A task-general connectivity model reveals variation in convergence of cortical inputs to functional regions of the cerebellum

Maedbh King[1,+,], Ladan Shahshahani[2,+], Richard Ivry[1,3] & Jörn Diedrichsen[2,4,5]*

1. Department of Psychology, University of California, Berkeley, CA, USA

2. Western Institute for Neuroscience, Western University, London, Ontario, Canada

3. Helen Wills Neuroscience Institute, University of California, Berkeley, CA, USA

4. Department of Statistical and Actuarial Sciences, Western University, London, Ontario, Canada

5. Department of Computer Science, Western University, London, Ontario, Canada

+ The first two authors contributed equally to the work.

* To whom correspondence should be addressed: Maedbh King, Department of Psychology, University of California, Berkeley, 2121 Berkeley Way West, Berkeley, CA, USA. Email: maedbhking@berkeley.edu.

Contributions: M.K. R.I. and J.D. originally conceived of the project. M.K. designed the experiment and collected the data. M.K. L.S. and J.D. performed the analyses. M.K., L.S., J.D., and R.I. wrote the manuscript. All authors discussed the results and contributed to the final manuscript.

Keywords: Cerebellum; Task-Based fMRI; Functional Connectivity; Encoding Models; Multi-Domain Task Battery (MDTB)

Acknowledgements: This work was supported by the Canadian Institutes of Health Research (PJT 159520 to J.D.), the Canada First Research Excellence Fund (BrainsCAN to Western University), and the National Institute of Health (NS092079, NS105839 to R.B.I.). Thanks to Da Zhi for providing cerebral cortical parcellations.

Competing Interests: R.I. is a co-founder with equity in Magnetic Tides, Inc. The other authors declare no competing interests.

Abstract

While resting-state fMRI studies have provided a broad picture of the connectivity between human neocortex and cerebellum, the degree of convergence of cortical inputs onto cerebellar circuits remains unknown: Does each cerebellar region receive input from a single cortical area or convergent inputs from multiple cortical areas? Here we use task-based fMRI data to build a range of cortico-cerebellar connectivity models, each allowing for a different degree of convergence. We compared these models by their ability to predict cerebellar activity patterns for novel task sets. Models that allow some degree of convergence provided the best predictions, arguing for convergence of multiple cortical inputs onto single cerebellar voxels. Importantly, the degree of convergence varied across the cerebellum with the highest convergence observed in areas linked to language, working memory, and social cognition. These findings suggest important differences in the way that functional subdivisions of the cerebellum support motor and cognitive function.

Introduction

The last 30 years has witnessed a paradigm shift with regards to cerebellar function, with broad recognition that this subcortical structure is engaged in many aspects of human cognition. Since the first report of cerebellar activation in a semantic retrieval task[1], thousands of neuroimaging papers have reported cerebellar recruitment during a broad range of tasks that cannot be attributed to the motor demands of these tasks. Functional interpretations include hypotheses concerning how the cerebellum may facilitate attentional shifts[2], stimulus-response mapping[3], higher order rule processing[4], verbal working memory[5], language[6], and social cognition[7]. This body of work has produced functional maps of the cerebellum that depict the association of particular cognitive processes with different subregions of the cerebellum[8].

Given the relative uniform cytoarchitecture of the cerebellar cortex, it is assumed that differences in function mainly arise from variation in the input to the cerebellum. Trans-synaptic tracing methods employed in non-human primates studies have revealed extensive reciprocal connections between many frontal and parietal areas and the cerebellum[9,10]. These studies have highlighted the closed-loop nature of these connections, with each (neo-)cortical region projecting to a specific cerebellar region, and receiving input from the same area[11]. In humans, resting state functional connectivity analyses have revealed a set of cerebellar networks, each one associated with a specific cortical network[12–14].

An important unanswered question is whether each cerebellar region receives input from a restricted cortical region or whether it receives convergent input from multiple cortical regions. Providing an answer to this question has important implications for our understanding of cerebellar function. An architecture marked by a one-to-one relationship between cortical and cerebellar regions would suggest that the function of each cerebellar region is to fine-tune the dynamics in its cortical input. In contrast, a convergent architecture would suggest that subregions within the cerebellum integrate information across disparate cortical regions and may coordinate their interactions. Indeed, recent work in the rodent brain has suggested convergence of mossy fibers from diverse sources onto the same cerebellar region[15,16], or even onto the same granule cells[17]. Furthermore, the pattern of cortico-cerebellar convergence may vary across the cerebellar cortex, similar to how cortical areas

show considerable variation in the degree to which they serve as points of convergence from other cortical regions[18–20].

In the current study we introduce a novel approach to study the pattern of connectivity between the cerebral cortex and cerebellum. We employed an fMRI data set obtained while human participants completed a large task battery that was designed to engage cognitive processes across a broad range of functional domains (e.g., visual cognition, memory, attention, cognitive control). Using this rich data set, we built models that predicted the cerebellar activity based solely on cortical activity, with the models varying in terms of the degree of convergence of cortical inputs onto each cerebellar area. To evaluate the models, we examined how well each model predicted data obtained from different tasks and/or different participants. These analyses reveal a novel picture of cortico-cerebellar connectivity, one in which the degree of convergence was higher in cerebellar regions associated with more complex cognitive functions.

Fig. 1 | Connectivity model training and evaluation. *(A) Models were trained on task set A of the multi-domain task battery (MDTB[8]). Model hyperparameters were tuned using 4-fold cross validation. Three types of models were tested (WTA, Lasso, Ridge), each with 7 cortical parcellations of different granularity. (B) Models were used to predict the cerebellar activity solely from cortical activity patterns when participants performed task set B (including both novel and common tasks). To avoid the influence of shared noise correlations across cortex and cerebellum, the models were trained and tested by using cortical and cerebellar activity patterns from different sessions (see methods). (C) As a*

test of generalization, the models were used to predict cerebellar activity from cortical data obtained from a separate experiment (King et al., unpublished data).

Results

Overview

We compared three models of cortico-cerebellar connectivity by using a region-to-region predictive modeling approach[21,22]. For each model, the cortical activity pattern for each task was used to predict the activity pattern in the cerebellum (Fig. 1). We chose this direction of prediction, as the cerebellar BOLD signal overwhelmingly reflects mossy-fiber input, with minimal contribution from cerebellar output neurons, the Purkinje cells[23,24]. As a model of sparse connectivity, we used a Winner-Take-All (WTA) model that imposed the strong constraint that only a single cortical region could be used to predict the activity in each cerebellar voxel. The other two models allowed for some degree of convergence. Models estimated with Lasso Regression (L1 regularization) found sparse solutions, minimizing the number of cortical inputs by setting to zero those weights that make negligible contribution to the predicted activity. Models estimated with Ridge Regression (L2 regularization) allowed for a wide distribution of inputs, keeping each weight as small as possible. For the input to each model, we used a set of cortical parcellations that varied in terms of their level of granularity.

