# Reconfigurations of cortical manifold structure during reward-based motor learning

Qasem Nick[1,2], Daniel J. Gale[1], Corson Areshenkoff[1,2], Anouk De Brouwer[1], Joseph Nashed[1], Jeffrey Wammes[1,2], Randy Flanagan[1,2], Jonny Smallwood[1,2] & Jason Gallivan*[1,2,3]

[1]Centre for Neuroscience Studies, [2]Department of Psychology, and [3]Department of Biomedical and Molecular Sciences, Queen's University, Kingston, Ontario, Canada. *Correspondence: Jason Gallivan, PhD; Department of Psychology, and Department of Biomedical and Molecular Sciences, Queen's University, gallivan@queensu.ca

## Abstract

Adaptive motor behavior depends on the coordinated activity of multiple neural systems distributed across cortex and subcortex. While the role of sensorimotor cortex in motor learning has been well-established, how higher-order brain systems interact with sensorimotor cortex to guide learning is less well understood. Using functional MRI, we examined human brain activity during a reward-based motor task where subjects learned to shape their hand trajectories through reinforcement feedback. We projected patterns of cortical and subcortical functional connectivity onto a low-dimensional manifold space and examined how regions expanded and contracted along the manifold during learning. During early learning, we found that several sensorimotor areas in the Dorsal Attention Network exhibited increased covariance with areas of the salience/ventral attention network and reduced covariance with areas of the default mode network (DMN). During late learning, these effects reversed, with sensorimotor areas now exhibiting increased covariance with DMN areas. However, areas in posteromedial cortex showed the opposite pattern across learning phases, with its connectivity suggesting a role in coordinating activity across different networks over time. Our results establish the whole-brain neural changes that support reward-based motor learning, and identify distinct transitions in the functional coupling of sensorimotor to transmodal cortex when adapting behavior.

**Keywords:** Motor learning, reward-guided learning, manifold, brain connectivity, motor control, reaching

## INTRODUCTION

Organizing our behaviors so that they match the demands of a given situation depends on establishing contingencies between specific features of an action and whether they lead to the desired outcome. In many real-world tasks this is a challenging endeavor, as the brain must learn how to modify its actions based on a single measure of performance feedback that reflects overall task success [1–4]. Moreover, this feedback must be communicated to multiple discrete neural systems distributed across cortex and subcortex, many of which are topographically segregated from one another [5]. Numerous studies have shown that neural systems anchored within the medial prefrontal cortex and striatum are important in evaluating whether the results of behavior are in line with expectations [5–8]. Specifically, when discrepancies arise between the expected versus actual results of an action — termed a "prediction error" — this information serves as the teaching signal that can be used to update behavior directly [9–12]. How exactly this information is communicated in a coherent manner across the multiple, distributed neural systems that guide behavior remains poorly understood. Our study addresses this gap in our knowledge using state-of-the-art manifold learning techniques to describe how the landscape of brain activity changes during reward-guided motor learning.

Contemporary systems and cognitive neuroscience has identified many large-scale neural systems that have each been linked to different components of effective behavior. For example, areas in sensory cortex provide representations of the external environment, whereas areas in motor cortex are involved in generating the final motor commands required for action. At higher levels in the cortical hierarchy, regions within the frontoparietal system, along with those that make up the brain's attention-orienting systems (dorsal and ventral attention systems), are important in the selection of sensory inputs and the guidance of rule-driven behavior [13,14]. Yet, how the activity of these various brain systems is coordinated during the learning process is unclear. An emerging literature suggests that this coordination may depend, in part, on functional activity in several key regions of higher-order association cortex, known collectively as the default mode network (DMN) [15–19].

Initially identified through its tendency to deactivate during cognitively demanding tasks, the DMN has traditionally been implicated in largely introspective, abstract cognitive functions such as autobiographical memory and internal mentation [20–24]. In recent years, however, this characterization of DMN activity has been difficult to reconcile with an emerging body of fMRI and neurophysiological evidence showing that areas of this network are activated during demanding

decision-making and working-memory tasks [25–31]. One hypothesis concerning this system's function pertains to its unique topographic positioning on the cortical mantle [17]: Each core region of the DMN is located in regions of association cortex that are equidistant between different primary systems; e.g., posteromedial cortex, a key node of the DMN, is located precisely at the midpoint between the calcarine (visual) and central (motor) sulci [19]. This unique topographic location is hypothesized to allow DMN regions broad oversight over distributed brain functions, enabling them to play a role in the coordination of activity across cortex [17].

Consistent with this contemporary perspective, recent work implicates several regions of the DMN in organizing different modes of behavior over time. For instance, DMN areas such as medial frontal cortex and posteromedial cortex, appear to play an important role in shifting between information gathering versus information exploitation during reward-guided decision-making tasks [28,31–35] — the so-called explore/exploit trade-off [36,37]. Consistent with this, recent studies have argued that broad features of the DMN's activity can be explained under the auspices that it supports behavior under conditions in which performance depends on knowledge accrued across several trials rather than by immediate sensory inputs [25–27,38]. Extending these ideas, we recently showed that the DMN plays a role in motor adaptation, showing that connectivity between DMN regions and sensorimotor cortex are altered when normal visual-motor contingencies governing behavior are interrupted over the course of several trials [39]. Given that the DMN is hypothesized to exert influence on functional brain activity via its topographic positioning on cortex [17], understanding how the DMN supports task behavior likely requires analytical techniques that allow for a characterization of whole-brain changes in functional architecture.

In the current study, we explored changes in the landscape of cortical and subcortical activity during a reward-based motor task in which human participants learned to produce, through purely reinforcement feedback, a specific movement trajectory that was initially unknown to them. To characterize learning-related changes at the neural level, we leveraged advanced manifold learning approaches that provide a low-dimensional description of whole-brain activity [19,40,41]. This approach builds on recent electrophysiological studies in macaques demonstrating that high-dimensional neural population activity can be described along a low-dimensional subspace or manifold [42–45], reflecting covariance patterns across the entire population. This same organizational structure also appears to govern the macroscale activity of cortex, with this manifold approach having recently provided key insights into the overarching structural and functional architecture of the human brain [41,46–49]. Here we applied this manifold approach to understand how brain activity across widely

distributed cortical and subcortical systems is coordinated during reward-based motor learning. We were particularly interested in understanding how connectivity between regions within the DMN and the rest of the brain changes as participants shift from learning the relationship between motor commands and reward feedback, during early learning, to subsequently using this information, during late learning. We were also interested in examining whether learning-dependent changes in manifold structure relate to variation in subject motor performance.

## RESULTS

Prior studies examining the neural processes underlying reward-based learning have typically used tasks requiring simple motor responses, such as button presses or lever movements [5–8,50,51]. The simplicity of these motor responses is intended to isolate participants' choice behavior by eliminating any variability related to movement execution (i.e., motor implementation of the choice) as a potential confounding factor in the learning process [52,53]. However, recent theories on learning, supported by both human and animal studies [1,54,55], have highlighted the crucial role of movement variability as a key ingredient for effective learning [55–57]. In order to incorporate this aspect to learning, and inspired by recent work in the field [55,57], we developed an MRI-compatible reward-based motor task in which human participants (N=36) learned to shape their hand trajectories purely through reinforcement feedback.

In this task, subjects used their right finger on an MRI-compatible touchpad to trace, without visual feedback of their finger, a rightward-curved path displayed on a screen (Fig. 1A,B). Participants began the MRI study by performing a *Baseline* block of 70 trials, wherein they did not receive any feedback about their performance. Following this, subjects began a separate *Learning* block of 200 trials in which they were told that they would now receive score feedback (from 0 to 100 points), presented at the end of each trial, based on how accurately they traced the visual path displayed on the screen. However, unbeknownst to subjects, the score they actually received was based on how well they traced a *hidden* mirror-image path (the 'reward' path, which was reflected across the vertical axis; Fig. 1C). Importantly, because subjects received no visual feedback about their actual finger trajectory and could not see their own hand, they could only use the score feedback — and thus only reward-based learning mechanisms — to modify their movements from one trial to the next [55,57].

Fig. 1C shows an example of a single subject's finger trajectories across trials. Initially, the subject begins by tracing the visual path displayed on the screen (as instructed), albeit with some expected motor noise due to the absence of any visual feedback about their finger paths (see cyan trajectories). However, over time, the subject learns to gradually trace a path more similar to the rewarded, mirror-image path (dark pink trajectories). As can be seen in Fig. 1D, subjects on average were able to use the reward-based feedback to increase their score, and thus produce a trajectory more similar to the hidden path, over the 200 learning trials.

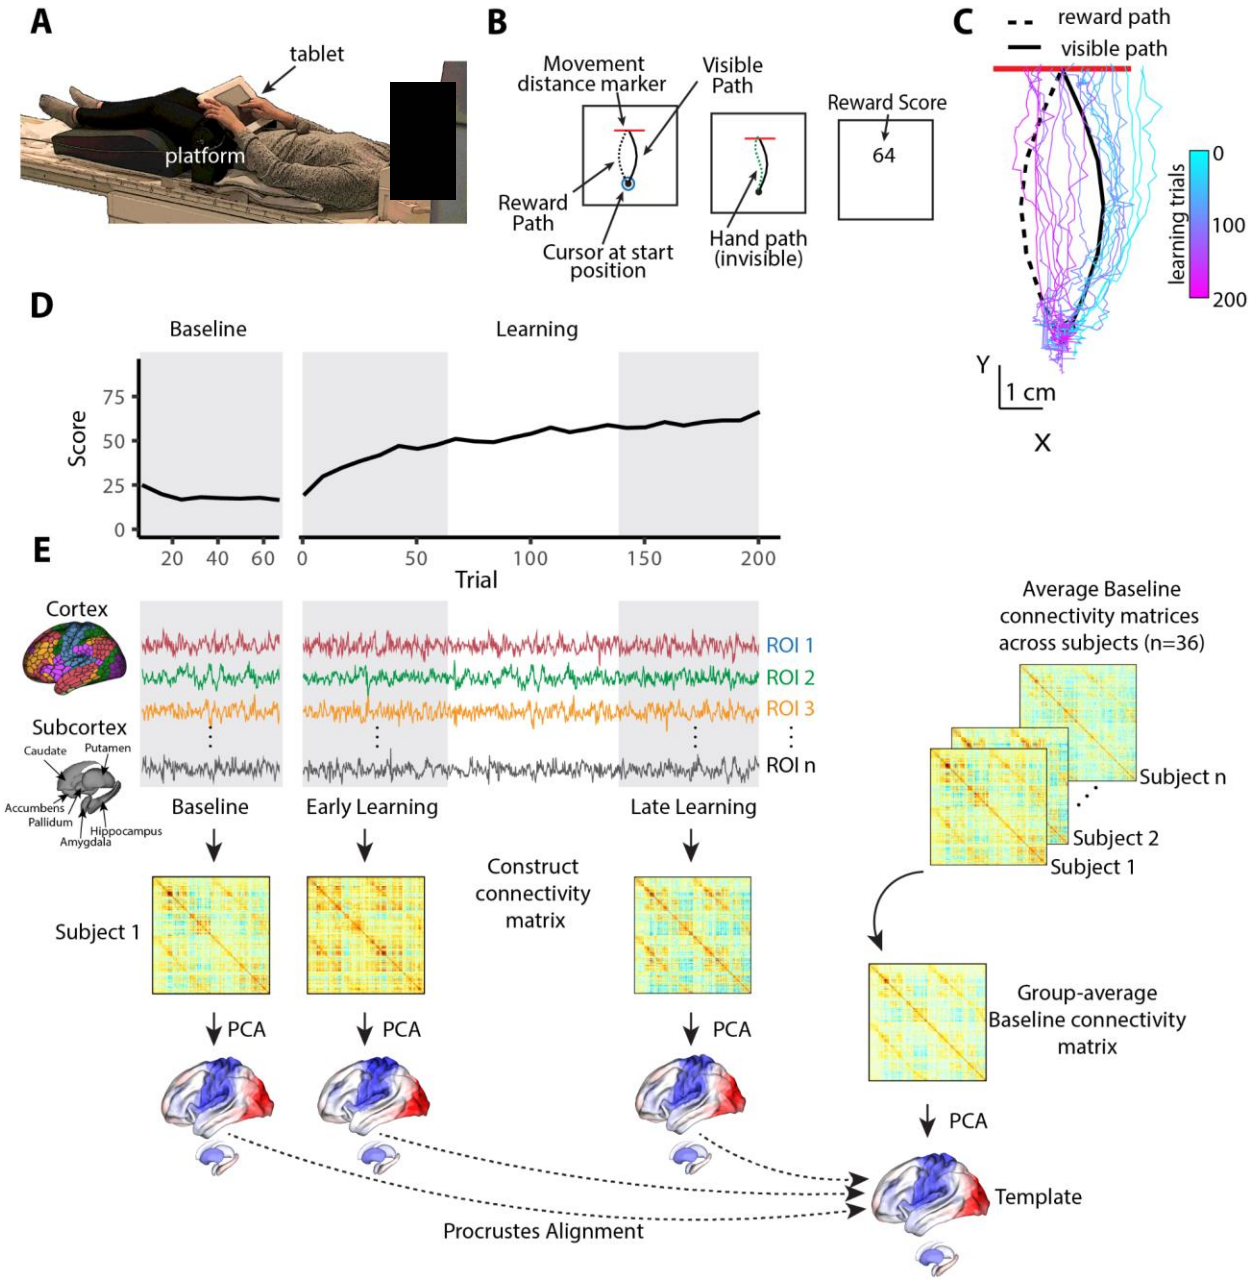

**Figure 1: Task structure and overview of fMRI analysis.** (A) Subject setup in the MRI scanner. (B) Trial structure of the reward-based motor learning task. On each trial, subjects were required to trace a curved (Visible) path from a start location to a target line (in red), without visual feedback of their finger location. Following a baseline block of trials, subjects were instructed that they would receive score feedback, presented at the end of the trial, based on their accuracy in tracing the visible path. However, unbeknownst to subjects, the score they received was actually based on how accurately they traced the mirror-image path (Reward Path), which was invisible to participants. (C) Example subject data from learning trials in the task. Coloured traces show individual trials over time (each trace is separated by ten trials to give a sense of the trajectory changes over time; 20 trials shown in total). (D) Average participant performance throughout the learning task. Three equal-length task epochs for subsequent neural analyses are indicated by the gray shaded boxes. (E) Neural analysis approach. For each participant and each task epoch (Baseline, Early and Late learning), we estimated functional connectivity matrices using region-wise timeseries extracted from the Schaefer 1000 cortical parcellation and the Harvard-Oxford subcortical parcellation. We estimated functional connectivity manifolds for each task epoch using PCA with centered and thresholded connectivity matrices (see Methods, as well as Fig. 2). All manifolds were aligned to a common template manifold created from a group-average Baseline connectivity matrix (right) using Proscrustes alignment. This allowed us to assess learning-related changes in manifold structure from this Baseline architecture.

