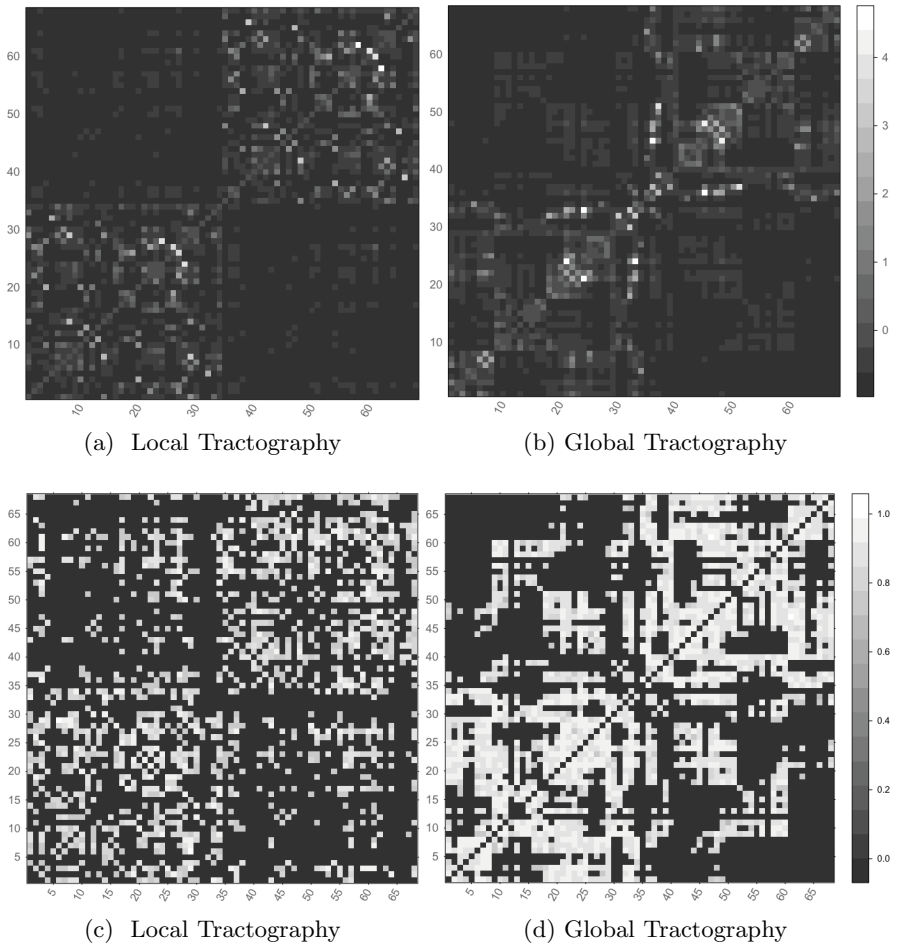

(a) Local Tractography (b) Global Tractography

(c) Local Tractography (d) Global Tractography

Fig. 2. The top row shows the z-scores along all structural scans for (a) local and (b) global tractography. The bottom row depicts the bootstrap results from (a) local and (b) global tractography.

voxel size of 2.5×2.5×2.5 mm and TR/TE=7300/81 ms. T2*-weighted gradient-echo EPI sequence of 125 volumes was also acquired with TR/TE=3320/50 ms, 36 slices with thickness 3.0 mm, voxel size 3.0×3.0×3.0 mm, flip angle 90° , FOV 192×192×108 mm, voxel matrix 64×64×36. High resolution T1-weighted whole-brain structural images were also obtained in all subjects.