We trained the models to predict the activity in each voxel of the human cerebellum as a linear function of cortical activity (Fig. 1A), using cortical and cerebellar fMRI data obtained from 24 participants who performed a large task battery over four scanning sessions. The battery was composed of 26 tasks that could be further subdivided into 47 conditions. Set A consisted of seventeen tasks presented in an interleaved fashion in the first two sessions. Set B consisted of 17 tasks, 8 of which had been part of Set A, presented in an interleaved fashion in the second two sessions. We used Set A to estimate the connectivity weights $\hat{W}$ for each model. These weights were then used to make predictions for the cerebellar activity patterns from Set B (Fig. 1B). Importantly, the model predictions relied on a single set of connectivity weights across all tasks, and therefore were based only on the cortical activity patterns without reference to any features of the tasks themselves. In a final generalization test, the models were used to predict cerebellar activity patterns from new participants who were tested on a set of novel tasks (Fig. 1C). Overall, the training and evaluation was designed to determine the best task-general model of cortico-cerebellar connectivity.

To validate that this approach enables us to distinguish between different forms of cortico-cerebellar connectivity, we conducted a range of model recovery simulations. Using the measured cortical activity for each participant, we generated artificial sets of cerebellar data imposing either a one-to-one or a many-to-one mapping. Following the procedure used with the real data, we trained the different models on Task Set A and tested them on Task Set B. The simulations (Extended Data Fig. 1) showed that the WTA performs best if each cerebellar voxel is connected to only one cortical parcel, whereas ridge regression performs better if there is substantial convergence, with Lasso performing in between the two methods. Thus, despite the considerable covariance between cortical parcels, the modeling approach was able to distinguish between different forms of connectivity.

Cortico-cerebellar connectivity is best captured by models with convergence

We first compared the different models, asking how well they predicted activity patterns obtained when the same participants were tested on task set B (Fig. 2A). Models allowing for some degree of convergence outperformed the WTA model. This advantage was observed across all levels of cortical

granularity; indeed, the prediction performance for the Ridge, Lasso, and WTA models was relatively independent of granularity. Averaged across all levels of granularity, the Ridge models outperformed the WTA models ($F_{1,23}$ = 47.122, p < .001). Post-hoc tests revealed that this advantage for the Ridge model was consistent across all levels of granularity, starting with a parcellation of 80 regions (t_{23} = 5.172, p < .001). We also found a significant difference in predictive accuracy between the Ridge and Lasso model ($F_{1,23}$ = 15.055, p < .001). Post-hoc tests showed that there was no significant difference in predictive accuracy at the lowest level of granularity (t_{23} = 1.279, p = 0.213), whereas the difference was significant for the the finer parcellations (all t_{23} > 10.937, p < .001).

The mean predictive accuracy of the Lasso and Ridge models was 0.257 (Fig. 2A). There are two issues of note here. First, the models are being used to predict activity of individual voxels, without any smoothing across cerebellar voxels. Thus, the upper bound for the predictive accuracy is limited by the reliability of the cerebellar test data. The correlation of the measured cerebellar activity patterns across the two sessions of Task Set B was, on average, r = 0.51 (SD = .102). Reliability was fairly consistent across the cerebellum (Fig. 2C), with some decreases in lobules I-IV. This dropoff likely reflects the fact that our battery only included hand and eye movements, and did not activate the lower body representation that is prominent in this area. Second, predictive accuracy is also limited by the reliability of the cortical data that was used to make the prediction. This is especially a limiting factor for WTA models that use fine granularity since the predictions will be based on data obtained from a small cortical region.

Given these issues, we calculated a noise ceiling for each model, taking into account the reliability of the cerebellar data, the reliability of the cortical data, and the effect of granularity (see Methods, Extended Data Fig. 2). As an unbiased comparison of the model predictions, Fig. 2B re-plots the predictive accuracy of each model normalized by its noise ceiling. As with the original analysis, the Ridge model significantly outperformed the WTA model ($F_{1,23}$ = 16.49, p < .01), and in post-hoc tests, the advantage was especially pronounced for finer parcellations (1848 regions; t_{23} = 5.073, p < .001). Overall, the noise ceiling calculation showed that the Ridge model was able to predict approximately 45% of the systematic variance of the cerebellar activity patterns across tasks (average $R^2=0.67^2$). While the predictive accuracy (Fig. 2D) was best in anterior motor regions, it was reasonably high across the entire cerebellar surface.

The predicted and observed activity patterns for two exemplary tasks (Fig. 2E) demonstrate the quality of these predictions at the group level. Both for the spatial working memory and social cognition tasks, the connectivity model predicted the pattern of task activity with a high degree of fidelity. This is especially compelling for the spatial working memory task as the training set did not include a task with similar characteristics.

To ensure that our results were not biased by the use of an essentially arbitrary parcellation of the neocortex, we repeated the analysis using a range of published functional parcellations. For example, we used a 7-network cortical parcellation based on resting state fMRI data[25] to train and test the three models. Here, too, the WTA model was inferior to the Ridge models (t_{23} = 2.956, p < .01), with no performance difference between Ridge and Lasso models, t_{23} = -1.235, p = 0.229. The same pattern held when we used other common functional parcellations of the neocortex (Extended Data Fig. 3).

In summary, these results demonstrate that models which entail some degree of convergence from the neocortex to the cerebellum outperform a model in which cerebellar activity is based on input from a single cortical region. This conclusion holds across a broad range of cortical parcellations[26].

Fig. 2 | Performance of cortico-cerebellar connectivity models. (*A*) *Predictive accuracy (Pearson correlation) of the Ridge, Lasso, and WTA regression models for the test data of task set B. (**B**) Predictive accuracy normalized to the noise ceiling based on reliability of both cerebellar and cortical data (see methods). (**C**) Voxelwise map of inter-session reliability of the test data. (**D**) Voxelwise map of predictive accuracy of the Ridge model (1848 parcels), normalized to the noise ceiling. (**E**) Observed and predicted activity for a novel task involving spatial working memory (Spatial Map) and a social cognition task (Theory of Mind). The spatial map task was not included in Task Set A.*

Lobules VII receives more distributed input from the cerebral cortex

Given that the pattern of cortico-cerebellar convergence has important consequences for understanding the function of a cerebellar region, we were interested in knowing whether the pattern of convergence varied across the cerebellar cortex. As a first step, we tallied the number of non-zero regression weights for each cerebellar voxel as a measure of its input surface area. For this calculation, we used the Lasso model (80 cortical regions) since it uses the minimal set of cortical areas necessary for predicting each cerebellar voxel. By setting all other weights to zero, the Lasso model allows for a quantification of the important input weights without applying an arbitrary threshold. This calculation revealed a substantial variation in the degree of convergence across the cerebellar cortex (Fig. 3a). For example, predicting the activity pattern of voxels within Crus I required inclusion of up to 10% of the cortical surface whereas predicting the activity of voxels in the anterior lobe required less than 5% of the neocortex.