In order to study the changes in functional cortical and subcortical organization during the learning task, we used three distinct, equal-length epochs over the time course of the study. Specifically, in addition to the task Baseline epoch (70 trials), we defined Early and Late learning epochs as the subsequent initial and final 70 trials, respectively, following the presentation onset of reward feedback. For each participant, we extracted mean blood oxygenation level-dependent (BOLD) timeseries data for each cortical region defined by the Shaefer 1000 cortical parcellation [58] and for striatal regions defined by the Harvard-Oxford parcellation [59,60], and then estimated covariance (functional connectivity) matrices for each epoch (Baseline, Early and Late; Fig. 1E)[for a similar approach, see 39].

Because prior work [61,62], including our own [39,63,64], suggests that individual differences in functional connectivity can obscure any task-related effects, we centered the connectivity matrices using the Riemmanian manifold approach [39,63–65](See Supplemental Figure 1 for an overview of the approach). To illustrate the effects of this centering procedure, and why it is important for elucidating task-related effects in the data, we projected participants' covariance matrices both prior to, and after the centering procedure, using uniform manifold approximation (UMAP; [66]). As shown in Fig. 2A, prior to the centering procedure the covariance matrices mainly cluster according to subject identity, consistent with prior findings showing that this subject-level structure explains the majority of the variance in functional connectivity data [61]. Clearly, this subject-level clustering could impact the ability to detect task-related effects in the data. However, after applying the centering procedure (Fig.

2B), this subject-level clustering is abolished, potentially allowing for the differentiation of the three task-related epochs.

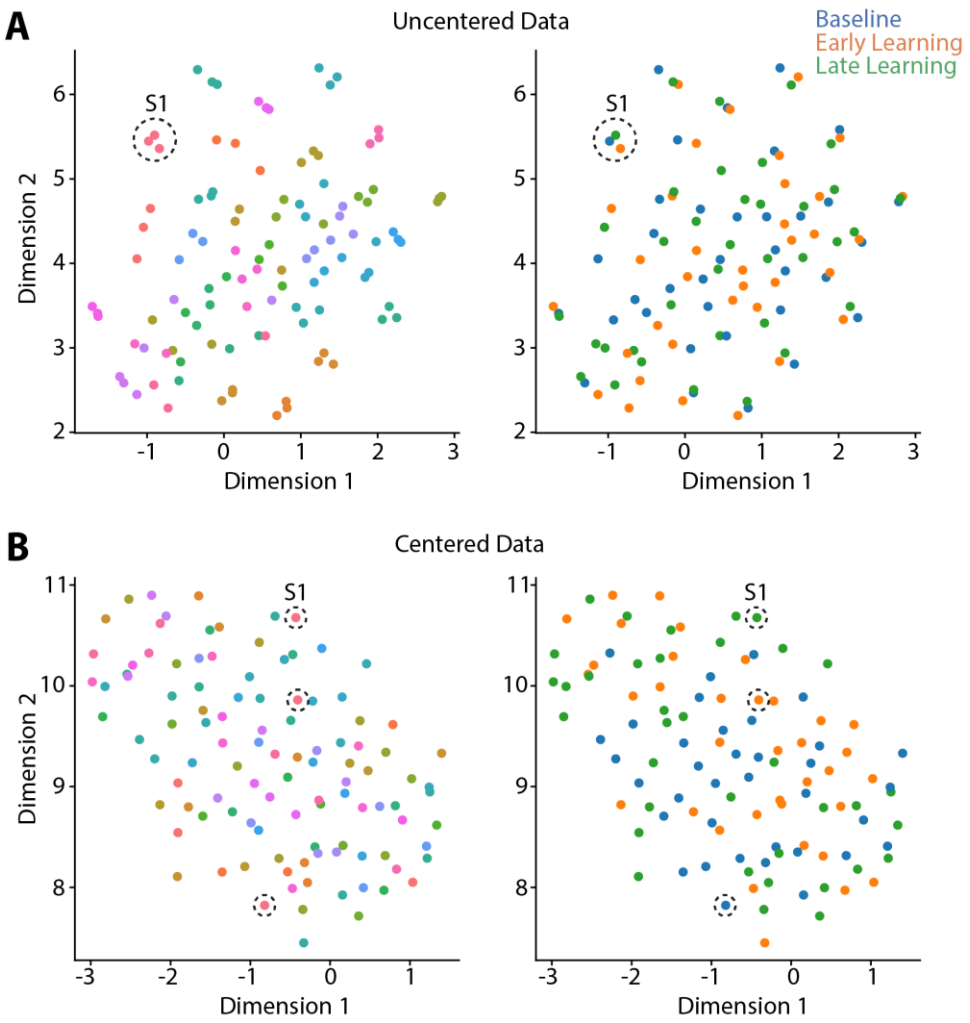


**Figure 2: Riemmanian centering removes subject-level clustering.** UMAP visualization of the similarity of connectivity matrices, both before centering (A) and after (B) centering. In these plots, each point represents a single functional connectivity matrix, color-coded either to subject identity (left panels) or task epoch (right panels), with its location in the multidimensional space based on the similarity between matrices. Note that the uncentered connectivity matrices in A show a high-degree of subject-level clustering, thus obscuring any differences in task structure. By contrast, the Riemmanian manifold centering approach (in B) abolishes this subject-level clustering. To help illustrate this point, the dashed circles in both A and B indicate the functional connectivity matrices belonging to the same single subject (Subject 1; S1).

To examine changes in cortical and subcortical connectivity during the reward-based motor learning task, we used the centered matrices from Fig. 2B to estimate separate cortical-subcortical connectivity manifolds for each participant's Baseline, Early and Late covariance matrices [see also 39]. Following from prior work [40,49,67], we transformed each matrix into an affinity matrix and then applied Principal Components Analysis (PCA) to obtain a set of principal components (PCs) that provides a low-dimensional representation of cortical-subcortical functional organization (i.e., a cortical-subcortical manifold). Next, we aligned the manifolds from each participant to a template Baseline manifold, which we constructed using the mean of all Baseline connectivity matrices across participants (Fig. 1E). We did this for two reasons: (1) the Baseline manifold provided a common target for manifold alignment [40] so that all subjects could be directly compared in a common task-based neural space, and (2) it allowed us to selectively detect deviations *from this* Baseline manifold architecture; i.e., observe the changes to this manifold structure that occur as a function of learning during the task (when subjects begin receiving reward feedback about their performance).

***Cortical-subcortical manifold structure during Baseline trials***

The top three principal components (PCs) of the template Baseline manifold (Fig. 3A) describe the cortical-subcortical functional organization during Baseline trials. As can be seen in Fig. 3A, PC1 distinguishes visual regions (positive loadings in red) from somatomotor regions (negative loadings in blue). Meanwhile, PC2 distinguishes visual and somatomotor regions (in red) from the remaining cortical areas (in blue), most prominently high-order association regions within the default mode network (DMN). Finally, PC3 mainly constitutes a gradient of frontoparietal areas of the Dorsal Attention Network (DAN) and Frontoparietal Control Network (FCN) versus DMN regions. Collectively, these top three PCs explain ~70% of the total variance (Fig. 3B).

When we mapped the brain regions onto their assigned intrinsic functional network architecture [68], we confirmed that PCs 1 and 2 jointly differentiate visual, DMN and somatomotor regions, replicating the tripartite structure of the brain's intrinsic functional architecture [19,41](Fig. 3D). Others have argued that this tripartite structure is a fundamental feature of functional brain organization, whereby the transition from unimodal cortex (visual and somatomotor networks) to transmodal cortex (the DMN) reflects a global processing hierarchy from lower- to higher-order brain systems [17,19,41].

We next sought to characterize the relative positions of cortical and subcortical brain regions along the Baseline connectivity-derived manifold space, thus providing a basis to examine future changes in the positioning of these regions during Early and Late learning. To this aim, and following from previous methods [39,69], we computed the manifold eccentricity of each region by taking its Euclidean distance from manifold centroid (coordinates (0,0,0), see Fig. 3E). This eccentricity measure provides a multivariate index of each brain region's embedding in the three-dimensional manifold space, whereby distal regions located at the extremes of the manifold have greater eccentricity than proximal regions located near the manifold center (Fig. 3E). Under this framework, regions with higher eccentricity are interpreted as having higher functional segregation from other networks in the rest of the brain, whereas regions with lower eccentricity are interpreted as having higher integration (lower segregation) with other networks in the rest of the brain [69–71]. Consistent with this interpretation, we find that our eccentricity measure strongly correlates with various graph theoretical measures of integration and segregation. For instance, we find that Baseline eccentricity is positively related to cortical node strength ($r=0.88$, two-tailed $p<0.001$) and within-manifold degree z-score ($r=0.45$, two-tailed $p<0.001$), consistent with the notion that more eccentric regions are more strongly functionally coupled with other members of the the same functional network (i.e., higher segregation, see Supplementary Fig 2). Likewise, we find that Baseline eccentricity is negatively related to a region's participation coefficient ($r=-0.74$, two-tailed $p<0.001$), which is a measure of a region's degree of cross-network integration. Thus, taken together, changes in a brain region's eccentricity can provide us with a multivariate measure of changes in that region's functional segregation versus integration during Early and Late learning.

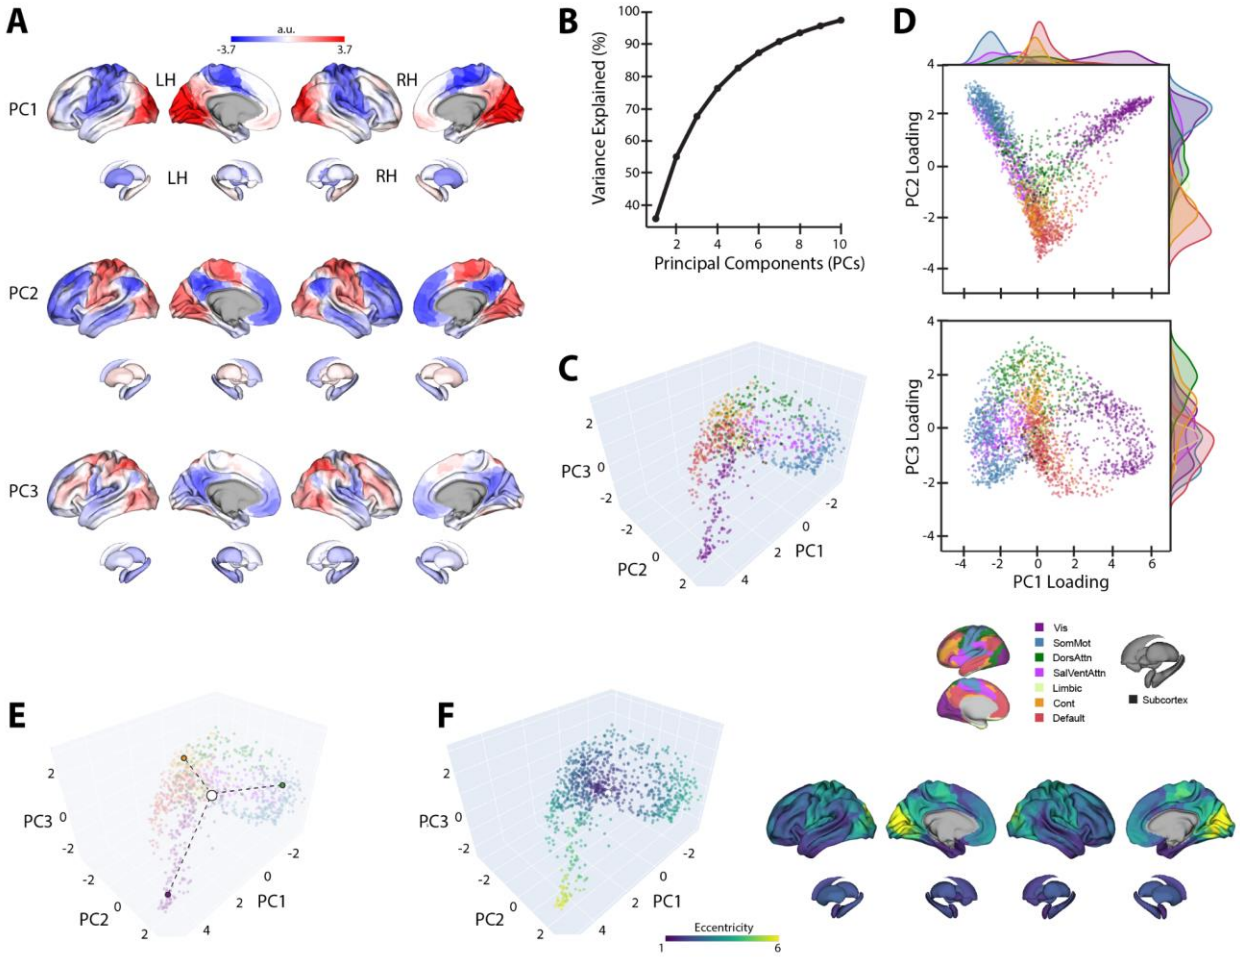


**Figure 3: Baseline manifold structure and eccentricity.** (A) Region loadings for the top three PCs. (B) Percent variance explained for the first 10 PCs. (C) The Baseline (template) manifold in low-dimensional space, with regions colored according to functional network assignment [58,68]. (D) Scatter plots showing the embedding of each region along the top three PCs. Probability density histograms at top and at right show the distribution of each functional network along each PC. Vis: Visual. SomMot: Somatomotor. DorsAttn: Dorsal attention. SalVentAttn: Salience/Ventral attention. Cont: Control. (E) Illustration of how eccentricity is calculated. Region eccentricity along the manifold is computed as the Euclidean distance (dashed line) from manifold centroid (white circle). The eccentricity of three example brain regions is highlighted (bordered coloured circles). (F) Regional eccentricity during Baseline. Each brain region's eccentricity is color-coded in the low-dimensional manifold space (left) and on the cortical and subcortical surfaces (right). White circle with black bordering denotes the center of the manifold (manifold centroid).

