

Tractometry-based Anomaly Detection for Single-subject White Matter Analysis

Maxime Chamberland¹, Sila Genc¹, Erika P. Raven¹, Greg D. Parker¹, Adam Cunningham², Joanne Doherty^{1,2}, Marianne van den Bree², Chantal M. W. Tax¹, Derek K. Jones¹

¹ Cardiff University Brain Research Imaging Centre (CUBRIC), Cardiff, United Kingdom

² MRC Centre for Neuropsychiatric Genetics and Genomics, Cardiff, United Kingdom

Abstract

There is an urgent need for a paradigm shift from group-wise comparisons to individual diagnosis in diffusion MRI (dMRI) to enable the analysis of rare cases and clinically-heterogeneous groups. Deep autoencoders have shown great potential to detect anomalies in neuroimaging data. We present a framework that operates on the manifold of white matter (WM) pathways to learn normative microstructural features, and discriminate those at genetic risk from controls in a paediatric population.

Keywords: Diffusion MRI, Tractography, Tractometry, Anomaly Detection, Autoencoder

1. Introduction

Considerable effort has gone into designing methods for group comparisons in dMRI (Jones and Cercignani, 2010) (i.e., N patients vs M controls) and as such, single-subject analysis frameworks (i.e., 1 patient vs M controls) are currently lacking. For clinically-heterogeneous groups, normative modeling (Marquand et al., 2016) has been proposed, but often relies on voxel-based methods, and hence is suboptimal for dMRI since WM tracts offer a more suitable manifold. In this work, we investigate individual differences in WM microstructure in children with copy number variants (CNVs) at high genetic risk of neurodevelopmental and psychiatric disorders (Chawner et al., 2019), which are relatively *rare* and challenging to recruit for research imaging studies (Villalón-Reina et al., 2019). We propose the following unsupervised framework for anomaly detection: First, we learn a normative set of features derived from microstructural tract profiles obtained from typically developing (TD) children. Second, we apply the framework to unseen subjects, to determine whether these deviate from controls (based on the hypothesis that deviations will stand out from the normative distribution).

2. Methods

2.1. Data acquisition & preprocessing

Diffusion data from 90 TD (age 8-18 years) and 8 children with a CNV and no apparent WM lesions (age 8-15 years) were acquired on a Siemens 3T Connectom MRI scanner with 14 b0 images, 30 directions at $b = 500, 1200$ s/mm², 60 directions at $b = 2400, 4000, 6000$

s/mm^2 and $2 \times 2 \times 2 \text{mm}^3$ voxels. Data were preprocessed as in Chamberland et al. (2019) and rotationally-invariant spherical harmonic (RISH0, Mirzaalian et al. (2015)) features were derived for each subject using the $b = 6000 \text{ s/mm}^2$ shell. Automated WM tract segmentation was performed using TractSeg (Wasserthal et al., 2018) to obtain 20 bundles of interest (Fig. 1, left). For each bundle, Tractometry (Bells et al., 2011; Cousineau et al., 2017) was performed (sampling at 20 locations, Chamberland et al. (2019)) and the resulting 20 tract profiles were concatenated to form a feature vector ($n = 20 \text{ tracts} \times 20 \text{ locations} = 400 \text{ features}$) for each subject.




Figure 1: Left: RISH0 feature mapped over 20 WM bundles with Tractometry. Right: The proposed unsupervised anomaly detection framework trained on healthy data only.

2.2. Anomaly detection

Our autoencoder architecture consists of five fully connected symmetric layers. The input and output layers have the same number of nodes as the length of the tract features vector. The inner layers consecutively apply a compression ratio of 2 by reducing the number of nodes by half, up to the bottleneck hidden layer (activation: ReLU, epochs: 25, batch size: 24, learning rate: 1.0×10^{-3} , optimiser: Adam, loss: mean squared error).

A validation set ($n = 16$) was generated and held-out by combining the individuals with a CNV ($n = 8$) with a random subset of TD ($n = 8$). The rest of the TD ($n = 82$) data was used to establish a normative distribution (Fig. 1, right). 10% of the TD data was held out for testing during the training phase where the goal was to generate an output (\hat{x}) similar to the input (x) by minimising the reconstruction error and computing the mean absolute error over all features as anomaly score. Age regression and feature normalization were performed on the training set and subsequently applied to the validation set. We repeated this entire process 50 times to assess variations within the TD population and derived a mean anomaly score for each subject. Using the subject labels, we report the mean ROC area under the curve (AUC) across the iterations. In addition, we compared our results with two previously-reported approaches: 1) univariate z-score distribution (Yeatman et al., 2012); and 2) multivariate PCA combined with the Mahalanobis distance (Yeatman et al., 2018; Sarica et al., 2017; Taylor et al., 2020) using the aforementioned bootstrapped approach.

