

Reconfigurations of cortical manifold structure during reward-based motor learning

Reviewed Preprint

v2 • April 8, 2024

Revised by authors

Reviewed Preprint

v1 • November 9, 2023

Qasem Nick, Daniel J. Gale, Corson Areshenkoff, Anouk De Brouwer, Joseph Nashed, Jeffrey Wammes, Tianyao Zhu, Randy Flanagan, Jonny Smallwood, Jason Gallivan 

Centre for Neuroscience Studies, Queen's University, Kingston, Ontario, Canada • Department of Psychology, Queen's University, Kingston, Ontario, Canada • Department of Biomedical and Molecular Sciences, Queen's University, Kingston, Ontario, Canada

 https://en.wikipedia.org/wiki/Open_access

 Copyright information

Abstract

Adaptive motor behavior depends on the coordinated activity of multiple neural systems distributed across the brain. 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 striatal 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 neural changes that support reward-based motor learning and identify distinct transitions in the functional coupling of sensorimotor to transmodal cortex when adapting behavior.

eLife assessment

This **valuable** study uses **convincing** state-of-the-art neuroimaging analyses to characterise whole-brain networks during reward-based motor learning. This work motivates future research to dissociate why the observed changes in neural connectivity occur and how they support reward-based motor learning. The study is highly relevant for researchers at the intersection of decision-making and sensorimotor learning.

<https://doi.org/10.7554/eLife.91928.2.sa3>

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 (Berniker and Kording, 2008 [↗](#); Dhawale et al., 2017 [↗](#); Houk et al., 1996 [↗](#); Wolpert et al., 2001 [↗](#)). Moreover, this feedback must be communicated to multiple discrete neural systems distributed across the cortex and striatum, many of which are topographically segregated from one another (Averbeck and O'Doherty, 2022 [↗](#)). 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 (Averbeck and O'Doherty, 2022 [↗](#); Klein-Flügge et al., 2022 [↗](#); Lee et al., 2012 [↗](#); O'Doherty et al., 2017 [↗](#)). 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 (Bayer and Glimcher, 2005 [↗](#); O'Doherty et al., 2003 [↗](#); Samejima et al., 2005 [↗](#); Schultz et al., 1997 [↗](#)). 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 (Corbetta et al., 2008 [↗](#); Corbetta and Shulman, 2002 [↗](#)). 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) (Buckner et al., 2008 [↗](#); Fox and Raichle, 2007 [↗](#); Margulies et al., 2016 [↗](#); Raichle, 2015 [↗](#); Smallwood et al., 2021 [↗](#)).

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 (Andrews-Hanna et al., 2014 [↗](#); Binder et al., 2009 [↗](#); Christoff et al., 2016 [↗](#); Schacter et al., 2012 [↗](#); Spreng et al., 2009 [↗](#)). 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 (Foster et al., 2023 [↗](#); Hayden et al., 2009 [↗](#); Heilbronner and Platt, 2013 [↗](#); Murphy et al., 2019 [↗](#), 2018 [↗](#); Pearson et al., 2009 [↗](#); Vatansever et al., 2017 [↗](#)). One hypothesis concerning this system's function pertains to its unique topographic positioning on the cortical mantle (Smallwood et al., 2021 [↗](#)): 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 (Margulies et al., 2016 [↗](#)). 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 (Smallwood et al., 2021 [↗](#)).

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 (Barack et al., 2017 [DOI](#); Foster et al., 2023 [DOI](#); Pearson et al., 2011 [DOI](#), 2009 [DOI](#); Schuck et al., 2015 [DOI](#); Trudel et al., 2020 [DOI](#)) — the so-called explore/exploit trade-off (Frank et al., 2009 [DOI](#); Sutton and Barto, 2018 [DOI](#)). 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 (Hayden et al., 2008 [DOI](#); Murphy et al., 2019 [DOI](#), 2018 [DOI](#); Vatansever et al., 2017 [DOI](#)). 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 (Gale et al., 2022 [DOI](#)). Given that the DMN is hypothesized to exert influence on functional brain activity via its topographic positioning on cortex (Smallwood et al., 2021 [DOI](#)), 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 striatal 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 cortical activity (Huntenburg et al., 2018 [DOI](#); Margulies et al., 2016 [DOI](#); Vos de Wael et al., 2020 [DOI](#)). 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 (Cunningham and Yu, 2014 [DOI](#); Gallego et al., 2017 [DOI](#); Shenoy et al., 2013 [DOI](#); Vyas et al., 2020 [DOI](#)), 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 (Huntenburg et al., 2018 [DOI](#); Paquola et al., 2019 [DOI](#); Shine et al., 2019a [DOI](#), 2019b [DOI](#); Vázquez-Rodríguez et al., 2019 [DOI](#)). Here we applied this manifold approach to explore how brain activity across widely distributed cortical and striatal systems is coordinated during reward-based motor learning. We were particularly interested in characterizing 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 exploring 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 (Averbeck and O'Doherty, 2022 [DOI](#); Daw et al., 2006 [DOI](#); Klein-Flügge et al., 2022 [DOI](#); Lee et al., 2012 [DOI](#); O'Doherty et al., 2017 [DOI](#); Rushworth et al., 2011 [DOI](#)). 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 (McDougle et al., 2019 [DOI](#), 2016 [DOI](#)). However, recent theories on learning, supported by both human and animal studies (Dhawale et al., 2019 [DOI](#), 2017 [DOI](#); Wu et al., 2014 [DOI](#)), have highlighted the crucial role of movement variability — and motor exploration in particular — as a key ingredient for effective learning (Dam et al., 2013 [DOI](#); Wilson et al., 2021 [DOI](#); Wu et al., 2014 [DOI](#)). In order to incorporate this aspect to learning, and inspired by recent work in the field

(Dam et al., 2013 [↗](#); Wu et al., 2014 [↗](#)), 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 (Dam et al., 2013 [↗](#); Wu et al., 2014 [↗](#)). That is, subjects could not use error-based learning mechanisms to achieve learning in our study, as this form of learning requires sensory errors that convey both the change in direction and magnitude needed to correct the movement.

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 (see Supplemental Fig. 1 for other measures of changes in subject motor behavior throughout learning, and Supplemental Fig. 2 for examples of all subjects’ finger trajectories across trials).

In order to study the changes in functional cortical and striatal 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 (Schaefer et al., 2018 [↗](#)) and for striatal regions defined by the Harvard-Oxford parcellation (Avants et al., 2008 [↗](#); Frazier et al., 2005 [↗](#)), and then estimated covariance (functional connectivity) matrices for each epoch (Baseline, Early and Late; **Fig. 1E** [↗](#)) (for a similar approach, see Gale et al., 2022 [↗](#)).

Because prior work (Gordon et al., 2017 [↗](#); Gratton et al., 2018 [↗](#)), including our own (Areshenkoff et al., 2022 [↗](#), 2021 [↗](#); Gale et al., 2022 [↗](#)), suggests that individual differences in functional connectivity can obscure any task-related effects, we centered the connectivity matrices using the Riemmanian manifold approach (Areshenkoff et al., 2022 [↗](#), 2021 [↗](#); Gale et al., 2022 [↗](#); Zhao et al., 2018 [↗](#)) (See Supplemental Figure 3 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; (McInnes et al., 2018 [↗](#))). 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 (Gratton et al., 2018 [↗](#)). 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 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. Black line denotes the mean across participants and whereas the gray banding denotes ± 1 SEM. 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 striatal 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 Procrustes alignment. This allowed us to assess learning-related changes in manifold structure from this Baseline architecture.




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 striatal connectivity during the reward-based motor learning task, we used the centered matrices from [Fig. 2B](#) to estimate separate cortical-striatal connectivity manifolds for each participant's Baseline, Early and Late covariance matrices (see also [Gale et al., 2022](#)). Following from prior work ([Hong et al., 2020](#); [Paguola et al., 2019](#); [Vos de Wael et al., 2020](#)), 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-striatal functional organization (i.e., a cortical-striatal 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 ([Vos de Wael et al., 2020](#)) 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-striatal manifold structure during Baseline trials

The top three principal components (PCs) of the template Baseline manifold ([Fig. 3A](#)) describe the cortical-striatal 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 ([Yeo et al., 2011](#)), 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 ([Huntenburg et al., 2018](#); [Margulies et al., 2016](#)) ([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 ([Huntenburg et al., 2018](#); [Margulies et al., 2016](#); [Smallwood et al., 2021](#)).

We next sought to characterize the relative positions of cortical and striatal 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 ([Gale et al., 2022](#); [Park et al., 2021a](#)), 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 ([Park et al., 2021a](#); [2021b](#); [Valk et al., 2021](#)). 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 4). Likewise, we find that Baseline eccentricity is negatively related to a region's participation coefficient ($r=-0.74$,




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 (Schaefer et al., 2018; Yeo et al., 2011). (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 striatal surfaces (right). White circle with black bordering denotes the center of the manifold (manifold centroid).

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.

Changes in cortical-striatal 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 striatal regions from Baseline to Early learning (however, for interested readers, Supplementary Fig 5 shows the unthresholded data from both cortex and striatum 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 striatal 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).

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 6). 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




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, (Yeo et al., 2011)). 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).

AG in particular — begin to integrate with regions outside of their respective networks. By contrast, during late learning, there is a clear 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 7). 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 ([Schaefer et al., 2018](#); [Yeo et al., 2011](#))].

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.




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 (Yeo et al., 2011)). 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).

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 (de Brouwer et al., 2022; 2018; Standage et al., 2022; Wu et al., 2014). 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 Supplementary Fig. 2). 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 in a manner that accounted the auto-correlation in learning performance over time (i.e., subjects who learned more quickly tend to exhibit better performance by the end of learning), we opted for a pure data-driven approach and performed functional principal component analysis (fPCA; (Shang, 2014)) on subjects’ learning curves. This approach allowed us to isolate the dominant patterns of variability in subject’s learning curves over time (see Methods for further details; see also Areshenkoff et al., 2022). 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. 6C**). 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 (see single data points in **Fig. 6C** at right).

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. 7A**). 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




Figure 6

Individual differences in subject learning performance.

(A) Examples of a good learner (bordered in green) and poor learner (bordered in red). (B) Individual subject learning curves for the task. Solid black line denotes the mean across all subjects whereas light gray lines denote individual participants. The green and red traces denote the learning curves for the example good and poor learners denoted in A. (C) Derivation of subject learning scores. We performed functional principal component analysis (fPCA) 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. The plot at right denotes the loading for each participant (open circles) onto this Learning score component.

effects across brain areas, which may correspond with known functional networks. Indeed, the full cortex and striatal correlation map in **Fig. 7A** 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 (Yeo et al., 2011) 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 (Markello et al., 2022; Váša et al., 2018) (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. 7C**), 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. 7C**, 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. 7D**). We should note that this pattern of performance effects for the DAN-A — i.e., greater contraction during Early learning and greater expansion during Late learning being associated with better learning — appears at odds with the group-level effects described in **Fig. 4A** and **B**, where we generally find the opposite pattern for the entire DAN network (composed of the DAN-A and DAN-B subnetworks). However, this potential discrepancy can be explained when examining the changes in eccentricity using the 17-network parcellation (see Supplementary Figure 8). At this higher resolution level we find that these group-level effects for the entire DAN network are being largely driven by eccentricity changes in the DAN-B network (areas in anterior superior parietal cortex and premotor cortex), and not by mean changes in the DAN-A network. By contrast, our present results suggest that it is the contraction and expansion of areas of the DAN-A network (and not DAN-B network) that 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. 7E**), 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. 7E**, 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.