Each of the 13 subjects' acquisition includes three structural scans that results in three structural connectivity matrices per subject and two rs-fMRI scans, which produce two functional connectivity matrices. These results in six combinations of structural-functional data per subject: $(S_i : F_j)$. From each leave-one-out cross-validation, we use data only from 12 subjects (a total of 72 samples

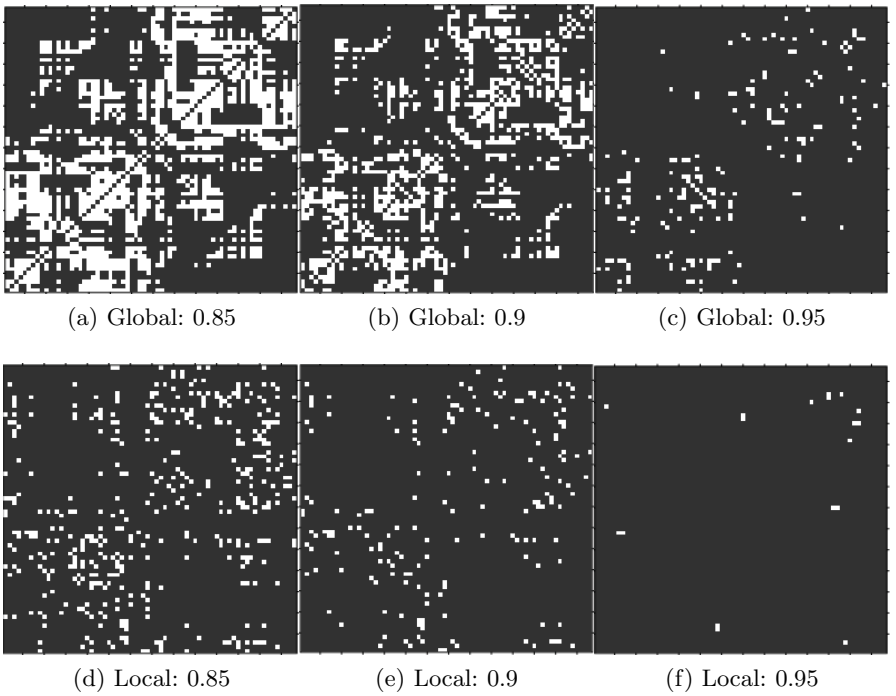

(a) Global: 0.85 (b) Global: 0.9 (c) Global: 0.95

(d) Local: 0.85 (e) Local: 0.9 (f) Local: 0.95

Fig. 3. Binary matrices derived by thresholding the bootstrap results shown in figure 2c-2d. The top and bottom rows show the structural connections that they are selected in more than 85%, 90% and 95% bootstrap iterations with structural networks derived based on global and local tractography, respectively.

across all connections) and we test the prediction performance of each structural scan of the left-out subject ($\hat{F}\{S_i\}$) according to equation 5.

Figure 1 shows how global tractography compares with local tractography for each subject and structural scan. Figure 1a demonstrates that the distance d_{AI} between the predicted and the observed functional connectivity is lower for global tractography than for local tractography. d_{AI} is a distance metric, with smaller values representing better performance. The inter-subject variability is a magnitude of order higher than the difference between the two algorithms. The prediction performance varies considerably across functional scans.

We also plot the difference in d_{AI} between global and local tractography per scan, figure 1b:

$$d_{AI}(\hat{F}\{{S_i}^G\}, F_j) - d_{AI}(\hat{F}\{{S_i}^L\}, F_j) \tag{7}$$

${S_i}^G$ and ${S_i}^L$ corresponds to the structural brain networks derived from global and local tractography, respectively, for the same structural scan S_i. Note that in equation 7 the effect of variability across functional scans is cancelled.

