





## RESEARCH ARTICLE

# Deep multimodal saliency parcellation of cerebellar pathways: Linking microstructure and individual function through explainable multitask learning

Ari Tchetchenian<sup>1</sup>  | Leo Zekelman<sup>2,3</sup> | Yuqian Chen<sup>2</sup> | Jarrett Rushmore<sup>4,5,6,7</sup> | Fan Zhang<sup>8</sup>  | Edward H. Yeterian<sup>9</sup> | Nikos Makris<sup>4,5,6,10</sup> | Yogesh Rathi<sup>2,10</sup>  | Erik Meijering<sup>1</sup> | Yang Song<sup>1</sup> | Lauren J. O'Donnell<sup>2</sup> 

<sup>1</sup>Biomedical Image Computing Group, School of Computer Science and Engineering, University of New South Wales (UNSW), Sydney, New South Wales, Australia

<sup>2</sup>Department of Radiology, Brigham and Women's Hospital, Harvard Medical School, Boston, Massachusetts, USA

<sup>3</sup>Harvard University, Cambridge, Massachusetts, USA

<sup>4</sup>Department of Psychiatry, Massachusetts General Hospital, Harvard Medical School, Boston, Massachusetts, USA

<sup>5</sup>Department of Neurology, Massachusetts General Hospital, Harvard Medical School, Boston, Massachusetts, USA

<sup>6</sup>Department of Radiology, Massachusetts General Hospital, Harvard Medical School, Boston, Massachusetts, USA

<sup>7</sup>Department of Anatomy and Neurobiology, Boston University School of Medicine, Boston, Massachusetts, USA

<sup>8</sup>School of Information and Communication Engineering, University of Electronic Science and Technology of China, Chengdu, China

<sup>9</sup>Department of Psychology, Colby College, Waterville, Maine, USA

<sup>10</sup>Department of Psychiatry, Brigham and Women's Hospital, Harvard Medical School, Boston, Massachusetts, USA

## Correspondence

Ari Tchetchenian and Yang Song, Biomedical Image Computing Group, School of Computer Science and Engineering, University of New South Wales (UNSW), Sydney, NSW, Australia. Email: [ari.tchetchenian@gmail.com](mailto:ari.tchetchenian@gmail.com) and [yang.song1@unsw.edu.au](mailto:yang.song1@unsw.edu.au)

## Funding information

National Key Research and Development Program of China, Grant/Award Number: 2023YFE0118600; National Natural Science Foundation of China, Grant/Award Number: 62371107; UNSW-USA Networks of Excellence Grant; National Institutes of Health, Grant/Award Numbers: K24MH116366, P41EB015902, R01AG042512, R01MH111917, R01MH112748, R01MH119222, R01MH125860, R01MH132610, R01NS125307, R01NS125781, R21DA042271, 1U54MH091657

## Abstract

Parcellation of human cerebellar pathways is essential for advancing our understanding of the human brain. Existing diffusion magnetic resonance imaging tractography parcellation methods have been successful in defining major cerebellar fibre tracts, while relying solely on fibre tract structure. However, each fibre tract may relay information related to multiple cognitive and motor functions of the cerebellum. Hence, it may be beneficial for parcellation to consider the potential importance of the fibre tracts for individual motor and cognitive functional performance measures. In this work, we propose a multimodal data-driven method for cerebellar pathway parcellation, which incorporates both measures of microstructure and connectivity, and measures of individual functional performance. Our method involves first training a multitask deep network to predict various cognitive and motor measures from a set of fibre tract structural features. The importance of each structural feature for predicting each functional measure is then computed, resulting in a set of structure-function saliency values that are clustered to parcellate cerebellar pathways. We refer to our method as *Deep Multimodal Saliency Parcellation (DeepMSP)*, as it computes the saliency of structural measures for predicting cognitive and motor functional

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2024 The Author(s). *Human Brain Mapping* published by Wiley Periodicals LLC.

performance, with these saliencies being applied to the task of parcellation. Applying DeepMSP to a large-scale dataset from the Human Connectome Project Young Adult study ( $n = 1065$ ), we found that it was feasible to identify multiple cerebellar pathway parcels with unique structure–function saliency patterns that were stable across training folds. We thoroughly experimented with all stages of the DeepMSP pipeline, including network selection, structure–function saliency representation, clustering algorithm, and cluster count. We found that a 1D convolutional neural network architecture and a transformer network architecture both performed comparably for the multitask prediction of endurance, strength, reading decoding, and vocabulary comprehension, with both architectures outperforming a fully connected network architecture. Quantitative experiments demonstrated that a proposed low-dimensional saliency representation with an explicit measure of motor versus cognitive category bias achieved the best parcellation results, while a parcel count of four was most successful according to standard cluster quality metrics. Our results suggested that motor and cognitive saliencies are distributed across the cerebellar white matter pathways. Inspection of the final  $k = 4$  parcellation revealed that the highest-saliency parcel was most salient for the prediction of both motor and cognitive performance scores and included parts of the middle and superior cerebellar peduncles. Our proposed saliency-based parcellation framework, DeepMSP, enables multimodal, data-driven tractography parcellation. Through utilising both structural features and functional performance measures, this parcellation strategy may have the potential to enhance the study of structure–function relationships of the cerebellar pathways.

#### KEYWORDS

cerebellar pathways, deep learning, diffusion MRI, explainable AI, multitask learning, tractography, white matter parcellation

#### Practitioner Points

- We propose a deep learning-based cerebellar pathway parcellation approach that integrates diffusion magnetic resonance imaging tractography with individual cognitive and motor functional performance measures.
- We show that our method can identify unique cerebellar parcels with unique structure–function saliency patterns.
- We publish a four-parcel parcellation and analyse its motor and cognitive saliency patterns, as well as the way in which it subdivides major cerebellar white matter tracts.

## 1 | INTRODUCTION

The cerebellum is increasingly understood as both a motor and cognitive structure (Koziol et al., 2014). Many parcellation approaches exist for the anatomy of the human cerebellum (Carass et al., 2018; Diedrichsen et al., 2009; Makris et al., 2005) and its cortical functional networks that often highlight the cerebellum's role in motor and cognitive performance (Buckner et al., 2011; Nettekoven et al., 2023). However, parcellations of the cerebellar pathways as reconstructed by diffusion magnetic resonance imaging (dMRI) tractography have received less attention. dMRI tractography is the only method that allows in vivo, non-invasive mapping of the human brain's white

matter connections (Basser et al., 2000). Existing tractography-based parcellations define major cerebellar fibre tracts, including the inferior, middle, and superior peduncles (Tang et al., 2018; van Baarsen et al., 2016; Yeh et al., 2018; Zhang, Wu, et al., 2018), which can be subdivided more finely using fibre clustering (Zhang, Wu, et al., 2018) or according to cerebellar cortical terminations of streamlines (Rousseau et al., 2022). These parcellations are powerful for enabling quantification of tractography (Zhang et al., 2022) to study the cerebellar pathways in health and disease (Beez et al., 2022; Filippi et al., 2018; Gupta et al., 2021; Habas & Manto, 2018; Harrison et al., 2021; Mittal et al., 2014; Phillips et al., 2015; Quartarone et al., 2020). As the cerebellum contains 80% of the total number of

neurons in the brain (Azevedo et al., 2009), and models of cerebellar pathways are helpful for the understanding of various conditions such as Parkinson's disease (Haghshomar et al., 2022), Alzheimer's disease (Toniolo et al., 2020), and autism spectrum disorder (Jeong et al., 2014), it can be clinically beneficial to develop and improve in vivo methods of cerebellar pathways parcellation.

However, current cerebellar pathway parcellation methods utilise fibre tract structure, but do not leverage the relationships between fibre tract structure and individual functional performance measures, such as motor or cognitive performance. In the field of brain parcellation, state-of-the-art methods frequently rely on multimodal sources of information, including anatomical, functional, connective, cytoarchitectonic, and/or microstructural neuroimage data to define parcels within the brain's grey or white matter (Arslan et al., 2018; Eickhoff et al., 2015; Eickhoff et al., 2018; Glasser et al., 2016; Zhang et al., 2022). In the cerebellum, functional information is known to be important for parcellating the cerebellar cortex (Nettekoven et al., 2023). Recent work has demonstrated relationships between the structure of cerebellar pathways and individual cognitive performance in health and disease (Chang et al., 2022; Chen et al., 2020; Fritz et al., 2022; Kim et al., 2021; Zekelman et al., 2023). Therefore, in this work we explore a parcellation method that can leverage the relationship between two sources of multimodal information, including individual structural (tractography) and functional (from individual functional performance measure data) information.

By training a network to predict functional performance measures from structural features, deep learning provides a data-driven approach that can encode this relationship between individual cerebellar pathway structure and individual functional performance. In the domain of medical imaging, there have been a variety of successful deep learning architectures. These range from convolutional neural networks [CNN; Sarvamangala & Kulkarni, 2022], to graph neural networks (Zhang et al., 2023) and transformers (Shamshad et al., 2023). Despite being the most recent of these architectures, transformers have already achieved state-of-the-art results across a variety of imaging (Chen et al., 2021; Matsoukas et al., 2021; Usman et al., 2022) and non-imaging (Deng et al., 2022; Peng et al., 2021; Rao et al., 2022) tasks in the medical field. Transformers (Vaswani et al., 2017) utilise an attention mechanism that allows for better handling of long sequences of data by assigning importance to different sections of the input (Lin et al., 2022). This makes transformers appealing for our task of neuroimaging-based individual functional performance prediction, where a large sequence of features (1940 dMRI microstructural and connectivity measures) is input to the model. Additionally, the paradigm of multitask learning (Caruana, 1997) has gained traction as a model-agnostic method for enhancing the performance of deep models by having the model learn to predict multiple tasks simultaneously. Multitask learning enables a model to learn features that are relevant across the multiple tasks, reducing the chance of learning spurious or irrelevant features (Ruder, 2017), and has been shown to improve the performance of deep networks across many neuroimaging tasks (He et al., 2020; Liang et al., 2021; Liu et al., 2019).

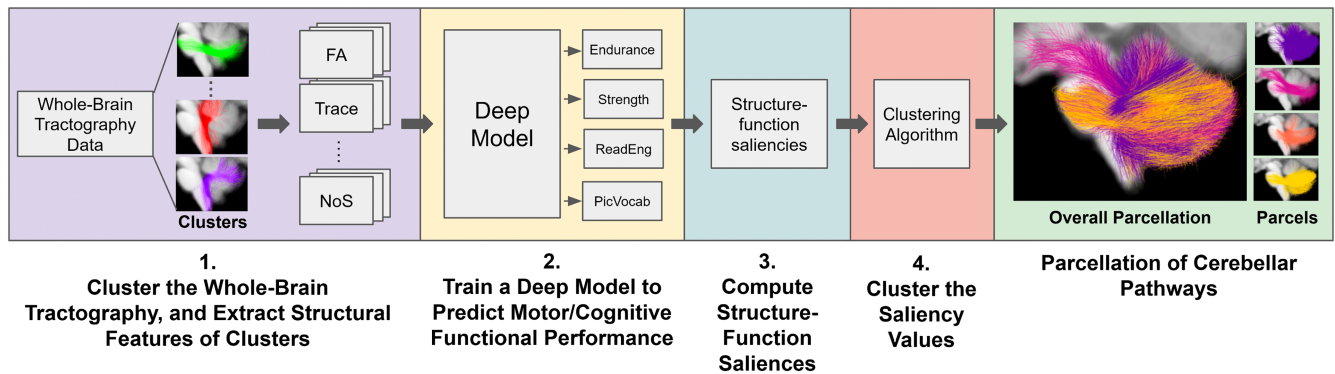
To quantify the structure–function relationships learned by a deep model, we leverage the computer vision concept of *saliency*, which measures the importance of input features for a model's predictions. In a deep learning context, a saliency estimation method assigns a value to each input feature of the network, which represents the importance of that feature to the network's final prediction (Fan et al., 2021). There exist a variety of saliency estimation methods, including gradient-based (Selvaraju et al., 2020; Simonyan et al., 2013; Springenberg et al., 2014), perturbation-based (Fong & Vedaldi, 2017; Ribeiro et al., 2016), and activation-based (Shrikumar et al., 2017; Zhou et al., 2016) methods. In this work, we use the gradient-based 'saliency maps' (Simonyan et al., 2013) approach, where saliency values for a particular target output are the partial derivative of the network's target output with respect to the input. In comparison to competing methods, the saliency maps approach has been shown to be invariant to constant shifts in input (Kindermans et al., 2019) and is sensitive to model parameters and output labels (Adebayo et al., 2018), both of which are desirable properties for saliencies.