### *Changes in cortical-subcortical manifold structure during learning*

To examine which regions exhibited significant changes in manifold eccentricity from (1) Baseline to Early Learning and then from (2) Early to Late learning, we performed two sets of paired t-tests, and corrected for multiple comparisons using a false-discovery rate correction (FDR; $q < 0.05$). To

directly test how regional eccentricity changes at the onset of learning (when subjects begin receiving reward feedback), we performed a contrast of Early > Baseline (Fig. 4A). This contrast primarily revealed a pattern of increasing eccentricity, i.e., manifold expansion, across several brain regions, indicating that these regions became segregated from the rest of the brain (red areas in Fig. 4A). This included areas located throughout the cortical sensorimotor system and Dorsal Attention Network (DAN), including bilateral superior-parietal, somatomotor, supplementary motor and premotor cortex, as well as regions in lateral visual cortex. In addition, this contrast identified many key areas of the DMN, including bilateral medial frontal gyrus (MFG), medial prefrontal cortex (MPFC), inferior frontal gyrus (IFG), and middle temporal cortex (MTC; for a network-level summary of these general effects, see the spider plot in Fig. 4A). In contrast to this general pattern of expansion-related effects, we also found that a small subset of areas in the posterior medial cortex (PMC) and posterior angular gyrus (AG) instead exhibited a decrease in eccentricity, i.e., manifold contraction (regions in blue in Fig. 4A), indicating that these areas increased their integration with other areas of the brain. Notably, we did not observe any significant changes in subcortical regions from Baseline to Early learning (however, for interested readers, Supplementary Fig 3 shows the unthresholded data from both cortex and subcortex to indicate any trends).

Next, to examine how regional eccentricity changes over the course of learning, we performed a direct contrast of Late > Early learning (Fig. 4B). This contrast mainly revealed a reversal in the general pattern of effects observed in the DMN during early learning. Specifically, during late learning, several regions in bilateral MFG, MPFC, IFG and MTC now exhibited contraction along the manifold, indicating an increased integration of these areas with other regions of the brain (this reversal can be easily observed by comparing the red areas in Fig. 4A to the blue areas in Fig. 4B). By contrast, areas in PMC and posterior AG now exhibited expansion, indicating an increased segregation of these areas from the brain. In addition, during late learning, we observed manifold expansion in several areas of the Salience/Ventral Attention Network (SalVentAttn), including the dorsal anterior cingulate cortex (dACC) and the anterior insula (AI), as well as higher-order lateral occipital cortical areas, and areas in retrosplenial cortex and medial ventro-temporal cortex (a network-level summary of these effects can be found in the Fig. 4A/B spider plot). Again, as in the Early > Baseline contrast, we did not observe any significant changes in subcortical regions from Early to Late learning. The only region that came close to reaching statistical significance in the striatum was the right pallidum ($p = 0.01$), but this region did not pass whole-brain FDR correction (corrected alpha = 0.086; note that Supplementary Fig 3 shows the unthresholded maps for this

contrast to demonstrate the strong reversal in the pattern of effects during Late learning, as well as indicate any trends).

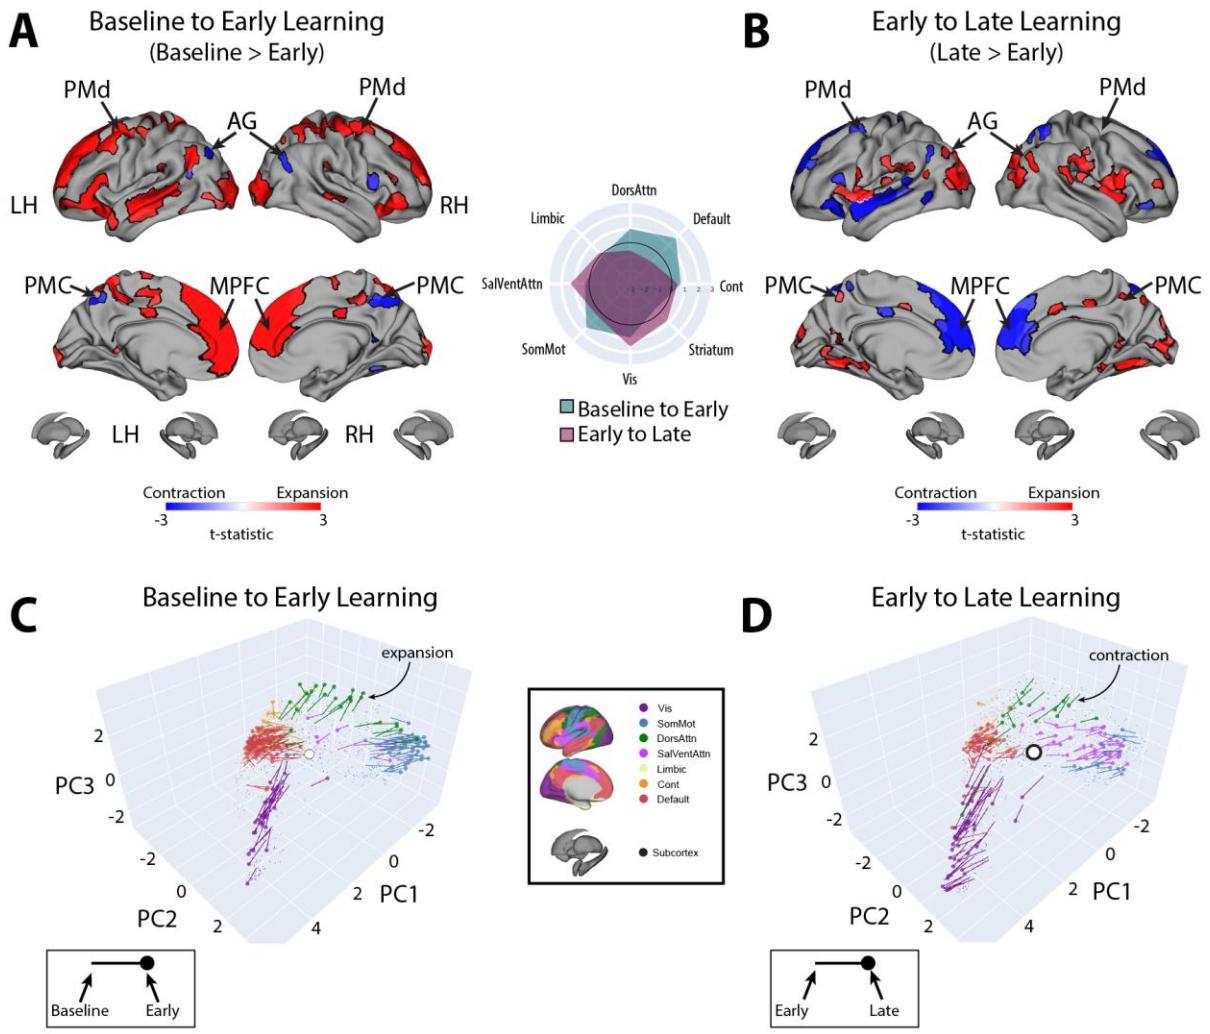


**Figure 4: Changes in manifold structure during reward-based motor learning.** (A & B) Pairwise contrasts of eccentricity between task epochs. Positive (red) and negative (blue) values show significant increases and decreases in eccentricity (i.e., expansion and contraction along the manifold), respectively, following FDR correction for region-wise paired t-tests (at $q<0.05$). The spider plot, at center, summarizes these patterns of changes in connectivity at the network-level (according to the Yeo networks, [68]). Note that the black circle in the spider plot denotes t=0 (i.e., no change in eccentricity between the epochs being compared). Radial axis values indicate t-values for the associated contrast (see color legend). (C & D) Temporal trajectories of statistically significant regions from A and B, shown in the low-dimensional manifold space. Traces show the displacement of each region for the relevant contrast and filled coloured circles indicate each region's final position along the manifold for a given contrast (see insets for legends). Each region is coloured according to its functional network assignment (middle). Nonsignificant regions are denoted by the gray point cloud. White circle with black bordering denotes the center of the manifold (manifold centroid).

Finally, for completeness, we also examined the contrast of Late > Baseline, which solely revealed a pattern of cortical expansion across several regions — in particular in areas of the SalVentAttn and visual network (see Supplementary Fig 4). This indicates a continuing expansion (and segregation) of these regions as learning progresses.

Taken together, the above pattern of results suggest that, during early learning, transmodal areas of the DMN, as well as several areas of the sensorimotor system (including areas of the DAN), begin to segregate from other brain networks, whereas a subset of areas — the PMC and posterior AG in particular — begin to integrate with regions outside of their respective networks. By contrast, during late learning, there is a reversal in these patterns, with regions within the DMN and DAN beginning to integrate with areas belonging to other brain networks. In the next section, we directly examine these interpretations of manifold expansion and contraction during early and late learning.

***Changes in connectivity that underlie patterns of manifold reconfiguration***

Given that eccentricity provides a multivariate index of a region's overall profile of connectivity (i.e., its relative positioning on the manifold), we next performed seed connectivity analyses to further characterize the patterns of effects that underlie the expansions and contractions of manifold structure during learning. To this aim, we selected several representative regions, distributed throughout the cortex, that epitomize the main changes in eccentricity that we observed during early learning (shown in Fig. 4A). These regions included the left (contralateral) MPFC, PMd and PMC, allowing us to characterize the patterns of connectivity changes across prefrontal, premotor and parietal cortex, respectively (For seed-connectivity analyses of their right hemisphere homologues, see Supplementary Fig 5). For each region, we contrasted seed connectivity maps between both the Early Learning vs. Baseline epochs (Early > Baseline) and the Late vs. Early Learning epochs (Late > Early) by computing region-wise paired t-tests, thus producing contrast maps associated with the connectivity change of each representative seed region (Fig. 5). Note that in Fig. 5 we display the unthresholded voxel-wise contrast maps (two leftmost panels), the region's corresponding change in eccentricity across epochs (second from rightmost panel), and the corresponding spider plots depicting network-level changes (rightmost panel), thus allowing for a complete visualization of the collective changes in connectivity that contribute to the changes in regional eccentricity. [Note that, for the spider plots, we used the 17-network mapping in order to capitalize on the improved spatial precision compared to the 7-network mapping [58,68]].