To statistically analyze these data, we opted to bin the cerebellar data using a functional 10-region parcellation of the cerebellum[8]. This analysis confirmed that the size of the estimated cortical input area differed across functional regions ($F_{9,207}$ = 7.244, p < .001, Fig. 3b). The lowest level of convergence was observed for regions 1 and 2, the anterior hand regions of the cerebellum. The highest levels of convergence were found in regions of the cerebellum associated with more cognitive

functions (e.g., regions 7 and 8, areas engaged during language comprehension). Noteworthy, this pattern holds for cortical parcellations of different levels of granularity (Extended Data Fig. 4), as well as for the thresholded coefficients from the Ridge regression models (Extended Data Fig. 5a, b).

As an independent method to quantify convergence, we determined the spatial spread of the inputs across the cortical surface. For this measure of dispersion, we calculated the spherical variance of the non-zero connectivity weights on an inflated model for each cerebral hemisphere (see Methods). This analysis revealed a similar pattern as seen in the surface area measure (Fig. 3c,d, $F_{9,207}$ = 18.322, p < .001). For example, the cortical inputs to the hand motor regions of the cerebellum were more concentrated whereas the cortical inputs to lobule VII were more dispersed. Again, this pattern also holds for the Ridge models (Extended Data Fig. 5c, d).

The preceding analyses provide a summary of convergence from the cerebral cortex to the cerebellum. To gain insight into the spatial profile of the cortical inputs, we visualized all of the cortical weights for each of the 10 functional regions of the cerebellum (Fig. 4). As expected given the crossed connectivity between M1 and lobules IV/V [9,27,28], the input to the hand regions of the cerebellum (regions 1 and 2) was centered around the contralateral primary sensorimotor cortex with some additional input from premotor and parietal cortex. Regions 3 and 4 are the other two cerebellar regions associated with motor function. Activity in region 3, the oculomotor vermis, is predicted by a bilateral set of cortical regions including the frontal eye fields (FEF), regions in the intraparietal sulcus, and extrastriate visual regions. Region 4, an area strongly activated during action observation, is predicted by a bilateral network of regions including pre-motor cortex, supplementary motor cortex (SMA) and parietal cortex. Cerebellar regions 5-10, the regions associated with more cognitive processes are predicted by a distributed set of cortical regions, with stronger input coming from the contralateral cortical hemisphere. For example, cerebellar region 5 is restricted to the left cerebellum and receives much stronger input from the right cerebral hemisphere. Overall, these findings point to significant variation in the patterns of cortico-cerebellar connectivity, with cerebellar regions in lobule VII, associated with cognitive processes such as language and working memory, exhibiting higher degrees of convergence as compared to cerebellar regions in lobule V and VIII.

Fig. 3 *Cortico-cerebellar convergence measures using the Lasso model.* (A) *Map of the cerebellum showing percentage of cortical parcels with non-zero weights for the Lasso model (n=80 parcels).* **(B)** *Percentage of parcels with non-zero weights for functional subregions of the cerebellum.* **(C)** *Spherical dispersion of the connectivity weights on the cortical surface for each cerebellar voxel.* **(D)** *Average cortical dispersion for each functional subregion of the cerebellum. Error bars indicate standard error of the mean across participants. See Extended Data Fig. 5 for the same results using Ridge regression.*

Fig. 4 | *Cortical connectivity weight maps* *for the Ridge regression model with 1848 cortical parcels for each of 10 functional cerebellar regions. Each region is denoted by the most important functional term[8]. Results are averaged across participants. Regression weights are in arbitrary units. See Extended Data Fig. 6. for the corresponding analysis using Lasso regression.*

Cortico-cerebellar connectivity model predicts new cerebellar data

A strong test of a model is how well it predicts novel data obtained in contexts distinct from that data used to develop the model. We conducted a generalization test using data from a separate experiment involving novel tasks and new participants. The data came from a study in which participants were trained over multiple sessions on five tasks, selected to span a set of distinct cognitive domains.

All three model types (WTA, Lasso, Ridge) were evaluated, using the connectivity weights estimated from Task Set A of the main study. Because the new study had different participants, we averaged the weights for each model across the individuals in the training set. These group-averaged weights were then used to predict cerebellar activity patterns for the new data.

As shown in Figure 5, the connectivity models were generally successful in predicting cerebellar activity patterns for the new dataset. The overall pattern is quite similar to that seen in the initial model tests (in which we had used data from different tasks and sessions, but involved training and test data from the same participants). Predictive accuracy was stable across levels of cortical granularity and the Ridge model provided the best overall predictive accuracy (r = .657) and the WTA (r = .352) the worst predictive accuracy. These results provide evidence that our cortico-cerebellar connectivity models capture important aspects of connectivity that are stable both across tasks and participants. Moreover, in accord with the earlier analyses based on predictions at the individual level, this generalization analysis again suggests that approximately 43% of the variation of cerebellar activity across tasks can be predicted by cortical activity alone.

Fig. 5 | *Generalization to new dataset*. *Models of cortico-cerebellar connectivity are tested on a new experiment. Each model is tested across different levels of cortical granularity. Predictive accuracy is the Pearson correlation between observed and predicted activity patterns, normalized to the noise ceiling.*

Discussion

To date, models of connectivity between the human neocortex and cerebellum have been based on fMRI resting-state data[12–14]. This work demonstrates that each region of the cerebellum receives input from a distinct set of cortical regions. For example, anterior and posterior cerebellar motor regions show correlated activity with contral-lateral sensorimotor cortex and non-motor or "cognitive" cerebellar regions show correlated activity with specific parietal and frontal association areas.