Figure 7

Relationship between learning performance and regional changes in eccentricity.

(A & B) Whole-brain correlation map between subject Learning score and the change in regional eccentricity from Baseline to Early learning (A) and Early to Late learning (B). Black bordering denotes regions that are significant at $p < 0.05$. (C & D) Results of the spin-test permutation procedure, assessing whether the topography of correlations in A and B are specific to individual functional brain networks. Single points indicate the real correlation value for each of the 17 Yeo et al. networks (Yeo et al., 2011), whereas the boxplots represent the parameters of a null distribution of correlations derived from 1000 iterations of a spatial autocorrelation-preserving null model (Markello et al., 2022; Váša et al., 2018). 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. (E) 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.

Discussion

Complex behavior necessitates the coordinated activity of multiple specialized neural systems distributed across cortex and striatum. 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 cortical and striatal 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 (Botvinick et al., 2004 [↗](#); Holroyd and Yeung, 2012 [↗](#); Pearson et al., 2009 [↗](#); Shenhav et al., 2014 [↗](#)), 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 (Corbetta et al., 2008 [↗](#)). In addition, recent work (Badre et al., 2012 [↗](#); Blanchard and Gershman, 2018 [↗](#); Boorman et al., 2009 [↗](#); Daw et al., 2006 [↗](#); Hogeveen et al., 2022 [↗](#); Kolling et al., 2016 [↗](#)) implicates several SalVentAttn areas in information gathering functions, so as to optimize reward outcomes. This neural perspective is consistent with both others' (Dam et al., 2013 [↗](#)) and our own behavioral findings (de Brouwer et al., 2022 [↗](#)) that, during the early phases of reward-guided learning, performance is more variable, presumably as individuals explore 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 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 (Buckner et al., 2008 [↗](#); Schacter et al., 2012 [↗](#); Smallwood et al., 2021 [↗](#)). This interpretation is consistent with prior work showing that connectivity between the DMN and premotor cortex tends to increase once rules have been learned (Shamlou and Helie, 2016 [↗](#)), 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 (Murphy et al., 2019; 2018; Vatansever et al., 2017). Our analysis, therefore, provides additional evidence, albeit 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 (Smallwood et al., 2021) — in this case, the history of reward information accrued across previous movement trajectories. Finally, it is important to note that the reversal pattern of effects noted above suggests that our findings during learning cannot be simply attributed to the introduction of reward feedback and/or the perturbation during Early learning, as both of these task-related features are also present during Late learning. In addition, these results cannot be simply explained due to the passage of time or increasing subject fatigue, as this would predict a consistent directional change in eccentricity across the Baseline, Early and Late learning epochs.

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 performed at the task. Our connectivity analysis revealed that this change in DAN-A activity was mainly driven by increases in connectivity with SalVentAttn subnetworks during early learning and by increases in connectivity 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 (Luo et al., 2010). 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. It is also possible that some of these task-related shifts in connectivity relate to shifts in task-general processes, such as changes in the allocation of attentional resources (Bédard and Song, 2013; Rosenberg et al., 2016) or overall cognitive engagement (Aben et al., 2020), which themselves play critical roles in shaping learning (Codol et al., 2018; Holland et al., 2018; Song, 2019; Taylor and Thoroughman, 2008, 2007; for a review of these topics, see Tsay et al., 2023). Such processes are particularly important during the earlier phases of learning when sensorimotor contingencies need to be established. While these remain questions for future work, our data nevertheless suggest that this shift in connectivity may be enabled through the PMC.

Although traditionally considered a member of the DMN (Yeo et al., 2011), studies have established that the PMC contains echoes of neural signals originating from across the cortex (Leech et al., 2012). 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 members of the DMN. This pattern of changes differed from other areas of the DMN, indicating 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 (Kable and Glimcher, 2007; Knutson and Bossaerts, 2007; McCoy and Platt, 2005) and the selection of response strategies that attempt to optimize reward outcomes (Barack et al., 2017; Pearson et al., 2009; Wan et al., 2015). Consistent with this, recent research shows that areas within the PMC are able to integrate information over particularly long periods of time (Hasson et al., 2015; Heilbronner and Platt, 2013; Lerner et al., 2011). 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 (Hagmann et al., 2008 [↗](#); Hutchison et al., 2015 [↗](#); Margulies et al., 2009 [↗](#)), suggest an important role for this region in directing long-term behavioral adaptation in accordance with higher-order task objectives (Braga et al., 2013 [↗](#); Pearson et al., 2011 [↗](#)).

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 the evidence indicates that this region plays a key role in reward-guided behavior (Averbeck and O'Doherty, 2022 [↗](#); O'Doherty et al., 2017 [↗](#)), 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 analysis 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 striatum 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 (Tian et al., 2020 [↗](#); Vos de Wael et al., 2021 [↗](#)). 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 striatal 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 striatal regions were included in the analysis (see Supplementary Fig. 9). As such, selective changes in the patterns of functional connectivity at the level of the striatum may be obscured using our cortex x striatum 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)(Tian et al., 2020 [↗](#)) and by focusing specifically on changes in the interactions *between* the striatum and cortex during learning. The latter can 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 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 striatum 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 cortical 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.

Materials and methods

Participants

Forty-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 (Oldfield, 1971 [↗](#)) 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 0.5° 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 (Pekny et al., 2015 [↗](#); Wolpert et al., 2011 [↗](#)). This task was a modification on the motor tasks developed by (Dam et al., 2013 [↗](#); Wu et al., 2014 [↗](#)).

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, other than this color change in the distance marker, the visible curved path remained constant and 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 (during which the screen was blanked), 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. Note that, at the end of testing, we did not assess participants' awareness of the manipulation (i.e., that they were, in fact, being rewarded based on a mirror image path of the visible path).

Behavioral data analysis

Data Preprocessing

Each movement trajectory was first re-sampled to 10 equally spaced points along the y (vertical) axis, between the starting position and the target distance marker. We defined subjects' reaction time (RT) as the time between trial onset and the cursor reaching 10% of the distance from the starting location, and defined subjects' movement time (MT) as the remaining time until reaching the target distance marker. To quantify the variability of the drawn path as a function of trial number, we used the following method. First, we calculated the average path at each trial by applying a moving average smoother (window size 7 trials) to the sequence of paths drawn by each participant. Then, for each trial, path variability was measured by the mean absolute x position difference between the actual path drawn and the average path in that trial (across the 10 sample points equally spaced along the y direction). Specifically, path variability was quantified in the following manner:

Path trial t : X_t^k ($k = 1, 2, \dots, 10$, i.e., 10 sample points)

For each t , $\overline{X}_t^k = \frac{1}{7} \sum_{i=t-3}^{t+3} X_i^k$ (sliding window mean, half window width 3)

Variability of path in trial t : $\frac{1}{10} \sum_{k=1}^{10} |X_t^k - \overline{X}_t^k|$

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 (Härdle, 1990 [\[1\]](#)). We then performed *functional PCA* (Ramsay and Silverman, 2013 [\[2\]](#)), 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* (Ramsay, J., Wickham, H., Ramsay, M. J., and deSolve, S., 2022 [\[3\]](#)), 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 and cerebellum. Note that for the current study, we did not examine changes in cerebellar activity during learning. 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 (Esteban et al., 2019, n.d.) (RRID:SCR_016216) which is based on Nipype 1.5.0 (Gorgolewski et al., 2011, 2018) (RRID:SCR_002502). Many internal operations of fMRIPrep use Nilearn 0.6.2 (Abraham et al., 2014) (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 (Tustison et al., 2010), distributed with ANTs 2.2.0 (Avants et al., 2008) (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) (Zhang et al., 2001). A T1w-reference map was computed after registration of the T1w images (after INU-correction) using mri_robust_template (FreeSurfer 6.0.1) (Reuter et al., 2010). Brain surfaces were reconstructed using recon-all (FreeSurfer 6.0.1, RRID:SCR_001847) (Dale et al., 1999), 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) (Klein et al., 2017). 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) (Jenkinson et al., 2002). BOLD runs were slice-time corrected using 3dTshift from AFNI 20160207 (Cox and Hyde, 1997) (RRID:SCR_005927). The BOLD reference was then co-registered to the T1w reference using bbregister (FreeSurfer) which implements boundary-based registration (Greve and Fischl, 2009). 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 (Lanczos, 1964). Striatal data in standard space was combined with resampled BOLD time series on the fsaverage surface to produce Grayordinates files (Glasser et al., 2013) 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 (Satterthwaite et al., 2013). Ten component-based physiological regressors were estimated using the aCompCor approach (Behzadi et al., 2007; Muschelli et al., 2014), 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 (Schaefer et al., 2018); two regions are removed from the parcellation due to their small parcel size) and, (2) each of the 12 striatal regions defined according to the Harvard Oxford atlas (Frazier et al., 2005; Makris et al., 2006), 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 (Ledoit and Wolf, 2004). 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 (Zhao et al., 2018), which leverages the natural geometry of the space of covariance matrices (Areshenkoff et al., 2022). 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 $\bar{S}_i$ 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 (Zhao et al., 2018 [2](#)), obtaining a centered tangent vector

$$T_{ij}^c = GT_{ij}G^T$$

where $G = \bar{S}_{gm}^{-1/2} \bar{S}_i^{-1/2}$. 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

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

where $\exp$ denotes the matrix exponential. For the benefits of this centering approach, see **Fig. 2** [2](#), and for an additional overview, see (Areshenkoff et al., 2022 [2](#)).

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> [2](#).

Manifold Construction

Connectivity manifolds for all centered functional connectivity matrices were derived in the following steps. First, consistent with previous studies (Gale et al., 2022 [2](#); Hong et al., 2020 [2](#); Margulies et al., 2016 [2](#); Vos de Wael et al., 2020 [2](#)), 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 (Hong et al., 2020 [2](#))).

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 $\times$ 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 $\sim 70\%$ 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 (Bethlehem et al., 2020 [2](#); Park et al., 2021a [2](#), 2021b [2](#); Valk et al., 2021 [2](#)). “Eccentricity” refers to the Euclidean distance between a single region and the manifold centroid (Park et al., 2021a [2](#)), 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 (Park et al., 2021a [↗](#), 2021b [↗](#); Valk et al., 2021 [↗](#)). 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 coefficient*, which measures the network diversity of a region's connectivity distribution (Rubinov and Sporns, 2010 [↗](#)). Regions were assigned to their respective intrinsic functional networks (Schaefer et al., 2018 [↗](#); Yeo et al., 2011 [↗](#)) 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 cortical and striatal 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 (Yeo et al., 2011 [↗](#)). 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 (Schaefer et al., 2018 [↗](#)). 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 (Markello et al., 2022 [↗](#); Váša et al., 2018 [↗](#)). 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/↗>).

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/↗>.

Acknowledgements

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 are employees of Voxel AI Inc. The other authors report no conflicts of interest. These funding sources had no role in the design, management, data analysis, presentation, or interpretation and write-up of the data.