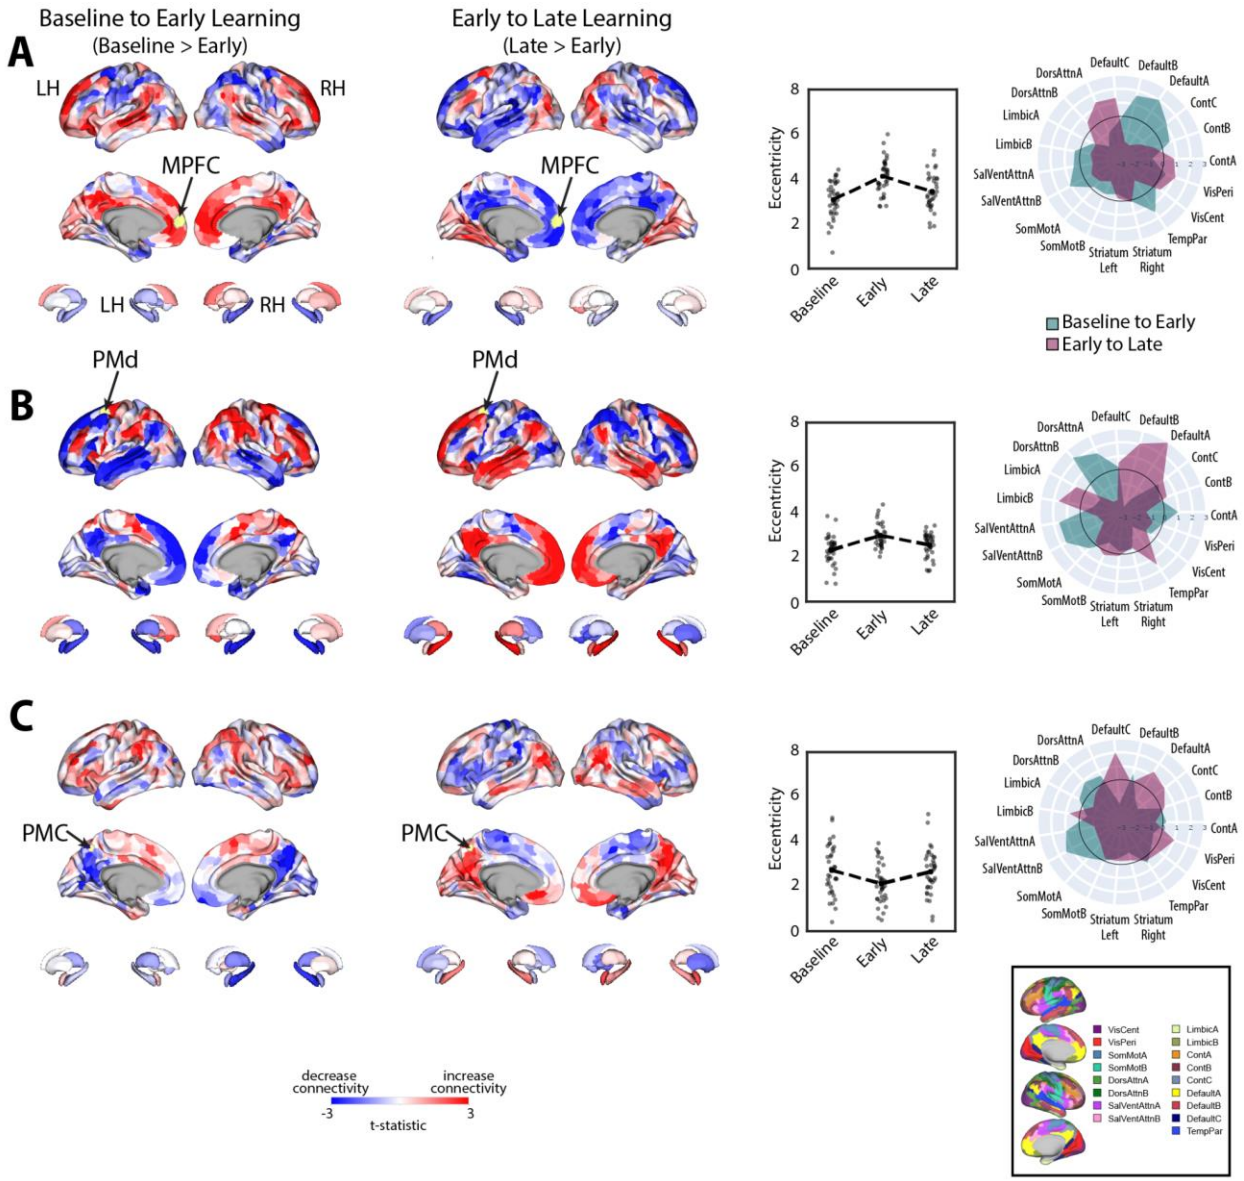


**Figure 5: Main patterns of connectivity changes that underlie manifold expansions and contractions.** (A-C) Connectivity changes for each seed region. Selected seed regions are shown in yellow and are also indicated by arrows. Positive (red) and negative (blue) values show increases and decreases in connectivity, respectively, from Baseline to Early learning (leftmost panel) and Early to Late learning (second from leftmost panel). Second from the rightmost panel shows the eccentricity of each region for each participant, with the black dashed line showing the group mean across task epochs. Rightmost panel contains spider plots, which summarize these patterns of changes in connectivity at the network-level (according to the Yeo 17-networks parcellation [68]). Note that the black circle in the spider plot denotes t=0 (i.e., zero change in eccentricity between the epochs being compared). Radial axis values indicate t-values for associated contrast (see color legend).

During early learning, we found that the left MPFC seed region, associated with the DMN network, exhibited increased connectivity with other DMN subregions and reduced connectivity with superior

parietal and premotor areas in the DAN (Fig. 5A). By contrast, during late learning, we observed a reversal in this pattern of connectivity changes, with the MPFC now exhibiting increased connectivity with the same regions of the DAN but reduced connectivity with other DMN areas. Together, these results suggest that the manifold expansion of the MPFC during early learning arises from its increased connectivity with other DMN areas (i.e., segregation of the DMN) whereas the manifold contraction of this region during late learning arises from its increased connectivity with areas *outside* of the DMN, such as sensorimotor areas of the DAN.

Notably, for the left PMd seed region (Fig. 5B), associated with the DAN, we observed an inverse pattern of results from that observed for the MPFC region above. Specifically, during early learning, we observed increased connectivity of the left PMd with other areas of the DAN, as well as areas belonging to the SalVentAttn network, such as the anterior insula/IFG, dACC, and inferior parietal cortex. Notably, this was coupled with its decreased connectivity to DMN areas and the hippocampus (Fig. 5B). By contrast, during late learning, we observed a reversal in this pattern of effects, whereby connectivity with DMN areas and the hippocampus now increased whereas connectivity with the DAN and SalVentAttn areas decreased. In this case, the pattern of manifold expansion and contraction of PMd during early and late learning, respectively, likely arises from its increased connectivity with attention networks in brain (the DAN and SalVentAttn) during early learning and an increase in between-network connectivity (i.e., integration) with DMN areas during late learning.

Finally, for the left PMC seed region, located at the border of the DMN-A and Control-C networks — and that was one of the few regions that exhibited *contraction* during early learning — we found that this region exhibited decreased connectivity during early learning with other DMN-A and Control-C subregions in bilateral PMC, as well as decreased connectivity with bilateral hippocampus and MPFC (Fig. 5C). Instead, this PMC region exhibited increased connectivity with DAN areas in superior parietal cortex and premotor cortex, and most prominently, with areas in the anterior insula/IFG, dACC, and inferior parietal cortex (belonging to the SalVentAttn networks). By contrast, during late learning, we again observed a reversal in this pattern of connectivity changes, with the PMC seed region now exhibiting increased connectivity with other bilateral PMC areas, MPFC and the hippocampus, as well as reduced connectivity with the same DAN and SalVentAttn areas. Together, these results suggest that manifold contractions of the PMC during early learning arise from its increased integration with regions outside of the DMN-A and Control-C networks, such as the DAN and SalVentAttn, whereas the manifold expansions of this region during late learning arise from its increased within-network connectivity with other DMN and Control areas (i.e., segregation).

Taken together, the results of our seed connectivity analyses above are broadly consistent with our interpretation of the patterns of manifold expansion as reflecting increases in *within*-network connectivity (segregation) and the patterns of manifold contraction as reflecting increases in *between*-network connectivity (integration). More generally, however, these findings point to changes in the landscape of communication between regions of the DAN, SalVentAttn and DMN in particular, as being associated with reward-based motor learning. Specifically, we find that, during early learning, there is increased functional coupling between several sensorimotor areas of the DAN with areas of the SalVentAttn network, whereas during late learning, these DAN sensorimotor areas switch their connectivity to DMN areas.

***Changes in eccentricity relate to learning performance***

In the previous sections we characterized the patterns of mean changes in manifold structure during learning across all participants. However, it is well-established that subjects exhibit significant variation in motor learning ability [55,72–74]. Indeed, while the learning curve in Fig. 1D shows that subjects, on average, improved their scores during the task, this group-level result obscures the fact that individuals differed greatly in their rates of learning (see individual subject learning curves in Fig. 6A). To emphasize this fact, Fig. 6A highlights the learning curves for two example subjects: An individual who learned the hidden shape quite rapidly (a 'fast learner' in green) and a second individual who only gradually learned to trace the hidden shape (a 'slow learner' in red). To quantify this variation in subject performance, we opted for a pure data-driven approach and performed functional principal component analysis (fPCA; [75]) on subjects' learning curves, which allowed us to isolate the dominant patterns of subject variability [see Methods for further details; see also 63]. Using this fPCA approach, we found that a single component — encoding overall learning — captured the majority (~75%) of the variability in subjects' learning curves (Fig. 6B). We thus used each subjects' loading on this dominant component as a single scalar measure of subjects' overall learning performance: Individuals who tended to learn the task more quickly had higher values on this 'Learning Score' component whereas individuals who tended to learn the task more slowly had lower values on this component.

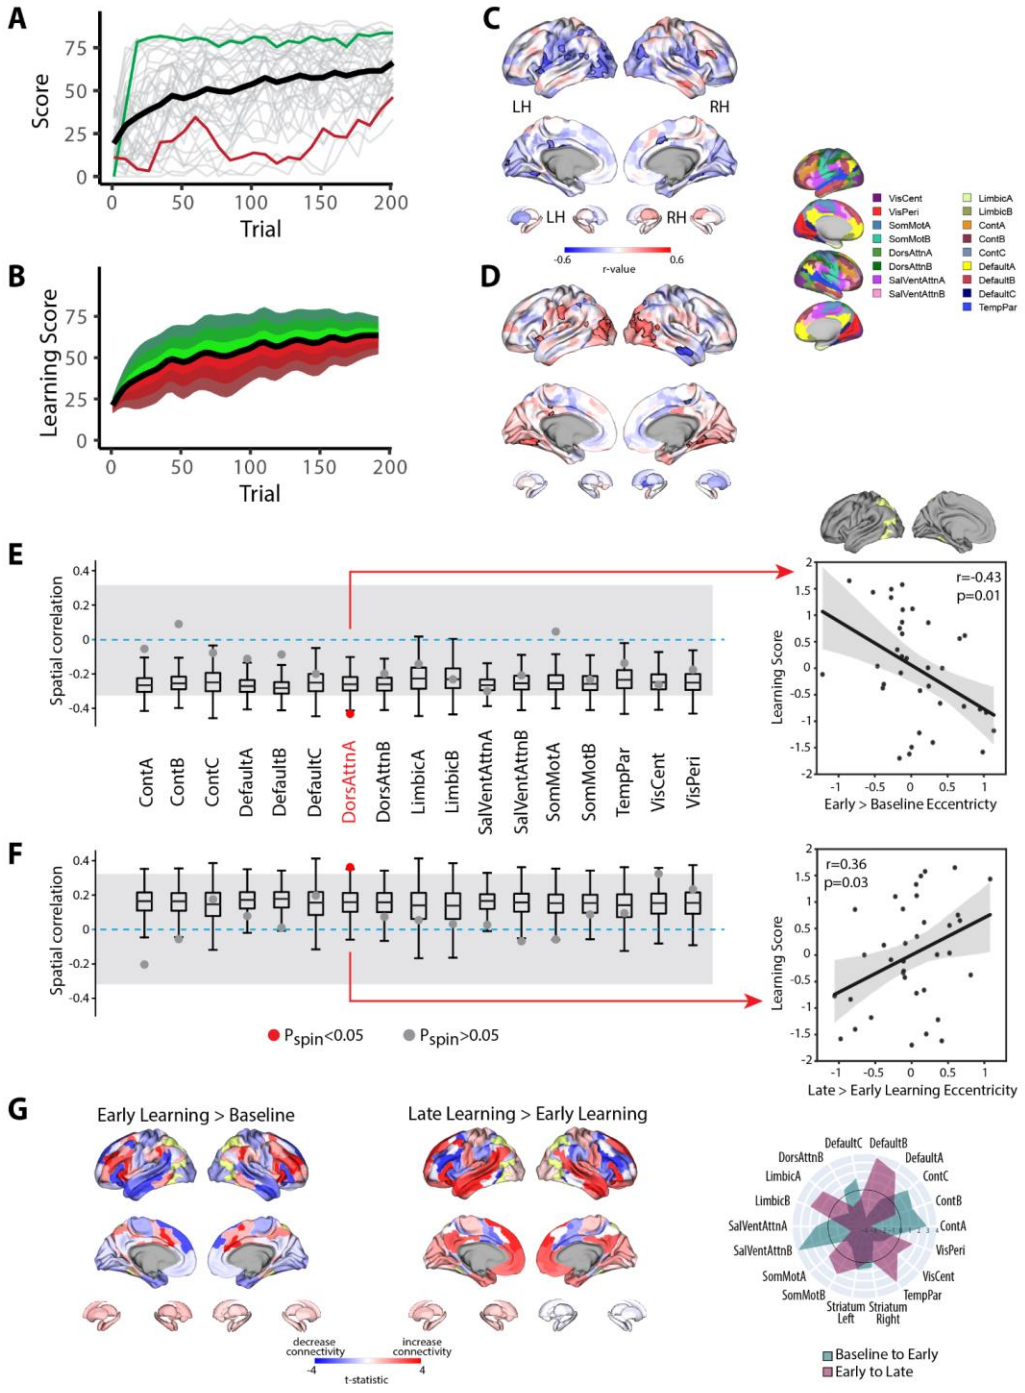


**Figure 6: Relationship between learning performance and regional changes in eccentricity.** (A) Individual subject learning curves for the learning task. Solid black line denotes the mean across all subjects whereas light gray lines denote individual participants. The green trace denotes an example fast learner and the red trace denotes an example slow learner (see text). (B) Derivation of subject learning scores. We performed functional principal component analysis on subjects' learning curves in order to identify the dominant patterns of variability during learning. The top component, which encodes overall learning, explained the majority of the observed variance (~75%). The green and red bands denote the effect of positive and negative component scores, respectively, relative to mean performance. Thus, subjects who learned more

quickly than average have a higher loading (in green) on this 'Learning score' component than subjects who learned more slowly (in red) than average. (C & D) Whole-brain correlation map between subject Learning score and the change in regional eccentricity from Baseline to Early learning (C) and Early to Late learning (D). Black bordering denotes regions that are significant at $p<0.05$. (E & F) Results of the spin-test permutation procedure, assessing whether the topography of correlations in C and D are specific to individual functional brain networks. Single points indicate the real correlation value for each of the 17 Yeo et al. networks [68], whereas the boxplots represent the parameters of a null distribution of correlations derived from 1000 iterations of a spatial autocorrelation-preserving null model [76,77]. In the boxplots, the ends of the boxes represent the first (25%) and third (75%) quartiles, the center line represents the median, and the whiskers represent the min-max range of the null distribution. All correlations were corrected for multiple comparisons ($q<0.05$). The dashed horizontal blue line indicates a correlation value of zero and the gray shading encompasses correlation values that do not significantly differ from zero ($p>0.05$). [Note that in the spin-test procedure, due to the sign of the correlations, it is possible for some networks to be significantly different from the null distribution, and yet not significantly different from zero. Thus, to be considered significant in our analyses, a brain network must satisfy both constraints; i.e., show a correlation that is significantly different from zero *and* from the spatial null distribution]. Right, scatterplots show the relationships between subject Learning score and the change in eccentricity from Baseline to Early learning (top) and Early to Late learning (bottom) for the DAN-A network (depicted in yellow on the cortical surface at top), the only brain network to satisfy the two constraints of our statistical testing procedure. Black line denotes the best-fit regression line, with shading indicating +1/- standard error of the mean. Dots indicate single participants. (G) Connectivity changes for the DAN-A network (highlighted in yellow) across epochs. Positive (red) and negative (blue) values show increases and decreases in connectivity, respectively, from Baseline to Early learning (left panel) and Early to Late learning (right panel). Spider plot, at right, summarizes the patterns of changes in connectivity at the network-level. Note that the black circle in the spider plot denotes t=0 (i.e., no change in eccentricity between the epochs being compared). Radial axis values indicate t-values for associated contrast (see color legend). VisCent: Visual Central. VisPer: Visual Peripheral. SomMotA: Somatomotor A. SomMotB: Somatomotor B. TempPar: Temporal Parietal. DorsAttnA: Dorsal Attention A. DorsAttnB: Dorsal Attention B. SalVentAttnA: Salience/Ventral Attention A. SalVentAttnB: Salience/Ventral Attention B. ContA: Control A. ContB: Control B. ContC: Control C.