3. Results

The autoencoder approach was better at identifying CNV subjects as outliers (AUC: 0.86 ± 0.06) compared with z-score (AUC: 0.53 ± 0.06) and PCA (AUC: 0.61 ± 0.09). Fig. 2 shows the reconstructed features of a random CNV subject which highlights significant discrepancies along various association pathways, suggesting subject-specific differences in microstructural attributes.

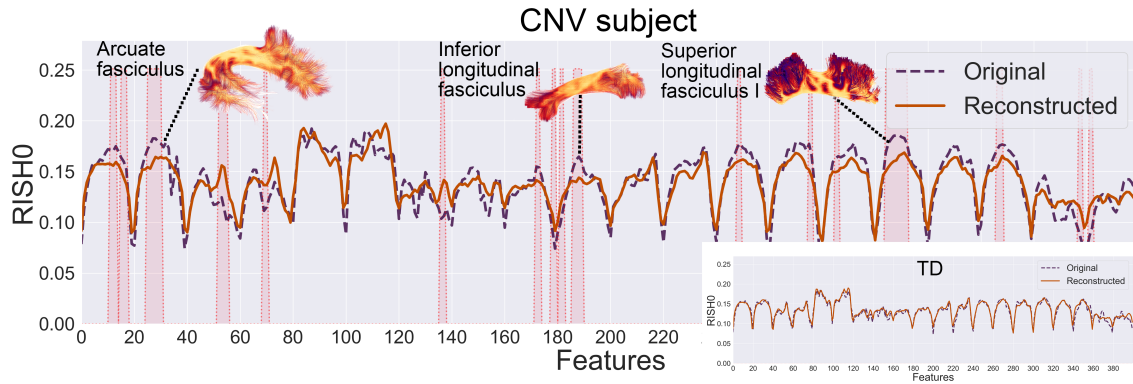


Figure 2: RISH0 features ($20 \text{ sections} \times 20 \text{ tracts}$) of a CNV subject before (x) and after (\hat{x}) reconstruction. Discrepancies were statistically identified along major association pathways using bootstrapped permutations (shaded area, $p < 0.01$). In comparison, features from a representative TD show no significant anomalies (bottom right).

4. Discussion & future work

The framework enabled *subject-* and *tract-specific* characterisation of WM microstructure. By training on only healthy data, our findings revealed that clinical cases (CNVs) can be classified as outliers. Furthermore, inspection of discrepancies in tract profiles allow the identification of clinically relevant differences in microstructural attributes. The framework also outperformed traditional multivariate outlier detection mostly due to its ability to handle high-dimensional data non-linearly. This extends the possibility of using anomaly detection in clinically-heterogeneous groups (Kia et al., 2018; Wolfers et al., 2020) and extremely rare cases (as little as $N=1$), where group comparisons are otherwise impossible. Future work will include additional cohorts to better assess the generalizability of the framework and its application to other pathology. Importantly, anomalies should be linked to clinical findings (Taylor et al., 2020). Our Tractometry-based anomaly detection framework paves the way to progress from the traditional paradigm of group-based comparison of patients against controls, to a personalised medicine approach, and takes us a step closer in transitioning microstructural MRI from the bench to the bedside. The browser-based tool will be made freely-available to the community via Github.

Acknowledgments

The authors thank Prof. Maxime Descoteaux and Jean-Christophe Houde (Sherbrooke Connectivity Imaging Lab) for their helpful discussions and code sharing. This work was supported by a Wellcome Trust Investigator Award (096646/Z/11/Z), a Wellcome Trust Strategic Award (104943/Z/14/Z), an EPSRC equipment grant (EP/M029778/1) and a Sir Henry Wellcome Fellowship (215944/Z/19/Z).