Despite these important insights, previous work has not been designed to examine the patterns of convergence between the neocortex and cerebellum. Resting state connectivity maps are generally produced by assigning each cerebellar voxel to a single cortical network, a *de facto* winner-take-all model. In the present study, we quantified and compared models of cortico-cerebellar connectivity. The models differed in the degree of convergence of cortical inputs to each cerebellar region, ranging from an architecture constrained to a strict one-to-one mapping to architectures that allowed for distributed inputs. We evaluated these models in terms of how well they could predict cerebellar activity patterns on a novel task set. For nearly the entire cerebellum, models that allowed for some convergence predicted cerebellar activity better than the WTA model.

Convergence differs across Cerebellar Circuits

Importantly, the amount of convergence differed across the cerebellar cortex. Specifically, regions in anterior (lobules I-V) and inferior cerebellum (lobules VIII-X) were well predicted by a relatively small and concentrated area of the cerebral cortex. In contrast, regions in lobule VI and especially lobule VII required input from a larger and spatially more dispersed set of cortical regions that were primarily located in association areas of prefrontal, parietal, and cingulate cortex. This finding underscores the heterogeneity of cortico-cerebellar connectivity with some cerebellar areas functioning in nearly a 1:1 relationship with a single cortical region, whereas other areas integrate input from a more diverse set of cortical regions.

This variation bears some resemblance to a motor/cognitive gradient identified from resting state data[29]. However, there are a few notable exceptions. First, based on our evaluation metrics, Region 3 (oculomotor vermis) was best explained by a large and relatively dispersed set of cortical regions, including intraparietal sulcus, the frontal eye fields (FEF), and extrastriate visual areas (Fig. 4). Second, sub-regions of lobules IX, functionally associated with non-motor tasks[12], are best explained by a relatively restricted set of cortical regions (Fig. 4a,c). Thus, rather than following a functional distinction, it appears that regions with strong convergence are anatomically restricted to lobules VI and VII, two areas that have disproportionately increased in size during primate evolution[30].

Variation in cortico-cerebellar connectivity has important implications for theories of cerebellar function[15,31,32]. A cerebellar region that forms a closed-loop circuit with a single cortical region can only fine-tune the computational processes that are ongoing in its connected cortical region. In contrast, a cerebellar region that receives convergent input would be able to coordinate the interactions between dispersed cortical regions. Indeed, using a rodent model, Pisano et al. have identified cerebellar areas that project back to multiple cortical areas, providing a substrate by which the cerebellum can influence neuronal dynamics not only in a focal cortical area but also interactions within networks of cortical areas[16].

Task-based vs. resting-state connectivity analyses

While our modeling approach could be applied to resting-state data, we opted to use the data from the multi-domain task battery for a variety of reasons. First, the breadth of tasks included in the battery assured that there would be reliable variation in activity across most of the neocortex and cerebellum, a necessary precondition for building a complete connectivity model. Second, the dataset allowed us to avoid biases that arise from correlated noise across abutting cortical and cerebellar regions[12]. We used the cortical activity patterns from one session to predict cerebellar activity patterns in a different session, relying on the fact that measurement noise is independent across sessions. Third, the MDTB data set allowed us to test the model across a broad set of tasks and mental states. While resting-state correlations are predictive of task-based activation patterns[8,26,33], resting in an fMRI scanner is arguably a relatively restricted situation. Our genarlization test shows that the connectivity model can robustly predict activity patterns for new tasks and participants.

Methodological limitations

Inter-region correlations of fMRI data can of course only provide indirect evidence of true functional connectivity. As such, it is important to consider methodological limitations that may influence the validity of our conclusions. From a statistical perspective, it was not clear, *a priori*, that we would have the power to distinguish between models of connectivity given that there can be substantial collinearity between different cortical regions. The model-recovery simulations (Extended Data Fig. 1a,b) suggest that the present dataset was suitable to make such inferences: Activity patterns in different cortical regions were sufficiently de-correlated, likely reflecting the use of a broad task battery. Thus, we were able to recover the correct model used to simulate the data, regardless of whether we assumed one-to-one connectivity or different degrees of convergence.

However, the simulations also indicated that the approach (and data) was not sufficient to determine the absolute degree of convergence with high confidence. For example, the size of the cortical input area for each cerebellar region differed substantially between the Ridge and Lasso regression models (Fig. 3 vs. Extended Data Fig. 5). Nonetheless, the two models result in similar predictive accuracy when using real data. Therefore, the true extent of the cortical input to a cerebellar

region likely lies somewhere between the extremes provided by the Ridge and Lasso models. Importantly, this ambiguity does not impact the core observation that the degree of convergence varies systematically across the cerebellar cortex. Whether using measures based on cortical surface area or dispersion, the general picture of variation in connectivity holds for both the Ridge and Lasso models, as well as when using cortical parcellations of different granularity. Thus, we are confident that the variation in convergence reflects a stable, method-independent characteristic of the cortico-cerebellar system.

A full understanding of how the cerebral cortex and cerebellum work as a functional network will also require a detailed analysis of cerebello-cortical connections. Our decision to focus here on modeling the connectivity from cortex to cerebellum comes, in large part, from studies of cerebellar blood flow in rodents. Increases in blood flow (a major contributor to the BOLD signal) are the result of increases in mossy fiber activity, one of the primary inputs to the cerebellar cortex[23]. In contrast, even a dramatic increase in the firing rate of Purkinje cells does not produce a measurable change in the BOLD signal from the cerebellar cortex[24,34]. Thus, the cerebellar BOLD signal provides a relatively clear image of the cortical inputs to the cerebellum, but not the output of the cerebellar cortex. As such, correlations in the fMRI signal between the neocortex and cerebellum should more strongly reflect the flow of information from the cortex to the cerebellum than the reverse. A complete analysis of the entire cortico-cerebellar circuit will therefore require methods such as viral tracing techniques[9,16] or functional activation measures that target key nodes in the ascending pathway (deep cerebellar nuclei, thalamus).

Future directions

Our approach was designed to identify the best task-general model of cortico-cerebellar connectivity. The connectivity model was able to take cortical activity patterns to make fairly accurate predictions of cerebellar activity, even when the data were obtained in a completely separate experiment from that used to build the model. This finding suggests that a substantial proportion of the cortico-cerebellar communication can be described by fixed information flow[21]; that is, each area of the cerebellum simply receives, independent of task, a copy of the neural activity occurring in the corresponding cortical areas.

However, the task-general model did not provide a perfect prediction of cerebellar activity. While these predictions may improve with more training data, we believe that there will likely be some systematic failures of the model. Such task-dependent deviations from the model prediction may offer important insights into cerebellar function, indicating that the cerebellum is more active than predicted by the cortical activity for some tasks and less for others. This pattern would suggest a "gating" of cortical input to the cerebellum, perhaps within the pontine nuclei - the main relay station of cortical inputs to the cerebellum. Rather than just transmitting cortical input to the cerebellar cortex, these nuclei may serve as an adaptive gate[35], amplifying or attenuating information as a function of the relative importance of cerebellar processing for the current task.