Next, to examine the neural correlates of intersubject differences in Learning scores, we calculated, for each region, the association between participants' scores and the change in eccentricity between the Baseline and Early learning epochs (Early > Baseline; Fig. 6C). This analysis did not reveal any brain regions that survived FDR-corrections for multiple comparisons ($q<0.05$). However, the FDR approach is completely agnostic to any topographical patterns of effects across brain areas, which may correspond with known functional networks. Indeed, the full whole-brain correlation map in Fig. 6C indicates that many region-level correlations exhibit a high degree of spatial contiguity, with many statistically significant regions (at $p<0.05$, bordered in black) being situated within much larger clusters of regions that exhibit a similar pattern of effects (i.e., areas in blue, denoting a negative correlation between Learning score versus the change in eccentricity from Baseline to Early learning, tend to lie adjacent to other regions exhibiting a similar negative correlation). This is because topographically adjacent regions are likely to have similar connectivity profiles, and thus project onto similar locations along the manifold, resulting in similar brain-behavior relationships. This spatial

topography suggests that the association between the eccentricity of certain brain regions and learning performance are likely to be better characterized at the level of distributed functional networks.

To examine this, we mapped each region onto its respective functional network using the Yeo et al. 17-network parcellation [68] and, for each participant, computed the mean manifold eccentricity for each network (i.e., network eccentricity). We then correlated the change in each brain network's eccentricity across epochs with subject Learning scores. We tested the statistical significance of these network-level correlations by building null models that account for the spatial autocorrelation in the brain maps [76,77](see Methods) and corrected for multiple comparisons (across all networks) using an FDR correction ($q<0.05$). Using this permutation testing approach, we found that it was only the change in eccentricity of the DAN-A network that correlated with Learning score (see Fig. 6E), such that the more the DAN-A network *decreased* in eccentricity from Baseline to Early learning (i.e., contracted along the manifold), the better subjects performed at the task (see Fig. 6E, scatterplot at right). Consistent with the notion that changes in the eccentricity of the DAN-A network are linked to learning performance, we also found the inverse pattern of effects during Late learning, whereby the more that this same network *increased* in eccentricity from Early to Late learning (i.e., expanded along the manifold), the better subjects performed at the task (Fig. 6F). These results suggest that the contraction and expansion of areas of the DAN-A network are selectively associated with subject learning performance.

To understand the global changes in connectivity that underlie these network eccentricity effects, we performed a network-level seed connectivity analyses (analogous to our single ROI seed connectivity analyses in the previous section) wherein constructed contrast maps, using the DAN-A as the seed network, for both the Early vs. Baseline epochs (Early > Baseline) and the Late vs. Early epochs (Late > Early). As before, we display the unthresholded voxel-wise contrast maps, along with corresponding spider plots depicting the network-level changes (Fig. 6G), to allow for a complete visualization of the collective changes in network-level connectivity that underlie the changes in eccentricity of the DAN-A. As can be observed in Fig. 6G, we find that, during early learning, DAN-A regions exhibited the largest increases in connectivity with one of the SalVentAttn subnetworks (SalVentAttn-B) whereas, during late learning, the DAN-A regions exhibited the largest increases in connectivity with one of the DMN subnetworks (DMN-B). These findings not only re-constitute the group-level effects reported above at the ROI level (in Fig. 5), but they also suggest that this general

transition in functional coupling — between the DAN and SalVentAttn areas during early learning, to DAN and DMN areas during late learning — are associated with differences in subject performance.

## DISCUSSION

Complex behavior necessitates the coordinated activity of multiple specialized neural systems distributed across cortex and subcortex. During motor learning, these systems must adapt their functional interactions to ensure appropriate behavior in response to changes in sensory feedback. While much research in motor learning has focused on understanding the role of sensorimotor cortex in isolation, our understanding of the contribution of higher-order brain systems, which play a role in the organization of behavior over time, remains incomplete. In the current study, we utilized state-of-the-art analytical methods that reconcile topographic and functional brain organization, enabling us to describe changes in the landscape of whole-brain activity during learning.

During early learning, when subjects were establishing the relationship between motor commands and reward feedback, we found that regions within both the DAN (e.g., premotor cortex) and DMN (e.g., MPFC) exhibited expansion along the manifold. Our connectivity analyses revealed that this expansion was largely driven by an increase in within-network communication in both the DAN and DMN networks. There were, however, two notable exceptions to this general pattern. First, we found that connectivity between regions of the DAN increased with regions in the SalVentAttn network (e.g., anterior insula/IFG and anterior cingulate cortex). Second, areas within the PMC, part of the posterior core of the DMN, showed a pattern of manifold contraction that was primarily driven by a decrease in covariance with other DMN regions and an increase in covariance with regions of the superior DAN and also with several regions within the SalVentAttn network. Together, these results suggest that functional interactions of the sensorimotor system with the SalVentAttn network are important during initial learning. In other work, areas within the SalVentAttn network have been implicated in several aspects of cognitive control and motivation [28,78–80], and in studies on reward-based decision-making, these regions are thought to support exploratory behavior. For instance, several neuroimaging studies have shown that SalVentAttn areas are activated in response to novel or salient stimuli in the environment, presumably reflecting the engagement of attentional resources for sampling new information [13]. In addition, recent work [50,81–85] implicates several SalVentAttn areas in information gathering functions, so as to optimize reward outcomes. This neural perspective is consistent with both others' [57] and our own behavioral findings [72] that, during the early phases of

motor learning, performance tends to be more variable, presumably reflecting an exploration of the relationship between motor commands and associated sensory feedback.

During late learning, we observed that many of the changes in manifold architecture observed during early learning reversed. For instance, areas within both the DAN and DMN now exhibited contraction along the manifold, whereas the PMC now exhibited expansion. Connectivity analyses showed that the contraction within the DAN and DMN was driven by increases in connectivity between these two networks — specifically, increases between premotor and superior parietal areas of the DAN with areas of the DMN. This suggests that once the mapping between motor commands and reward feedback have been largely learned, regions within the DAN and DMN become more integrated with one another. In the context of the current task, the shift in DAN connectivity from the SalVentAttn network during early learning to the DMN during late learning may reflect the hypothesized role of the DMN in supporting behavior using information from memory [15,17,22]. This interpretation is consistent with prior work showing that connectivity between the DMN and premotor cortex tends to increase once rules have been learned [86], as well as studies from other task domains showing that the DMN contributes to behavior when actions must be guided by information from memory and a knowledge of task structure [25–27]. Our analysis, therefore, provides new evidence, from the domain of human motor learning, that functional interactions between the DMN and brain regions involved in sensorimotor processes support a mode of action in which behavior must be guided by memory processes [17] — in this case, the history of reward information accrued across previous movement trajectories.

We also observed a relationship between changes in the manifold eccentricity of a subnetwork of the DAN (DAN-A) with subject learning performance. We found that the more this subnetwork contracted, and then subsequently expanded, along the manifold during early and late learning respectively, the better subjects tended to perform at the task. Our connectivity analysis revealed that this change in DAN-A activity was mainly driven by increases in covariance with SalVentAttn subnetworks during early learning and by increases in covariance with DMN subnetworks during late learning. Notably, areas of the DAN are not thought to generate top-down signals for response selection but instead transform the input signals they receive (e.g., related to reward, memory) into motor commands [87]. Our analysis thus suggests that this system's contribution to motor behavior may be facilitated through changes in its functional coupling to both the SalVentAttn and DMN over time. To speculate, this shift in functional coupling may reflect a shift from more exploratory to more

exploitative modes of behavior across early to late periods of motor learning, respectively. Moreover, our data suggest that this shift in functional coupling could be mediated through the PMC.

Although traditionally considered a member of the DMN [68], studies have established that the PMC contains echoes of neural signals originating from across the cortex [88]. In our study, this region initially became more segregated from the rest of the DMN and increased connectivity with the SalVentAttn network during early learning. However, during late learning, the PMC reduced its connectivity with the SalVentAttn network and became more integrated with other areas of the DMN. Notably, this pattern of changes differed significantly from other areas of the DMN, suggesting that the PMC may serve a different function during motor learning than other areas of this system. Prior studies have linked the activity of PMC areas to reward processing [89–91] and the selection of response strategies that attempt to optimize reward outcomes [28,35,92]. Consistent with this, recent research shows that areas within the PMC are able to integrate information over particularly long periods of time [30,93,94]. This characteristic positions the PMC as an ideal candidate region to orchestrate the neural transition from (1) exploring the relationship between motor commands and sensory feedback during early learning to (2) subsequently exploiting this learned relationship during late learning. Taken together, these distinctive functional properties of PMC activity, coupled with its diverse patterns of whole-brain connectivity [95–97], suggest an important role for this region in directing long-term behavioral adaptation in accordance with higher-order task objectives [32,98].

While we identified several changes in the cortical manifold that are associated with reward-based motor learning, it is noteworthy that we did not observe any significant changes in manifold eccentricity within the striatum. While clearly prior evidence indicates that this region plays a key role in reward-guided behavior [5,6], there are several possible reasons why our manifold approach did not identify this collection of brain areas. First, the relatively small size of the striatum may mean that our whole-brain approach was too coarse to identify changes in the connectivity of this region. Though we used a 3T scanner and employed a widely-used parcellation scheme that divided the subcortex into its constituent anatomical regions (e.g., hippocampus, caudate, etc.), both of these approaches may have obscured important differences in connectivity that exist *within* each of these regions. For example, areas such the hippocampus and caudate are not homogenous areas but themselves exhibit gradients of connectivity (e.g., head versus tail) that can only be revealed at the voxel level [99,100]. Second, while our dimension reduction approach, by design, aims to identify gradients of functional connectivity that account for the largest amounts of variance, the limited number of subcortical regions (as compared to cortex) necessitates that their contribution to the

total whole-brain variance is relatively small. Consistent with this perspective, we found that the low-dimensional manifold architecture in cortex did not strongly depend on whether or not subcortical regions were included in the analysis (see Supplementary Fig. 6). As such, selective changes in the patterns of functional connectivity occurring at the level of subcortex may be obscured using our whole-brain dimension reduction approach. Future work can help address some of these limitations by using both finer parcellations of striatal cortex (perhaps even down to the voxel level)[100] and/or by focusing solely on changes in the interactions *between* subcortex and cortex during learning. The latter could be accomplished by selectively performing dimension reduction on the slice of the functional connectivity matrix that corresponds to functional coupling between striatum and cortex.

**Conclusions**

Our study set out to characterize changes in the landscape of whole-brain activity that underlies reward-based motor learning. We used dimensionality reduction techniques to build a manifold that describes changes in the functional organization of the cortex and subcortex during different phases of learning. During early learning, we found that regions within the DAN and DMN became relatively segregated from each other, with the DAN becoming more integrated with the SalVentAttn network. This pattern reversed during later learning, with regions within the DAN now becoming more integrated with the DMN. Notably, regions of the PMC, within the posterior core of the DMN, showed the reverse pattern, exhibiting coupling with the SalVentAttn early during learning and with other regions of the DMN later during learning. Together, these findings provide a unique whole-brain perspective into the neural changes that underlie reward-based motor learning, and point to marked transitions in the activity of transmodal cortical regions in organizing behavior over time.

# METHODS

## *Participants*

Fourty-six right-handed individuals (27 females, aged 18-28 years) participated in the MRI study. Of these 46 participants, ten individuals were removed from the final analysis, either due to excessive head motion in the MRI scanner, incomplete scans, poor task compliance (i.e., > 25% of trials not being completed within the maximal trial duration), or missing data (i.e., > 20% of trials being missed due to insufficient pressure of the fingertip on the MRI-compatible tablet). We assessed right-handedness using the Edinburgh handedness questionnaire [101] and obtained informed consent before beginning the experimental protocol. The Queen's University Research Ethics Board approved the study and it was conducted in coherence to the principles outlined in the Canadian Tri-Council Policy Statement on Ethical Conduct for Research Involving Humans and the principles of the Declaration of Helsinki (1964).

## *Procedure*

Prior to MRI testing, participants first took part in an MRI training session inside a mock (0 T) scanner, made to look and sound like a real MRI scanner. This training session served multiple purposes. First, it introduced participants to key features of the motor task that was subsequently performed in the MRI scanner. Second, it allowed us to screen for subjects that could obtain baseline performance levels on the task. Third, it allowed us to screen for subjects that could remain still for a long period of time without feeling claustrophobic. With respect to the latter, we monitored subjects' head movement in the mock scanner while they performed practice task trials and during simulated anatomical scans. This monitoring was done by attaching, via medical tape, a Polhemus sensor to each subject's forehead (Polhemus, Colchester, Vermont), which allowed a real-time read-out of subject head displacement in each of the three axes of translation and rotation (6 dimensions total). Whenever subjects' head translation and/or rotation reached 0.5 mm or 0.5º rotation (within a pre-specified velocity criterion), subjects received an unpleasant auditory tone, delivered through a speaker system located near the head. All subjects learned to constrain their head movement via this auditory feedback. Following this first training session, subjects that met our criteria were invited to subsequently participate in the reward-based motor learning task (see below for details), approximately 1-week later.

### *Apparatus*

During testing in the mock (0 T) scanner, subjects performed hand movements that were directed towards a target by applying fingertip pressure on a digitizing touchscreen tablet (Wacom Intuos Pro M tablet). During the actual MRI testing sessions, subjects used an MRI-compatible digitizing tablet (Hybridmojo LLC, CA, USA). In both the mock and real MRI scanner, the visual stimuli were rear-projected with an LCD projector (NEC LT265 DLP projector, 1024 x 768 resolution, 60 Hz refresh rate) onto a screen mounted behind the participant. The stimuli on the screen were viewed through a mirror fixated on the MRI coil directly above the participants' eyes, thus preventing the participant from being able to see their hand.