References

Aben B, Buc Calderon C, Van den Bussche E, Verguts T. (2020) **Cognitive Effort Modulates Connectivity between Dorsal Anterior Cingulate Cortex and Task-Relevant Cortical Areas** *J Neurosci* **40**:3838–3848

Abraham A, Pedregosa F, Eickenberg M, Gervais P, Mueller A, Kossaifi J, Gramfort A, Thirion B, Varoquaux G (2014) **Machine learning for neuroimaging with scikit-learn** *Front Neuroinform* **8**

Andrews-Hanna JR, Smallwood J, Spreng RN (2014) **The default network and self-generated thought: component processes, dynamic control, and clinical relevance** *Ann N Y Acad Sci* **1316**:29–52

Areshenkoff C, Gale DJ, Standage D, Nashed JY, Flanagan JR, Gallivan JP (2022) **Neural excursions from manifold structure explain patterns of learning during human sensorimotor adaptation** *Elife* **11** <https://doi.org/10.7554/eLife.74591>

Areshenkoff CN, Nashed JY, Hutchison RM, Hutchison M, Levy R, Cook DJ, Menon RS, Everling S, Gallivan JP (2021) **Muting, not fragmentation, of functional brain networks under general anesthesia** *Neuroimage* **231**

Avants BB, Epstein CL, Grossman M, Gee JC (2008) **Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain** *Med Image Anal* **12**:26–41

Averbeck B, O'Doherty JP (2022) **Reinforcement-learning in fronto-striatal circuits** *Neuropsychopharmacology* **47**:147–162

Badre D, Doll BB, Long NM, Frank MJ (2012) **Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration** *Neuron* **73**:595–607

Barack DL, Chang SWC, Platt ML (2017) **Posterior Cingulate Neurons Dynamically Signal Decisions to Disengage during Foraging** *Neuron* **96**:339–347

Bayer HM, Glimcher PW (2005) **Midbrain Dopamine Neurons Encode a Quantitative Reward Prediction Error Signal** *Neuron* <https://doi.org/10.1016/j.neuron.2005.05.020>

Bédard P, Song J-H (2013) **Attention modulates generalization of visuomotor adaptation** *J Vis* **13**

Behzadi Y, Restom K, Liao J, Liu TT (2007) **A component based noise correction method (CompCor) for BOLD and perfusion based fMRI** *Neuroimage* **37**:90–101

Berniker M, Kording K (2008) **Estimating the sources of motor errors for adaptation and generalization** *Nat Neurosci* **11**:1454–1461

Bethlehem RAI, Paquola C, Seidlitz J, Ronan L, Bernhardt B, Consortium C-C, Tsvetanov KA (2020) **Dispersion of functional gradients across the adult lifespan** *Neuroimage* **222**

Binder JR, Desai RH, Graves WW, Conant LL (2009) **Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies** *Cereb Cortex* **19**:2767–2796

Blanchard TC, Gershman SJ (2018) **Pure correlates of exploration and exploitation in the human brain** *Cogn Affect Behav Neurosci* **18**:117–126

Boorman ED, Behrens TEJ, Woolrich MW, Rushworth MFS (2009) **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

- Botvinick MM, Cohen JD, Carter CS (2004) **Conflict monitoring and anterior cingulate cortex: an update** *Trends Cogn Sci* **8**:539–546
- Braga RM, Sharp DJ, Leeson C, Wise RJS, Leech R (2013) **Echoes of the brain within default mode, association, and heteromodal cortices** *Journal of Neuroscience* **33**:14031–14039
- Buckner RL, Andrews-Hanna JR, Schacter DL (2008) **The Brain's Default Network** *Annals of the New York Academy of Sciences* <https://doi.org/10.1196/annals.1440.011>
- Christoff K, Irving ZC, Fox KCR, Spreng RN, Andrews-Hanna JR (2016) **Mind-wandering as spontaneous thought: a dynamic framework** *Nat Rev Neurosci* **17**:718–731
- Codol O, Holland PJ, Galea JM (2018) **The relationship between reinforcement and explicit control during visuomotor adaptation** *Sci Rep* **8**
- Corbetta M, Patel G, Shulman GL (2008) **The reorienting system of the human brain: from environment to theory of mind** *Neuron* **58**:306–324
- Corbetta M, Shulman GL (2002) **Control of goal-directed and stimulus-driven attention in the brain** *Nat Rev Neurosci* **3**:201–215
- Cox RW, Hyde JS (1997) **Software tools for analysis and visualization of fMRI data** *NMR Biomed* **10**:171–178
- Cunningham JP, Yu BM (2014) **Dimensionality reduction for large-scale neural recordings** *Nat Neurosci* **17**:1500–1509
- Dale AM, Fischl B, Sereno MI (1999) **Cortical Surface-Based Analysis** *NeuroImage* <https://doi.org/10.1006/nimg.1998.0395>
- Dam G, Kording K, Wei K (2013) **Credit assignment during movement reinforcement learning** *PLoS One* **8**
- Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ (2006) **Cortical substrates for exploratory decisions in humans** *Nature* **441**:876–879
- de Brouwer AJ, Albaghdadi M, Flanagan JR, Gallivan JP. (2018) **Using gaze behavior to parcellate the explicit and implicit contributions to visuomotor learning** *J Neurophysiol* **120**:1602–1615
- de Brouwer AJ, Areshenkoff CN, Rashid MR, Flanagan JR, Poppenk J, Gallivan JP. (2022) **Human Variation in Error-Based and Reinforcement Motor Learning Is Associated With Entorhinal Volume** *Cereb Cortex* **32**:3423–3440
- Dhawale AK, Miyamoto YR, Smith MA, Ölveczky BP (2019) **Adaptive Regulation of Motor Variability** *Curr Biol* **29**:3551–3562
- Dhawale AK, Smith MA, Ölveczky BP (2017) **The Role of Variability in Motor Learning** *Annu Rev Neurosci* **40**:479–498
- Esteban O, Blair R, Markiewicz CJ, Berleant SL. **fMRIPrep. Software**
- Esteban O *et al.* (2019) **fMRIPrep: a robust preprocessing pipeline for functional MRI** *Nat Methods* **16**:111–116

- Foster BL, Koslov SR, Aponik-Gremillion L, Monko ME, Hayden BY, Heilbronner SR (2023) **A tripartite view of the posterior cingulate cortex** *Nature Reviews Neuroscience* <https://doi.org/10.1038/s41583-022-00661-x>
- Fox MD, Raichle ME (2007) **Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging** *Nat Rev Neurosci* 8:700–711
- Frank MJ, Doll BB, Oas-Terpstra J, Moreno F (2009) **Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation** *Nat Neurosci* 12:1062–1068
- Frazier JA *et al.* (2005) **Structural brain magnetic resonance imaging of limbic and thalamic volumes in pediatric bipolar disorder** *Am J Psychiatry* 162:1256–1265
- Gale DJ, Areshenkoff CN, Standage DI, Nashed JY, Markello RD, Flanagan JR, Smallwood J, Gallivan JP (2022) **Distinct patterns of cortical manifold expansion and contraction underlie human sensorimotor adaptation** *Proceedings of the National Academy of Sciences* 119
- Gallego JA, Perich MG, Miller LE, Solla SA (2017) **Neural Manifolds for the Control of Movement** *Neuron* 94:978–984
- Glasser MF *et al.* (2013) **The minimal preprocessing pipelines for the Human Connectome Project** *NeuroImage* <https://doi.org/10.1016/j.neuroimage.2013.04.127>
- Gordon EM *et al.* (2017) **Precision Functional Mapping of Individual Human Brains** *Neuron* 95:791–807
- Gorgolewski K, Burns CD, Madison C, Clark D, Halchenko YO, Waskom ML, Ghosh SS (2011) **Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python** *Front Neuroinform* 5
- Gorgolewski KJ, Esteban O, Markiewicz CJ, Ziegler E, Ellis DG, Notter MP, Jarecka D, Johnson H, Burns C, Manhães-Savio A (2018) **Nipype [Software]** *Zenodo*
- Gratton C *et al.* (2018) **Functional Brain Networks Are Dominated by Stable Group and Individual Factors, Not Cognitive or Daily Variation** *Neuron* 98:439–452
- Greve DN, Fischl B (2009) **Accurate and robust brain image alignment using boundary-based registration** *Neuroimage* 48:63–72
- Hagmann P, Cammoun L, Gigandet X, Meuli R, Honey CJ, Wedeen VJ, Sporns O (2008) **Mapping the structural core of human cerebral cortex** *PLoS Biol* 6
- Härdle W (1990) **Härdle W. 1990. Applied Nonparametric Regression.** [doi:10.1017/ccol0521382483](https://doi.org/10.1017/ccol0521382483) *Applied Nonparametric Regression* <https://doi.org/10.1017/ccol0521382483>
- Hasson U, Chen J, Honey CJ (2015) **Hierarchical process memory: memory as an integral component of information processing** *Trends Cogn Sci* 19:304–313
- Hayden BY, Nair AC, McCoy AN, Platt ML (2008) **Posterior Cingulate Cortex Mediates Outcome-Contingent Allocation of Behavior** *Neuron* <https://doi.org/10.1016/j.neuron.2008.09.012>

- Hayden BY, Smith DV, Platt ML (2009) **Electrophysiological correlates of default-mode processing in macaque posterior cingulate cortex** *Proceedings of the National Academy of Sciences* <https://doi.org/10.1073/pnas.0812035106>
- Heilbronner SR, Platt ML (2013) **Causal Evidence of Performance Monitoring by Neurons in Posterior Cingulate Cortex during Learning** *Neuron* <https://doi.org/10.1016/j.neuron.2013.09.028>
- Hogeveen J, Mullins TS, Romero JD, Eversole E, Rogge-Obando K, Mayer AR, Costa VD. (2022) **The neurocomputational bases of explore-exploit decision-making** *Neuron* 110:1869–1879
- Holland P, Codol O, Galea JM (2018) **Contribution of explicit processes to reinforcement-based motor learning** *J Neurophysiol* 119:2241–2255
- Holroyd CB, Yeung N (2012) **Motivation of extended behaviors by anterior cingulate cortex** *Trends Cogn Sci* 16:122–128
- Hong S-J, Xu T, Nikolaidis A, Smallwood J, Margulies DS, Bernhardt B, Vogelstein J, Milham MP (2020) **Toward a connectivity gradient-based framework for reproducible biomarker discovery** *NeuroImage* <https://doi.org/10.1016/j.neuroimage.2020.117322>
- Houk JC, Buckingham JT, Barto AG (1996) **Models of the cerebellum and motor learning** *Behav Brain Sci* 19:368–383
- Huntenburg JM, Bazin P-L, Margulies DS (2018) **Large-Scale Gradients in Human Cortical Organization** *Trends Cogn Sci* 22:21–31
- Hutchison RM, Culham JC, Flanagan JR, Everling S, Gallivan JP (2015) **Functional subdivisions of medial parieto-occipital cortex in humans and nonhuman primates using resting-state fMRI** *Neuroimage* 116:10–29
- Jenkinson M, Bannister P, Brady M, Smith S (2002) **Improved optimization for the robust and accurate linear registration and motion correction of brain images** *Neuroimage* 17:825–841
- Kable JW, Glimcher PW (2007) **The neural correlates of subjective value during intertemporal choice** *Nat Neurosci* 10:1625–1633
- Klein A *et al.* (2017) **Mindboggling morphometry of human brains** *PLoS Comput Biol* 13
- Klein-Flügge MC, Bongioanni A, Rushworth MFS (2022) **Medial and orbital frontal cortex in decision-making and flexible behavior** *Neuron* 110:2743–2770
- Knutson B, Bossaerts P (2007) **Neural antecedents of financial decisions** *J Neurosci* 27:8174–8177
- Kolling N, Wittmann MK, Behrens TEJ, Boorman ED, Mars RB, Rushworth MFS (2016) **Value, search, persistence and model updating in anterior cingulate cortex** *Nat Neurosci* 19:1280–1285
- Lanczos C (1964) **Evaluation of Noisy Data** *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis* 1:76–85
- Ledoit O, Wolf M (2004) **A well-conditioned estimator for large-dimensional covariance matrices** *J Multivar Anal* 88:365–411