While such task-dependent gating implies that cortico-cerebellar connectivity may be better modeled using task-1s such, this approach provides a potent new tool to test hypotheses concerning cerebellar function.

Methods

Multi-domain task battery

To build models of cortico-cerebellar connectivity, we used the publicly available multi-domain task battery dataset (MDTB[8]). The MDTB includes fMRI data from two independent task sets (A & B). Each set consists of 17 tasks, eight of which were common to both sets. The 26 unique tasks were designed to sample activity during a broad range of task domains including motor (e.g., sequence production), working memory (e.g., 2-back task), language (e.g., reading), social (e.g., theory of mind), cognitive control (no-go, stroop), and emotion (e.g., facial expression) (see Supplemental Table 1 in King et al. [8]).

24 participants (16 females, 8 males, mean age = 23.8) were scanned during four sessions, with task set A employed in sessions 1 and 2 and task set B in sessions 3 and 4. Within a session, there were eight 10-minute runs, with each run composed of 17 blocks of 35 s each. Each block involved a unique task and included an initial 5 s instruction period. Most tasks contained multiple conditions (i.e., different levels of difficulty), resulting in a total of 47 unique task conditions.

Image acquisition and preprocessing

All fMRI and MRI data were collected on a 3T Siemens Prisma located at the Center for Functional and Metabolic Mapping at Western University, Canada. The protocol used the following parameters: 1 s repetition time; field-of-view measuring 20.8 cm; P-to-A phase encoding direction; 48 slices; 3 mm thickness; in-plane resolution 2.5×2.5 mm^2 . In order to localize and normalize the functional data, a high-resolution anatomical scan (T1-weighted MPRAGE, 1 mm isotropic resolution) of the whole brain was acquired during the first session.

Functional data were realigned for head motion artifacts within each session and different head positions across sessions using a six-parameter rigid body transformation. The mean functional image was co-registered to the anatomical image and this rigid-body transformation was applied to all functional images. No smoothing or group normalization was applied. SPM12 (www.fil.ion.ucl.ac.uk/ spm/doc/spm12_manual.pdf) and custom-written scripts in MATLAB were used to conduct data pre-processing.

General Linear Model (GLM)

To generate estimates of the activity related to each task condition, a general linear model (GLM) was fitted to the time series data of each voxel. This was done separately for each imaging run and task set (A & B). The 47 task conditions were modeled as separate regressors. Each of the 5-second task-specific instructions was modeled as a separate regressor in the GLM. For each task set, the beta weight for each task condition was averaged across the 8 runs within a session. Task set A was based on 92 activity estimates (2 sessions x (17 Instructions + 29 condition regressors)) and task set B was based on 98 activity estimates (2 sessions x (17 instructions + 32 conditions regressors)).

Cerebral Cortex Surface Reconstruction

For each of the 24 participants, the anatomical surfaces of the cortical hemispheres were reconstructed using the standard recon-all pipeline of the FreeSurfer package[39] (v. 5.0). The pipeline included brain extraction, generation of white and pial surfaces, inflation, and spherical alignment to the

symmetric fsLR-32k template[40]. Individual surfaces were re-sampled to this standard grid, resulting in surfaces with 32,492 vertices per hemisphere.

Spatial Normalization of Cerebellar Data

The cerebellum was isolated and normalized to the high-resolution Spatially Unbiased Infratentorial Template of the cerebellum using the SUIT toolbox[41]. This non-linear transformation was applied to both the anatomical and functional data. Task condition activity estimates (i.e., the beta weights) were resampled to a resolution of 3 mm isotropic and resliced into SUIT space. The cerebellar isolation mask was edited to remove voxels in the superior cerebellum that abutted voxels in the primary visual cortex. Functional images were masked with the cerebellar isolation mask resulting in activation signals that originate only from the cerebellar cortex. The cerebellar data were visualized using a surface-based representation of the cerebellar gray matter included in the SUIT toolbox[42]. This flat map is not intended to represent a true unfolding of the cerebellar cortex, but rather provides a convenient way to visualize volume-averaged cerebellar imaging data in a 2d representation.

Connectivity Models

Model training and testing was performed on fitted time series rather than the beta weights. To obtain the fitted time series (length T = 9568), we premultiplied the activity estimates with the first-level design matrix. This procedure reweights the activity estimates in an optimal manner, accounting for the fact that regressors for the instructions were based on 5 s of fMRI data, while regressors for the conditions were based on 10-30 s of data. This method also mitigates problems that may arise in event-related designs due to correlations of regressors[43].

We normalized the time series by dividing them by the standard deviation of the residuals from the first-level GLM. The time series were then averaged for all Q cortical parcels (see next section), resulting in a TxQ matrix ($\mathbf{X}$) of cortical features. For the cerebellum, we used all P=6937 cerebellar voxels (SUIT atlas space with 3 mm resolution), resulting in a TxP matrix ($\mathbf{Y}$).

Each model was trained to predict the cerebellar data $\mathbf{Y}$ from the cortical data $\mathbf{X}$, using a linear model, where $\mathbf{W}$ is the QxP matrix of connectivity weights, and $\mathbf{E}$ is the matrix of measurement errors.

$$\mathbf{Y} = \mathbf{XW} + \mathbf{E}$$

When estimating connectivity models, correlated fMRI noise across cortex and cerebellum (e.g., from head movement, physiological artifacts, etc) can lead to the erroneous detection of connectivity between these two structures. This is especially problematic for the superior cerebellum and the directly abutting regions of the occipital and inferior temporal lobes. To reduce the influence of these noise correlations, we used a "crossed" approach to train the models: The cerebellar time series for the first session was predicted by the cortical time series from the second session, and vice-versa (see Fig. 1).This procedure effectively negates the influence of noise processes, given that noise processes are uncorrelated across sessions.

Cortical parcels

We created a set of models that used different levels of granularity to subdivide the cortex. Specifically, each hemisphere was divided into a set of regular hexagonal parcels, using icosahedrons with 42, 162, 362, 642, or 1002 parcels per hemisphere (see[26] for details). Parcels that were completely contained within the medial wall were excluded, resulting in 80, 304, 670, 1190, or 1848 parcels

combined for the left and right hemisphere. Activity within a parcel was averaged across voxels for each condition, defining the Q regressors for each model.

The regular icosahedron parcellations are arbitrary in the sense that they do not align to functional boundaries in the human neocortex more than expected by chance[26]. We therefore also employed functionally-defined parcellations, repeating the main analyses using 12 different parcellations derived from resting state fMRI data (see Extended Data Fig. 3).