References

- Sonya Bells, Mara Cercignani, Sean Deoni, Yaniv Assaf, Ofer Pasternak, C John Evans, Alexander Leemans, and Derek K Jones. Tractometry–comprehensive multi-modal quantitative assessment of white matter along specific tracts. In *Proc. ISMRM*, volume 678, 2011.
- Maxime Chamberland, Erika P. Raven, Sila Genc, Kate Duffy, Maxime Descoteaux, Greg D. Parker, Chantal M.W. Tax, and Derek K. Jones. Dimensionality reduction of diffusion MRI measures for improved tractometry of the human brain. *NeuroImage*, 200:89–100, oct 2019. doi: 10.1016/j.neuroimage.2019.06.020.
- Samuel JRA Chawner, Michael J Owen, Peter Holmans, F Lucy Raymond, David Skuse, Jeremy Hall, and Marianne BM van den Bree. Genotype–phenotype associations in children with copy number variants associated with high neuropsychiatric risk in the uk (imagine-id): a case-control cohort study. *The Lancet Psychiatry*, 6(6):493–505, 2019.
- Martin Cousineau, Pierre-Marc Jodoin, Eleftherios Garyfallidis, Marc-Alexandre Cote, Felix C Morency, Verena Rozanski, Marilyn GrandMaison, Barry J Bedell, and Maxime Descoteaux. A test-retest study on Parkinson’s PPMI dataset yields statistically significant white matter fascicles. *NeuroImage: Clinical*, 16:222–233, 2017.
- Derek K. Jones and Mara Cercignani. Twenty-five pitfalls in the analysis of diffusion MRI data. *NMR in Biomedicine*, 23(7):803–820, sep 2010.
- Seyed Mostafa Kia, Christian F Beckmann, and Andre F Marquand. Scalable multi-task gaussian process tensor regression for normative modeling of structured variation in neuroimaging data. *arXiv preprint arXiv:1808.00036*, 2018.
- Andre F. Marquand, Iead Rezek, Jan Buitelaar, and Christian F. Beckmann. Understanding Heterogeneity in Clinical Cohorts Using Normative Models: Beyond Case-Control Studies. *Biological Psychiatry*, 80(7):552–561, oct 2016.
- Hengameh Mirzaalian, Amicie de Pierrefeu, Peter Savadjiev, Ofer Pasternak, Sylvain Bouix, Marek Kubicki, Carl-Fredrik Westin, Martha E Shenton, and Yogesh Rathi. Harmonizing diffusion MRI data across multiple sites and scanners. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 12–19. Springer, 2015.
- Alessia Sarica, Antonio Cerasa, Paola Valentino, Jason Yeatman, Maria Trotta, Stefania Barone, Alfredo Granata, Rita Nisticò, Paolo Perrotta, Franco Pucci, and Aldo Quattrone.

- The corticospinal tract profile in amyotrophic lateral sclerosis. *Human Brain Mapping*, 38(2):727–739, feb 2017.
- Peter Neal Taylor, Nádia Moreira da Silva, Andrew Blamire, Yujiang Wang, and Rob Forsyth. Early deviation from normal structural connectivity: A novel intrinsic severity score for mild tbi. *Neurology*, 2020.
- Julio E Villalón-Reina, Kenia Martínez, Xiaoping Qu, Christopher RK Ching, Talia M Nir, Deydeep Kothapalli, Conor Corbin, Daqiang Sun, Amy Lin, Jennifer K Forsyth, et al. Altered white matter microstructure in 22q11. 2 deletion syndrome: a multisite diffusion tensor imaging study. *Molecular psychiatry*, pages 1–14, 2019.
- Jakob Wasserthal, Peter Neher, and Klaus H. Maier-Hein. TractSeg - Fast and accurate white matter tract segmentation. *NeuroImage*, 183:239–253, dec 2018.
- Thomas Wolfers, Christian F Beckmann, Martine Hoogman, Jan K Buitelaar, Barbara Franke, and Andre F Marquand. Individual differences v. the average patient: mapping the heterogeneity in adhd using normative models. *Psychological Medicine*, 50(2):314–323, 2020.
- Jason D Yeatman, Robert F Dougherty, Nathaniel J Myall, Brian A Wandell, and Heidi M Feldman. Tract profiles of white matter properties: automating fiber-tract quantification. *PloS one*, 7(11):e49790, 2012.
- Jason D. Yeatman, Adam Richie-Halford, Josh K. Smith, Anisha Keshavan, and Ariel Rokem. A browser-based tool for visualization and analysis of diffusion MRI data. *Nature Communications*, 9(1), dec 2018.