- Leech R, Braga R, Sharp DJ (2012) **Echoes of the brain within the posterior cingulate cortex** *Neurosci* **32**:215–222
- Lee D, Seo H, Jung MW (2012) **Neural basis of reinforcement learning and decision making** *Annu Rev Neurosci* **35**:287–308
- Lerner Y, Honey CJ, Silbert LJ, Hasson U (2011) **Topographic mapping of a hierarchy of temporal receptive windows using a narrated story** *Journal of Neuroscience* **31**:2906–2915
- Luo L, Rodriguez E, Jerbi K, Lachaux J-P, Martinerie J, Corbetta M, Shulman GL, Piomelli D, Turrigiano GG, Nelson SB (2010) **Ten years of Nature Reviews Neuroscience: insights from the highly cited** *Nat Rev Neurosci* **11**:718–726
- Makris N, Goldstein JM, Kennedy D, Hodge SM, Caviness VS, Faraone SV, Tsuang MT, Seidman LJ (2006) **Decreased volume of left and total anterior insular lobule in schizophrenia** *Schizophrenia Research* <https://doi.org/10.1016/j.schres.2005.11.020>
- Margulies DS *et al.* (2016) **Situating the default-mode network along a principal gradient of macroscale cortical organization** *Proceedings of the National Academy of Sciences* **113**:12574–12579
- Margulies DS, Vincent JL, Kelly C, Lohmann G, Uddin LQ, Biswal BB, Villringer A, Castellanos FX, Milham MP, Petrides M (2009) **Precuneus shares intrinsic functional architecture in humans and monkeys** *Proc Natl Acad Sci U S A* **106**:20069–20074
- Markello RD *et al.* (2022) **. neuromaps: structural and functional interpretation of brain maps** *Nat Methods* **19**:1472–1479
- McCoy AN, Platt ML (2005) **Risk-sensitive neurons in macaque posterior cingulate cortex** *Nature Neuroscience* <https://doi.org/10.1038/nn1523>
- McDougle SD, Boggess MJ, Crossley MJ, Parvin D, Ivry RB, Taylor JA (2016) **Credit assignment in movement-dependent reinforcement learning** *Proc Natl Acad Sci U S A* **113**:6797–6802
- McDougle SD, Butcher PA, Parvin DE, Mushtaq F, Niv Y, Ivry RB, Taylor JA (2019) **Neural Signatures of Prediction Errors in a Decision-Making Task Are Modulated by Action Execution Failures** *Curr Biol* **29**:1606–1613
- McInnes L, Healy J, Melville J. (2018) **UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction** *arXiv [statML]*
- Murphy C, Jefferies E, Rueschemeyer S-A, Sormaz M, Wang H-T, Margulies DS, Smallwood J (2018) **Distant from input: Evidence of regions within the default mode network supporting perceptually-decoupled and conceptually-guided cognition** *Neuroimage* **171**:393–401
- Murphy C, Wang H-T, Konu D, Lowndes R, Margulies DS, Jefferies E, Smallwood J. (2019) **Modes of operation: A topographic neural gradient supporting stimulus dependent and independent cognition** *Neuroimage* **186**:487–496
- Muschelli J, Nebel MB, Caffo BS, Barber AD, Pekar JJ, Mostofsky SH (2014) **Reduction of motion-related artifacts in resting state fMRI using aCompCor** *Neuroimage* **96**:22–35

O'Doherty JP, Cockburn J, Pauli WM (2017) **Learning, Reward, and Decision Making** *Annu Rev Psychol* **68**:73–100

O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2003) **Temporal difference models and reward-related learning in the human brain** *Neuron* **38**:329–337

Oldfield RC (1971) **The assessment and analysis of handedness: the Edinburgh inventory** *Neuropsychologia* **9**:97–113

Paquola C *et al.* (2019) **Microstructural and functional gradients are increasingly dissociated in transmodal cortices** *PLOS Biology* <https://doi.org/10.1371/journal.pbio.3000284>

Park B-Y, Bethlehem RA, Paquola C, Larivière S, Rodríguez-Cruces R, Vos de Wael R, Neuroscience in Psychiatry Network (NSPN) Consortium, Bullmore ET, Bernhardt BC. (2021) **An expanding manifold in transmodal regions characterizes adolescent reconfiguration of structural connectome organization** *Elife* **10** <https://doi.org/10.7554/eLife.64694>

Park B-Y, Park H, Morys F, Kim M, Byeon K, Lee H, Kim S-H, Valk SL, Dagher A, Bernhardt BC (2021) **Inter-individual body mass variations relate to fractionated functional brain hierarchies** *Communications Biology* <https://doi.org/10.1038/s42003-021-02268-x>

Pearson JM, Hayden BY, Raghavachari S, Platt ML (2009) **Neurons in posterior cingulate cortex signal exploratory decisions in a dynamic multioption choice task** *Curr Biol* **19**:1532–1537

Pearson JM, Heilbronner SR, Barack DL, Hayden BY, Platt ML (2011) **Posterior cingulate cortex: adapting behavior to a changing world** *Trends Cogn Sci* **15**:143–151

Pekny SE, Izawa J, Shadmehr R (2015) **Reward-dependent modulation of movement variability** *J Neurosci* **35**:4015–4024

Raichle ME (2015) **The brain's default mode network** *Annu Rev Neurosci* **38**:433–447

Ramsay J, Silverman BW. (2013) **Functional Data Analysis**. Springer Science & Business Media

Ramsay J., Wickham H., Ramsay M. J., deSolve S. (2022) **Ramsay, J., Wickham, H., Ramsay, M. J., and deSolve, S. 2022. Package "fda."** *Package "fda."*

Reuter M, Rosas HD, Fischl B (2010) **Highly accurate inverse consistent registration: a robust approach** *Neuroimage* **53**:1181–1196

Rosenberg MD, Finn ES, Scheinost D, Papademetris X, Shen X, Constable RT, Chun MM (2016) **A neuromarker of sustained attention from whole-brain functional connectivity** *Nat Neurosci* **19**:165–171

Rubinov M, Sporns O (2010) **Complex network measures of brain connectivity: uses and interpretations** *Neuroimage* **52**:1059–1069

Rushworth MFS, Noonan MP, Boorman ED, Walton ME, Behrens TE (2011) **Frontal cortex and reward-guided learning and decision-making** *Neuron* **70**:1054–1069

- Samejima K, Ueda Y, Doya K, Kimura M (2005) **Representation of action-specific reward values in the striatum** *Science* **310**:1337–1340
- Satterthwaite TD *et al.* (2013) **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
- Schacter DL, Addis DR, Hassabis D, Martin VC, Spreng RN, Szpunar KK (2012) **The future of memory: remembering, imagining, and the brain** *Neuron* **76**:677–694
- Schaefer A, Kong R, Gordon EM, Laumann TO, Zuo X-N, Holmes AJ, Eickhoff SB, Yeo BTT (2018) **Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI** *Cereb Cortex* **28**:3095–3114
- Schuck NW, Gaschler R, Wenke D, Heinzle J, Frensch PA, Haynes J-D, Reverberi C (2015) **Medial Prefrontal Cortex Predicts Internally Driven Strategy Shifts** *Neuron* <https://doi.org/10.1016/j.neuron.2015.03.015>
- Schultz W, Dayan P, Montague PR (1997) **A neural substrate of prediction and reward** *Science* **275**:1593–1599
- Shamloo F, Helie S (2016) **Changes in default mode network as automaticity develops in a categorization task** *Behav Brain Res* **313**:324–333
- Shang HL (2014) **A survey of functional principal component analysis** *Adv Stat Anal* **98**:121–142
- Shenhav A, Straccia MA, Cohen JD, Botvinick MM (2014) **Anterior cingulate engagement in a foraging context reflects choice difficulty, not foraging value** *Nat Neurosci* **17**:1249–1254
- Shenoy KV, Sahani M, Churchland MM (2013) **Cortical control of arm movements: a dynamical systems perspective** *Annu Rev Neurosci* **36**:337–359
- Shine JM, Breakspear M, Bell PT, Ehgoetz Martens KA, Shine R, Koyejo O, Sporns O, Poldrack RA (2019) **Human cognition involves the dynamic integration of neural activity and neuromodulatory systems** *Nat Neurosci* **22**:289–296
- Shine JM, Hearne LJ, Breakspear M, Hwang K, Müller EJ, Sporns O, Poldrack RA, Mattingley JB, Cocchi L (2019) **The Low-Dimensional Neural Architecture of Cognitive Complexity Is Related to Activity in Medial Thalamic Nuclei** *Neuron* **104**:849–855
- Smallwood J, Bernhardt BC, Leech R, Bzdok D, Jefferies E, Margulies DS (2021) **The default mode network in cognition: a topographical perspective** *Nat Rev Neurosci* **22**:503–513
- Song J-H (2019) **The role of attention in motor control and learning** *Curr Opin Psychol* **29**:261–265
- Spreng RN, Mar RA, Kim ASN (2009) **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
- Standage DI, Areshenkoff CN, Gale DJ, Nashed JY, Flanagan JR, Gallivan JP (2022) **Whole-brain dynamics of human sensorimotor adaptation** *Cereb Cortex* <https://doi.org/10.1093/cercor/bhac378>