### *Reward-based motor learning task*

In the motor task participants were trained, through reward-based feedback, to produce finger movement trajectories for an unseen shape. Specifically, subjects were instructed to repeatedly trace, without visual feedback of their actual finger paths, a subtly curved path displayed on the screen (the visible path, Fig. 1B,C). Participants were told that, following each trial, they would receive a score based on how 'accurately' they traced the visible path. However, unbeknownst to them, they actually received points based on how well they traced the mirror-image path (the reward path, Fig. 1B,C). Critically, because participants received no visual feedback about their actual finger trajectories or the 'rewarded' shape, they could not use error-based learning mechanisms to guide learning [102,103]. This task was a modification on the motor tasks developed by [55,57].

Each trial started with the participant moving a cursor (3 mm radius cyan circle), which represented their finger position, into the start position (4 mm radius white circle) at the bottom of the screen (by sliding the index finger on the tablet). The cursor was only visible when it was within 30 mm of the start position. After the cursor was held within the start position for 0.5 s, the cursor disappeared and a rightward-curved path (Visible Path) and a movement distance marker appeared on the screen (Fig. 1B). The movement distance marker was a horizontal red line (30 x 1 mm) that appeared 60 mm above the start position. The visible path connected the start position and movement distance marker, and had the shape of a half sine wave with an amplitude of 0.15 times the marker distance. Participants were instructed to trace the curved path. When the cursor reached the target distance, the target changed color from red to green to indicate that the trial was completed. Importantly,

beyond this color change in the distance marker, participants never received any feedback about the position of their cursor.

In the baseline block, participants did not receive any feedback about their performance. In the learning block, participants were rewarded 0 to 100 points after reaching the movement distance marker, and participants were instructed to do their best to maximize this score across trials (following the movement, the points were displayed as text centrally on the screen). Each trial was terminated after 4.5 s, independent of whether the cursor had reached the target. After a delay of 1.5 s, allowing time to save the data and the subject to return to the starting location, the next trial started with the presentation of the start position.

To calculate the reward score on each trial in the learning block, the *x* position of the cursor was interpolated at each cm displacement from the start position in the *y* direction (i.e., at exactly 10, 20, 30, 40, 50 and 60 mm). For each of the six *y* positions, the absolute distance between the interpolated *x* position of the cursor and the *x* position of the rewarded path was calculated. The sum of these errors was scaled by dividing it by the sum of errors obtained for a half cycle sine-shaped path with an amplitude of 0.5 times the target distance, and then multiplied by 100 to obtain a score ranging between 0 and 100. The scaling worked out such that a perfectly traced visible path would result in an imperfect score of 40 points. This scaling was chosen on the basis of extensive pilot testing in order to: (1) encourage motor exploration across trials (in search of higher scores), (2) achieve variation across subjects in overall performance (i.e., individual differences in learning curves), and (3) to ensure that subjects still received informative score feedback when tracing in the vicinity of the visible trajectory.

During the training session in the mock MRI scanner (i.e., ~1-week prior to the MRI testing session), participants performed only a practice block in which they traced a straight line, first with (40 trials) and then without (40 trials), visual feedback of the position of the cursor during the movement (80 trials total). This training session exposed participants to several key features of the task (e.g., use of the touchscreen tablet, trial timing, removal of cursor feedback) and allowed us to establish adequate performance levels. Importantly, however, subjects did not encounter any reward-based feedback (reward scores) during this initial training session.

At the beginning of the MRI testing session, but prior to the first scan being collected, participants re-acquainted themselves with the motor task by first performing a practice block in which they traced

a straight line with (40 trials) and then without (40 trials) visual feedback of the position of the cursor. Next, we collected an anatomical scan, followed by a DTI scan, followed by a resting-state fMRI scan. During the latter resting-state scan, participants were instructed to rest with their eyes open while fixating on a central cross location presented on the screen. [Note that the DTI and resting-state fMRI data are not the focus on the present study]. Next, participants performed the motor task, which consisted of two separate experimental runs without visual feedback of the cursor: (1) a baseline block of 70 trials in which they attempted to trace the curved path and no score feedback was provided, and (2) a separate learning block of 200 trials in which participants were instructed to maximize their score shown at the end of each trial.

## *Behavioral data analysis*

**Data Preprocessing.** Trials in which the cursor did not reach the target within the time limit were excluded from the offline analysis of hand movements (~1% of trials). As insufficient pressure on the touchpad resulted in a default state in which the cursor was reported as lying in the top left corner of the screen, we excluded trials in which the cursor jumped to this position before reaching the target region (~2% of trials). We then applied a conservative threshold on the movement and reaction times, removing the top 0.05% of trials across all subjects. As the motor task did not involve response discrimination, we did not set a lower threshold on these variables.

**Functional PCA of subject behavioral data.** All subject behavioral data were averaged over 8 trial bins. We represented individual learning curves as functional data using a cubic spline basis with smoothing penalty estimated by generalized cross-validation [104]. We then performed *functional PCA* [105], which allowed us to extract components capturing the dominant patterns of variability in subject performance. Using this analysis, we found that the top component, which describes overall learning, explained a majority of the variability (~75%) in performance. Spline smoothing and fPCA were performed using the R package fda [106], and example code is provided in a publically available repository (see Data and Code Availability statement).

## *MRI Acquisition*

Participants were scanned using a 3-Tesla Siemens TIM MAGNETOM Trio MRI scanner located at the Centre for Neuroscience Studies, Queen's University (Kingston, Ontario, Canada). Subject anatomicals were acquired using a 32-channel head coil and a T1-weighted ADNI MPRAGE sequence

(TR = 1760 ms, TE = 2.98 ms, field of view = 192 mm x 240 mm x 256 mm, matrix size = 192 x 240 x 256, flip angle = 9°, 1 mm isotropic voxels). This was followed by a series of Diffusion-Weighted scans, and a resting-state scan (which are not the focus of the present investigation). Next, we acquired functional MRI volumes using a T2*-weighted single-shot gradient-echo echo-planar imaging (EPI) acquisition sequence (time to repetition (TR) = 2000 ms, slice thickness = 4 mm, in-plane resolution = 3 mm x 3 mm, time to echo (TE) = 30 ms, field of view = 240 mm x 240 mm, matrix size = 80 x 80, flip angle = 90°, and acceleration factor (integrated parallel acquisition technologies, iPAT) = 2 with generalized auto-calibrating partially parallel acquisitions (GRAPPA) reconstruction. Each volume comprised 34 contiguous (no gap) oblique slices acquired at a ~30° caudal tilt with respect to the plane of the anterior and posterior commissure (AC-PC), providing whole-brain coverage of the cerebrum. For the baseline and learning scans, we acquired 222 and 612 imaging volumes, respectively. Each of these task-related scans included an additional 6 imaging volumes at both the beginning and end of the scan. On average, the total MRI scanning session lasted ~2 hrs. (including setup time and image acquisition).

### *fMRI Preprocessing*

Preprocessing of anatomical and functional MRI data was performed using fMRIPrep 20.1.1 [107,108](RRID:SCR_016216) which is based on Nipype 1.5.0 [109,110](RRID:SCR_002502). Many internal operations of fMRIPrep use Nilearn 0.6.2 [111](RRID:SCR_001362), mostly within the functional processing workflow. For more details of the pipeline, see the section corresponding to workflows in fMRIPrep's documentation. Below we provide a condensed description of the preprocessing steps.

T1w images were corrected for intensity non-uniformity (INU) with N4BiasFieldCorrection [112], distributed with ANTs 2.2.0 [59](RRID:SCR_004757). The T1w-reference was then skull-stripped with a Nipype implementation of the antsBrainExtraction.sh workflow (from ANTs), using OASIS30ANTs as target template. Brain tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM) and gray-matter (GM) was performed on the brain-extracted T1w using fast (FSL 5.0.9, RRID:SCR_002823)[113]. A T1w-reference map was computed after registration of the T1w images (after INU-correction) using mri_robust_template (FreeSurfer 6.0.1)[114]. Brain surfaces were reconstructed using recon-all (FreeSurfer 6.0.1, RRID:SCR_001847)[115], and the brain mask estimated previously was refined with a custom variation of the method to reconcile ANTs-derived and FreeSurfer-derived segmentations of the cortical gray-matter of Mindboggle (RRID:SCR_002438)[116]. Volume-based spatial normalization to standard space (MNI152NLin6Asym) was performed through

nonlinear registration with antsRegistration (ANTs 2.2.0), using brain-extracted versions of both T1w reference and the T1w template.

For each BOLD run, the following preprocessing was performed. First, a reference volume and its skull-stripped version were generated using a custom methodology of fMRIPrep. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding rotation and translation parameters) are estimated before any spatiotemporal filtering using mcflirt (FSL 5.0.9)[117]. BOLD runs were slice-time corrected using 3dTshift from AFNI 20160207 [118](RRID:SCR_005927). The BOLD reference was then co-registered to the T1w reference using bbregister (FreeSurfer) which implements boundary-based registration [119]. Co-registration was configured with six degrees of freedom. The BOLD time series were resampled with a single interpolation step by composing all the pertinent transformations (i.e. head-motion transform matrices, and co-registrations to anatomical and output spaces). BOLD time series were resampled onto their original, native space, as well as standard space (MNI152NLin6Asym), using antsApplyTransforms (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels [120]. Subcortical data in standard space was combined with resampled BOLD time series on the fsaverage surface to produce Grayordinates files [121] containing 91k samples, using fsaverage as the intermediate standardized surface space. Resampling onto fsaverage was performed using mri_vol2surf (FreeSurfer).

A set of 34 motion and physiological regressors were extracted in order to mitigate the impact of head motion and physiological noise. The six head-motion estimates calculated in the correction step were expanded to include temporal derivatives and quadratic terms of each of the original and derivative regressors, totalling 24 head-motion parameters [122]. Ten component-based physiological regressors were estimated using the aCompCor approach [123,124], where the top five principal components were separately extracted from WM and CSF masks. Principal components were estimated after high-pass filtering the preprocessed BOLD time series (using a discrete cosine filter with 128s cut-off).

**Regional time series extraction.** For each participant and scan, the average BOLD time series were computed from the grayordinate time series for (1) each of the 998 regions defined according to the Schaefer 1000 parcellation [58]; two regions are removed from the parcellation due to their small parcel size) and, (2) each of the 12 subcortical regions defined according to the Harvard Oxford atlas [60,125], which included the caudate, putamen, accumbens, pallidum, hippocampus and amygdala.

Region timeseries were denoised using the above-mentioned confound regressors in conjunction with the discrete cosine regressors (128s cut-off for high-pass filtering) produced from fMRIprep and low-pass filtering using a Butterworth filter (100s cut-off) implemented in Nilearn. Finally, all region timeseries were z-scored. Regional time series extraction was performed using the nixtract package (https://github.com/danjgale/nixtract).

## *Neuroimaging data analysis*

**Covariance estimation and centering.** For every participant, region timeseries from the task scans were spliced into three equal-lengthed task epochs (210 imaging volumes each), after having discarded the first 6 imaging volumes (thus avoiding scanner equilibrium effects). This allowed us to estimate functional connectivity from continuous brain activity over the corresponding 70 trials for each epoch; Baseline comprised of the initial 70 trials in which subjects performed the motor task in the absence of any reward feedback, whereas the Early and Late learning epochs consisted of the first and last 70 trials after the onset of reward feedback, respectively. Then, we separately estimated functional connectivity matrices for each epoch by computing the region-wise covariance matrices using the Ledoit-Wolf estimator [126]. Note that our use of equal-length epochs for the three phases ensured that no biases in covariance estimation were introduced due to differences in time series length.

Next, we centered the connectivity matrices using the approach advocated by [65], which leverages the natural geometry of the space of covariance matrices [63,64]. In brief, this involved adjusting the covariance matrices of each participant to have a common mean, which was equivalent to the overall mean covariance, thus removing subject-specific variations in functional connectivity. First, a grand mean covariance matrix, $\bar{S}_{gm}$, was computed by taking the geometric mean covariance matrix across all $i$ participants and $j$ epochs. Then, for each participant we computed the geometric mean covariance matrix across task epochs, $\bar{S}_i$, and each task epoch covariance matrix $S_{ij}$ was projected onto the tangent space at this mean participant covariance matrix $S_i$ to obtain a tangent vector

$$T_{ij} = \bar{S}_i^{1/2} \log\left(\bar{S}_i^{-1/2} S_{ij} \bar{S}_i^{-1/2}\right) \bar{S}_i^{1/2}$$

where log denotes the matrix logarithm. We then transported each tangent vector to the grand mean $\bar{S}_{gm}$ using the transport proposed by [65], obtaining a centered tangent vector

$$T_{ij}^c = G T_{ij} G^\top$$

where $G = \bar{S}_{gm}^{1/2} S_i^{-1/2}$. This centered tangent vector now encodes the same difference in covariance, but now expressed relative to the grand mean. Finally, we projected each centered tangent vector back onto the space of covariance matrices, to obtain the centered covariance matrix

$$\mathrm{S}_{\mathrm{ij}}^{\mathrm{c}} = \bar{\mathrm{S}}_{\mathrm{gm}}^{1/2} \exp\left(\bar{\mathrm{S}}_{\mathrm{gm}}^{-1/2} \mathrm{T}_{\mathrm{ij}}^{\mathrm{c}} \bar{\mathrm{S}}_{\mathrm{gm}}^{-1/2}\right) \bar{\mathrm{S}}_{\mathrm{gm}}^{1/2}$$

where *exp* denotes the matrix exponential. For the benefits (and general necessity) of this centering approach, see Fig. 2, and for an additional overview, see [63].

Note that we have implemented many of the computations required to replicate the analysis in an publically available R package ***spdm***, which is freely available from a Git repository at https://github.com/areshenk-rpackages/spdm.

**Manifold Construction.** Connectivity manifolds for all centered functional connectivity matrices were derived in the following steps. First, consistent with previous studies [19,39,40,67], we applied row-wise thresholding to retain the top 10% connections in each row, and then computed cosine similarity between each row to produce an affinity matrix that describes the similarity of each region's connectivity profiles. Second, we applied principal components analysis (PCA) to obtain a set of principal components (PCs) that provide a low-dimensional representation of connectivity structure (i.e. connectivity gradients). We selected PCA as our dimension reduction technique based on recent work demonstrating the improved reliability of PCA over non-linear dimensionality reduction techniques (e.g., diffusion map embedding [67]).