Hence, in this work, we propose *Deep Multimodal Saliency Parcellation (DeepMSP)*, a multimodal data-driven approach for human cerebellar pathway parcellation, which leverages the relationship between the structural and functional information of individuals. We train a multitask learning transformer to take a set of high-dimensional dMRI measures describing the microstructure (fractional anisotropy, trace of tractography tensor) and connectivity (number of streamlines [NoS], number of points) of various fibre clusters, and predict performance on multiple individual functional performance measures from the NIH Toolbox. By investigating the resulting trained network, we can estimate the saliency of each dMRI feature and the fibre cluster from which it was measured, for the prediction of individual functional performance measures. We refer to this measure of importance as the *structure–function saliency*, as it encodes the importance of a structural measure for functional prediction. By then grouping fibre clusters with similar structure–function saliencies into parcels, we aim to discover meaningful parcellations of the cerebellar pathways. This overall strategy can enable the usage of high-dimensional, multimodal, structural and functional information to enhance the parcellation of white matter connections. Overall, in this work we obtain new insights into the predictive importance of various dMRI measures and the potential relevance of particular parcels within pathways for studying structure–function relationships of the human cerebellar pathways.

## 2 | MATERIALS AND METHODS

### 2.1 | Overview

DeepMSP (Figure 1) includes four major steps. First, input tractography data are processed to extract finely parcellated fibre clusters from which structural features are extracted to describe microstructure and connectivity. These parcels are based only on tractography streamline trajectory information and therefore have no inherent functional



**FIGURE 1** An overview of DeepMSP. ‘FA’ and ‘NoS’ refer to fractional anisotropy and NoS. Endurance, Strength, ReadEng (reading decoding), and PicVocab (vocabulary comprehension) are examples of individual functional performance measures.

relevance, but can serve as a substrate to enable the definition of parcels with potential functional relevance. Second, the proposed multi-task deep learning network is trained to predict individual functional performance measures from the structural features. Third, during inference on testing data, saliencies are measured to describe the impact of each structural feature from each fibre cluster on the prediction of individual functional performance measures. We call these *structure–function saliencies*. Fourth, saliency values are used as input features to a clustering algorithm, which discovers parcels with common saliency patterns. The final result is a parcellation of cerebellar pathways.

## 2.2 | Dataset

We studied NIH Toolbox and dMRI data from the 1065 subjects of the Human Connectome Project minimally preprocessed young adult dataset (HCP-YA) (Glasser et al., 2013). The HCP-YA dataset consisted of 46% male and 54% female subjects, from 22 to 35 years old, where all dMRI scans were preprocessed according to the HCP’s minimal preprocessing pipeline. This pipeline included normalisation of b0 image intensity across runs, removal of various distortions (EPI distortions, eddy-current-induced distortions, subject motion), correction for gradient nonlinearities, registration of dMRI data to structural scans, bringing the data into 1.25 mm structural space, and masking the dMRI data with a brain mask (Glasser et al., 2013). Regarding the NIH Toolbox measures, we studied all age-adjusted motor and cognitive NIH Toolbox measures (Table 1) (Gershon et al., 2013; Weintraub et al., 2013). We studied all measures due to the known importance of the cerebellum for both motor and cognitive processing (Schmahmann, 2019) and the recently reported relationships between cerebellar input and output pathways (peduncles) and individual performance across cognitive domains in many diseases (Chang et al., 2022; Chen et al., 2020; Fritz et al., 2022; Kim et al., 2021). In healthy individuals, our group recently showed that cerebellar pathway microstructure significantly covaries with multiple NIH Toolbox cognitive performance assessments in the HCP-YA dataset (Zekelman et al., 2023).

The HCP-YA dMRI data were processed to perform whole brain tractography in each subject using the two-tensor unscented Kalman filter (UKF) method (Reddy & Rathi, 2016) as implemented in the ukf-tractography package (<https://github.com/pnlbwh/ukftractography>) following published methods (He et al., 2022; Zekelman et al., 2022). In contrast to other tractography methods that fit a model to the diffusion signal independently at each voxel, UKF employs prior information from the previous tracking step to stabilise model fitting. The two-tensor model can reconstruct fibre crossings that are prevalent in the brain and brainstem (Ford et al., 2013; Jeurissen et al., 2013). The first tensor is associated with the tract being traced, while the second tensor models fibres that cross through the tract. For each subject, the whole brain tractogram was divided into fibre clusters with a robust machine-learning approach that consistently extracts white matter connections across datasets, acquisitions, and the human lifespan (Zhang, Wu, et al., 2018) as implemented in the WMA package (<https://github.com/SlicerDMRI/whitematteranalysis>). This fibre clustering approach has been previously demonstrated to be successful for tractography approaches beyond the UKF method used in this article, for instance diffusion tensor and constrained spherical deconvolution tractography (Zhang, Wu, et al., 2018). The approach first involves the registration of the subject’s tractogram to an atlas tractogram (O’Donnell et al., 2012), followed by a spectral embedding to represent each streamline using its similarity to thousands of streamlines in the atlas (O’Donnell & Westin, 2007). Next, spectral embedding clustering is performed to extract subject-specific anatomical fibre clusters according to the atlas. We used an anatomically curated white matter tract atlas, the ORG tractography atlas (Zhang, Wu, et al., 2018) (<http://dmri.slicer.org/atlas/>) provided by SlicerDMRI (Norton et al., 2017; Zhang et al., 2020). The atlas defines 97 cerebellum fibre clusters, which are anatomically categorised into five tract categories, including the inferior cerebellar peduncle (ICP), the middle cerebellar peduncle (MCP), the superior cerebellar peduncle (SCP), the input and Purkinje cell fibres (IP, including streamlines in the medullary white matter core of the cerebellum), and the putative parallel fibres (PF, including intra-cortical streamlines in the cerebellar cortex). We note that these fibre clusters describe a fine parcellation of the cerebellar pathways. In this work, our goal is to go beyond this structural

**TABLE 1** A description for each NIH Toolbox motor and cognitive measure used in this article. Abbreviations for NIH Toolbox measures are in brackets.

Category	NIH Toolbox measure	Description
Motor	Endurance	Measures cardiovascular endurance by the distance a participant walks in 2 minutes on a 50-foot course.
	Locomotion (GaitSpeed)	Assesses walking speed by timing how quickly a participant walks a 4-m distance at their usual pace.
	Dexterity	Evaluates manual dexterity by timing how quickly a participant can insert and remove 9 pegs into a pegboard with their dominant hand.
	Strength	Measures grip strength by having participants squeeze a Jamar Plus Digital dynamometer, which provides a digital reading of force in pounds.
Cognitive	Episodic memory (PicSeq)	Measures memory retention through the recall of a sequence of pictures shown on a computer screen.
	Cognitive flexibility (CardSort)	Assesses the ability to switch between tasks by asking participants to match pictures based on varying dimensions like shape and colour.
	Inhibition (Flanker)	Measures attention and inhibitory control by requiring focus on a central stimulus while ignoring surrounding stimuli.
	Reading decoding (ReadEng)	Evaluates reading skills in English or Spanish by asking participants to read and pronounce letters and words accurately.
	Vocabulary comprehension (PicVocab)	Tests understanding of vocabulary by presenting a spoken word and asking participants to select the corresponding picture.
	Processing speed (ProcSpeed)	Measures the speed at which a participant can determine if two side-by-side pictures are the same or different.
	Working memory (ListSort)	Assesses the ability to remember and sequence visually- and orally-presented items, such as foods and animals, in specific conditions.

parcellation to investigate parcellations that have the potential to be more meaningful for studying structure–function relationships of the cerebellar pathways. We used these input fibre clusters as a fine parcellation to enable the initial measurement of quantitative dMRI features in all subjects.

For a given subject, from the fibre clusters extracted using the previously described spectral embedding approach, we further

extracted 20 quantitative dMRI structural features from the subset of fibre clusters belonging to the cerebellum. These features included microstructural and connectional geometric information measured within each fibre cluster. The microstructural features computed from the multi-tensor model included the minimum, maximum, median, mean, and variance of the fractional anisotropy of tensors 1 (FA1) and 2 (FA2), and the minimum, maximum, median, and mean trace of tensors 1 (Trace1) and 2 (Trace2). FA is a measure of the amount of directionality of the diffusion, while trace is a mathematical property derived from a diffusion tensor that measures the total amount of diffusion in a voxel (O'Donnell & Westin, 2011). Rather than focusing on the traditional mean value (Zhang et al., 2022), employing multiple summary statistics to describe tissue microstructure in this way can benefit machine-learning applications (He et al., 2022; Zhang, Savadjiev, et al., 2018). In addition to these microstructural features, connectional geometric features described the ‘connectivity’ and geometry of each cluster and included the NoS and the number of points (NoP) along streamlines. We note that while the NoS is popularly considered to be a measure of ‘connectivity’ and is used in a large number of studies, it is only indirectly related to structural connectivity ‘strength’ (Zhang et al., 2022). Thus NoS and NoP can also be considered to be measures of the geometry of connections. In total, 20 dMRI features were extracted within each of the 97 fibre clusters, resulting in 1940 input features per subject. Each input feature was standardised by subtracting the training set mean and dividing by the training set standard deviation before being fed into the model.

### 2.3 | Model architecture

In order to perform parcellation, we required a network that was trained to predict individual functional performance measures so that we could investigate the fibre clusters that were most salient for its predictions.

We trained three multitask learning regression neural networks, each designed to predict 11 continuous individual functional performance measures from the NIH Toolbox, from an input of a 1D vector with 1940 dMRI features. A 1D vector representation was chosen, as it was shown to outperform a 2D representation (He et al., 2022). We chose to force the network to predict multiple NIH Toolbox measures simultaneously (i.e., multitask learning), as multitask learning has been demonstrated to improve network performance (He et al., 2020; Liang et al., 2021; Liu et al., 2019) by learning shared feature representations for all tasks that reduces the chance of learning irrelevant features correlated with only a single task (Ruder, 2017).

Our first architecture was a fully connected model comprised of a sequence of fully connected layers. This is the simplest possible deep network architecture where each neuron in one layer is connected to every neuron in the next layer (supplementary Figure S1), and was chosen to establish a baseline. The second architecture was a 1D CNN (1DCNN) (supplementary Figure S2), which was found to perform well in He et al. (2022) for age and sex prediction. A 1DCNN

utilises specialised layers called 1D convolutional layers, which integrate information from adjacent input elements via a learned convolution operation. The final model was a multi-head attention transformer (Vaswani et al., 2017) (Figure 2), which excels in capturing long-range dependencies in the data by focusing attention on different regions of input (Lin et al., 2022). Each model was balanced to contain approximately 20 million parameters, ensuring a fair comparison across architectures by isolating the effectiveness of the architecture design, independent of model capacity.

## 2.4 | Model training and evaluation details

All networks were trained using fivefold cross-validation with a 60/20/20 train/validation/test split of 639/213/213 subjects, where each subject was randomly allocated to a single fold. We utilised a mean squared error loss function, which was optimised over 50 epochs via the AdamW optimiser (Loshchilov & Hutter, 2017), with  $\beta_1 = .9$ ,  $\beta_2 = .999$ , and a weight decay of 0.01. The learning rate was reduced by a factor of 0.1 when the validation loss had not decreased for 15 epochs. All models were implemented and trained using PyTorch on a National Computational Infrastructure (NCI Australia) node containing 16 GB of memory, 12 24-core Intel Xeon Cascade Lake processors, and an Nvidia V100 GPU with 32 GB of memory.

Each model's hyperparameters were optimised via a 50-epoch grid search on a single fold of cross-validation (learning rate [1e-3, 1e-4, 1e-5], batch size [1, 5, 10, 20], and dropout percentage [0, 10%, 50%, 90%]). The set of possible values for each hyperparameter was chosen to cover a broad range of values, while being limited to three values to allow for the completion of the grid search within a reasonable timeframe. Additionally, the upper end of the batch size values (20) was chosen based on GPU memory limits. The hyperparameter values resulting in the lowest validation set loss across the 50 epochs were as follows: the fully connected network had batch size 20, learning rate 1e-3, and no dropout; the 1DCNN had batch size 1, learning rate 1e-3, and no dropout; and the transformer had batch size 10, learning rate 1e-4, and 10% dropout. Models with these

hyperparameters then underwent proper fivefold cross-validation training and evaluation.