Model Estimation

We used three regression methods to estimate the connectivity matrix $\hat{\mathbf{W}}$ at the individual participant level. In the core analyses, each of these methods was combined with the five arbitrary cortical parcellations, resulting in 15 connectivity models. Each method was selected to favor a specific pattern of cortico-cerebellar connectivity. For the **winner-take-all** models, we assumed that the time series of each cerebellar voxel is explained by one and only one cortical region. To build this model, we simply chose the cortical region with the highest correlation with the cerebellar voxel in question. The connectivity weight corresponding to the rest of the cortical regions were set to 0.

The other two methods allowed for some degree of convergence in the cortical input to each cerebellar voxel. **Lasso regression** (least absolute shrinkage and selection operator[44]), seeks to explain activity in each cerebellar voxel by a restricted set of cortical features. Specifically, Lasso minimizes the squared-error, plus an additional penalty calculated as the sum of the absolute values (L1-norm) of the regression coefficients:

$$\hat{\mathbf{W}} = \underset{\mathbf{W}}{\operatorname{argmin}} \|\mathbf{Y} - \mathbf{X}\mathbf{W}\|_2^2 + \lambda\|\mathbf{W}\|_1$$

In contrast, **Ridge regression** penalizes large connectivity weights, attempting to find a broad set of cortical regions to explain each cerebellar voxel. Ridge regression minimizes the following loss function:

$$\hat{\mathbf{W}} = \underset{\mathbf{W}}{\operatorname{argmin}} \|\mathbf{Y} - \mathbf{X}\mathbf{W}\|_2^2 + \lambda\|\mathbf{W}\|_2^2$$

Both Lasso and Ridge models include a hyperparameter λ that must be tuned in model estimation. The value of these regularization parameters was determined using gridsearch with 4-fold cross-validation on the training set. We divided the conditions of the training set into four non-overlapping subsets, reconstructed the time series using tasks from three of the four subsets and evaluated the model using the left-out subset. Extended Data Fig. 7 depicts the average predictive accuracy for the training set and the optimal value of λ for the Lasso and Ridge models.

Model Testing

After model training with task set A, task set B was used as the test set. We applied the same procedure to construct the **X** and **Y** matrices. We then used the estimated weights for each model to generate predicted voxel-wise cerebellar activity patterns from the observed cortical time series. These predictions were generated at each level of granularity for each participant. Model performance was measured by correlating predicted and observed cerebellar time series for each voxel. As with the procedure for model training, the cortical and cerebellar time series were crossed: Cortical time series

from session 3 were used to predict cerebellar time series from session 4 and vice versa. Model predictions were calculated separately for each cerebellar voxel and visualized on voxelwise maps of the cerebellar cortex.

In sum, we evaluated 15 models in which the cortical parcels were based on an arbitrary icosahedron, based on the combination of three regression methods and five levels of granularity. For the 12 functionally-defined parcellations (Extended Data Fig. 3), we limited the evaluation to the WTA and Ridge models.

Noise Ceiling

For each model evaluated in the current study, the noise ceiling quantifies the expected performance of each model, under the assumption that the estimated weights reflect the true connectivity between the cerebral cortex and cerebellum. When predicting the cerebellar activity patterns, our models are limited by two factors: Measurement noise associated with the cerebellar data and measurement noise associated with the cortical data.

To estimate these sources of noise, we used the fitted cerebellar time-series for the two sessions, assuming that each session's data are composed of the true time series $\mathbf{Y}^*$ plus noise $\boldsymbol{\epsilon}$: $\mathbf{Y}_i = \mathbf{Y}^* + \boldsymbol{\epsilon}$. If σ_Y^2 signifies the variance of the true time series , and σ_ϵ^2 the variance of the noise, then the correlation between the two sessions, i.e. the split-half reliability, is:

$$r(\mathbf{Y}_1, \mathbf{Y}_2) = \frac{\sigma_Y^2}{\sigma_Y^2 + \sigma_\epsilon^2}.$$

Similarly, the reliability of the prediction, even if we knew the true connectivity weights $\mathbf{W}$, is limited by the noise on the cortical data $\mathbf{X}$. The prediction $\mathbf{X}_i\mathbf{W}$ has noise variance σ_{XW}^2, which can be estimated by calculating the split-half reliability across the two sessions:

$$r(\mathbf{X}_1\mathbf{W}, \mathbf{X}_2\mathbf{W}) = \frac{\sigma_Y^2}{\sigma_Y^2 + \sigma_{XW}^2}.$$

Thus, for the true model the expected correlation would be:

$$r_{ceil} = r(\mathbf{X}_i\mathbf{W}, \mathbf{Y}_i) = \frac{\sigma_Y^2}{\sqrt{\sigma_Y^2 + \sigma_{XW}^2}\sqrt{\sigma_Y^2 + \sigma_\epsilon^2}} = \sqrt{r(\mathbf{Y}_1, \mathbf{Y}_2)r(\mathbf{X}_1\mathbf{W}, \mathbf{X}_2\mathbf{W})}$$

Because we do not know the true weights, we use the estimated weights from each model to estimate the noise ceiling. Thus, the noise ceiling is model dependent, specifying what the correlation would be if the current model was the true model.

Cortico-Cerebellar Convergence

Two measures were used to assess the amount of convergence of cortical inputs to each cerebellar voxel. For the first measure, we calculated the percentage of the cortical surface contributing to the prediction in each cerebellar voxel. We used the estimated weights from the Lasso regression model with the best unnormalized performance (80 cortical parcels). We opted to focus on the Lasso model since it sets weights that make a negligible contribution to the predictions to zero, thus identifying the sparsest cortical input that explains the cerebellar data. To determine the likely input area, we

determined the number of non-zero coefficients for each cerebellar voxel, and expressed this as a percentage of the cortical surface. We also repeated these analyses using Ridge regression (Extended Data Fig. S5). Because Ridge does not result in zero coefficients, we applied a threshold, counting the number of connectivity weights with a value one standard above the mean of the weight matrix.

For the second measure, we calculated the dispersion (spherical variance) of the input weights. Using the spherical representation of the cerebral hemispheres from FreeSurfer, we defined unit vectors $\mathbf{v}_i$, pointing from the center of the sphere to the center of each of the cortical parcels. Each vector was then weighted by the cortico-cerebellar connectivity weight for that region w_i. All negative connectivity weights were set to zero. The spherical variance for hemisphere h across all Q_h parcels of this hemisphere was defined as:

$$var_h = 1 - \left| \frac{1}{Q_h} \sum_{Q_h} w_i \mathbf{v}_i \right|.$$

To obtain a composite measure, we averaged the variances for the two hemispheres, with each weighted by the size of the summed (non-negative) connectivity weights. For summary and statistical testing, we then averaged the size of the input area, as well as the dispersion, across cerebellar voxels within each of the 10 functionally defined MDTB regions.