To provide a basis for comparing changes in functional network architecture that arise during learning specifically, we constructed a template manifold, using the same aforementioned manifold construction procedures, from a group-average Baseline connectivity matrix that was derived from the geometric mean (across participants) of all centered Baseline connectivity matrices. We aligned all individual manifolds (36 participants × 3 epochs; 108 total) to this Baseline template manifold using Procrustes alignment. All analyses on the aligned manifolds were performed using the top three PCs, which cumulatively explained ~70% of the total variance in the template manifold. Across participants and epochs, the top three PCs, following Procrustes alignment, had an average correlation of r = .88 with their respective PCs in the template manifold, thus demonstrating good overall reliability and alignment across participants and epochs. Together, this approach enabled us to uniquely examine the learning-related changes in manifold structure (during Early and Late learning), and specifically how these deviate from the Baseline task functional architecture.

**Manifold Eccentricity and Analyses.** Recent work has quantified the embedding of regions and networks in low-dimensional spaces using Euclidean distance as a measure [69–71,127]. "Eccentricity" refers to the Euclidean distance between a single region and the manifold centroid [69], which, in the case of PCA, is equivalent to a region's magnitude, or vector length. Note that eccentricity provides a scalar index of network integration and segregation, in which distal regions with greater eccentricity are more segregated than proximal regions that integrate more broadly across functional networks [69–71]. To validate this interpretation with our own data, we correlated the Baseline template manifold eccentricity with three graph theoretical measures of functional integration and segregation. These measures were calculated on the row-wise thresholded template connectivity matrix and included *node strength*, which is the sum of a region's connectivity weights; *within-module degree z-score*, which measures the degree centrality of a region within its respective network; and *participation coefficien*t, which measures the network diversity of a region's connectivity distribution [128]. Regions were assigned to their respective intrinsic functional networks [58,68] for calculations of within-module degree z-score and participation coefficient.

We computed eccentricity for each brain region for all individual manifolds (each participant and epoch). This allowed us to observe manifold expansions (increases in eccentricity) and contractions (decreases in eccentricity) throughout early and late learning, thereby probing learning-related changes in functional segregation and integration (e.g., see Fig. 4). We compared region eccentricity between the Baseline, Early and Late learning epochs by performing a series of region-wise paired t-tests between these three key epochs. We applied FDR correction ($q<0.05$) across regions for each contrast.

**Seed Connectivity Analyses.** In order to understand the underlying changes in regional covariance that ultimately give rise to the observed changes in manifold eccentricity, we performed seed connectivity contrasts between the different task epochs. To this end, we selected several seed regions that were statistically significant in the Early Learning > Baseline contrast, which included areas in the left MPFC, left Premotor cortex (PMd), and left PMC , thereby allowing us to characterize some of the whole-brain changes that are associated reward-based learning. For completeness, we also selected homologous regions in the right hemisphere (see Supplemental Fig 5). For each seed region, we generated functional connectivity maps for the epochs of interest in every participant and computed region-wise paired t-tests for both the Early > Baseline and Late > Early contrasts. For all contrasts, we opted to show the unthresholded t-maps so as to visualize the complete multivariate pattern of connectivity changes that drive changes in eccentricity (a multivariate measure). In

addition, we constructed spider plots further characterizing these changes at the network-level by averaging the t-values across individual regions according to their network assignment [68]. Note that these analyses are mainly intended to provide characterization (and interpretation) of the connectivity changes of representative regions from our main eccentricity analyses.

**Behavioural Correlation Analyses.** To investigate the relationship between changes in manifold structure and individual differences in learning performance, we computed a correlation, across participants, between learning score and each region's change in eccentricity from Baseline to Early Learning (Fig. 6C) and from Early to Late Learning (Fig. 6D). This produced two correlation maps, one for each contrast. We found that the spatial specificity of significant regions in these correlation maps overlapped substantially with the Dorsal Attention A (DAN-A) network, from the 17-network Schaefer 1000 assignments [58]. This was determined by evaluating the mean correlation (across regions) for each of the 17-network assignments against a null distribution generated by projecting each brain region's correlation onto their respective Schaefer 1000 parcels on the 32k fsLR spherical mesh and performing 1000 iterations of the Váša spin-testing permutation procedure [76,77]. This allowed us to generate, for each brain network and pairwise comparison (Baseline to Early Learning and Early to Late learning), a topographical distribution of correlations that could be expected simply due to chance from spatial autocorrelations in the brain maps (see Fig. 6E and F). We then empirically assessed the statistical significance of our real correlation values against this spatial null distribution for each brain network. Because the spin-testing procedure assesses only the probability of having observed a given correlation value, and not whether that correlation value itself differs from zero, we incorporated the additional stipulation that an effect would be deemed significant only if the real correlation value was also statistically different from zero (at $p<0.05$).

To explore the underlying changes in functional connectivity that give way to these brain-behaviour correlations, we performed seed connectivity contrasts, using paired t-tests on the 17-network parcellation, between the different task epochs (Early > Baseline and Late > Early). As with the 'Seed Connectivity Analyses' section above, we opted to show the unthresholded t-maps so as to visualize the complete multivariate pattern of connectivity that underlies the brain-behaviour correlations (Fig. 6G). Together, these complementary approaches enabled us to explore how individual differences in performance relate to changes in manifold structure at the region- and network-levels.

***Data availability***

Upon acceptance for publication, all behavioral and imaging data (including T1w and functional scans) will be posted in a repository at OpenNeuro (https://openneuro.org/). We will also upload processed behavioral data and ROI time courses in a Dryad repository (https://datadryad.org/).

***Code availability***

Imaging data were preprocessed using fmriPrep, which is open source and freely available. Operations on covariance matrices, including estimation and centering, were performed using the R package spdm, which is freely available in a repository at https://github.com/areshenk-rpackages/spdm. Regional timeseries extraction was performed using Nixtract, which is freely available in a repository at https://github.com/danjgale/nixtract.

***Acknowledgments***

This work was supported by operating grants from the Canadian Institutes of Health Research Grant (MOP126158), the Natural Sciences and Engineering Research Council (RGPIN-2017-04684), and Botterell Foundation Award, as well as funding from the Canadian Foundation for Innovation (35559). The authors would like to thank Martin York and Sean Hickman for technical assistance, and Don O'Brien for assistance with data collection.

***Disclosures***

Jason Gallivan and Daniel Gale serve as part-time employees for Voxel AI Inc.

## REFERENCES


1. Dhawale, A. K., Smith, M. A. & Ölveczky, B. P. The Role of Variability in Motor Learning. *Annu. Rev. Neurosci.* **40**, 479–498 (2017).
2. Houk, J. C., Buckingham, J. T. & Barto, A. G. Models of the cerebellum and motor learning. *Behav. Brain Sci.* **19**, 368–383 (1996).
3. Berniker, M. & Kording, K. Estimating the sources of motor errors for adaptation and generalization. *Nat. Neurosci.* **11**, 1454–1461 (2008).
4. Wolpert, D. M., Ghahramani, Z. & Flanagan, J. R. Perspectives and problems in motor learning. *Trends Cogn. Sci.* **5**, 487–494 (2001).
5. Averbeck, B. & O'Doherty, J. P. Reinforcement-learning in fronto-striatal circuits. *Neuropsychopharmacology* **47**, 147–162 (2022).
6. O'Doherty, J. P., Cockburn, J. & Pauli, W. M. Learning, Reward, and Decision Making. *Annu. Rev. Psychol.* **68**, 73–100 (2017).
7. Lee, D., Seo, H. & Jung, M. W. Neural basis of reinforcement learning and decision making. *Annu. Rev. Neurosci.* **35**, 287–308 (2012).
8. Klein-Flügge, M. C., Bongioanni, A. & Rushworth, M. F. S. Medial and orbital frontal cortex in decision-making and flexible behavior. *Neuron* **110**, 2743–2770 (2022).
9. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
10. Samejima, K., Ueda, Y., Doya, K. & Kimura, M. Representation of action-specific reward values in the striatum. *Science* **310**, 1337–1340 (2005).
11. O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H. & Dolan, R. J. Temporal difference models and reward-related learning in the human brain. *Neuron* **38**, 329–337 (2003).
12. Bayer, H. M. & Glimcher, P. W. Midbrain Dopamine Neurons Encode a Quantitative Reward Prediction Error Signal. *Neuron* vol. 47 129–141 Preprint at https://doi.org/10.1016/j.neuron.2005.05.020 (2005).
13. Corbetta, M., Patel, G. & Shulman, G. L. The reorienting system of the human brain: from environment to theory of mind. *Neuron* **58**, 306–324 (2008).
14. Corbetta, M. & Shulman, G. L. Control of goal-directed and stimulus-driven attention in the brain. *Nat. Rev. Neurosci.* **3**, 201–215 (2002).
15. Buckner, R. L., Andrews-Hanna, J. R. & Schacter, D. L. *The Brain's Default Network*. *Annals of the New York Academy of Sciences* vol. 1124 1–38 Preprint at https://doi.org/10.1196/annals.1440.011 (2008).
16. Raichle, M. E. The brain's default mode network. *Annu. Rev. Neurosci.* **38**, 433–447 (2015).
17. Smallwood, J. *et al.* The default mode network in cognition: a topographical perspective. *Nat. Rev. Neurosci.* **22**, 503–513 (2021).
18. Fox, M. D. & Raichle, M. E. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat. Rev. Neurosci.* **8**, 700–711 (2007).
19. Margulies, D. S. *et al.* Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences* **113**, 12574–12579 (2016).
20. Binder, J. R., Desai, R. H., Graves, W. W. & Conant, L. L. Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cereb. Cortex* **19**, 2767–2796 (2009).
21. Christoff, K., Irving, Z. C., Fox, K. C. R., Spreng, R. N. & Andrews-Hanna, J. R. Mind-wandering as spontaneous thought: a dynamic framework. *Nat. Rev. Neurosci.* **17**, 718–731 (2016).
22. Schacter, D. L. *et al.* The future of memory: remembering, imagining, and the brain. *Neuron* **76**,

677–694 (2012).
23. Andrews-Hanna, J. R., Smallwood, J. & Spreng, R. N. The default network and self-generated thought: component processes, dynamic control, and clinical relevance. *Ann. N. Y. Acad. Sci.* **1316**, 29–52 (2014).
24. Spreng, R. N., Mar, R. A. & Kim, A. S. N. The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: a quantitative meta-analysis. *J. Cogn. Neurosci.* **21**, 489–510 (2009).
25. Vatansever, D., Menon, D. K. & Stamatakis, E. A. Default mode contributions to automated information processing. *Proc. Natl. Acad. Sci. U. S. A.* **114**, 12821–12826 (2017).
26. Murphy, C. *et al.* Distant from input: Evidence of regions within the default mode network supporting perceptually-decoupled and conceptually-guided cognition. *Neuroimage* **171**, 393–401 (2018).
27. Murphy, C. *et al.* Modes of operation: A topographic neural gradient supporting stimulus dependent and independent cognition. *Neuroimage* **186**, 487–496 (2019).
28. Pearson, J. M., Hayden, B. Y., Raghavachari, S. & Platt, M. L. Neurons in posterior cingulate cortex signal exploratory decisions in a dynamic multioption choice task. *Curr. Biol.* **19**, 1532–1537 (2009).
29. Hayden, B. Y., Smith, D. V. & Platt, M. L. Electrophysiological correlates of default-mode processing in macaque posterior cingulate cortex. *Proceedings of the National Academy of Sciences* vol. 106 5948–5953 Preprint at https://doi.org/10.1073/pnas.0812035106 (2009).
30. Heilbronner, S. R. & Platt, M. L. Causal Evidence of Performance Monitoring by Neurons in Posterior Cingulate Cortex during Learning. *Neuron* vol. 80 1384–1391 Preprint at https://doi.org/10.1016/j.neuron.2013.09.028 (2013).
31. Foster, B. L. *et al.* A tripartite view of the posterior cingulate cortex. *Nature Reviews Neuroscience* vol. 24 173–189 Preprint at https://doi.org/10.1038/s41583-022-00661-x (2023).
32. Pearson, J. M., Heilbronner, S. R., Barack, D. L., Hayden, B. Y. & Platt, M. L. Posterior cingulate cortex: adapting behavior to a changing world. *Trends Cogn. Sci.* **15**, 143–151 (2011).
33. Schuck, N. W. *et al.* Medial Prefrontal Cortex Predicts Internally Driven Strategy Shifts. *Neuron* vol. 86 331–340 Preprint at https://doi.org/10.1016/j.neuron.2015.03.015 (2015).
34. Trudel, N. *et al.* Polarity of uncertainty representation during exploration and exploitation in ventromedial prefrontal cortex. *Nature Human Behaviour* vol. 5 83–98 Preprint at https://doi.org/10.1038/s41562-020-0929-3 (2020).
35. Barack, D. L., Chang, S. W. C. & Platt, M. L. Posterior Cingulate Neurons Dynamically Signal Decisions to Disengage during Foraging. *Neuron* **96**, 339–347.e5 (2017).
36. Sutton, R. S. & Barto, A. G. *Reinforcement Learning, second edition: An Introduction*. (MIT Press, 2018).
37. Frank, M. J., Doll, B. B., Oas-Terpstra, J. & Moreno, F. Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nat. Neurosci.* **12**, 1062–1068 (2009).
38. Hayden, B. Y., Nair, A. C., McCoy, A. N. & Platt, M. L. Posterior Cingulate Cortex Mediates Outcome-Contingent Allocation of Behavior. *Neuron* vol. 60 19–25 Preprint at https://doi.org/10.1016/j.neuron.2008.09.012 (2008).
39. Gale, D. J. *et al.* Distinct patterns of cortical manifold expansion and contraction underlie human sensorimotor adaptation. *Proceedings of the National Academy of Sciences* **119**, e2209960119 (2022).
40. Vos de Wael, R. *et al.* BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. *Commun Biol* **3**, 103 (2020).
41. Huntenburg, J. M., Bazin, P.-L. & Margulies, D. S. Large-Scale Gradients in Human Cortical