The prediction performance of each deep model was evaluated using two metrics: mean absolute error (MAE) and Pearson correlation coefficient ( $r$ ) (Freedman et al., 2007). MAE calculates the average magnitude of errors between predicted ( $y_i$ ) and true ( $x_i$ ) values, with lower values indicating better performance:

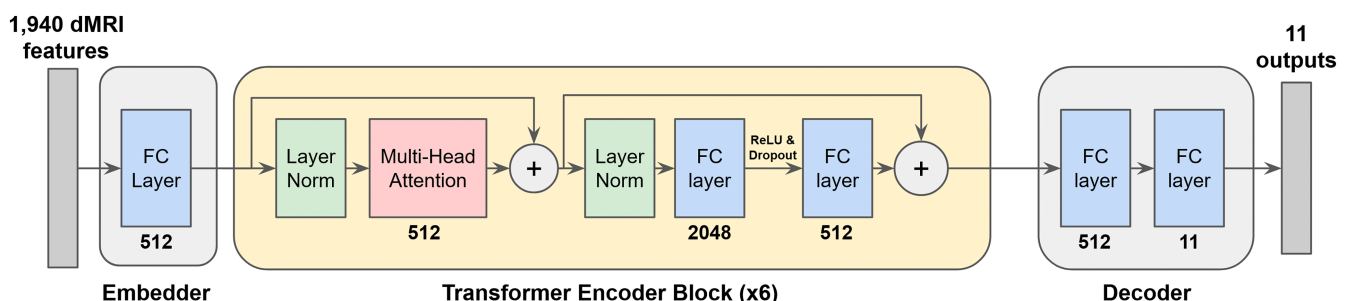
$$MAE = \frac{\sum_{i=1}^n |y_i - x_i|}{n} \quad (1)$$

The Pearson correlation coefficient measures the linear correlation between the predicted ( $y_i$ ) and true ( $x_i$ ) values, where values closer to 1 and  $-1$  indicate strong positive and negative correlations respectively, while a value of 0 indicates no correlation. It is computed as follows, where  $\bar{y}$  and  $\bar{x}$  indicate the mean of the predicted and true values:

$$r = \frac{\sum (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum (x_i - \bar{x})^2 \sum (y_i - \bar{y})^2}} \quad (2)$$

The mean and standard deviation for these two metrics across all cross-validation test sets were reported for each NIH Toolbox measure independently. NIH Toolbox measures for which no model achieved  $r > .1$  after subtracting the standard deviation, were determined to have performed insufficiently, and were discarded from further analysis (Gong et al., 2021). The model that performed the best for the remaining NIH Toolbox measures was then selected for saliency computation and parcellation.

The statistical significance of performance differences for Pearson correlation coefficient values between models was determined via a repeated measures ANOVA, followed by post hoc pairwise Tukey tests. The normality of residuals was verified using the Shapiro-Wilk test, and homogeneity of variances was assessed via Levene's test. Pairwise comparisons were deemed statistically significant if the post hoc Tukey test resulted in a  $p$ -value below .05. The



**FIGURE 2** Architecture diagram of our multi-head attention transformer based on (Vaswani et al., 2017). ‘FC layer’ indicates a fully connected layer, and the number under each layer indicates the output size of that layer. Multi-head attention used eight heads, and the entire transformer encoder block was repeated six times. The output of the penultimate fully connected layer in each encoder block is followed by an ReLU activation function and a dropout layer. Note that layer normalisation is placed before the multi-head attention and fully connected layers rather than after, as this was shown to improve model training time in Xiong et al. (2020).

statistical significance of performance differences for MAE between models was determined as follows. Because the Shapiro–Wilk test indicated non-normality of MAE residuals, for MAE value comparisons we applied the non-parametric Kruskal–Wallis test to determine the statistical significance of MAE value differences, followed by a pairwise post hoc Dunn's test with Bonferroni correction for multiple comparisons.

## 2.5 | Computing structure–function saliencies as feature vectors for parcellation

A set of structure–function saliency values were generated for each NIH Toolbox measure that the network was trained to predict. This resulted in a 1940-element vector (97 fibre clusters times 20 dMRI features per cluster) for each task of each subject, corresponding to the importance of each dMRI feature for the network's output prediction. To enable symmetric parcellation of cerebellar structures (Nettekoven et al., 2023), the saliency values of bilaterally defined fibre clusters (Zhang, Wu, et al., 2018) were averaged, resulting in a 1060 element structure–function saliency vector, corresponding to 53 bilaterally defined fibre clusters. Symmetric averaging of saliency vectors across hemispheres can enhance stability to benefit this initial investigation. A benefit of symmetric tractography parcellation is that it provides corresponding structures across hemispheres, thus enabling the study of the structural and tissue microstructure lateralization of brain pathways (O'Donnell & Westin, 2007; Propper et al., 2010).

These saliency values were computed by taking the absolute value of task-specific saliency maps (Simonyan et al., 2013). Saliency maps assign a value to each element of a deep network's input vector, which is the partial derivative of the network's output with respect to the input at the particular point of the current input vector. Practically, this is achieved by forward propagating the input data, resulting in an output vector. The output for the target class is then backpropagated to the input layer. Each input element is thus assigned a signed saliency value, the magnitude of the value being proportional to the amount of impact the corresponding input element has on the network's prediction for the target class, while the sign of the value indicates whether the element contributes to increasing (positive sign) or decreasing (negative sign) the network's predicted value. In our experimentation, we are explicitly interested in the magnitude of the importance of each feature for predicting NIH Toolbox measures, regardless of the direction of influence on the predicted measures; hence, we disregard the sign of the saliency vectors by taking the absolute value.

When comparing saliency values across multiple NIH Toolbox measures, it is necessary to establish a consistent scale for valid comparison. Saliency values for different NIH Toolbox measures are produced at different scales due to being a gradient value tied to the scale of the NIH Toolbox measure. Hence, comparing the raw saliency values across different NIH Toolbox measures would be unproductive, as a particular feature's saliency value being larger or smaller between different measures could be due to it being more or less salient, but may also be due to the differing scales of the output spaces. We address this issue by applying min-max normalisation to each

generated saliency vector, resulting in all elements falling within the range [0,1], representing the relative saliency of input features.

Computing the saliency values for each NIH Toolbox task, and averaging across all subjects to get a single model-level representation of structure–function saliency, results in a 1060-element vector. This vector can be separated into a 20-element saliency vector for each of the 53 fibre clusters, where each value of the 20-element vector corresponds to one of the 20 dMRI features per fibre cluster. The simplest method of parcellation involves concatenating these 20-element vectors for each NIH Toolbox measure into a single 1D vector with  $N \times 20$  elements, where  $N$  is the number of NIH Toolbox measures that have been deemed as sufficiently predictive ( $r > .1$  after subtracting standard deviation) by the deep model. However, this is a very high-dimensional vector, which is not ideal, as clustering algorithms are known to struggle in high-dimensional spaces (Domingos, 2012). Hence, we experimented with a variety of dimensionality reduction approaches, including both automated approaches like principal component analysis (PCA), which reduces dimensionality by transforming the original elements of the vector into a new set of uncorrelated elements called principal components, and hand-crafted approaches (Table 2).

## 2.6 | Parcellating pathways using saliencies

Parcellation of the cerebellar pathways was achieved by applying clustering algorithms to the set of structure–function saliency vectors described in Section 2.5.

To choose a suitable clustering algorithm, we experimented with both k-means and agglomerative hierarchical clustering to parcellate the cerebellar pathways. These algorithms were chosen as they represent different clustering paradigms (Hastie et al., 2009): partitioning (k-means) and hierarchical (agglomerative) (Rokach, 2010). In k-means (MacQueen, 1967) clusters are formed by initialising  $k$  centroids, assigning clusters to their nearest centroid, and iteratively updating datapoint-centroid assignments until cluster centroids no longer significantly change. Alternatively, agglomerative hierarchical clustering (Hastie et al., 2009) is a bottom-up hierarchical clustering approach that iteratively merges nearest-neighbour datapoints into clusters until only a single cluster remains. The resulting dendrogram is then cut off at the desired height, where the number of branches at that height indicates the number of clusters.

Evaluation of brain parcellation is a known challenge: selection of a parcellation depends on the task at hand (Arslan et al., 2018) and there may not be an optimal number of parcels because of the multi-scale nature of the brain (Eickhoff et al., 2018). Therefore, we experimented with a [2,14] range for parcel counts, choosing values within that range for further downstream analysis based on performance across a variety of metrics. Since we did not have a ground-truth clustering dataset, we employed the following internal cluster quality metrics that evaluated the intrinsic properties of the clustering such as within-cluster dispersion and between-cluster distance: Silhouette coefficient (SC) (Rousseeuw, 1987), Davies–Bouldin index (DBI) (Davies & Bouldin, 1979). These metrics are often applied to assess

**TABLE 2** Descriptions of various saliency vector representations we use in our experiments. The dimensionality column indicates the number of elements in the saliency vector representing a fibre cluster when performing parcellation in Section 2.6.

Name	Description	Dimensionality
Original vector	The original saliency vector, which is a concatenation of the 20-element saliency vectors for each of the $N$ NIH Toolbox measures.	$20 \times N$
PCA-8	Applying PCA with eight components to the original vector.	8
PCA-4	Applying PCA with four components to the original vector.	4
Mean of NIH measures	Taking the average across the 20 dMRI saliencies for each NIH Toolbox measure to get a single average saliency value per NIH Toolbox measure.	4
Mean of categories	As per 'Mean of NIH measures', but further averaging saliencies for all motor and all cognitive measures to get a single average saliency value for motor and cognitive categories.	2
Category displacements	As per 'Mean of Category Saliencies' approach, but adding an element that explicitly measures the amount of bias towards one of the categories, calculated by: $\frac{\text{motor} - \text{cognitive}}{\sqrt{2}}$ This value is the signed perpendicular distance of a given <motor, cognitive> point from the motor = cognitive line in 2D space.	3

the quality of tractography parcellation results (Siless et al., 2013; Vázquez et al., 2020). See supplementary Table S1 for the full descriptions and equations for these metrics. For interpretation purposes, we note that in traditional cluster analysis an  $SC \geq 0.5$  indicates a high-quality clustering result (Mooi & Sarstedt, 2011), while for the challenging problem of tractography parcellation, SC values can range from 0.2 to 0.4 (Siless et al., 2013). DBI has also been used to evaluate tractography parcellation, with values around 0.7 indicating success (lower is better) (Vázquez et al., 2020).

### 3 | RESULTS

#### 3.1 | Choosing a deep network architecture

NIH Toolbox measures with satisfactory prediction performance ( $r > .1$  as in Gong et al. (2021), described in Section 2.4) included two motor (endurance, strength) and two cognitive (reading decoding,

vocabulary comprehension) measures. For each of these measures, both the transformer and 1DCNN outperformed the fully connected network in terms of both Pearson's  $r$  and MAE ( $p < .05$ ) (Table 3), with the exception of the MAE of the strength measure, where there was no statistically significant difference between the fully connected network and the 1DCNN ( $p = .071$ ). Additionally, there was no statistically significant difference in Pearson  $r$  or MAE between the 1DCNN and transformer for each NIH Toolbox measure. However, due to the mean scores for both Pearson  $r$  and MAE of the transformer being slightly better than the 1DCNN for the vast majority of metrics (Table 3), as well as the previously stated lack of statistically significant difference between the fully connected network and 1DCNN for the strength measure, we chose to proceed with the transformer model for subsequent experiments. For a complete report of the performance of all networks for all 11 NIH Toolbox measures, see supplementary Tables S2 and S3, and for a complete list of statistical  $p$ -values, see Supplementary Table S4.

#### 3.2 | Structure–function saliency values and consistency across folds

Examining the saliency values computed from the multitask transformer (Figure 3a), we found that the overall mean saliency for each dMRI feature category was relatively similar (0.24 for FA1, NoS, Trace1, and 0.25 for FA2, NoP, and Trace2). We also found that the 'min' statistic was the most salient regardless of feature category (mean saliency of 0.27), with more typical statistics used in dMRI prediction tasks: 'mean' and 'median', being the lowest 2 salient statistics for all categories (both with mean saliencies of 0.22). We also observe different patterns of saliencies across clusters when averaging cross motor and cognitive NIH Toolbox measures (Figure 3b), with differences in motor and cognitive scores being larger towards the right of Figure 3b. Since we have sorted clusters based on their mean motor and cognitive saliency, this indicates that the majority of the difference in saliency across motor/cognitive categories stems from the most salient clusters.

Averaging the 1060-element vectors across all subjects within a fold, we found very similar standard deviations across folds for each NIH Toolbox measure: 0.14 for strength, and 0.13 for endurance, reading decoding (ReadEng), and vocabulary comprehension (PicVocab). We also found that within each task, the 1060-element saliency vectors were very similar, achieving mean pairwise cosine similarity scores of: 0.96 ( $\pm 0.00$ ) for strength, and 0.93 ( $\pm 0.01$ ) for endurance, ReadEng, and PicVocab.