Generalization to new participants

Although the models were trained and tested on distinct datasets with more than half the tasks unique to each set, the data for training and testing were obtained from the same set of participants. As a stronger test of generalization, we evaluated the models with data from a new experiment involving naive participants (n=20, 11 females, 9 males, mean age=21.3). These participants were trained to perform a 5-task battery (cognitive control, social prediction, action prediction, visual search, semantic prediction). For each task, there was an easy and hard condition (e.g., 2-back vs 0-back for cognitive control, high vs low predictability), along with the task's associated instruction. Thus, there were a total of 16 conditions when including rest (fixation). Each participant completed five sessions, with fMRI data obtained in the first, third, and fifth sessions. Within a scanning session, there were six 11-minute runs, and each run included three repetitions of each task (35 s) with a 10 s rest period prior to each 5 s instruction period.

The participants were scanned on a Siemens MAGNETOM TrioTim syngo MR B17 located in the Henry Wheeler Jr. Brain Imaging Center at the University of California, Berkeley. The protocol used the following parameters: 1 s repetition time; field-of-view measuring 20.8 cm; A-to-P phase encoding direction; 48 slices; 3 mm thickness; in-plane resolution 2.5×2.5 mm^2. To localize and normalize the functional data, a high-resolution, whole-brain anatomical scan (T1-weighted MPRAGE, 1mm isotropic resolution) was acquired during the first session. fMRIPrep[45,46](https://fmriprep.org/en/stable/) was used to preprocess the anatomical and functional data, following the same analysis procedures as conducted for the main experiment.

To generate estimates of activity (beta weights), a General Linear Model (GLM) was fitted for each voxel to the time series data. Each of the 16 conditions was modeled as a separate regressor, with this repeated separately for each imaging run. Nipype (https://nipype.readthedocs.io/en/latest/) and custom-written scripts in Python were used to estimate the beta weights from a first-level GLM.

To evaluate model generalization to this new dataset, we used the full set of cortico-cerebellar connectivity models estimated with the original data set, limiting the analysis to the regular cortical parcellations. Since the participants were different from the original data set, we averaged the model weights for the original 24 participants (i.e., 3 model types x 5 levels of granularity). Each model was then used to predict cerebellar activity from the cortical activity obtained in the new study, with these predictions compared to actual cerebellar activity at the individual level. Given that the experiment was a training study, we expected that both the true activation patterns, as well as the signal-to-noise ratio would change across sessions. We therefore determined the noise ceiling for each session separately by calculating the reliability across odd and even runs within each session. The normalized predictive accuracy was determined separately for each of the three imaging sessions, with the results across sessions averaged.

References

1. Petersen, S. E., Fox, P. T., Posner, M. I., Mintun, M. & Raichle, M. E. Positron emission tomographic studies of the processing of singe words. *J. Cogn. Neurosci.* **1**, 153–170 (1989).
2. Allen, G., Buxton, R. B., Wong, E. C. & Courchesne, E. Attentional activation of the cerebellum independent of motor involvement. *Science* **275**, 1940–1943 (1997).
3. Bischoff-Grethe, A., Ivry, R. B. & Grafton, S. T. Cerebellar involvement in response reassignment rather than attention. *J. Neurosci.* **22**, 546–553 (2002).
4. Balsters, J. H., Whelan, C. D., Robertson, I. H. & Ramnani, N. Cerebellum and cognition: evidence for the encoding of higher order rules. *Cereb. Cortex* **23**, 1433–1443 (2013).
5. Marvel, C. L. & Desmond, J. E. Functional topography of the cerebellum in verbal working memory. *Neuropsychol. Rev.* **20**, 271–279 (2010).
6. Fiez, J. A. The cerebellum and language: Persistent themes and findings. *Brain Lang.* **161**, 1–3 (2016).
7. Van Overwalle, F., D'aes, T. & Mariën, P. Social cognition and the cerebellum: A meta-analytic connectivity analysis. *Hum. Brain Mapp.* **36**, 5137–5154 (2015).
8. King, M., Hernandez-Castillo, C. R., Poldrack, R. A., Ivry, R. B. & Diedrichsen, J. Functional boundaries in the human cerebellum revealed by a multi-domain task battery. *Nat. Neurosci.* **22**, 1371–1378 (2019).
9. Kelly, R. M. & Strick, P. L. Cerebellar loops with motor cortex and prefrontal cortex of a nonhuman primate. *J. Neurosci.* **23**, 8432–8444 (2003).
10. Dum, R. P. & Strick, P. L. An Unfolded Map of the Cerebellar Dentate Nucleus and its Projections to the Cerebral Cortex. *J. Neurophysiol.* **89**, 634–639 (2003).
11. Strick, P. L., Dum, R. P. & Fiez, J. A. Cerebellum and nonmotor function. *Annu. Rev. Neurosci.* **32**, 413–434 (2009).
12. Buckner, R. L., Krienen, F. M., Castellanos, A., Diaz, J. C. & Yeo, B. T. T. The organization of the human cerebellum estimated by intrinsic functional connectivity. *J. Neurophysiol.* **106**, 2322–2345 (2011).
13. Ji, J. L. *et al.* Mapping the human brain's cortical-subcortical functional network organization. *Neuroimage* **185**, 35–57 (2019).
14. Marek, S. *et al.* Spatial and Temporal Organization of the Individual Human Cerebellum. *Neuron* **100**, 977–993.e7 (2018).
15. Henschke, J. U. & Pakan, J. M. P. Disynaptic cerebrocerebellar pathways originating from multiple functionally distinct cortical areas. *Elife* **9**, e59148 (2020).
16. Pisano, T. J. *et al.* Homologous organization of cerebellar pathways to sensory, motor, and

associative forebrain. *Cell Rep.* **36**, 109721 (2021).

17. Huang, C.-C. *et al.* Convergence of pontine and proprioceptive streams onto multimodal cerebellar granule cells. *Elife* **2**, e00400 (2013).

18. Yeo, B. T. T. *et al.* Functional Specialization and Flexibility in Human Association Cortex. *Cereb. Cortex* **25**, 3654–3672 (2015).