Organization. *Trends Cogn. Sci.* **22**, 21–31 (2018).
42. Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation Through Neural Population Dynamics. (2020) doi:10.1146/annurev-neuro-092619-094115.
43. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: a dynamical systems perspective. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
44. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17**, 1500–1509 (2014).
45. Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A. Neural Manifolds for the Control of Movement. *Neuron* **94**, 978–984 (2017).
46. Shine, J. M. *et al.* Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nat. Neurosci.* **22**, 289–296 (2019).
47. Shine, J. M. *et al.* The Low-Dimensional Neural Architecture of Cognitive Complexity Is Related to Activity in Medial Thalamic Nuclei. *Neuron* **104**, 849–855.e3 (2019).
48. Vázquez-Rodríguez, B. *et al.* Gradients of structure–function tethering across neocortex. *Proceedings of the National Academy of Sciences* **116**, 21219–21227 (2019).
49. Paquola, C. *et al.* Microstructural and functional gradients are increasingly dissociated in transmodal cortices. *PLOS Biology* vol. 17 e3000284 Preprint at https://doi.org/10.1371/journal.pbio.3000284 (2019).
50. Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B. & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* **441**, 876–879 (2006).
51. Rushworth, M. F. S., Noonan, M. P., Boorman, E. D., Walton, M. E. & Behrens, T. E. Frontal cortex and reward-guided learning and decision-making. *Neuron* **70**, 1054–1069 (2011).
52. McDougle, S. D. *et al.* Credit assignment in movement-dependent reinforcement learning. *Proc. Natl. Acad. Sci. U. S. A.* **113**, 6797–6802 (2016).
53. McDougle, S. D. *et al.* Neural Signatures of Prediction Errors in a Decision-Making Task Are Modulated by Action Execution Failures. *Curr. Biol.* **29**, 1606–1613.e5 (2019).
54. Dhawale, A. K., Miyamoto, Y. R., Smith, M. A. & Ölveczky, B. P. Adaptive Regulation of Motor Variability. *Curr. Biol.* **29**, 3551–3562.e7 (2019).
55. Wu, H. G., Miyamoto, Y. R., Gonzalez Castro, L. N., Ölveczky, B. P. & Smith, M. A. Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat. Neurosci.* **17**, 312–321 (2014).
56. Wilson, R. C., Bonawitz, E., Costa, V. D. & Ebitz, R. B. Balancing exploration and exploitation with information and randomization. *Curr Opin Behav Sci* **38**, 49–56 (2021).
57. Dam, G., Kording, K. & Wei, K. Credit assignment during movement reinforcement learning. *PLoS One* **8**, e55352 (2013).
58. Schaefer, A. *et al.* Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cereb. Cortex* **28**, 3095–3114 (2018).
59. Avants, B. B., Epstein, C. L., Grossman, M. & Gee, J. C. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* **12**, 26–41 (2008).
60. Frazier, J. A. *et al.* Structural brain magnetic resonance imaging of limbic and thalamic volumes in pediatric bipolar disorder. *Am. J. Psychiatry* **162**, 1256–1265 (2005).
61. Gratton, C. *et al.* Functional Brain Networks Are Dominated by Stable Group and Individual Factors, Not Cognitive or Daily Variation. *Neuron* **98**, 439–452.e5 (2018).
62. Gordon, E. M. *et al.* Precision Functional Mapping of Individual Human Brains. *Neuron* **95**, 791–807.e7 (2017).
63. Areshenkoff, C. *et al.* Neural excursions from manifold structure explain patterns of learning during human sensorimotor adaptation. *Elife* **11**, (2022).
64. Areshenkoff, C. N. *et al.* Muting, not fragmentation, of functional brain networks under general

anesthesia. *Neuroimage* **231**, 117830 (2021).
65. Zhao, Q., Kwon, D. & Pohl, K. M. A Riemannian Framework for Longitudinal Analysis of Resting-State Functional Connectivity. *Med. Image Comput. Comput. Assist. Interv.* **11072**, 145–153 (2018).
66. McInnes, L., Healy, J. & Melville, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv [stat.ML]* (2018).
67. Hong, S.-J. *et al.* Toward a connectivity gradient-based framework for reproducible biomarker discovery. *NeuroImage* vol. 223 117322 Preprint at https://doi.org/10.1016/j.neuroimage.2020.117322 (2020).
68. Yeo, B. T. T. *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* **106**, 1125–1165 (2011).
69. Park, B.-Y. *et al.* An expanding manifold in transmodal regions characterizes adolescent reconfiguration of structural connectome organization. *Elife* **10**, (2021).
70. Park, B.-Y. *et al.* Inter-individual body mass variations relate to fractionated functional brain hierarchies. *Communications Biology* vol. 4 Preprint at https://doi.org/10.1038/s42003-021-02268-x (2021).
71. Valk, S. *et al. Changing the social brain: plasticity along macro-scale axes of functional connectivity following social mental training*. https://juser.fz-juelich.de/record/905774/ (2021).
72. de Brouwer, A. J. *et al.* Human Variation in Error-Based and Reinforcement Motor Learning Is Associated With Entorhinal Volume. *Cereb. Cortex* **32**, 3423–3440 (2022).
73. de Brouwer, A. J., Albaghdadi, M., Flanagan, J. R. & Gallivan, J. P. Using gaze behavior to parcellate the explicit and implicit contributions to visuomotor learning. *J. Neurophysiol.* **120**, 1602–1615 (2018).
74. Standage, D. I. *et al.* Whole-brain dynamics of human sensorimotor adaptation. *Cereb. Cortex* (2022) doi:10.1093/cercor/bhac378.
75. Shang, H. L. A survey of functional principal component analysis. *Adv. Stat. Anal.* **98**, 121–142 (2014).
76. Markello, R. D. *et al.* neuromaps: structural and functional interpretation of brain maps. *Nat. Methods* **19**, 1472–1479 (2022).
77. Váša, F. *et al.* Adolescent Tuning of Association Cortex in Human Structural Brain Networks. *Cereb. Cortex* **28**, 281–294 (2018).
78. Botvinick, M. M., Cohen, J. D. & Carter, C. S. Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn. Sci.* **8**, 539–546 (2004).
79. Shenhav, A., Straccia, M. A., Cohen, J. D. & Botvinick, M. M. Anterior cingulate engagement in a foraging context reflects choice difficulty, not foraging value. *Nat. Neurosci.* **17**, 1249–1254 (2014).
80. Holroyd, C. B. & Yeung, N. Motivation of extended behaviors by anterior cingulate cortex. *Trends Cogn. Sci.* **16**, 122–128 (2012).
81. Hogeveen, J. *et al.* The neurocomputational bases of explore-exploit decision-making. *Neuron* **110**, 1869–1879.e5 (2022).
82. Blanchard, T. C. & Gershman, S. J. Pure correlates of exploration and exploitation in the human brain. *Cogn. Affect. Behav. Neurosci.* **18**, 117–126 (2018).
83. Kolling, N. *et al.* Value, search, persistence and model updating in anterior cingulate cortex. *Nat. Neurosci.* **19**, 1280–1285 (2016).
84. Boorman, E. D., Behrens, T. E. J., Woolrich, M. W. & Rushworth, M. F. S. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron* **62**, 733–743 (2009).
85. Badre, D., Doll, B. B., Long, N. M. & Frank, M. J. Rostrolateral prefrontal cortex and individual

differences in uncertainty-driven exploration. *Neuron* **73**, 595–607 (2012).
86. Shamloo, F. & Helie, S. Changes in default mode network as automaticity develops in a categorization task. *Behav. Brain Res.* **313**, 324–333 (2016).
87. Luo, L. *et al.* Ten years of Nature Reviews Neuroscience: insights from the highly cited. *Nat. Rev. Neurosci.* **11**, 718–726 (2010).
88. Leech, R., Braga, R. & Sharp, D. J. Echoes of the brain within the posterior cingulate cortex. *J. Neurosci.* **32**, 215–222 (2012).
89. McCoy, A. N. & Platt, M. L. Risk-sensitive neurons in macaque posterior cingulate cortex. *Nature Neuroscience* vol. 8 1220–1227 Preprint at https://doi.org/10.1038/nn1523 (2005).
90. Kable, J. W. & Glimcher, P. W. The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* **10**, 1625–1633 (2007).
91. Knutson, B. & Bossaerts, P. Neural antecedents of financial decisions. *J. Neurosci.* **27**, 8174–8177 (2007).
92. Wan, X., Cheng, K. & Tanaka, K. Neural encoding of opposing strategy values in anterior and posterior cingulate cortex. *Nat. Neurosci.* **18**, 752–759 (2015).
93. Lerner, Y., Honey, C. J., Silbert, L. J. & Hasson, U. Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *Journal of Neuroscience* **31**, 2906–2915 (2011).
94. Hasson, U., Chen, J. & Honey, C. J. Hierarchical process memory: memory as an integral component of information processing. *Trends Cogn. Sci.* **19**, 304–313 (2015).
95. Hagmann, P. *et al.* Mapping the structural core of human cerebral cortex. *PLoS Biol.* **6**, e159 (2008).
96. Margulies, D. S. *et al.* Precuneus shares intrinsic functional architecture in humans and monkeys. *Proc. Natl. Acad. Sci. U. S. A.* **106**, 20069–20074 (2009).
97. Hutchison, R. M., Culham, J. C., Flanagan, J. R., Everling, S. & Gallivan, J. P. Functional subdivisions of medial parieto-occipital cortex in humans and nonhuman primates using resting-state fMRI. *Neuroimage* **116**, 10–29 (2015).
98. Braga, R. M., Sharp, D. J., Leeson, C., Wise, R. J. S. & Leech, R. Echoes of the brain within default mode, association, and heteromodal cortices. *Journal of Neuroscience* **33**, 14031–14039 (2013).
99. Vos de Wael, R. *et al.* Structural connectivity gradients of the temporal lobe serve as multiscale axes of brain organization and cortical evolution. *Cereb. Cortex* **31**, 5151–5164 (2021).
100. Tian, Y., Margulies, D. S., Breakspear, M. & Zalesky, A. Topographic organization of the human subcortex unveiled with functional connectivity gradients. *Nat. Neurosci.* **23**, 1421–1432 (2020).
101. Oldfield, R. C. The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia* **9**, 97–113 (1971).
102. Wolpert, D. M., Diedrichsen, J. & Flanagan, J. R. Principles of sensorimotor learning. *Nat. Rev. Neurosci.* **12**, 739–751 (2011).
103. Pekny, S. E., Izawa, J. & Shadmehr, R. Reward-dependent modulation of movement variability. *J. Neurosci.* **35**, 4015–4024 (2015).
104. Härdle, W. Applied Nonparametric Regression. Preprint at https://doi.org/10.1017/ccol0521382483 (1990).
105. Ramsay, J. & Silverman, B. W. *Functional Data Analysis*. (Springer Science & Business Media, 2013).
106. Ramsay, J., Wickham, H., Ramsay, M. J., and deSolve, S. Package 'fda'. (2022).
107. Esteban, O. *et al.* fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat. Methods* **16**, 111–116 (2019).
108. Esteban, O., Blair, R., Markiewicz, C. J. & Berleant, S. L. fMRIPrep. Software. *Zenodo*.
109. Gorgolewski, K. *et al.* Nipype: a flexible, lightweight and extensible neuroimaging data

processing framework in python. *Front. Neuroinform.* **5**, 13 (2011).
110. Gorgolewski, K. J. *et al.* Nipype [Software]. Zenodo. Preprint at (2018).
111. Abraham, A. *et al.* Machine learning for neuroimaging with scikit-learn. *Front. Neuroinform.* **8**, 14 (2014).
112. Tustison, N. J. *et al.* N4ITK: improved N3 bias correction. *IEEE Trans. Med. Imaging* **29**, 1310–1320 (2010).
113. Zhang, Y., Brady, M. & Smith, S. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans. Med. Imaging* **20**, 45–57 (2001).
114. Reuter, M., Rosas, H. D. & Fischl, B. Highly accurate inverse consistent registration: a robust approach. *Neuroimage* **53**, 1181–1196 (2010).
115. Dale, A. M., Fischl, B. & Sereno, M. I. Cortical Surface-Based Analysis. *NeuroImage* vol. 9 179–194 Preprint at https://doi.org/10.1006/nimg.1998.0395 (1999).
116. Klein, A. *et al.* Mindboggling morphometry of human brains. *PLoS Comput. Biol.* **13**, e1005350 (2017).
117. Jenkinson, M., Bannister, P., Brady, M. & Smith, S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* **17**, 825–841 (2002).
118. Cox, R. W. & Hyde, J. S. Software tools for analysis and visualization of fMRI data. *NMR Biomed.* **10**, 171–178 (1997).
119. Greve, D. N. & Fischl, B. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage* **48**, 63–72 (2009).
120. Lanczos, C. Evaluation of Noisy Data. *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis* **1**, 76–85 (1964).
121. Glasser, M. F. *et al.* The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* vol. 80 105–124 Preprint at https://doi.org/10.1016/j.neuroimage.2013.04.127 (2013).
122. Satterthwaite, T. D. *et al.* An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *Neuroimage* **64**, 240–256 (2013).
123. Behzadi, Y., Restom, K., Liau, J. & Liu, T. T. A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage* **37**, 90–101 (2007).
124. Muschelli, J. *et al.* Reduction of motion-related artifacts in resting state fMRI using aCompCor. *Neuroimage* **96**, 22–35 (2014).
125. Makris, N. *et al.* Decreased volume of left and total anterior insular lobule in schizophrenia. *Schizophrenia Research* vol. 83 155–171 Preprint at https://doi.org/10.1016/j.schres.2005.11.020 (2006).
126. Ledoit, O. & Wolf, M. A well-conditioned estimator for large-dimensional covariance matrices. *J. Multivar. Anal.* **88**, 365–411 (2004).
127. Bethlehem, R. A. I. *et al.* Dispersion of functional gradients across the adult lifespan. *Neuroimage* **222**, 117299 (2020).
128. Rubinov, M. & Sporns, O. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage* **52**, 1059–1069 (2010).