- Sutton RS, Barto AG (2018) **Reinforcement Learning, second edition: An Introduction**
- Taylor JA, Thoroughman KA (2008) **Motor adaptation scaled by the difficulty of a secondary cognitive task** *PLoS One* **3**
- Taylor JA, Thoroughman KA (2007) **Divided attention impairs human motor adaptation but not feedback control** *J Neurophysiol* **98**:317–326
- Tian Y, Margulies DS, Breakspear M, Zalesky A (2020) **Topographic organization of the human subcortex unveiled with functional connectivity gradients** *Nat Neurosci* **23**:1421–1432
- Trudel N, Scholl J, Klein-Flügge MC, Fouragnan E, Tankelevitch L, Wittmann MK, Rushworth MFS (2020) **Polarity of uncertainty representation during exploration and exploitation in ventromedial prefrontal cortex** *Nature Human Behaviour* <https://doi.org/10.1038/s41562-020-0929-3>
- Tsay JS, Kim HE, McDougale SD, Taylor J, Haith A, Avraham G, Krakauer JW, Collins A, Ivry R (2023) **Strategic Processes in Sensorimotor Learning: Reasoning Refinement, and Retrieval** <https://doi.org/10.31234/osf.io/x4652>
- Tustison NJ, Avants BB, Cook PA, Zheng Y, Egan A, Yushkevich PA, Gee JC (2010) **N4ITK: improved N3 bias correction** *IEEE Trans Med Imaging* **29**:1310–1320
- Valk S, Bernhardt BC, Hong SJ, Park B-Y, Trautwein F-M, Kanske P, Böckler-Raettig A, Singer T (2021) **Changing the social brain: plasticity along macro-scale axes of functional connectivity following social mental training (No. FZJ-2022-00997)** *Gehirn & Verhalten*
- Váša F *et al.* (2018) **Adolescent Tuning of Association Cortex in Human Structural Brain Networks** *Cereb Cortex* **28**:281–294
- Vatansever D, Menon DK, Stamatakis EA (2017) **Default mode contributions to automated information processing** *Proc Natl Acad Sci U S A* **114**:12821–12826
- Vázquez-Rodríguez B, Suárez LE, Markello RD, Shafiei G, Paquola C, Hagmann P, van den Heuvel MP, Bernhardt BC, Spreng RN, Misisic B. (2019) **Gradients of structure–function tethering across neocortex** *Proceedings of the National Academy of Sciences* **116**:21219–21227
- Vos de Wael R *et al.* (2020) **BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets** *Commun Biol* **3**
- Vos de Wael R, Royer J, Tavakoli S, Wang Y, Paquola C, Benkarim O, Eichert N, Larivière S, Xu T, Misisic B (2021) **Structural connectivity gradients of the temporal lobe serve as multiscale axes of brain organization and cortical evolution** *Cereb Cortex* **31**:5151–5164
- Vyas S, Golub MD, Sussillo D, Shenoy KV (2020) **Vyas S, Golub MD, Sussillo D, Shenoy KV. 2020. Computation Through Neural Population Dynamics.** doi:10.1146/annurev-neuro-092619-094115 *Computation Through Neural Population Dynamics* <https://doi.org/10.1146/annurev-neuro-092619-094115>
- Wan X, Cheng K, Tanaka K (2015) **Neural encoding of opposing strategy values in anterior and posterior cingulate cortex** *Nat Neurosci* **18**:752–759

- Wilson RC, Bonawitz E, Costa VD, Ebitz RB (2021) **Balancing exploration and exploitation with information and randomization** *Curr Opin Behav Sci* **38**:49–56
- Wolpert DM, Diedrichsen J, Flanagan JR (2011) **Principles of sensorimotor learning** *Nat Rev Neurosci* **12**:739–751
- Wolpert DM, Ghahramani Z, Flanagan JR (2001) **Perspectives and problems in motor learning** *Trends Cogn Sci* **5**:487–494
- Wu HG, Miyamoto YR, Gonzalez Castro LN, Ölveczky BP, Smith MA (2014) **Temporal structure of motor variability is dynamically regulated and predicts motor learning ability** *Nat Neurosci* **17**:312–321
- Yeo BTT *et al.* (2011) **The organization of the human cerebral cortex estimated by intrinsic functional connectivity** *J Neurophysiol* **106**:1125–1165
- Zhang Y, Brady M, Smith S (2001) **Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm** *IEEE Trans Med Imaging* **20**:45–57
- Zhao Q, Kwon D, Pohl KM (2018) **A Riemannian Framework for Longitudinal Analysis of Resting-State Functional Connectivity** *Med Image Comput Comput Assist Interv* **11072**:145–153

Editors

Reviewing Editor

Jonathan Tsay

Carnegie Mellon University, Pittsburgh, United States of America

Senior Editor

Michael Frank

Brown University, Providence, United States of America

Reviewer #1 (Public Review):

This important study uses a wide variety of convincing, state-of-the-art neuroimaging analyses to characterize whole-brain networks and relate them to reward-based motor learning. During early learning, the authors found increased covariance between the sensorimotor and dorsal attention networks, coupled with reduced covariance between the sensorimotor and default mode networks. During late learning, they observed the opposite pattern. It remains to be seen whether these changes reflect generic changes in task engagement during learning or are specific to reward-based motor learning. This study is highly relevant for researchers interested in reward-based motor learning and decision-making.

<https://doi.org/10.7554/eLife.91928.2.sa2>

Reviewer #2 (Public Review):

This useful investigation of learning-driven dynamics of cortical and some subcortical structures combines a novel in-scanner learning paradigm with interesting analysis approaches. The new task for reward-based motor learning is compelling and goes beyond the current state of the art. The results are of interest to neuroscientists working on motor control and reward-based learning.

Comments on revised version:

The revision has produced a stronger manuscript. Thank you for your thorough responses to the comments and concerns.

<https://doi.org/10.7554/eLife.91928.2.sa1>

Reviewer #3 (Public Review):

Summary:

The manuscript of Nick and colleagues addresses the intriguing question of how brain connectivity evolves during reward-based motor learning. The concept of quantifying connectivity through changes in extraction and contraction across lower-dimensional manifolds is both novel and interesting and the presented results are clear and well-presented. Overall, the manuscript is a valuable addition to the field.

Strengths:

This manuscript is written in a clear and comprehensible way. It introduces a rather novel technique of assessing connectivity across lower-dimensional manifold which has hitherto not been applied in this way to the question of reward-based motor learning. Thus, this presents a unique viewpoint on understanding how the brain changes with motor learning. I particularly enjoyed the combination of connectivity-based, followed by further scrutiny of seed-based connectivity analyses, thus providing a more comprehensive viewpoint. Now it also has added a more comprehensive report on the behavioural changes of learning, and the added statistical quantification, which is useful.

Weaknesses:

The main weakness of the manuscript is the lack of direct connection between the reported neural changes and behavioural learning. Namely, most of the results could also be explained by changes in attention allocation during the session, or changes in movement speed (independent of learning). The authors acknowledge some of these potential confounds and argue that factors like attention are important for learning. While this is true, it is nonetheless very limiting if one cannot ascertain whether the observed effects are due to attention (independent of learning) or attention allocated in the pursuit of learning. The only direct analysis linking behavioural changes to neural changes is based on individual differences in learning performance, where the DAN-A shows the opposite trend than group level effects, which they interpret as differences given the used higher-resolution parcellation. However, it could be that these learning effects are indeed much smaller and subtler compared to more dominant group-level attention effects during the task. The lack of a control condition in the task limits the interpretability of results as learning-related.

<https://doi.org/10.7554/eLife.91928.2.sa0>

Author Response

The following is the authors' response to the original reviews.

Recommendations

Recommendation #1: Address potential confounds in the experimental design:


(1a) Confounding factors between baseline to early learning. While the visual display of the curved line remains constant, there are at least three changes between these two phases: 1) the presence of reward feedback (the focus of the paper); 2) a perturbation introduced to draw a hidden, mirror-symmetric curved line; 3) instructions provided to use reward feedback to trace the line on the screen (intentionally deceitful). As such, it remains unclear which of these factors are driving the changes in both behavior and bold signals between the two phases. The absence of a veridical feedback phase in which participants received reward feedback associated with the shown trajectory seems like a major limitation.

(1b) Confounding Factors Between Early and Late Learning. While the authors have focused on interpreting changes from early to late due to the explore-exploit trade-off, there are three additional factors possibly at play: 1) increasing fatigue, 2) withdrawal of attention, specifically related to individuals who have either successfully learned the perturbation within the first few trials or those who have simply given up, or 3) increasing awareness of the perturbation (not clear if subjective reports about perturbation awareness were measured.). I understand that fMRI research is resource-intensive; however, it is not clear how to rule out these alternatives with their existing data without additional control groups. [Another reviewer added the following: Why did the authors not acquire data during a control condition? How can we be confident that the neural dynamics observed are not due to the simple passage of time? Or if these effects are due to the task, what drives them? The reward component, the movement execution, increased automaticity?]

We have opted to address both of these points above within a single reply, as together they suggest potential confounding factors across the three phases of the task. We would agree that, if the results of our pairwise comparisons (e.g., Early > Baseline or Late > Early) were considered in isolation from one another, then these critiques of the study would be problematic. However, when considering the pattern of effects across the three task phases, we believe most of these critiques can be dismissed. Below, we first describe our results in this context, and then discuss how they address the reviewers' various critiques.

Recall that from Baseline to Early learning, we observe an expansion of several cortical areas (e.g., core regions in the DMN) along the manifold (red areas in Fig. 4A, see manifold shifts in Fig. 4C) that subsequently exhibit contraction during Early to Late learning (blue areas in Fig. 4B, see manifold shifts in Fig. 4D). We show this overlap in brain areas in Author response image 1 below, panel A. Notably, several of these brain areas appear to contract back to their original, Baseline locations along the manifold during Late learning (compare Fig. 4C and D). This is evidenced by the fact that many of these same regions (e.g., DMN regions, in Author response image 1 panel A below) fail to show a significant difference between the Baseline and Late learning epochs (see Author response image 1 panel B below, which is taken from supplementary Fig 6). That is, the regions that show significant expansion and subsequent contraction (in Author response image 1 panel A below) tend not to overlap with the regions that significantly changed over the time course of the task (in Author response image 1 panel B below).

Author response image 1.




Overlap (in dark blue) between brain areas that exhibit significant *expansion* from Baseline to Early Learning and areas that exhibit significant *contraction* from Early to Late learning.

From Supplementary Figure 6

Note that this basic observation above is not only true of our regional manifold eccentricity data, but also in the underlying functional connectivity data associated with individual brain regions. To make this second point clearer, we have modified and annotated our Fig. 5 and included it below. Note the reversal in seed-based functional connectivity from Baseline to Early learning (leftmost brain plots) compared to Early to Late learning (rightmost brain plots). That is, it is generally the case that for each seed-region (A-C) the areas that increase in seed-connectivity with the seed region (in red; leftmost plot) are also the areas that decrease in seed-connectivity with the seed region (in blue; rightmost plot), and vice versa. [Also note that these connectivity reversals are conveyed through the eccentricity data — the horizontal red line in the rightmost plots denote the mean eccentricity of these brain regions during the Baseline phase, helping to highlight the fact that the eccentricity of the Late learning phase reverses back towards this Baseline level].

Author response image 2.



Critically, these reversals in brain connectivity noted above directly counter several of the critiques noted by the reviewers. For instance, this reversal pattern of effects argues against the idea that our results during Early Learning can be simply explained due to the (i) presence of reward feedback, (ii) presence of the perturbation or (iii) instructions to use reward feedback to trace the path on the screen. Indeed, all of these factors are also present during Late learning, and yet many of the patterns of brain activity during this time period revert back to the Baseline patterns of connectivity, where these factors are absent. Similarly, this reversal pattern strongly refutes the idea that the effects are simply due to the passage of time, increasing fatigue, or general awareness of the perturbation. Indeed, if any of these factors alone could explain the data, then we would have expected a gradual increase (or decrease) in eccentricity and connectivity from Baseline to Early to Late learning, which we do not observe. We believe these are all important points when interpreting the data, but which we failed to mention in our original manuscript when discussing our findings.

We have now rectified this in the revised paper, where we now write in our Discussion:

“Finally, it is important to note that the reversal pattern of effects noted above suggests that our findings during learning cannot be simply attributed to the introduction of reward feedback and/or the perturbation during Early learning, as both of these task-related features are also present during Late learning. In addition, these results cannot be simply explained due to the passage of time or increasing subject fatigue, as this would predict a consistent directional change in eccentricity across the Baseline, Early and Late learning epochs.”