19. Yeo, B. T. T., Krienen, F. M., Chee, M. W. L. & Buckner, R. L. Estimates of segregation and overlap of functional connectivity networks in the human cerebral cortex. *Neuroimage* **88**, 212–227 (2014).

20. Bertolero, M. A., Yeo, B. T. T. & D'Esposito, M. The modular and integrative functional architecture of the human brain. *Proc. Natl. Acad. Sci. U. S. A.* **112**, E6798–807 (2015).

21. Cole, M. W., Ito, T., Bassett, D. S. & Schultz, D. H. Activity flow over resting-state networks shapes cognitive task activations. *Nat. Neurosci.* **19**, 1718–1726 (2016).

22. Mell, M. M., St-Yves, G. & Naselaris, T. Voxel-to-voxel predictive models reveal unexpected structure in unexplained variance. *Neuroimage* **238**, 118266 (2021).

23. Mathiesen, C., Caesar, K. & Lauritzen, M. Temporal coupling between neuronal activity and blood flow in rat cerebellar cortex as indicated by field potential analysis. *J. Physiol.* **523 Pt 1**, 235–246 (2000).

24. Thomsen, K., Offenhauser, N. & Lauritzen, M. Principal neuron spiking: neither necessary nor sufficient for cerebral blood flow in rat cerebellum. *J. Physiol.* **560**, 181–189 (2004).

25. Yeo, B. T. T. *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* **106**, 1125–1165 (2011).

26. Zhi, D., King, M., Hernandez-Castillo, C. R. & Diedrichsen, J. Evaluating brain parcellations using the distance-controlled boundary coefficient. *Hum. Brain Mapp.* (2022) doi:10.1002/hbm.25878.

27. Krienen, F. M. & Buckner, R. L. Segregated fronto-cerebellar circuits revealed by intrinsic functional connectivity. *Cereb. Cortex* **19**, 2485–2497 (2009).

28. O'Reilly, J. X., Beckmann, C. F., Tomassini, V., Ramnani, N. & Johansen-Berg, H. Distinct and overlapping functional zones in the cerebellum defined by resting state functional connectivity. *Cereb. Cortex* **20**, 953–965 (2010).

29. Guell, X., Schmahmann, J. D., Gabrieli, J. D. E. & Ghosh, S. S. Functional gradients of the cerebellum. *Elife* **7**, e36652 (2018).

30. Balsters, J. H. *et al.* Evolution of the cerebellar cortex: the selective expansion of prefrontal-projecting cerebellar lobules. *Neuroimage* **49**, 2045–2052 (2010).

31. Diedrichsen, J., King, M., Hernandez-Castillo, C., Sereno, M. & Ivry, R. B. Universal Transform or Multiple Functionality? Understanding the Contribution of the Human Cerebellum across Task Domains. *Neuron* **102**, 918–928 (2019).

32. De Zeeuw, C. I., Lisberger, S. G. & Raymond, J. L. Diversity and dynamism in the cerebellum. *Nat. Neurosci.* **24**, 160–167 (2021).

33. Tavor, I. *et al.* Task-free MRI predicts individual differences in brain activity during task performance. *Science* **352**, 216–220 (2016).

34. Thomsen, K., Piilgaard, H., Gjedde, A., Bonvento, G. & Lauritzen, M. Principal cell spiking, postsynaptic excitation, and oxygen consumption in the rat cerebellar cortex. *J. Neurophysiol.* **102**, 1503–1512 (2009).

35. Schwarz, C. & Thier, P. Binding of signals relevant for action: towards a hypothesis of the functional role of the pontine nuclei. *Trends Neurosci.* **22**, 443–451 (1999).

36. Cole, M. W., Bassett, D. S., Power, J. D., Braver, T. S. & Petersen, S. E. Intrinsic and task-evoked network architectures of the human brain. *Neuron* **83**, 238–251 (2014).

37. Braun, U. *et al.* Dynamic reconfiguration of frontal brain networks during executive cognition in humans. *Proc. Natl. Acad. Sci. U. S. A.* **112**, 11678–11683 (2015).

38. Allen, E. A. *et al.* Tracking Whole-Brain Connectivity Dynamics in the Resting State. *Cereb. Cortex* **24**, 663–676 (2012).
39. Fischl, B. FreeSurfer. *Neuroimage* **62**, 774–781 (2012).
40. Van Essen, D. C., Glasser, M. F., Dierker, D. L., Harwell, J. & Coalson, T. Parcellations and Hemispheric Asymmetries of Human Cerebral Cortex Analyzed on Surface-Based Atlases. *Cereb. Cortex* **22**, 2241–2262 (2011).
41. Diedrichsen, J. A spatially unbiased atlas template of the human cerebellum. *Neuroimage* **33**, 127–138 (2006).
42. Diedrichsen, J. & Zotow, E. Surface-Based Display of Volume-Averaged Cerebellar Imaging Data. *PLoS One* **10**, e0133402 (2015).
43. Mumford, J. A., Turner, B. O., Ashby, F. G. & Poldrack, R. A. Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. *Neuroimage* **59**, 2636–2643 (2012).
44. Tibshirani, R. Regression Shrinkage and Selection via the Lasso. *J. R. Stat. Soc. Series B Stat. Methodol.* **58**, 267–288 (1996).
45. Esteban, O. *et al.* fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat. Methods* **16**, 111–116 (2019).
46. Esteban, O. *et al.* Analysis of task-based functional MRI data preprocessed with fMRIPrep. *Nat. Protoc.* **15**, 2186–2202 (2020).
47. Schaefer, A. *et al.* Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cereb. Cortex* **28**, 3095–3114 (2017).
48. Arslan, S., Parisot, S. & Rueckert, D. Joint Spectral Decomposition for the Parcellation of the Human Cerebral Cortex Using Resting-State fMRI. in *Information Processing in Medical Imaging* 85–97 (Springer International Publishing, 2015). doi:10.1007/978-3-319-19992-4_7.
49. Fan, L. *et al.* The Human Brainnetome Atlas: A New Brain Atlas Based on Connectional Architecture. *Cereb. Cortex* **26**, 3508–3526 (2016).
50. Gordon, E. M. *et al.* Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. *Cereb. Cortex* **26**, 288–303 (2016).
51. Shen, X., Tokoglu, F., Papademetris, X. & Constable, R. T. Groupwise whole-brain parcellation from resting-state fMRI data for network node identification. *Neuroimage* **82**, 403–415 (2013).
52. Zhi, D., King, M., Hernandez-Castillo, C. R. & Diedrichsen, J. Evaluating brain parcellations using the distance controlled boundary coefficient. *bioRxiv* 2021.05.11.443151 (2021) doi:10.1101/2021.05.11.443151.