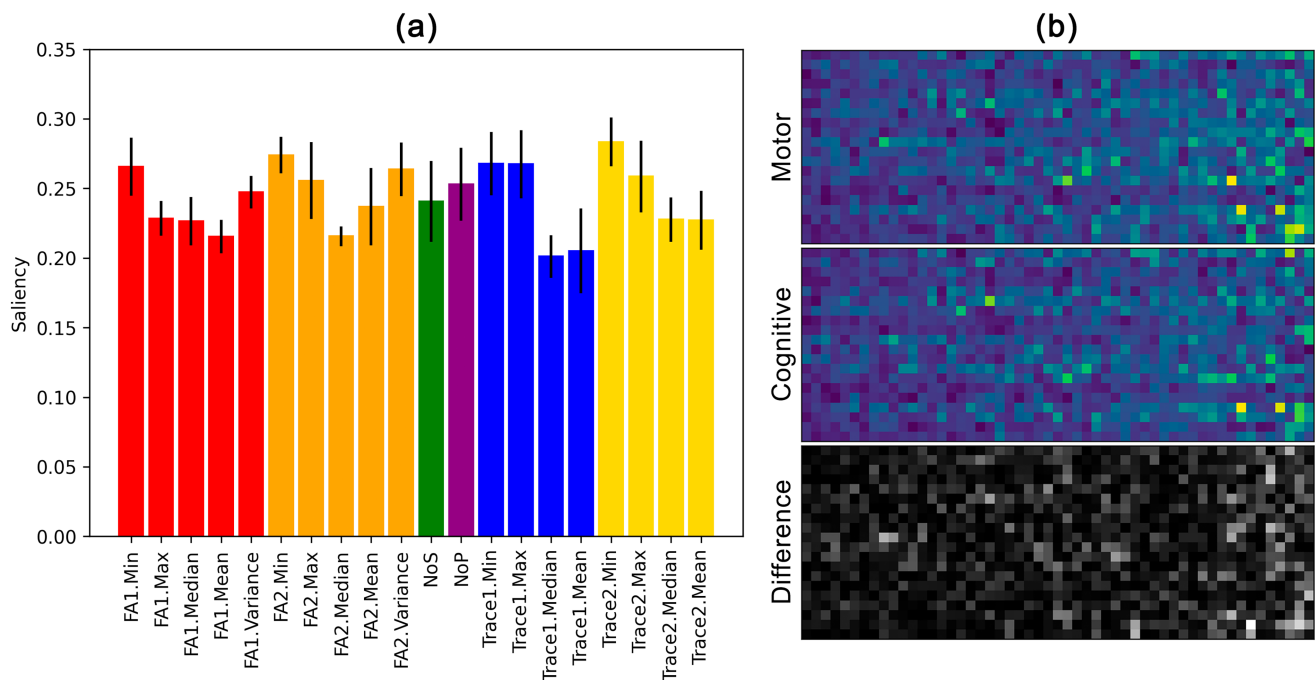
#### 3.3 | Selecting a saliency vector representation

We found (Figure 4) that reducing the dimensionality of the vectors that are being used for clustering results in higher SC values across a variety of cluster counts, for both k-means and agglomerative clustering. We selected the 'Category Displacements (three elements)



**TABLE 3** The Pearson correlation coefficients and mean absolute errors (MAE) of the fully connected, 1DCNN, and transformer models for endurance, strength, reading decoding (ReadEng), and vocabulary comprehension (PicVocab) NIH Toolbox measures. Values are the mean across the five folds of cross-validation, with the standard deviation indicated in brackets. The results with the highest means for each NIH Toolbox measure are bolded.

Category	NIH Toolbox measure	Pearson's <i>r</i>			MAE		
		Fully connected	1DCNN	Transformer	Fully connected	1DCNN	Transformer
Motor	Endurance	0.09 ( $\pm 0.07$ )	0.23 ( $\pm 0.06$ )	<b>0.25 (<math>\pm 0.02</math>)</b>	17.89 ( $\pm 0.41$ )	<b>11.24 (<math>\pm 0.76</math>)</b>	11.34 ( $\pm 0.76$ )
	Strength	0.22 ( $\pm 0.17$ )	<b>0.57 (<math>\pm 0.04</math>)</b>	<b>0.57 (<math>\pm 0.05</math>)</b>	19.88 ( $\pm 1.36$ )	14.51 ( $\pm 1.34$ )	<b>13.51 (<math>\pm 1.25</math>)</b>
Cognitive	ReadEng	0.02 ( $\pm 0.07$ )	0.20 ( $\pm 0.05$ )	<b>0.21 (<math>\pm 0.05</math>)</b>	18.64 ( $\pm 0.89$ )	12.49 ( $\pm 1.01$ )	<b>12.38 (<math>\pm 0.92</math>)</b>
	PicVocab	0.02 ( $\pm 0.06$ )	0.21 ( $\pm 0.07$ )	<b>0.25 (<math>\pm 0.07</math>)</b>	18.70 ( $\pm 0.55$ )	12.01 ( $\pm 0.75$ )	<b>11.81 (<math>\pm 0.94</math>)</b>



**FIGURE 3** Visualising mean saliencies. Depicted in (a) is the mean of each dMRI feature, where colour indicates a common dMRI measure, and error bars indicate the standard deviation over folds of cross validation. Depicted in (b) is the mean motor, cognitive, and absolute difference between motor and cognitive saliencies, where each row corresponds to a dMRI feature, and each column corresponds to a fibre cluster. Results are averaged across all 1065 subjects. Columns are sorted by the mean motor/cognitive score of the column.

representation for further analysis, as it exceeded the desired SC threshold of 0.5 (Mooi & Sarstedt, 2011) for both clustering methods, and achieved a higher median SC value (0.50 k-means, 0.49 agglomerative) compared to the 'Mean of Categories (two elements)' representation (0.46 k-means, 0.44 agglomerative) which also exceeded the threshold.

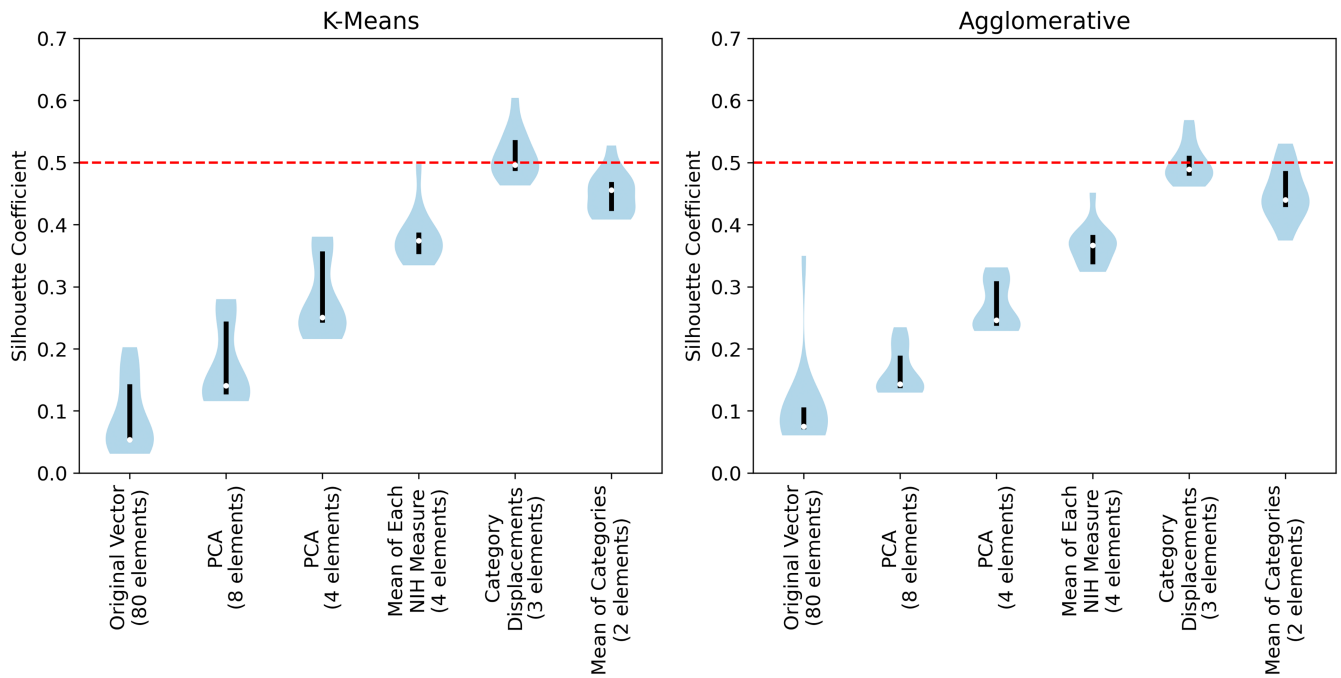
### 3.4 | Selecting the number of parcels and clustering algorithm

Investigating the three-element category displacement feature representation more closely (Figure 5), we found that a parcel count of 4 achieved the best internal cluster quality scores (k-means: 0.60 SC, 0.45 DBI; agglomerative: 0.57 SC, 0.46 DBI). Hence, we selected a


parcel count of 4 for parcellating the cerebellar pathways. In regard to choosing a clustering algorithm, there were very small differences in performance between k-means and agglomerative clustering for a parcel count of 4, with k-means achieving equal or slightly better (5.7% higher SC, 2.17% lower DBI) performance across every metric. Hence, k-means was selected as the clustering algorithm for creating the cerebellar pathway parcellation.

### 3.5 | Parcellation

Cerebellar fibre clusters were parcellated into four parcels (Figure 6a). Parcel 2 was found to have the highest mean saliency (Figure 6b) and was motor dominant (Figure 6c), while parcel 3 was cognitive dominant, and parcels 1 and 4 had very similar motor and cognitive



**FIGURE 4** Violin plots of the silhouette coefficient values for k-means and agglomerative clustering, where clustering was applied to each vector representation across cluster counts in the range [2,14]. Black bars indicate interquartile range, with the white dot indicating the median value. A red horizontal line is set at a silhouette coefficient value of 0.5, which is a standard criterion for a good quality SC result (Mooi & Sarstedt, 2011).

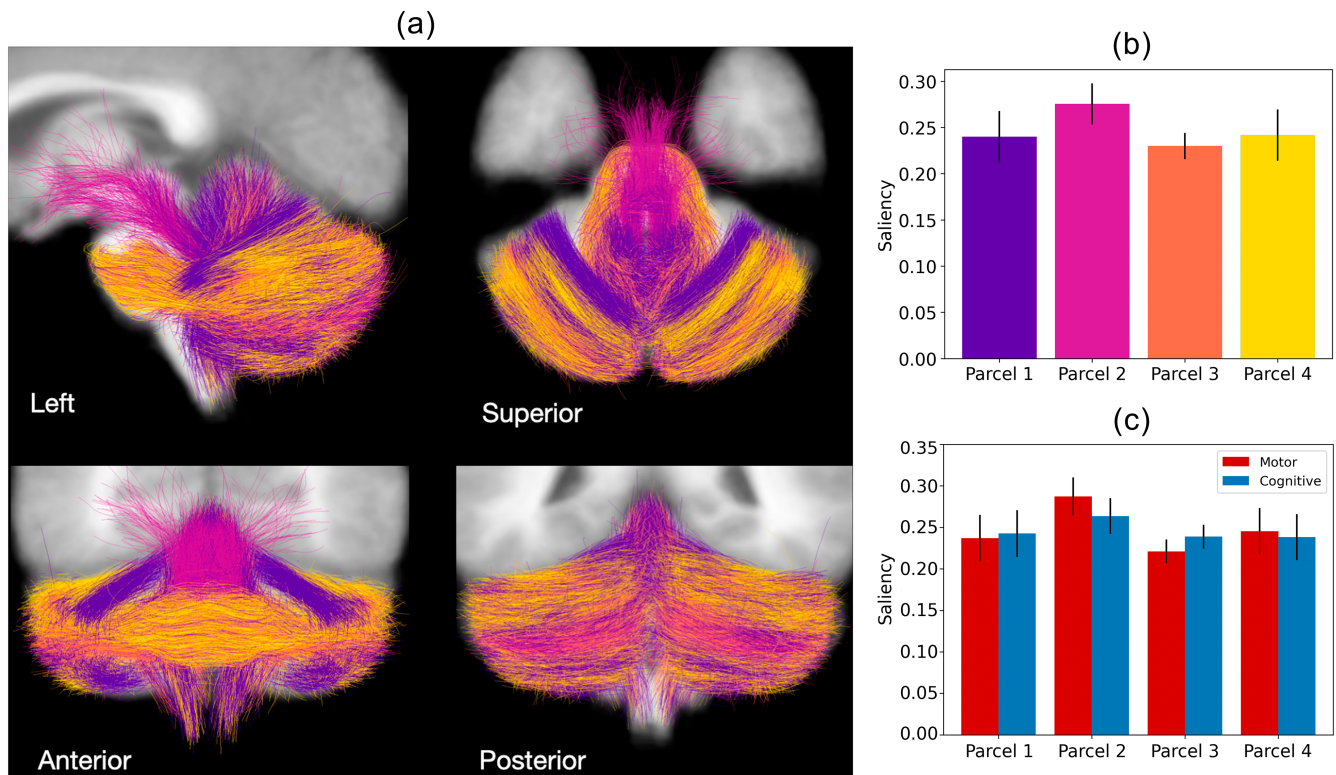


**FIGURE 5** Clustering quality for parcel counts from 2 to 14. (a) Silhouette coefficient (higher is better), (b) Davies-Bouldin index (lower is better). Clustering was performed on ‘Category Displacements’ saliency vectors.

saliencies. See supplementary Figure S3 for a more detailed breakdown of the mean saliencies for NIH Toolbox measures and dMRI features. Each parcel was also visually distinct (Figure 7), had a different distribution of white matter tracts, and connected different cerebellar regions (Table 4, Figure 8). As a complement to the preceding visualisations illustrating the parcels as structural connections, we include a visualisation of the parcels coloured by their relative motor versus cognitive saliency values (Figure 9).