However, having said the above, we acknowledge that one potential factor that our findings cannot exclude is that they are (at least partially) attributable to changes in subjects’ state of attention throughout the task. Indeed, one can certainly argue that Baseline trials in our study don’t require a great deal of attention (after all, subjects are simply tracing a curved path presented on the screen). Likewise, for subjects that have learned the hidden shape, the Late learning trials are also likely to require limited attentional resources (indeed, many subjects at this point are simply producing the same shape trial after trial). Consequently, the large shift in brain connectivity that we observe from Baseline to Early Learning, and the subsequent reversion back to Baseline-levels of connectivity during Late learning, could actually reflect a heightened allocation of attention as subjects are attempting to learn the (hidden) rewarded shape. However, we do not believe that this would reflect a ‘confound’ of our study per se — indeed, any subject who has participated in a motor learning study would agree that the early learning phase of a task is far more cognitively demanding than Baseline trials and Late learning trials. As such, it is difficult to disentangle this ‘attention’ factor from the learning process itself (and in fact, it is likely central to it).

Of course, one could have designed a ‘control’ task in which subjects must direct their attention to something other than the learning task itself (e.g., divided attention paradigm, e.g., Taylor & Thoroughman, 2007, 2008, and/or perform a secondary task concurrently (Codol et al., 2018; Holland et al., 2018), but we know that this type of manipulation impairs the learning process itself. Thus, in such a case, it wouldn’t be obvious to the experimenter what they are actually measuring in brain activity during such a task. And, to extend this argument even further, it is true that any sort of brain-based modulation can be argued to reflect some ‘attentional’ process, rather than modulations related to the specific task-based process under consideration (in our case, motor learning). In this regard, we are sympathetic to the views of Richard Andersen and colleagues who have eloquently stated that “The study of how attention interacts with other neural processing systems is a most important endeavor. However, we think that over-generalizing attention to encompass a large variety of different neural processes weakens the concept and undercuts the ability to develop a robust understanding of other cognitive functions.” (Andersen & Cui, 2007, *Neuron*). In short, it appears that different fields/researchers have alternate views on the usefulness of attention as an explanatory construct (see also articles from Hommel et al., 2019, “No one knows what attention is”, and Wu, 2023, “We know what attention is!”), and we personally don’t have a dog in this fight. We only highlight these issues to draw attention (no pun intended) that it is not trivial to separate these different neural processes during a motor learning study.

Nevertheless, we do believe these are important points worth flagging for the reader in our paper, as they might have similar questions. To this end, we have now included in our Discussion section the following text:

“It is also possible that some of these task-related shifts in connectivity relate to shifts in task-general processes, such as changes in the allocation of attentional resources (Bédard and Song, 2013; Rosenberg et al., 2016) or overall cognitive engagement (Aben et al., 2020), which themselves play critical roles in shaping learning (Codol et al., 2018; Holland et al., 2018; Song, 2019; Taylor and Thoroughman, 2008, 2007; for a review of these topics, see Tsay et al., 2023). Such processes are particularly important during the earlier phases of learning when

sensorimotor contingencies need to be established. While these remain questions for future work, our data nevertheless suggest that this shift in connectivity may be enabled through the PMC.”

Finally, we should note that, at the end of testing, we did not assess participants' awareness of the manipulation (i.e., that they were, in fact, being rewarded based on a mirror image path). In hindsight, this would have been a good idea and provided some value to the current project. Nevertheless, it seems clear that, based on several of the learning profiles observed (e.g., subjects who exhibited very rapid learning during the Early Learning phase, more on this below), that many individuals became aware of a shape approximating the rewarded path. Note that we have included new figures (see our responses below) that give a better example of what fast versus slower learning looks like. In addition, we now note in our Methods that we did not probe participants about their subjective awareness re: the perturbation:

“Note that, at the end of testing, we did not assess participants' awareness of the manipulation (i.e., that they were, in fact, being rewarded based on a mirror image path of the visible path).”

Recommendation #2: Provide more behavioral quantification.


(2a) The authors chose to only plot the average learning score in Figure 1D, without an indication of movement variability. I think this is quite important, to give the reader an impression of how variable the movements were at baseline, during early learning, and over the course of learning. There is evidence that baseline variability influences the 'detectability' of imposed rotations (in the case of adaptation learning), which could be relevant here. Shading the plots by movement variability would also be important to see if there was some refinement of the moment after participants performed at the ceiling (which seems to be the case ~ after trial 150). This is especially worrying given that in Fig 6A there is a clear indication that there is a large difference between subjects' solutions on the task. One subject exhibits almost a one-shot learning curve (reaching a score of 75 after one or two trials), whereas others don't seem to really learn until the near end. What does this between-subject variability mean for the authors' hypothesized neural processes?

In line with these recommendations, we have now provided much better behavioral quantification of subject-level performance in both the main manuscript and supplementary material. For instance, in a new supplemental Figure 1 (shown below), we now include mean subject ($\pm$ SE) reaction times (RTs), movement times (MTs) and movement path variability (our computing of these measures are now defined in our Methods section).

As can be seen in the figure, all three of these variables tended to decrease over the course of the study, though we note there was a noticeable uptick in both RTs and MTs from the Baseline to Early learning phase, once subjects started receiving trial-by-trial reward feedback based on their movements. With respect to path variability, it is not obvious that there was a significant refinement of the paths created during late learning (panel D below), though there was certainly a general trend for path variability to decrease over learning.

Author response image 3.

Behavioral measures of learning across the task. (A-D) shows average participant reward scores (A), reaction times (B), movement times (C) and path variability (D) over the course of the task. In each plot, the black line denotes the mean across participants and the gray banding denotes $\pm$ 1 SEM. The three equal-length task epochs for subsequent neural analyses are indicated by the gray shaded boxes.




In addition to these above results, we have also created a new Figure 6 in the main manuscript, which now solely focuses on individual differences in subject learning (see below). Hopefully, this figure clarifies key features of the task and its reward structure, and also depicts (in movement trajectory space) what fast versus slow learning looks like in the task. Specifically, we believe that this figure now clearly delineates for the reader the mapping between movement trajectory and the reward score feedback presented to participants, which appeared to be a source of confusion based on the reviewers' comments below. As can be clearly observed in this figure, trajectories that approximated the 'visible path' (black line) resulted in fairly mediocre scores (see score color legend at right), whereas trajectories that approximated the 'reward path' (dashed black line, see trials 191-200 of the

fast learner) resulted in fairly high scores. This figure also more clearly delineates how fPCA loadings derived from our functional data analysis were used to derive subject-level learning scores (panel C).

Author response image 4.

Individual differences in subject learning performance. (A) Examples of a good learner (bordered in green) and poor learner (bordered in red). (B) Individual subject learning curves for the task. Solid black line denotes the mean across all subjects whereas light gray lines denote individual participants. The green and red traces denote the learning curves for the example good and poor learners denoted in A. (C) Derivation of subject learning scores. We performed functional principal component analysis (fPCA) 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. The plot at right denotes the loading for each participant (open circles) onto this Learning score component.



The reviewers note that there are large individual differences in learning performance across the task. This was clearly our hope when designing the reward structure of this task, as it would allow us to further investigate the neural correlates of these individual differences (indeed, during pilot testing, we sought out a reward structure to the task that would allow for these intersubject differences). The subjects who learn early during the task end up having higher fPCA scores than the subjects who learn more gradually (or learn the task late). From our perspective, these differences are a feature, and not a bug, and they do not negate any of our original interpretations. That is, subjects who learn earlier on average tend to contract their DAN-A network during the early learning phase whereas subjects who learn more slowly on average (or learn late) instead tend to contract their DAN-A network during late learning (Fig. 7).

(2b) In the methods, the authors stated that they scaled the score such that even a perfectly traced visible path would always result in an imperfect score of 40 patients. What happens if a subject scores perfectly on the first try (which seemed to have happened for the green highlighted subject in Fig 6A), but is then permanently confronted with a score of 40 or below? Wouldn't this result in an error-clamp-like (error-based motor adaptation) design for this subject and all other high performers, which would vastly differ from the task demands for the other subjects? How did the authors factor in the wide between-subject variability?

We think the reviewers may have misinterpreted the reward structure of the task, and we apologize for not being clearer in our descriptions. The reward score that subjects received after each trial was based on how well they traced the mirror-image of the visible path. However, all the participant can see on the screen is the visible path. We hope that our inclusion of the new Figure 6 (shown above) makes the reward structure of the task, and its relationship to movement trajectories, much clearer. We should also note that, even for the

highest performing subject (denoted in Fig. 6), it still required approximately 20 trials for them to reach asymptote performance.

(2c) The study would benefit from a more detailed description of participants' behavioral performance during the task. Specifically, it is crucial to understand how participants' motor skills evolve over time. Information on changes in movement speed, accuracy, and other relevant behavioral metrics would enhance the understanding of the relationship between behavior and brain activity during the learning process. Additionally, please clarify whether the display on the screen was presented continuously throughout the entire trial or only during active movement periods. Differences in display duration could potentially impact the observed differences in brain activity during learning.

We hope that with our inclusion of the new Supplementary Figure 1 (shown above) this addresses the reviewers' recommendation. Generally, we find that RTs, MTs and path variability all decrease over the course of the task. We think this relates to the early learning phase being more attentionally demanding and requiring more conscious effort, than the later learning phases.

Also, yes, the visible path was displayed on the screen continuously throughout the trial, and only disappeared at the 4.5 second mark of each trial (when the screen was blanked and the data was saved off for 1.5 seconds prior to commencement of the next trial; 6 seconds total per trial). Thus, there were no differences in display duration across trials and phases of the task. We have now clarified this in the Methods section, where we now write the following:

"When the cursor reached the target distance, the target changed color from red to green to indicate that the trial was completed. Importantly, other than this color change in the distance marker, the visible curved path remained constant and participants never received any feedback about the position of their cursor."

(2d) It is unclear from plots 6A, 6B, and 1D how the scale of the behavioral data matches with the scaling of the scores. Are these the 'real' scores, meaning 100 on the y-axis would be equivalent to 40 in the task? Why then do all subjects reach an asymptote at 75? Or is 75 equivalent to 40 and the axis labels are wrong?

As indicated above, we clearly did a poor job of describing the reward structure of our task in our original paper, and we now hope that our inclusion of Figure 6 makes things clear. A '40' score on the y-axis would indicate that a subject has perfectly traced the visible path whereas a perfect '100' score would indicate that a subject has perfectly traced the (hidden) mirror image path.

The fact that several of the subjects reach asymptote around 75 is likely a byproduct of two factors. Firstly, the subjects performed their movements in the absence of any visual error feedback (they could not see the position of a cursor that represented their hand position), which had the effect of increasing motor variability in their actions from trial to trial. Secondly, there appears to be an underestimation among subjects regarding the curvature of the concealed, mirror-image path (i.e., that the rewarded path actually had an equal but opposite curvature to that of the visible path). This is particularly evident in the case of the top-performing subject (illustrated in Figure 6A) who, even during late learning, failed to produce a completely arched movement.