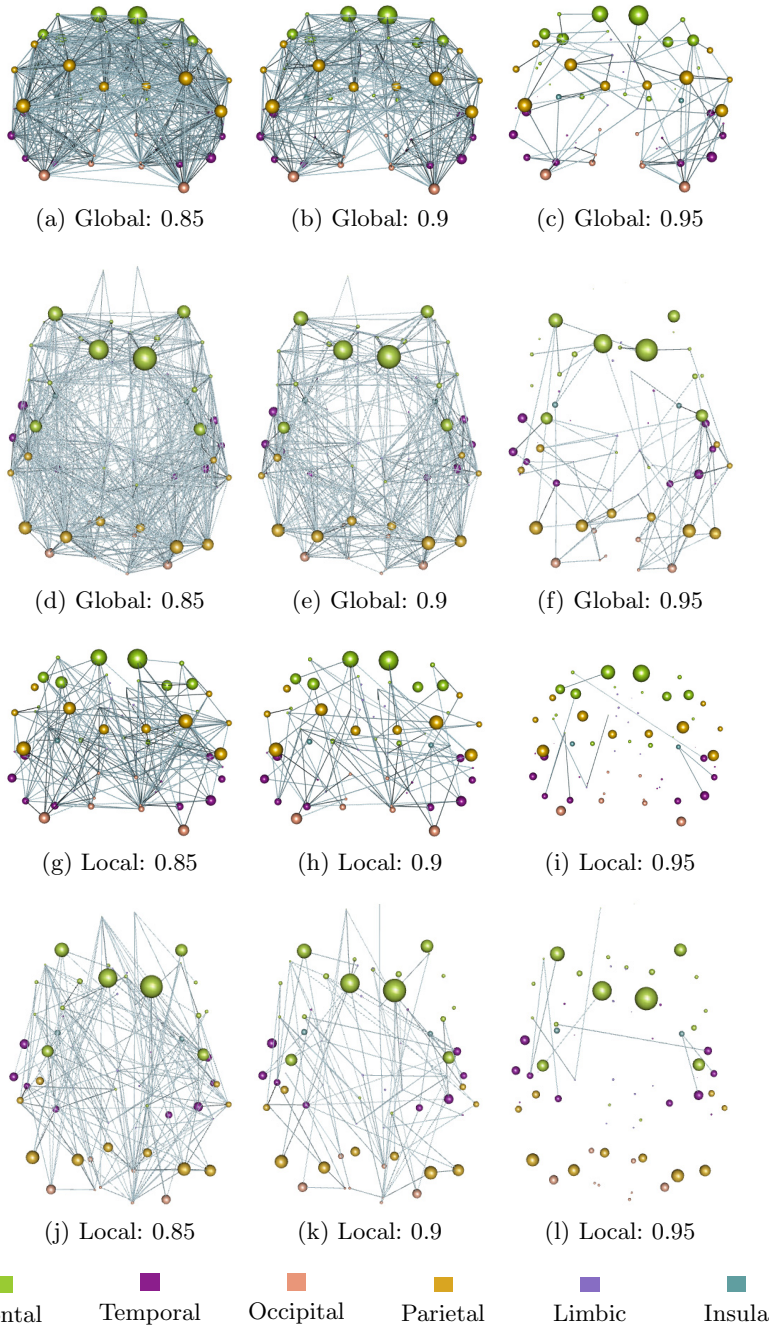

Fig. 4. Thresholded matrices in figure 3 are mapped in brain space. Brain regions are represented with spheres and their radius reflect the relative size of the region.

To further investigate which structural connections are more consistently selected in each tractography method, we use bootstrap and examine the recovered u vector. We resample with replacement all the available datasets and run SCCA 1000 times. The probability of a connection is estimated as the number of times the connection is selected divided by the number of bootstrap repetitions. Figure 2 summarises the results. The top row of the figure 2a-2b, demonstrates the z-scores across the original structural connectivity matrices for local and global tractography, respectively. Global tractography results in structural networks with stronger inter-hemispheric connections. The second row, figure 2c-2d shows the results of the bootstrap. Results indicate that global tractograpy leads to a higher number of connections more consistently selected across the bootstrap iterations. This is apparent when we threshold the bootstrapped connectivity matrices in figure 3.

Figure 3 shows the binary matrices derived by thresholding the bootstrap results shown in figure 2c-2d. These matrices are also mapped in brain space in figure 4. Brain regions are represented with spheres. Their centres are the center of masses of each underlying region and their radius reflect the relative size of each region. The color coding corresponds to different brain lobes: Temporal lobe (dark magenta), frontal lobe (yellow green), parietal lobe (golden road), occipital lobe (dark salmon), insula (cadet blue) and limbic (medium purple).