#### 4 | DISCUSSION

We have introduced DeepMSP, a novel deep learning approach for cerebellar pathway parcellation. We thoroughly experimented with all stages of the DeepMSP pipeline, including network selection, structure–function saliency representation, and clustering algorithm and cluster count. We found that the 1DCNN and transformer network architectures performed comparably for the multitask prediction




**FIGURE 6** Visualisation of the cerebellar pathway parcellation. Shown in (a) is the overall parcellation overlaid on a T1-weighted MRI scan, with each parcel being a different colour, (b) is the mean saliency for each parcel, with error bars indicating standard deviation across clusters within each parcel, and the colour of each bar corresponding to the parcel colour in (a), (c) is the mean motor and cognitive saliency of each parcel.

of endurance, strength, reading decoding, and vocabulary comprehension, as well as both significantly outperforming a simpler fully connected network architecture for these same tasks. Investigating the saliencies of this trained transformer, which achieved slightly better mean performance on the evaluation metrics compared to the 1DCNN, we found that the minimum statistic for all microstructural measures (FA1, FA2, Trace1, Trace2) was more salient than commonly used statistics such as mean and variance. This finding is in line with other machine learning work that has found that using minima or maxima statistics for dMRI structural fibre cluster measures achieved higher classification performance than summary statistic measures such as the mean (Zhang, Savadjiev, et al., 2018). It is well known that the mean value of a microstructure measure (such as FA) within a fibre tract can be less informative than studying individual microstructure values along the tract (Colby et al., 2012; O'Donnell et al., 2009; Yeatman et al., 2012). The minimum FA generally occurs at streamline endpoints, that is, the white matter grey matter interface, an anatomical region that is highly predictive of individual cognitive performance (Chen et al., 2024). In contrast, the minimum Trace value may occur in the central core (or 'skeleton') of a fibre tract, another region that is considered highly informative for the study of the white matter (Smith et al., 2006). Our findings suggest the importance of including summary statistics beyond the popular mean value for machine learning applications related to microstructure. In addition to this finding

concerning microstructural measure statistics, we also observed that the majority of differences when comparing motor and cognitive saliencies were for clusters which had a high saliency when averaging over both categories. Finally, we found that a saliency representation that explicitly encoded category bias ('Category Displacements (three elements)') resulted in better clustering performance, with a cluster count of 4 achieving the best internal cluster quality metric scores.

Applying DeepMSP, we identified four cerebellar pathway parcels that exhibited unique structure–function saliency patterns that were stable across training folds. We found that the parcel with the highest mean saliency was motor dominant, with a higher mean saliency for motor tasks compared to cognitive tasks. The opposite was true for the second most salient parcel, while the remaining two parcels had very similar saliencies for motor and cognitive tasks. However, the absolute differences in cognitive and motor saliencies were relatively small for each parcel. In the HCP-YA dataset, it can be observed that motor (strength/endurance) and cognitive (reading decoding/vocabulary comprehension) functional performance measures are correlated (average Pearson's  $r = .19$ ), making structure–function saliencies for motor and cognitive categories potentially difficult to disentangle. Hence, experimentation with categories beyond motor and cognitive function should be explored in future work. This parcellation also subdivided all cerebellar white matter tracts, where, for example, the nine fibre clusters of the MCP were split across all four



**FIGURE 7** Visualisations of each saliency-based cerebellar pathway parcel overlaid on a reference T1-weighted MRI scan.


**TABLE 4** The distribution of white matter tracts for each parcel, and the top cerebellar regions connected by each parcel according to the SUI parcellation (Diedrichsen et al., 2009; Diedrichsen et al., 2011). Abbreviations of white matter tracts are: Inferior cerebellar peduncle (ICP), middle cerebellar peduncle (MCP), superior cerebellar peduncle (SCP), Input and Purkinje (IP), parallel fibres (PF).

Parcel	Number of fibre clusters in each tract					Top 5 cerebellar regions connected
	ICP	MCP	SCP	IP	PF	
1	2	1	1	8	10	Crus II, Crus I, Lobule I_IV, Lobule VI, Dentate
2	0	2	3	0	1	Crus II, Dentate, Crus I, Lobule I_IV, Lobule VIIb
3	0	2	0	3	3	Crus II, Crus I, Lobule VIIb, Lobule I_IV, Lobule IX
4	1	4	0	1	11	Crus I, Crus II, Lobule VI, Lobule VIIb, Vermis VI


parcels. Each parcel also contained clusters from multiple white matter tracts, such as the most salient parcel that included clusters from the MCP, SCP and PF, and also connected multiple cerebellar regions, such as Crus II, Dentate, Crus I, Lobule I\_IV, and Lobule VIIb. With the cerebellum playing a role in various disorders such as Parkinson's disease (Haghshomar et al., 2022), Alzheimer's disease (Toniolo et al., 2020), and autism spectrum disorder (Jeong et al., 2014), findings such as these, which have been produced via DeepMSP, may impact the diagnosis and treatment of such cerebellar-related disorders.

Several detailed observations about the experimental results are discussed below.

We demonstrated that the proposed 1DCNN and transformer models significantly outperformed a simpler fully connected model in terms of predictive ability, and we showed that the saliencies derived from the transformer were consistent across the five test-retest folds. We also showed that different parcels had unique motor and cognitive saliency patterns, as expected based on the known complex functionality of the human cerebellum. Finally, the resulting parcels demonstrated coherent spatial/anatomical patterns, even though no



**FIGURE 8** Visualisations of each white matter tract, where colour indicates parcel. White matter tracts shown are: inferior cerebellar peduncle (ICP), middle cerebellar peduncle (MCP), superior cerebellar peduncle (SCP), Input and Purkinje (IP), and parallel fibres (PF).



**FIGURE 9** Parcel-specific tract visualisation of motor/cognitive distribution. Each fibre cluster for a given white matter tract is assigned the motor and cognitive mean saliency values across all fibre clusters that belong to the same tract and parcel. These motor and cognitive values are then combined into a single score by calculating:  $(\text{motor} - \text{cognitive}) / (\text{motor} + \text{cognitive})$ . These scores are visualised from blue (lower value = more cognitive) to red (higher value = more motor), with a white value indicating a middle neutral value of neither motor nor cognitive dominance. For visualisation purposes, the colours for each tract are normalised independently by dividing by the absolute maximum value across all fibre cluster scores within the tract.

spatial or tract label information was input to the network. Overall, these findings support the potentially meaningful nature of the saliencies employed in this project, and the parcellations resulting from these saliencies. However, one aspect of network selection that remains for future research is to investigate the stability of the parcellations formed by training different architectures for the prediction of the same NIH Toolbox measures. If relatively consistent parcellations are produced across a variety of network architectures, it would further validate the DeepMSP pipeline, while also allowing for smaller and less computationally expensive networks to be used, increasing the accessibility of DeepMSP for those with limited computational resources. This is particularly relevant to the neuroimage analysis field, where simpler models have been shown to perform comparably to more complex state of the art models (Kugelman et al., 2022).

We focused our experimental results on the four-parcel parcellation, which was the most successful according to the SC and DBI cluster quality metrics. These metrics are popularly used to evaluate the quality of tractography parcellation (Chen et al., 2023; Siless et al., 2013; Vázquez et al., 2020). We relied on the use of cluster quality metrics for the selection of the parcel count, though it is well known that there may not be an optimal number of parcels in neuroimaging (Eickhoff et al., 2018). This stems from the complex, multiscale organisation of the brain, where different applications that leverage a parcellation may require different levels of granularity (Eickhoff et al., 2018). Recent work in parcellation of the cerebellar cortex suggests that a hierarchical, multiscale parcellation is effective to allow analysis at multiple granularities, where four parcels are chosen for the coarsest scale (Nettekoven et al., 2023). Therefore, depending on the application of interest, a finer parcellation (more than four parcels) may be useful. For this reason, upon publication we will openly provide our four-parcel parcellation, as well as parcellations of other scales on GitHub (<https://github.com/SlicerDMRI/DeepMSP>). We note that because the proposed parcellations are based on a finer fibre cluster parcellation that can be automatically applied across the lifespan (Zhang, Wu, et al., 2018) using open-source software (<https://github.com/SlicerDMRI/whitematteranalysis>), it is straightforward to apply the proposed parcellations for new studies.

Human brain parcellation is a crucial challenge (Arslan et al., 2018; Bijsterbosch et al., 2020) and is increasingly important for data reduction in large-scale neuroimaging studies (Cetin-Karayumak et al., 2023; Eickhoff et al., 2018). There is a large body of work on cerebellar parcellation including cerebellar cortical parcellation, cerebellar deep nuclei segmentation, and atlas of major cerebellar tracts (Buckner et al., 2011; Carass et al., 2018; Diedrichsen et al., 2009; Makris et al., 2005; Nettekoven et al., 2023). We have aimed to situate our results in the context of this established work by using the well-known SUIT anatomical parcellation (Diedrichsen et al., 2009) and providing information about the cerebellar cortical connectivity of the proposed four-parcel parcellation. In our study, the most salient parcel was primarily connected to Crus II and Crus I, which are considered to form part of key cognitive sociolinguistic and demand networks (Nettekoven et al., 2023), as well as the dentate nucleus, which

participates in both motor and non-motor functions and is crucial in the communication between the cerebellum and the rest of the brain (Steele et al., 2017). Our results generally suggest that motor and cognitive functions are distributed across the cerebellar pathways. This finding is in line with previous research that finds microstructural measures of the cerebellar white matter tracts are associated with motor and cognitive function and dysfunction (Cao et al., 2021; Chang et al., 2022; Chen et al., 2020; Fritz et al., 2022; Hernandez-Castillo et al., 2015; Hernandez-Castillo et al., 2016; Kim et al., 2021; Travis et al., 2015; Wu et al., 2017; Zekelman et al., 2023), as well as recent research identifying both motor and cognitive regions within the cerebellar dentate nucleus (Kulkarni et al., 2023; Palesi et al., 2021).

In addition, our results are supported by anatomical and physiological studies that show that cerebellar neurons, modules and regions receive sensory and motor signals from diverse sources (e.g., Apps et al., 2018; Huang et al., 2013; Ishikawa et al., 2015; Schmahmann, 1996). Our results also suggest that motor and cognitive saliencies are spatially organised across cerebellar cortical and subcortical white matter pathways. Indeed, anatomical studies have found that many sensory, motor, and association regions of the cerebral cortex communicate with the cerebellum via the pontine nuclei, which indicates that the signals carried in the middle cerebellar peduncle are functionally diverse (e.g., Kelly & Strick, 2003; Schmahmann et al., 2004; Schmahmann & Pandya, 1997a; Schmahmann & Pandya, 1997b). Similarly, the inferior cerebellar peduncle carries sensory signals from the spinal cord, motor signals to vestibular nuclei, and cognitive information likely through the principal nuclei of the inferior olivary nuclear complex (e.g., Koziol et al., 2014; Schmahmann, 2010). Overall, our findings are in line with the notion that the functional organisations of the brain extend beyond the cortical surface and are reflected in the composition of white matter (Ghimire et al., 2021). Interrelating the diverse signals in cerebellar white matter with functional localisation in the cerebellar cortex and deep cerebellar nuclei is an important next step that this method may begin to address (Buckner et al., 2011; Makris et al., 2003; Makris et al., 2005; Stoodley & Schmahmann, 2009; Stoodley & Schmahmann, 2018). Looking forward, with the cerebellar pathways playing a role in various disorders such as Parkinson's disease (Haghshomar et al., 2022), Alzheimer's disease (Toniolo et al., 2020), and autism spectrum disorder (Jeong et al., 2014), future work in parcellation of cerebellar pathways may impact the study of such cerebellar-related disorders.

As this is a completely novel approach to cerebellar pathway parcellation, there are a few current limitations and potential directions for future work. In regards to the deep learning and saliency computation stages of DeepMSP, we focused on one high-performing network (multitask transformer) designed to provide saliencies for the prediction of multiple motor and cognitive individual functional performance measures. However, a different network could potentially focus on a different aspect of the high-dimensional input data, producing different saliency patterns. We believe we have somewhat ameliorated this concern by averaging the saliency across many dMRI features and across over 1000 testing subjects to base our final parcellation on

robust input saliency features representative of a large population. However, the investigation of different saliency measures (e.g., LIME (Ribeiro et al., 2016), Guided Backpropagation (Springenberg et al., 2014), Guided Grad-CAM (Selvaraju et al., 2020)) and the effect of different networks is of interest for future research, especially when considering additional individual functional performance measures that may be of interest to inform parcellation for future studies.