(2e) Labeling of Contrasts: There is a consistent issue with the labeling of contrasts in the presented figures, causing confusion. While the text refers to the difference as "baseline to early learning," the label used in figures, such as Figure 4, reads "baseline > early." It is essential to clarify whether the presented contrast is indeed "baseline > early" or "early > baseline" to avoid any misinterpretation.

We thank the reviewers for catching this error. Indeed, the intended label was Early > Baseline, and this has now been corrected throughout.


Recommendation #3. Clarify which motor learning mechanism(s) are at play.

(3a) Participants were performing at a relatively low level, achieving around 50-60 points by the end of learning. This outcome may not be that surprising, given that reward-based learning might have a substantial explicit component and may also heavily depend on reasoning processes, beyond reinforcement learning or contextual recall (Holland et al., 2018; Tsay et al., 2023). Even within our own data, where explicit processes are isolated, average performance is low and many individuals fail to learn (Brudner et al., 2016; Tsay et al., 2022). Given this, many participants in the current study may have simply given up. A potential indicator of giving up could be a subset of participants moving straight ahead in a rote manner (a heuristic to gain moderate points). Consequently, alterations in brain networks may not reflect exploration and exploitation strategies but instead indicate levels of engagement and disengagement. Could the authors plot the average trajectory and the average curvature changes throughout learning? Are individuals indeed defaulting to moving straight ahead in learning, corresponding to an average of 50-60 points? If so, the interpretation of brain activity may need to be tempered.

We can do one better, and actually give you a sense of the learning trajectories for every subject over time. In the figure below, which we now include as Supplementary Figure 2 in our revision, we have plotted, for each subject, a subset of their movement trajectories across learning trials (every 10 trials). As can be seen in the diversity of these trajectories, the average trajectory and average curvature would do a fairly poor job of describing the pattern of learning-related changes across subjects. Moreover, it is not obvious from looking at these plots the extent to which poor learning subjects (i.e., subjects who never converge on the reward path) actually ‘give up’ in the task — rather, many of these subjects still show some modulation (albeit minor) of their movement trajectories in the later trials (see the purple and pink traces). As an aside, we are also not entirely convinced that straight ahead movements, which we don’t find many of in our dataset, can be taken as direct evidence that the subject has given up.

Author response image 5

Variability in learning across subjects. Plots show representative trajectory data from each subject (n=36) over the course of the 200 learning trials. Coloured traces show individual trials over time (each trace is separated by ten trials, e.g., trial 1, 10, 20, 30, etc.) to give a sense of the trajectory changes throughout the task (20 trials in total are shown for each subject).



We should also note that we are not entirely opposed to the idea of describing aspects of our findings in terms of subject engagement versus disengagement over time, as such processes are related at some level to exploration (i.e., cognitive engagement in finding the best solution) and exploitation (i.e., cognitively disengaging and automating one's behavior). As noted in our reply to Recommendation #1 above, we now give some consideration of these explanations in our Discussion section, where we now write:

"It is also possible that these task-related shifts in connectivity relates to shifts in task-general processes, such as changes in the allocation of attentional resources (Bédard and Song, 2013; Rosenberg et al., 2016) or overall cognitive engagement (Aben et al., 2020), which themselves play critical roles in shaping learning (Codol et al., 2018; Holland et al., 2018; Song, 2019; Taylor and Thoroughman, 2008, 2007; for a review of these topics, see Tsay et al., 2023). Such processes are particularly important during the earlier phases of learning when sensorimotor contingencies need to be established. While these remain questions for future work, our data nevertheless suggest that this shift in connectivity may be enabled through the PMC."

(3b) The authors are mixing two commonly used paradigms, reward-based learning, and motor adaptation, but provide no discussion of the different learning processes at play here. Which processes were they attempting to probe? Making this explicit would help the reader understand which brain regions should be implicated based on previous literature. As it stands, the task is hard to interpret. Relatedly, there is a wealth of literature on explicit vs implicit learning mechanisms in adaptation tasks now. Given that the authors are specifically looking at brain structures in the cerebral cortex that are commonly associated with explicit and strategic learning rather than implicit adaptation,

how do the authors relate their findings to this literature? Are the learning processes probed in the task more explicit, more implicit, or is there a change in strategy usage over time? Did the authors acquire data on strategies used by the participants to solve the task? How does the baseline variability come into play here?

As noted in our paper, our task was directly inspired by the reward-based motor learning tasks developed by Dam et al., 2013 (Plos One) and Wu et al., 2014 (Nature Neuroscience). What drew us to these tasks is that they allowed us to study the neural bases of reward-based learning mechanisms in the absence of subjects also being able to exploit error-based mechanisms to achieve learning. Indeed, when first describing the task in the Results section of our paper we wrote the following:


“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 (Dam et al., 2013; Wu et al., 2014).”

If the reviewers are referring to ‘motor adaptation’ in the context in which that terminology is commonly used — i.e., the use of sensory prediction errors to support error-based learning — then we would argue that motor adaptation is not a feature of the current study. It is true that in our study subjects learn to ‘adapt’ their movements across trials, but this shaping of the movement trajectories must be supported through reinforcement learning mechanisms (and, of course, supplemented by the use of cognitive strategies as discussed in the nice review by Tsay et al., 2023). We apologize for not being clearer in our paper about this key distinction and we have now included new text in the introduction to our Results to directly address this:

“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 (Dam et al., 2013; Wu et al., 2014). That is, subjects could not use error-based learning mechanisms to achieve learning in our study, as this form of learning requires sensory errors that convey both the change in direction and magnitude needed to correct the movement.”

With this issue aside, we are well aware of the established framework for thinking about sensorimotor adaptation as being composed of a combination of explicit and implicit components (indeed, this has been a central feature of several of our other recent neuroimaging studies that have explored visuomotor rotation learning, e.g., Gale et al., 2022 PNAS, Areshenkoff et al., 2022 eLife, Standage et al., 2023 Cerebral Cortex). However, there has been comparably little work done on these parallel components within the domain of reinforcement learning tasks (though see Codol et al., 2018; Holland et al., 2018, van Mastrigt et al., 2023; see also the Tsay et al., 2023 review), and as far as we can tell, nothing has been done to date in the reward-based motor learning area using fMRI. By design, we avoided using descriptors of ‘explicit’ or ‘implicit’ in our study because our experimental paradigm did not allow a separate measurement of those two components to learning during the task. Nevertheless, it seems clear to us from examining the subjects’ learning curves (see supplementary figure 2 above), that individuals who learn very quickly are using strategic processes (such as action exploration to identify the best path) to enhance their learning. As we noted in an above response, we did not query subjects after the fact about their strategy use, which admittedly was a missed opportunity on our part.

Author response image 6.



With respect to the comment on baseline variability and its relationship to performance, this is an interesting idea and one that was explored in the Wu et al., 2014 Nature Neuroscience paper. Prompted by the reviewers, we have now explored this idea in the current data set by testing for a relationship between movement path variability during baseline trials (all 70 baseline trials, see Supplementary Figure 1D above for reference) and subjects' fPCA score on our learning task. However, when we performed this analysis, we did not observe a significant positive relationship between baseline variability and subject performance. Rather, we actually found a trend towards a negative relationship (though this was non-significant; $r=-0.2916$, $p=0.0844$). Admittedly, we are not sure what conclusions can be drawn from this analysis, and in any case, we believe it to be tangential to our main results. We provide the results (at right) for the reviewers if they are interested. This may be an interesting avenue for exploration in future work.

Recommendation #4: Provide stronger justification for brain imaging methods.


(4a) Observing how brain activity varies across these different networks is remarkable, especially how sensorimotor regions separate and then contract with other, more cognitive areas. However, does the signal-to-noise ratio in each area/network influence manifold eccentricity and limit the possible changes in eccentricity during learning?

Specifically, if a region has a low signal-to-noise ratio, it might exhibit minimal changes during learning (a phenomenon perhaps relevant to null manifold changes in the striatum due to low signal-to-noise); conversely, regions with higher signal-to-noise (e.g., motor cortex in this sensorimotor task) might exhibit changes more easily detected. As such, it is unclear how to interpret manifold changes without considering an area/network's signal-to-noise ratio.

We appreciate where these concerns are coming from. First, we should note that the timeseries data used in our analysis were z-transformed (mean zero, 1 std) to allow normalization of the signal both over time and across regions (and thus mitigate the possibility that the changes observed could simply reflect mean overall signal changes across different regions). Nevertheless, differences in signal intensity across brain regions — particularly between cortex and striatum — are well-known, though it is not obvious how these differences may manifest in terms of a task-based modulation of MR signals.

To examine this issue in the current data set, we extracted, for each subject and time epoch (Baseline, Early and Late learning) the raw scanner data (in MR arbitrary units, a.u.) for the cortical and striatal regions and computed the (1) mean signal intensity, (2) standard deviation of the signal (Std) and (3) temporal signal to noise ratio (tSNR; calculated by mean/Std). Note that in the fMRI connectivity literature tSNR is often the preferred SNR measure as it normalizes the mean signal based on the signal's variability over time, thus providing a general measure of overall 'signal quality'. The results of this analysis, averaged across subjects and regions, is shown below.

Author response image 7.



Note that, as expected, the overall signal intensity (left plot) of cortex is higher than in the striatum, reflecting the closer proximity of cortex to the receiver coils in the MR head coil. In fact, the signal intensity in cortex is approximately 38% higher than that in the striatum ($\sim 625 - 450/450$). However, the signal variation in cortex is also greater than striatum (middle plot), but in this case approximately 100% greater (i.e., $(\sim 5 - 2.5)/2.5$). The result of this is that the tSNR (mean/std) for our data set and the ROI parcellations we used is actually greater in the striatum than in cortex (right plot). Thus, all else being equal, there seems to have been sufficient tSNR in the striatum for us to have detected motor-learning related effects. As such, we suspect the null effects for the striatum in our study actually stem from two sources.

The first likely source is the relatively lower number of striatal regions (12) as compared to cortical regions (998) used in our analysis, coupled with our use of PCA on these data (which, by design, identifies the largest sources of variation in connectivity). In future studies, this unbalance could be rectified by using finer parcellations of the striatum (even down to the voxel level) while keeping the same parcellation of cortex (i.e., equate the number of 'regions' in each of striatum and cortex). The second likely source is our use of a striatal atlas (the Harvard-Oxford atlas) that divides brain regions based on their neuroanatomy rather than their function. In future work, we plan on addressing this latter concern by using finer, more functionally relevant parcellations of striatum (such as in Tian et al., 2020, Nature Neuroscience). Note that we sought to capture these interrelated possible explanations in our Discussion section, where we wrote the following:


“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 the evidence indicates that this region plays a key role in reward-guided behavior (Averbeck and O’Doherty, 2022; O’Doherty et al., 2017), 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 analysis 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 striatum 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 (Tian et al., 2020; Vos de Wael et al., 2021). 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 striatal 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 striatal regions were included in the analysis (see Supplementary Fig. 6). As such, selective changes in the patterns of functional connectivity at the level of the striatum may be obscured using our cortex x striatum 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)(Tian et al., 2020) and by focusing specifically on changes in the interactions between the striatum and cortex during learning. The latter can be accomplished by selectively performing dimension reduction on the slice of the functional connectivity matrix that corresponds to functional coupling between striatum and cortex.”