4 Conclusions

To fully understand how the brain works as a network, the physical connections through the white matter, that mediate information exchange must be accurately characterised. Global tracking may be more stable than local tractography in the presence of noise and imaging artifacts in the data. However, due to the lack of ground truth *in-vivo* tracing data direct evaluation is difficult. Here, we present a robust model selection framework to compare local and global tractography approaches in predicting functional brain networks from structural brain networks and show compelling results that global tractography outperforms local tractography. Structural connectivity only restrains functional connectivity, which is influenced from several other physiological factors. In fact, functional connectivity varies considerably across scans and it does not represent an absolute 'ground truth' for tractography. Nevertheless, we demonstrate evidence that the relationship between structure and function can capture systematic differences in tractography. Future work should aim to compare several tractography algorithms and investigate which brain connections play an important role in the prediction performance. This work is of paramount importance in understanding links between function and structure.

References

1. van den Heuvel, M.P., et al.: Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. Human Brain Mapping 30(10), 3127–3141 (2009)

2. Fillard, P., et al.: Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom. NeuroImage 56(1), 220–234 (2011)
3. Behrens, T., et al.: Characterization and propagation of uncertainty in diffusion-weighted MR imaging. Magnet. Reson. Med. 50, 1077–1088 (2003)
4. Reisert, M., et al.: Global fiber reconstruction becomes practical. NeuroImage 54(2), 955–962 (2011)
5. Li, L., et al.: Quantitative assessment of a framework for creating anatomical brain networks via global tractography. NeuroImage 61(4), 1017–1030 (2012)
6. Witten, D.M., Tibshirani, R.J.: Extensions of sparse canonical correlation analysis with applications to genomic data. Stat. Appl. Genet. Mol. Biol. 8, Article 28 (2009)
7. Vounou, M., Nichols, T.E., Montana, G.: Discovering genetic associations with high-dimensional neuroimaging phenotypes: A sparse reduced-rank regression approach. NeuroImage 53(3), 1147–1159 (2010)
8. Krämer, N., Schäfer, J., Boulesteix, A.-L.: Regularized estimation of large-scale gene association networks using graphical gaussian models. BMC Bioinformatics 10, 384 (2009)
9. Förstner, W., Moonen, B.: A metric for covariance matrices. Qua Vadis Geodesia, 113 (1999)
10. Mitteroecker, P., Bookstein, F.: The ontogenetic trajectory of the phenotypic covariance matrix, with examples from craniofacial shape in rats and humans. Evolution 63(3), 727–737 (2009)
11. Deligianni, F., Varoquaux, G., Thirion, B., Robinson, E., Sharp, D.J., Edwards, A.D., Rueckert, D.: A probabilistic framework to infer brain functional connectivity from anatomical connections. In: Székely, G., Hahn, H.K. (eds.) IPMI 2011. LNCS, vol. 6801, pp. 296–307. Springer, Heidelberg (2011)
12. Deligianni, F., et al.: Inference of functional connectivity from structural brain connectivity. In: ISBI, p. 1113 (2010)
13. Desikan, R.S., et al.: An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. NeuroImage 31(3), 968–980 (2006)
14. Modat, M., et al.: Fast free-form deformation using graphics processing units. Comput. Methods Programs Biomed. 98(3), 278–284 (2010)
15. Smith, S., et al.: Advances in functional and structural MR image analysis and implementation as FSL. NeuroImage 23, 208–219 (2004)
16. Clayden, J., et al.: Tractor: Magnetic resonance imaging and tractography with R. Journal of Statistical Software 44(8), 1–18 (2011)
17. Sporns, O., Tononi, G., Edelman, G.: Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices. Cereb. Cortex 10, 127–141 (2000)
18. Ledoit, O., Wolf, M.: A well-conditioned estimator for large-dimensional covariance matrices. J. Multivar. Anal. 88, 365–411 (2004)
19. Schäfer, J., Strimmer, K.: A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. Stat. Appl. Genet. Mol. Biol. 4, Article 32 (2005)