We also believe it may be of interest for future work to investigate measures of saliency that incorporate directional or other forms of information beyond purely the magnitude of predictive importance. Although we have successfully shown that saliency varies within and across anatomical tracts, a parcel containing fibre clusters of similar saliency suggests that these fibre clusters have similar predictive importance, and not necessarily that they have a similar relationship with an individual functional performance measure of interest. For example, a fibre cluster where FA is positively correlated with an individual functional performance measure, and another where it is negatively correlated, and yet another with a nonlinear relationship could be considered similarly important (i.e., similar saliencies) by a deep learning model. However, as described in Section 2.5, we have exclusively focused on the magnitude of the saliency values, while disregarding the value's sign. Future work could investigate this aspect more deeply.

Additionally, the proposed parcellation is based on a particular tractography algorithm (which is consistent across the lifespan and robust) that provides a particular set of fibre-specific microstructure measures that are estimated during fibre tracking based on a multi-tensor model (Malcolm et al., 2010; Reddy & Rathi, 2016; Zhang, Wu, et al., 2018). While the proposed parcellation can be applied straightforwardly to tractography computed using other models (via application of our fibre cluster atlas (Zhang, Wu, et al., 2018)), it is not known if microstructure measures from another model would have similar saliency patterns to those observed here. Additionally, we have investigated bilateral parcellation, a popular strategy for cerebellum parcellation (Nettekoven et al., 2023). However, as we did not explore an asymmetric approach to structure–function parcellation, the degree to which the structure–function saliency patterns are asymmetric remains unknown. Hence, future work may benefit from exploring a lateralized approach.

Finally, while the NIH Toolbox assessments utilised in this study are an established and standard approach to measuring cognitive and motor function, these assessments measure limited and broad aspects of cognitive behaviour and motor performance (Hodes et al., 2013; Reuben et al., 2013; Weintraub et al., 2013). We anticipate that more refined and controlled outcome measures may reveal a stronger relationship of sensory, motor, and cognitive functions with fibre clusters in the cerebellum. Additionally, the tractography and fibre-cluster-based quantitative measurements in this study were performed using high-quality HCP-YA data. However, it is important to note that the fine-grained anatomy of the cerebellum requires higher resolution scanning to depict, for example, the fine folds of the cerebellar cortical folia (Serenio et al., 2020). Partial volume effects in the current study were ameliorated by using measurements (e.g., FA or

NoS) performed at the level of a fine fibre cluster parcellation, where each cluster occupied a relatively small region of white or grey matter, and where fibre-specific tissue microstructure measures were computed from a multi-tensor model. However, future work using higher resolution diffusion MRI (Ramos-Llordén et al., 2020) may enable more detailed insights into finer scale cerebellar pathway parcellation. We also note the fibre clusters in the employed tractography atlas are biased towards the longer connections in the posterior lobe of the cerebellum, the neocerebellum. In this way, our investigation may be biased towards phylogenetically newer functions of the cerebellum. Thus, future investigations may continue to investigate across the entire cerebellum.

Looking forward, there are many future research directions for building on DeepMSP. First, the quality of the saliency values used in DeepMSP depends on both the prediction performance of the deep network, and the saliency calculation algorithm. Hence, it could be beneficial to experiment with more architectures, larger datasets or augmentation strategies, and other well-known saliency computation methods such as LIME (Ribeiro et al., 2016), Guided Backpropagation (Springenberg et al., 2014), and Guided Grad-CAM (Selvaraju et al., 2020). Second, our saliencies encoded the magnitude of the predictive importance of a particular fibre cluster; hence, future work could explore directional structure–function relationships.

## 5 | CONCLUSION

In this work, we proposed DeepMSP, a novel saliency-based parcellation framework for multimodal, data-driven tractography parcellation. Our method generates parcellations that aim to reflect the importance of pathway microstructure and connectivity for predicting individual cognitive and motor functional performance measures. Through utilising both structural features and functional performance measures, this parcellation strategy may have the potential to enhance the study of structure–function relationships of the cerebellar pathways.

## ACKNOWLEDGEMENTS

This research was undertaken with the assistance of resources from the National Computational Infrastructure (NCI Australia), an NCRIS enabled capability supported by the Australian Government. This research has been supported by an Australian Government Research Training Program (RTP) Scholarship. Data were provided (in part) by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centres that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. The authors gratefully acknowledge the funding provided by the following National Institutes of Health (NIH) grants: R01MH132610, R01MH125860, R01MH119222, P41EB015902, R01NS125307, R01NS125781, R01MH112748, R01MH111917, R01AG042512, R21DA042271, and K24MH116366. Fan Zhang is supported by the National Key R&D Program of China (No. 2023YFE0118600) and

the National Natural Science Foundation of China (No. 62371107). Additionally, this study was made possible by the support of the UNSW-USA Networks of Excellence Grant. Open access publishing facilitated by University of New South Wales, as part of the Wiley - University of New South Wales agreement via the Council of Australian University Librarians.

## CONFLICT OF INTEREST STATEMENT

The authors declare no conflict of interest.

## DATA AVAILABILITY STATEMENT

The Human Connectome Project minimally preprocessed young adult dataset imaging data and associated NIH Toolbox measures are publicly available at <https://db.humanconnectome.org/>. The ORG tractography atlas is publicly available at <http://dmri.slicer.org/atlasses/> and code to apply the atlas is publicly available at <https://github.com/SlicerDMRI/whitematteranalysis>. All code developed for our experiments, as well our cerebellar parcellations will be publicly available at <https://github.com/SlicerDMRI/DeepMSP>.

## ORCID

Ari Tchetchenian  <https://orcid.org/0000-0001-8682-7516>

Fan Zhang  <https://orcid.org/0000-0002-5032-6039>

Yogesh Rathi  <https://orcid.org/0000-0002-9946-2314>

Lauren J. O'Donnell  <https://orcid.org/0000-0003-0197-7801>

## REFERENCES

- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., & Kim, B. (2018). Sanity checks for saliency maps. *Advances in Neural Information Processing Systems* 31. Retrieved from <https://proceedings.neurips.cc/paper/8160-sanity-checks-for-saliency-maps>
- Apps, R., Hawkes, R., Aoki, S., Bengtsson, F., Brown, A. M., Chen, G., Ebner, T. J., Isope, P., Jörntell, H., Lackey, E. P., Lawrenson, C., Lumb, B., Schonewille, M., Sillitoe, R. V., Spaeth, L., Sugihara, I., Valera, A., Voogd, J., Wylie, D. R., & Ruigrok, T. J. H. (2018). Cerebellar modules and their role as operational cerebellar processing units. *Cerebellum*, 17, 654–682.
- Arslan, S., Ktena, S. I., Makropoulos, A., Robinson, E. C., Rueckert, D., & Parisot, S. (2018). Human brain mapping: A systematic comparison of parcellation methods for the human cerebral cortex. *NeuroImage*, 170, 5–30.
- Azevedo, F. A. C., Carvalho, L. R. B., Grinberg, L. T., Farfel, J. M., Ferretti, R. E. L., Leite, R. E. P., Jacob Filho, W., Lent, R., & Herculano-Houzel, S. (2009). Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *The Journal of Comparative Neurology*, 513, 532–541.
- Basser, P. J., Pajevic, S., Pierpaoli, C., Duda, J., & Aldroubi, A. (2000). In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44(4), 625–632.
- Beez, T., Munoz-Bendix, C., Mijderwijk, H.-J., Remke, M., & Hänggi, D. (2022). Structural damage burden and hypertrophic olivary degeneration in pediatric postoperative cerebellar mutism syndrome. *Neurosurgical Review*, 45, 2757–2765. <https://doi.org/10.1007/s10143-022-01791-7>
- Bijsterbosch, J., Harrison, S. J., Jbabdi, S., Woolrich, M., Beckmann, C., Smith, S., & Duff, E. P. (2020). Challenges and future directions for representations of functional brain organization. *Nature Neuroscience*, 23, 1484–1495.
- Buckner, R. L., Krienen, F. M., Castellanos, A., Diaz, J. C., & Thomas Yeo, B. T. (2011). The organization of the human cerebellum estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106, 2322–2345. <https://doi.org/10.1152/jn.00339.2011>
- Cao, S., Nie, J., Zhang, J., Chen, C., Wang, X., Liu, Y., Mo, Y., Du, B., Hu, Y., Tian, Y., Wei, Q., & Wang, K. (2021). The cerebellum is related to cognitive dysfunction in white matter hyperintensities. *Frontiers in Aging Neuroscience*, 13, 670463.
- Carass, A., Cuzzocreo, J. L., Han, S., Hernandez-Castillo, C. R., Rasser, P. E., Ganz, M., Beliveau, V., Dolz, J., Ben Ayed, I., Desrosiers, C., Thyreau, B., Romero, J. E., Coupé, P., Manjón, J. V., Fonov, V. S., Collins, D. L., Ying, S. H., Onyike, C. U., Crocetti, D., ... Prince, J. L. (2018). Comparing fully automated state-of-the-art cerebellum parcellation from magnetic resonance images. *NeuroImage*, 183, 150–172.
- Caruana, R. (1997). Multitask learning. *Machine Learning*, 28, 41–75.
- Cetin-Karayumak, S., Zhang, F., Billah, T., Zekelman, L., Makris, N., Pieper, S., O'Donnell, L. J., & Rathi, Y. (2023). Harmonized diffusion MRI data and white matter measures from the adolescent brain cognitive development study. *bioRxiv*. <https://doi.org/10.1101/2023.04.04.535587>
- Chang, X., Jia, X., Wang, Y., & Dong, D. (2022). Alterations of cerebellar white matter integrity and associations with cognitive impairments in schizophrenia. *Frontiers in Psychiatry*, 13, 993866.
- Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., Lu, L., Yuille, A. L., & Zhou, Y. (2021). TransUNet: Transformers make strong encoders for medical image segmentation. *arXiv [cs.CV]*. *arXiv*. Retrieved from <http://arxiv.org/abs/2102.04306>
- Chen, Y., Landin-Romero, R., Kumfor, F., Irish, M., Hodges, J. R., & Piguet, O. (2020). Cerebellar structural connectivity and contributions to cognition in frontotemporal dementias. *Cortex*, 129, 57–67.
- Chen, Y., Zekelman, L. R., Zhang, C., Xue, T., Song, Y., Makris, N., Rathi, Y., Golby, A. J., Cai, W., Zhang, F., & O'Donnell, L. J. (2024). TractGeoNet: A geometric deep learning framework for pointwise analysis of tract microstructure to predict language assessment performance. *Medical Image Analysis*, 94, 103120.
- Chen, Y., Zhang, C., Xue, T., Song, Y., Makris, N., Rathi, Y., Cai, W., Zhang, F., & O'Donnell, L. J. (2023). Deep fiber clustering: Anatomically informed fiber clustering with self-supervised deep learning for fast and effective tractography parcellation. *NeuroImage*, 273, 120086.
- Colby, J. B., Soderberg, L., Lebel, C., Dinov, I. D., Thompson, P. M., & Sowell, E. R. (2012). Along-tract statistics allow for enhanced tractography analysis. *NeuroImage*, 59, 3227–3242.
- Davies, D. L., & Bouldin, D. W. (1979). A cluster separation measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1, 224–227.
- Deng, X., Zhang, J., Liu, R., & Liu, K. (2022). Classifying ASD based on time-series fMRI using spatial-temporal transformer. *Computers in Biology and Medicine*, 151, 106320.
- Diedrichsen, J., Balsters, J. H., Flavell, J., Cussans, E., & Ramnani, N. (2009). A probabilistic MR atlas of the human cerebellum. *NeuroImage*, 46, 39–46.
- Diedrichsen, J., Maderwald, S., Küper, M., Thürling, M., Rabe, K., Gizewski, E. R., Ladd, M. E., & Timmann, D. (2011). Imaging the deep cerebellar nuclei: A probabilistic atlas and normalization procedure. *NeuroImage*, 54, 1786–1794.
- Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, 55, 78–87.
- Eickhoff, S. B., Thirion, B., Varoquaux, G., & Bzdok, D. (2015). Connectivity-based parcellation: Critique and implications. *Human Brain Mapping*, 36, 4771–4792.
- Eickhoff, S. B., Yeo, B. T. T., & Genon, S. (2018). Imaging-based parcellations of the human brain. *Nature Reviews. Neuroscience*, 19, 672–686.
- Fan, F.-L., Xiong, J., Li, M., & Wang, G. (2021). On interpretability of artificial neural networks: A survey. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 5, 741–760.



- Filippi, M., Basaia, S., Canu, E., Imperiale, F., Magnani, G., Falautano, M., Comi, G., Falini, A., & Agosta, F. (2018). Changes in functional and structural brain connectome along the Alzheimer's disease continuum. *Molecular Psychiatry*, 25, 230–239.
- Fong, R., & Vedaldi, A. (2017). Interpretable explanations of black boxes by meaningful perturbation. In Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp 3429–3437.
- Ford, A. A., Colon-Perez, L., Triplett, W. T., Gullett, J. M., Mareci, T. H., & Fitzgerald, D. B. (2013). Imaging white matter in human brainstem. *Frontiers in Human Neuroscience*, 7, 400.
- Freedman, D., Pisani, R., & Purves, R. (2007). *Statistics*. Norton.
- Fritz, N. E., Edwards, E. M., Ye, C., Prince, J., Yang, Z., Gressett, T., Keller, J., Myers, E., Calabresi, P. A., & Zackowski, K. M. (2022). Cerebellar contributions to motor and cognitive control in multiple sclerosis☆☆☆. *Archives of Physical Medicine and Rehabilitation*, 103, 1592–1599.
- Gershon, R. C., Wagster, M. V., Hendrie, H. C., Fox, N. A., Cook, K. F., & Nowinski, C. J. (2013). NIH Toolbox for assessment of neurological and behavioral function. *Neurology*, 80, S2–S6.
- Ghimire, P., Lavrador, J. P., Baig Mirza, A., Pereira, N., Keeble, H., Borri, M., Furlanetti, L., Brogna, C., Jarosz, J., Gullan, R., Vergani, F., Bhangoo, R., & Ashkan, K. (2021). Intraoperative mapping of precentral motor cortex and subcortex: A proposal for supplemental cortical and novel subcortical maps to Penfield's motor homunculus. *Brain Structure & Function*, 226, 1601–1611.
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C. F., Jenkinson, M., Smith, S. M., & Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, 536, 171–178.
- Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., Van Essen, D. C., Jenkinson, M., & WU-Minn HCP Consortium. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80, 105–124.
- Gong, W., Beckmann, C. F., & Smith, S. M. (2021). Phenotype discovery from population brain imaging. *Medical Image Analysis*, 71, 102050.
- Gupta, V., Booth, S., & Ko, J. H. (2021). Hypermetabolic cerebellar connectome in Alzheimer's disease. *Brain Connectivity*, 13, 356–366. <https://doi.org/10.1089/brain.2020.0937>
- Habas, C., & Manto, M. (2018). Probing the neuroanatomy of the cerebellum using tractography. *Handbook of Clinical Neurology*, 154, 235–249.
- Haghsomar, M., Shobeiri, P., Seyedi, S. A., Abbasi-Feijani, F., Poopak, A., Sotoudeh, H., Kamali, A., & Aarabi, M. H. (2022). Cerebellar microstructural abnormalities in Parkinson's disease: A systematic review of diffusion tensor imaging studies. *Cerebellum*, 21, 545–571.
- Harrison, O. K., Guell, X., Klein-Flügge, M. C., & Barry, R. L. (2021). Structural and resting state functional connectivity beyond the cortex. *NeuroImage*, 240, 118379.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer.
- He, H., Zhang, F., Pieper, S., Makris, N., Rathi, Y., Wells, W., & O'Donnell, L. J. (2022). Model and predict age and sex in healthy subjects using brain white matter features: A deep learning approach. In 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI). IEEE. pp 1–5.
- He, L., Li, H., Wang, J., Chen, M., Gozdas, E., Dillman, J. R., & Parikh, N. A. (2020). A multi-task, multi-stage deep transfer learning model for early prediction of neurodevelopment in very preterm infants. *Scientific Reports*, 10, 15072.
- Hernandez-Castillo, C. R., Galvez, V., Mercadillo, R., Diaz, R., Campos-Romo, A., & Fernandez-Ruiz, J. (2015). Extensive white matter alterations and its correlations with ataxia severity in SCA 2 patients. *PLoS One*, 10, e0135449.
- Hernandez-Castillo, C. R., Vaca-Palomares, I., Barrios, F., Martinez, L., Boll, M.-C., & Fernandez-Ruiz, J. (2016). Ataxia severity correlates with white matter degeneration in spinocerebellar ataxia type 7. *AJNR. American Journal of Neuroradiology*, 37, 2050–2054.
- Hodes, R. J., Insel, T. R., Landis, S. C., & NIH Blueprint for Neuroscience Research. (2013). The NIH Toolbox: Setting a standard for biomedical research. *Neurology*, 80, S1.
- Huang, C.-C., Sugino, K., Shima, Y., Guo, C., Bai, S., Mensh, B. D., Nelson, S. B., & Hantman, A. W. (2013). Convergence of pontine and proprioceptive streams onto multimodal cerebellar granule cells. *eLife*, 2, e00400.
- Ishikawa, T., Shimuta, M., & Häusser, M. (2015). Multimodal sensory integration in single cerebellar granule cells in vivo. *eLife*, 4, e12916.
- Jeong, J.-W., Tiwari, V. N., Behen, M. E., Chugani, H. T., & Chugani, D. C. (2014). In vivo detection of reduced Purkinje cell fibers with diffusion MRI tractography in children with autistic spectrum disorders. *Frontiers in Human Neuroscience*, 8, 110.
- Jeurissen, B., Leemans, A., Tournier, J.-D., Jones, D. K., & Sijbers, J. (2013). Investigating the prevalence of complex fiber configurations in white matter tissue with diffusion magnetic resonance imaging. *Human Brain Mapping*, 34, 2747–2766.
- Kelly, R. M., & Strick, P. L. (2003). Cerebellar loops with motor cortex and prefrontal cortex of a nonhuman primate. *The Journal of Neuroscience*, 23, 8432–8444.
- Kim, S. E., Jung, S., Sung, G., Bang, M., & Lee, S.-H. (2021). Impaired cerebro-cerebellar white matter connectivity and its associations with cognitive function in patients with schizophrenia. *NPJ Schizophrenia*, 7, 38.
- Kindermans, P.-J., Hooker, S., Adebayo, J., Alber, M., Schütt, K. T., Dähne, S., Erhan, D., & Kim, B. (2019). The (un)reliability of saliency methods. In W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, & K.-R. Müller (Eds.), *Explainable AI: Interpreting, explaining and visualizing deep learning* (pp. 267–280). Springer International Publishing.
- Kozioł, L. F., Budding, D., Andreasen, N., D'Arrigo, S., Bulgheroni, S., Imamizu, H., Ito, M., Manto, M., Marvel, C., Parker, K., Pezzulo, G., Ramnani, N., Riva, D., Schmahmann, J., Vandervert, L., & Yamazaki, T. (2014). Consensus paper: The cerebellum's role in movement and cognition. *Cerebellum*, 13, 151–177.
- Kugelman, J., Allman, J., Read, S. A., Vincent, S. J., Tong, J., Kalloniatis, M., Chen, F. K., Collins, M. J., & Alonso-Caneiro, D. (2022). A comparison of deep learning U-net architectures for posterior segment OCT retinal layer segmentation. *Scientific Reports*, 12, 14888.
- Kulkarni, M., Kent, J. S., Park, K., Guell, X., & Anteraper, S. (2023). Resting-state functional connectivity-based parcellation of the human dentate nucleus: New findings and clinical relevance. *Brain Structure & Function*, 228, 1799–1810.
- Liang, W., Zhang, K., Cao, P., Liu, X., Yang, J., & Zaiane, O. (2021). Rethinking modeling Alzheimer's disease progression from a multi-task learning perspective with deep recurrent neural network. *Computers in Biology and Medicine*, 138, 104935.
- Lin, T., Wang, Y., Liu, X., & Qiu, X. (2022). A survey of transformers. *AI Open*, 3, 111–132.
- Liu, M., Zhang, J., Adeli, E., & Shen, D. (2019). Joint classification and regression via deep multi-task multi-channel learning for Alzheimer's disease diagnosis. *IEEE Transactions on Biomedical Engineering*, 66, 1195–1206.
- Loshchilov, I., & Hutter, F. (2017). Decoupled weight decay regularization. arXiv [cs.LG]. arXiv. Retrieved from <http://arxiv.org/abs/1711.05101>
- MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics. University of California Press. Vol. 5.1, pp 281–298.
- Makris, N., Hodge, S. M., Haselgrove, C., Kennedy, D. N., Dale, A., Fischl, B., Rosen, B. R., Harris, G., Caviness, V. S., Jr., & Schmahmann, J. D. (2003). Human cerebellum: Surface-assisted cortical parcellation and volumetry with magnetic resonance imaging. *Journal of Cognitive Neuroscience*, 15, 584–599.

- Makris, N., Schlerf, J. E., Hodge, S. M., Haselgrove, C., Albaugh, M. D., Seidman, L. J., Rauch, S. L., Harris, G., Biederman, J., Caviness, V. S., Jr., Kennedy, D. N., & Schmahmann, J. D. (2005). MRI-based surface-assisted parcellation of human cerebellar cortex: An anatomically specified method with estimate of reliability. *NeuroImage*, 25, 1146–1160.
- Malcolm, J. G., Shenton, M. E., & Rath, Y. (2010). Filtered multitensor tractography. *IEEE Transactions on Medical Imaging*, 29, 1664–1675.
- Matsoukas, C., Haslum, J. F., Söderberg, M., & Smith, K. (2021). Is it time to replace CNNs with transformers for medical images? arXiv [cs.CV]. arXiv. Retrieved from <http://arxiv.org/abs/2108.09038>
- Mittal, V. A., Dean, D. J., Bernard, J. A., Orr, J. M., Pelletier-Baldelli, A., Carol, E. E., Gupta, T., Turner, J., Leopold, D. R., Robustelli, B. L., & Millman, Z. B. (2014). Neurological soft signs predict abnormal cerebellar-thalamic tract development and negative symptoms in adolescents at high risk for psychosis: A longitudinal perspective. *Schizophrenia Bulletin*, 40, 1204–1215.
- Mooi, E., & Sarstedt, M. (2011). *A concise guide to market research: The process, data, and methods using IBM SPSS statistics*. Springer Science & Business Media.
- Nettekoven, C., Zhi, D., Shahshahani, L., Pinho, A. L., Saadon-Grosman, N., & Diedrichsen, J. (2023). A hierarchical atlas of the human cerebellum for functional precision mapping. bioRxiv. <https://doi.org/10.1101/2023.09.14.557689v1.abstract>
- Norton, I., Essayed, W. I., Zhang, F., Pujol, S., Yarmarkovich, A., Golby, A. J., Kindlmann, G., Wassermann, D., Estepar, R. S. J., Rath, Y., Pieper, S., Kikinis, R., Johnson, H. J., Westin, C.-F., & O'Donnell, L. J. (2017). SlicerDMRI: Open source diffusion MRI software for brain cancer research. *Cancer Research*, 77, e101–e103.
- O'Donnell, L. J., Wells, W. M., 3rd, Golby, A. J., & Westin, C.-F. (2012). Unbiased groupwise registration of white matter tractography. *Medical Image Computing and Computer Assisted Interventions*, 15, 123–130.
- O'Donnell, L. J., & Westin, C.-F. (2007). Automatic tractography segmentation using a high-dimensional white matter atlas. *IEEE Transactions on Medical Imaging*, 26, 1562–1575.
- O'Donnell, L. J., & Westin, C.-F. (2011). An introduction to diffusion tensor image analysis. *Neurosurgery Clinics of North America*, 22, 185–196.
- O'Donnell, L. J., Westin, C.-F., & Golby, A. J. (2009). Tract-based morphometry for white matter group analysis. *NeuroImage*, 45, 832–844.
- Palesi, F., Ferrante, M., Gaviraghi, M., Misiti, A., Savini, G., Lascialfari, A., D'Angelo, E., & Gandini Wheeler-Kingshott, C. A. M. (2021). Motor and higher-order functions topography of the human dentate nuclei identified with tractography and clustering methods. *Human Brain Mapping*, 42, 4348–4361.
- Peng, X., Long, G., Shen, T., Wang, S., & Jiang, J. (2021). Sequential diagnosis prediction with transformer and ontological representation. In 2021 IEEE International Conference on Data Mining (ICDM). IEEE. pp 489–498.
- Phillips, J. R., Hewedi, D. H., Eissa, A. M., & Moustafa, A. A. (2015). The cerebellum and psychiatric disorders. *Frontiers in Public Health*, 3, 66.
- Propper, R. E., O'Donnell, L. J., Whalen, S., Tie, Y., Norton, I. H., Suarez, R. O., Zollei, L., Radmanesh, A., & Golby, A. J. (2010). A combined fMRI and DTI examination of functional language lateralization and arcuate fasciculus structure: Effects of degree versus direction of hand preference. *Brain and Cognition*, 73, 85–92.
- Quartarone, A., Cacciola, A., Milardi, D., Ghilardi, M. F., Calamuneri, A., Chillemi, G., Anastasi, G., & Rothwell, J. (2020). New insights into cortico-basal-cerebellar connectome: Clinical and physiological considerations. *Brain*, 143, 396–406.
- Ramos-Llordén, G., Ning, L., Liao, C., Mukhometzianov, R., Michailovich, O., Setsompop, K., & Rath, Y. (2020). High-fidelity, accelerated whole-brain submillimeter in vivo diffusion MRI using gSlider-spherical ridgelets (gSlider-SR). *Magnetic Resonance in Medicine*, 84, 1781–1795.
- Rao, S., Li, Y., Ramakrishnan, R., Hassaine, A., Canoy, D., Cleland, J., Lukasiewicz, T., Salimi-Khorshidi, G., & Rahimi, K. (2022). An explainable transformer-based deep learning model for the prediction of incident heart failure. *IEEE Journal of Biomedical and Health Informatics*, 26, 3362–3372.
- Reddy, C. P., & Rath, Y. (2016). Joint multi-fiber NODDI parameter estimation and tractography using the unscented information filter. *Frontiers in Neuroscience*, 10, 166.
- Reuben, D. B., Magasi, S., McCreath, H. E., Bohannon, R. W., Wang, Y.-C., Bubela, D. J., Rymer, W. Z., Beaumont, J., Rine, R. M., Lai, J.-S., & Gershon, R. C. (2013). Motor assessment using the NIH Toolbox. *Neurology*, 80, S65–S75.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. New York, NY, USA: Association for Computing Machinery. KDD '16. pp 1135–1144.
- Rokach, L. (2010). A survey of clustering algorithms. In O. Maimon & L. Rokach (Eds.), *Data mining and knowledge discovery handbook* (pp. 269–298). Springer US.
- Rousseau, P.-N., Chakravarty, M. M., & Steele, C. J. (2022). Mapping pontocerebellar connectivity with diffusion MRI. *NeuroImage*, 264, 119684.
- Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20, 53–65.
- Ruder, S. (2017). An overview of multi-task learning in deep neural networks. arXiv [cs.LG]. arXiv. Retrieved from <http://arxiv.org/abs/1706.05098>
- Sarvamangala, D. R., & Kulkarni, R. V. (2022). Convolutional neural networks in medical image understanding: A survey. *Evolutionary Intelligence*, 15, 1–22.
- Schmahmann, J. D. (1996). From movement to thought: Anatomic substrates of the cerebellar contribution to cognitive processing. *Human Brain Mapping*, 4, 174–198.
- Schmahmann, J. D. (2010). The role of the cerebellum in cognition and emotion: Personal reflections since 1982 on the dysmetria of thought hypothesis, and its historical evolution from theory to therapy. *Neuropsychology Review*, 20, 236–260.
- Schmahmann, J. D. (2019). The cerebellum and cognition. *Neuroscience Letters*, 688, 62–75.
- Schmahmann, J. D., & Pandya, D. N. (1997a). The cerebrocerebellar system. *International Review of Neurobiology*, 41, 31–60.
- Schmahmann, J. D., & Pandya, D. N. (1997b). Anatomic organization of the basilar pontine projections from prefrontal cortices in rhesus monkey. *The Journal of Neuroscience*, 17, 438–458.
- Schmahmann, J. D., Rosene, D. L., & Pandya, D. N. (2004). Motor projections to the basis pontis in rhesus monkey. *The Journal of Comparative Neurology*, 478, 248–268.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2020). Grad-CAM: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision*, 128, 336–359.
- Sereno, M. I., Diedrichsen, J., Tachrount, M., Testa-Silva, G., d'Arceuil, H., & De Zeeuw, C. (2020). The human cerebellum has almost 80% of the surface area of the neocortex. *Proceedings of the National Academy of Sciences of the United States of America*, 117, 19538–19543.
- Shamshad, F., Khan, S., Zamir, S. W., Khan, M. H., Hayat, M., Khan, F. S., & Fu, H. (2023). Transformers in medical imaging: A survey. *Medical Image Analysis*, 88, 102802.
- Shrikumar, A., Greenside, P., & Kundaje, A. (2017). Learning important features through propagating activation differences. In: Precup, D., Teh, Y. W., eds. Proceedings of the 34th International Conference on Machine Learning. PMLR. Proceedings of Machine Learning Research Vol. 70, pp. 3145–3153.
- Siless, V., Medina, S., Varoquaux, G., & Thirion, B. (2013). A comparison of metrics and algorithms for fiber clustering. In 2013 International

- Workshop on Pattern Recognition in Neuroimaging. IEEE. pp 190–193.
- Simonyan, K., Vedaldi, A., & Zisserman, A. (2013). Deep inside convolutional networks: Visualising image classification models and saliency maps. arXiv [cs.CV]. arXiv. Retrieved from <http://arxiv.org/abs/1312.6034>
- Smith, S. M., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T. E., Mackay, C. E., Watkins, K. E., Ciccarelli, O., Cader, M. Z., Matthews, P. M., & Behrens, T. E. J. (2006). Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data. *NeuroImage*, 31, 1487–1505.
- Springenberg, J. T., Dosovitskiy, A., Brox, T., & Riedmiller, M. (2014). Striving for simplicity: The all convolutional net. arXiv [cs.LG]. arXiv. Retrieved from <http://arxiv.org/abs/1412.6806>
- Steele, C. J., Anwander, A., Bazin, P.-L., Trampel, R., Schaefer, A., Turner, R., Ramnani, N., & Villringer, A. (2017). Human cerebellar sub-millimeter diffusion imaging reveals the motor and non-motor topography of the dentate nucleus. *Cerebral Cortex*, 27, 4537–4548.
- Stoodley, C. J., & Schmahmann, J. D. (2009). Functional topography in the human cerebellum: A meta-analysis of neuroimaging studies. *NeuroImage*, 44, 489–501.
- Stoodley, C. J., & Schmahmann, J. D. (2018). Functional topography of the human cerebellum. *Handbook of Clinical Neurology*, 154, 59–70.
- Tang, Y., Sun, W., Toga, A. W., Ringman, J. M., & Shi, Y. (2018). A probabilistic atlas of human brainstem pathways based on connectome imaging data. *NeuroImage*, 169, 227–239.
- Toniolo, S., Serra, L., Olivito, G., Caltagirone, C., Mercuri, N. B., Marra, C., Cercignani, M., & Bozzali, M. (2020). Cerebellar white matter disruption in Alzheimer's disease patients: A diffusion tensor imaging study. *Journal of Alzheimer's Disease*, 74, 615–624.
- Travis, K. E., Leitner, Y., Feldman, H. M., & Ben-Shachar, M. (2015). Cerebellar white matter pathways are associated with reading skills in children and adolescents. *Human Brain Mapping*, 36, 1536–1553.
- Usman, M., Zia, T., & Tariq, A. (2022). Analyzing transfer learning of vision transformers for interpreting chest radiography. *Journal of Digital Imaging*, 35, 1445–1462.
- van Baarsen, K. M., Kleinnijenhuis, M., Jbabdi, S., Sotiropoulos, S. N., Grotenhuis, J. A., & van Cappellen van Walsum, A. M. (2016). A probabilistic atlas of the cerebellar white matter. *NeuroImage*, 124, 724–732.
- Vaswani, A., Shazeer, N., & Parmar, N. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*. Retrieved from <https://proceedings.neurips.cc/paper/7181-attention-is-all>
- Vázquez, A., López-López, N., Sánchez, A., Houenou, J., Poupon, C., Mangin, J.-F., Hernández, C., & Guevara, P. (2020). FFClust: Fast fiber clustering for large tractography datasets for a detailed study of brain connectivity. *NeuroImage*, 220, 117070.
- Weintraub, S., Dikmen, S. S., Heaton, R. K., Tulsky, D. S., Zelazo, P. D., Bauer, P. J., Carlozzi, N. E., Slotkin, J., Blitz, D., Wallner-Allen, K., Fox, N. A., Beaumont, J. L., Mungas, D., Nowinski, C. J., Richler, J., Deocampo, J. A., Anderson, J. E., Manly, J. J., Borosh, B., ... Gershon, R. C. (2013). Cognition assessment using the NIH Toolbox. *Neurology*, 80, S54–S64.
- Wu, X., Liao, X., Zhan, Y., Cheng, C., Shen, W., Huang, M., Zhou, Z., Wang, Z., Qiu, Z., Xing, W., Liao, W., Tang, B., & Shen, L. (2017). Microstructural alterations in asymptomatic and symptomatic patients with spinocerebellar ataxia type 3: A tract-based spatial statistics study. *Frontiers in Neurology*, 8, 714.
- Xiong, R., Yang, Y., He, D., Zheng, K., Zheng, S., Xing, C., Zhang, H., Lan, Y., Wang, L., & Liu, T. (2020). On layer normalization in the transformer architecture. In: Iii, H. D., Singh, A., eds. *Proceedings of the 37th International Conference on Machine Learning*. PMLR. *Proceedings of Machine Learning Research* Vol. 119, pp. 10524–10533.
- Yeatman, J. D., Dougherty, R. F., Myall, N. J., Wandell, B. A., & Feldman, H. M. (2012). Tract profiles of white matter properties: Automating fiber-tract quantification. *PLoS One*, 7, e49790.
- Yeh, F.-C., Panesar, S., Fernandes, D., Meola, A., Yoshino, M., Fernandez-Miranda, J. C., Vettel, J. M., & Verstynen, T. (2018). Population-averaged atlas of the macroscale human structural connectome and its network topology. *NeuroImage*, 178, 57–68.
- Zekelman, L., Chen, Y., Zhang, F., Cetin Karayumak, S., Almeida, M., Pieper, S., Makris, N., Rath, Y., Golby, A. J., & O'Donnell, L. J. (2023). The effect of cerebellar pathway microstructure on cognition: A diffusion MRI investigation of >9000 pre-adolescents & >1000 young adults presented at the the Organization for Human Brain Mapping (OHBM) 2023 Annual Meeting, July 26, Montréal, QC. Retrieved from <https://event.fourwaves.com/ohbm2023/abstracts/26a836cb-1054-411c-9ecf-c260ab9f9b7a/poster>
- Zekelman, L. R., Zhang, F., Makris, N., He, J., Chen, Y., Xue, T., Liera, D., Drane, D. L., Rath, Y., Golby, A. J., & O'Donnell, L. J. (2022). White matter association tracts underlying language and theory of mind: An investigation of 809 brains from the Human Connectome Project. *NeuroImage*, 246, 118739.
- Zhang, F., Daducci, A., He, Y., Schiavi, S., Seguin, C., Smith, R. E., Yeh, C.-H., Zhao, T., & O'Donnell, L. J. (2022). Quantitative mapping of the brain's structural connectivity using diffusion MRI tractography: A review. *NeuroImage*, 249, 118870.
- Zhang, F., Noh, T., Juvekar, P., Frisken, S. F., Rigolo, L., Norton, I., Kapur, T., Pujol, S., Wells, W., 3rd, Yarmarkovich, A., Kindlmann, G., Wassermann, D., San Jose Estepar, R., Rath, Y., Kikinis, R., Johnson, H. J., Westin, C.-F., Pieper, S., Golby, A. J., & O'Donnell, L. J. (2020). SlicerDMRI: Diffusion MRI and tractography research software for brain cancer surgery planning and visualization. *JCO Clinical Cancer Informatics*, 4, 299–309.
- Zhang, F., Savadjiev, P., Cai, W., Song, Y., Rath, Y., Tunç, B., Parker, D., Kapur, T., Schultz, R. T., Makris, N., Verma, R., & O'Donnell, L. J. (2018). Whole brain white matter connectivity analysis using machine learning: An application to autism. *NeuroImage*, 172, 826–837.
- Zhang, F., Wu, Y., Norton, I., Rigolo, L., Rath, Y., Makris, N., & O'Donnell, L. J. (2018). An anatomically curated fiber clustering white matter atlas for consistent white matter tract parcellation across the lifespan. *NeuroImage*, 179, 429–447.
- Zhang, L., Zhao, Y., Che, T., Li, S., & Wang, X. (2023). Graph neural networks for image-guided disease diagnosis: A review. *iRADIOLOGY*, 1, 151–166.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., & Torralba, A. (2016). Learning deep features for discriminative localization. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE. pp 2921–2929.

## SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

**How to cite this article:** Tchetchenian, A., Zekelman, L., Chen, Y., Rushmore, J., Zhang, F., Yeterian, E. H., Makris, N., Rath, Y., Meijering, E., Song, Y., & O'Donnell, L. J. (2024). Deep multimodal saliency parcellation of cerebellar pathways: Linking microstructure and individual function through explainable multitask learning. *Human Brain Mapping*, 45(12), e70008. <https://doi.org/10.1002/hbm.70008>