(4b) Could the authors clarify how activity in the dorsal attention network (DAN) changes throughout learning, and how these changes also relate to individual differences in learning performance? Specifically, on average, the DAN seems to expand early and contract late, relative to the baseline. This is interpreted to signify that the DAN exhibits lesser connectivity followed by greater connectivity with other brain regions. However, in terms of how these changes relate to behavior, participants who go against the average trend (DAN exhibits more contraction early in learning, and expansion from early to late) seem to exhibit better learning performance. This finding is quite puzzling. Does this mean that the average trend of expansion and contraction is not facilitative, but rather detrimental, to learning? [Another reviewer added: The authors do not state any explicit hypotheses, but only establish that DMN coordinates activity among several regions. What predictions can we derive from this? What are the authors looking for in the data? The work seems more descriptive than hypothesis-driven. This is fine but should be clarified in the introduction.]

These are good questions, and we are glad the reviewers appreciated the subtlety here. The reviewers are indeed correct that the relationship of the DAN-A network to behavioral performance appears to go against the grain of the group-level results that we found for the entire DAN network (which we note is composed of both the DAN-A and DAN-B networks). That is, subjects who exhibited greater contraction from Baseline to Early learning and likewise, greater expansion from Early to Late learning, tended to perform better in the task (according to our fPCA scores). However, on this point it is worth noting that it was mainly the DAN-B network which exhibited group-level expansion from Baseline to Early Learning whereas the DAN-A network exhibited negligible expansion. This can be seen in Author response image 8 below, which shows the pattern of expansion and contraction (as in Fig. 4), but instead broken down into the 17-network parcellation. The red asterisk denotes the

expansion from Baseline to Early learning for the DAN-B network, which is much greater than that observed for the DAN-A network (which is basically around the zero difference line).

Author response image 8.



Thus, it appears that the DAN-A and DAN-B networks are modulated to a different extent during the task, which likely contributes to the perceived discrepancy between the group-level effects (reported using the 7-network parcellation) and the individual differences effects (reported using the finer 17-network parcellation). Based on the reviewers' comments, this seems like an important distinction to clarify in the manuscript, and we have now described this nuance in our Results section where we now write:

"...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. 7C), 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. 7C, 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. 7D). We should note that this pattern of performance effects for the DAN-A — i.e., greater contraction during Early learning and greater expansion during Late learning being associated with better learning — appears at odds with the group-level effects described in Fig. 4A and B, where we generally find the opposite pattern for the entire DAN network (composed of the DAN-A and DAN-B subnetworks). However, this potential discrepancy can be explained when examining the changes in eccentricity using the 17-network parcellation (see Supplementary Figure 8). At this higher resolution level we find that these group-level effects for the entire DAN network are being largely driven by eccentricity changes in the DAN-B network (areas in anterior superior parietal cortex and premotor cortex), and not by mean changes in the DAN-A network. By contrast, our present results suggest that it is the contraction and expansion of areas of the DAN-A network (and not DAN-B network) that are selectively associated with differences in subject learning performance."

Finally, re: the reviewers' comments that we do not state any explicit hypotheses etc., we acknowledge that, beyond our general hypothesis stated at the outset about the DMN being involved in reward-based motor learning, our study is quite descriptive and exploratory in nature. Such little work has been done in this research area (i.e., using manifold learning

approaches to study motor learning with fMRI) that it would be disingenuous to have any stronger hypotheses than those stated in our Introduction. Thus, to make the exploratory nature of our study clear to the reader, we have added the following text (in red) to our Introduction:

“Here we applied this manifold approach to explore how brain activity across widely distributed cortical and striatal systems is coordinated during reward-based motor learning. We were particularly interested in characterizing 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 exploring whether learning-dependent changes in manifold structure relate to variation in subject motor performance.”

We hope these changes now make it obvious the intention of our study.

(4c) The paper examines a type of motor adaptation task with a reward-based learning component. This, to me, strongly implicates the cerebellum, given that it has a long-established crucial role in adaptation and has recently been implicated in reward-based learning (see work by Wagner & Galea). Why is there no mention of the cerebellum and why it was left out of this study? Especially given that the authors state in the abstract they examine cortical and subcortical structures. It's evident from the methods that the authors did not acquire data from the cerebellum or had too small a FOV to fully cover it (34 slices at 4 mm thickness 136 mm which is likely a bit short to fully cover the cerebellum in many participants). What was the rationale behind this methodological choice? It would be good to clarify this for the reader. Related to this, the authors need to rephrase their statements on 'whole-brain' connectivity matrices or analyses - it is not whole-brain when it excludes the cerebellum.

As we noted above, we do not believe this task to be a motor adaptation task, in the sense that subjects are not able to use sensory prediction errors (and thus error-based learning mechanisms) to improve their performance. Rather, by denying subjects this sensory error feedback they are only able to use reinforcement learning processes, along with cognitive strategies (nicely covered in Tsay et al., 2023), to improve performance. Nevertheless, we recognize that the cerebellum has been increasingly implicated in facets of reward-based learning, particularly within the rodent domain (e.g., Wagner et al., 2017; Heffley et al., 2018; Kostadinov et al., 2019, etc.). In our study, we did indeed collect data from the cerebellum but did not include it in our original analyses, as we wanted (1) the current paper to build on prior work in the human and macaque reward-learning domain (which focuses solely on striatum and cortex, and which rarely discusses cerebellum, see Averbeck & O'Doherty, 2022 & Klein-Flugge et al., 2022 for recent reviews), and, (2) allow this to be a more targeted focus of future work (specifically we plan on focusing on striatal-cerebellar interactions during learning, which are hypothesized based on the neuroanatomical tract tracing work of Bostan and Strick, etc.). We hope the reviewers respect our decisions in this regard.

Nevertheless, we acknowledge that based on our statements about ‘whole-brain’ connectivity and vagueness about what we mean by ‘subcortex,’ that this may be confusing for the reader. We have now removed and/or corrected such references throughout the paper (however, note that in some cases it is difficult to avoid reference to “whole-brain” — e.g., “whole-brain correlation map” or “whole-brain false discovery rate correction”, which is standard terminology in the field).

In addition, we are now explicit in our Methods section that the cerebellum was not included in our analyses.

“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 and cerebellum. Note that for the current study, we did not examine changes in cerebellar activity during learning.”

(4d) The authors centered the matrices before further analyses to remove variance associated with the subject. Why not run a PCA on the connectivity matrices and remove the PC that is associated with subject variance? What is the advantage of first centering the connectivity matrices? Is this standard practice in the field?

Centering in some form has become reasonably common in the functional connectivity literature, as there is considerable evidence that task-related (or cognitive) changes in whole-brain connectivity are dwarfed by static, subject-level differences (e.g., Gratton, et al, 2018, Neuron). If covariance matrices were ordinary scalar values, then isolating task-related changes could be accomplished simply by subtracting a baseline scan or mean score; but because the space of covariance matrices is non-Euclidean, the actual computations involved in this subtraction are more complex (see our Methods). However, fundamentally (and conceptually) our procedure is simply ordinary mean-centering, but adapted to this non-Euclidean space. Despite the added complexity, there is considerable evidence that such computations — adapted directly to the geometry of the space of covariance matrices — outperform simpler methods, which treat covariance matrices as arrays of real numbers (e.g. naive subtraction, see Dodero et al. & Ng et al., references below). Moreover, our previous work has found that this procedure works quite well to isolate changes associated with different task conditions (Areshenkoff et al., 2021, Neuroimage; Areshenkoff et al., 2022, eLife).

Although PCA can be adapted to work well with covariance matrix valued data, it would at best be a less direct solution than simply subtracting subjects' mean connectivity. This is because the top components from applying PCA would be dominated by both subject-specific effects (not of interest here), and by the large-scale connectivity structure typically observed in component based analyses of whole-brain connectivity (i.e. the principal gradient), whereas changes associated with task-condition (the thing of interest here) would be buried among the less reliable components. By contrast, our procedure directly isolates these task changes.

References cited above:

Dodero, L., Minh, H. Q., San Biagio, M., Murino, V., & Sona, D. (2015, April). Kernel-based classification for brain connectivity graphs on the Riemannian manifold of positive definite matrices. In 2015 IEEE 12th international symposium on biomedical imaging (ISBI) (pp. 42-45). IEEE.

Ng, B., Dressler, M., Varoquaux, G., Poline, J. B., Greicius, M., & Thirion, B. (2014). Transport on Riemannian manifold for functional connectivity-based classification. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014: 17th International Conference, Boston, MA, USA, September 14-18, 2014, Proceedings, Part II 17 (pp. 405-412). Springer International Publishing.

(4e) Seems like a missed opportunity that the authors just use a single, PCA-derived measure to quantify learning, where multiple measures could have been of interest, especially given that the introduction established some interesting learning-related concepts related to exploration and exploitation, which could be conceptualized as movement variability and movement accuracy. It is unclear why the authors designed a task that was this novel and interesting, drawing on several psychological concepts, but then chose to ignore these concepts in the analysis.

We were disappointed to hear that the reviewers did not appreciate our functional PCA-derived measure to quantify subject learning. This is a novel data-driven analysis approach that we have previously used with success in recent work (e.g., Areshenkoff et al., 2022, eLife) and, from our perspective, we thought it was quite elegant that we were able to describe the entire trajectory of learning across all participants along a single axis that explained the majority (~75%) of the variance in the patterns of behavioral learning data. Moreover, the creation of a single behavioral measure per participant (what we call a ‘Learning score’, see Fig. 6C) helped simplify our brain-behavior correlation analyses considerably, as it provided a single measure that accounts for the natural auto-correlation in subjects’ learning curves (i.e., that subjects who learn quickly also tend to be better overall learners by the end of the learning phase). It also avoids the difficulty (and sometimes arbitrariness) of having to select specific trial bins for behavioral analysis (e.g., choosing the first 5, 10, 20 or 25 trials as a measure of ‘early learning’, and so on). Of course, one of the major alternatives to our approach would have involved fitting an exponential to each subject’s learning curves and taking measures like learning rate etc., but in our experience we have found that these types of models don’t always fit well, or derive robust/reliable parameters at the individual subject level. To strengthen the motivation for our approach, we have now included the following text in our Results:

“To quantify this variation in subject performance in a manner that accounted the auto-correlation in learning performance over time (i.e., subjects who learned more quickly tend to exhibit better performance by the end of learning), we opted for a pure data-driven approach and performed functional principal component analysis (fPCA; (Shang, 2014)) on subjects’ learning curves. This approach allowed us to isolate the dominant patterns of variability in subject’s learning curves over time (see Methods for further details; see also Areshenkoff et al., 2022).”

In any case, the reviewers may be pleased to hear that in current work in the lab we are using more model-based approaches to attempt to derive sets of parameters (per participant) that relate to some of the variables of interest described by the reviewers, but that we relate to much more dynamical (shorter-term) changes in brain activity.

(4f) Overall Changes in Activity: The manuscript should delve into the potential influence of overall changes in brain activity on the results. The choice of using Euclidean distance as a metric for quantifying changes in connectivity is sensitive to scaling in overall activity. Therefore, it is crucial to discuss whether activity in task-relevant areas increases from baseline to early learning and decreases from early to late learning, or if other patterns emerge. A comprehensive analysis of overall activity changes will provide a more complete understanding of the findings.

These are good questions and we are happy to explore this in the data. However, as mentioned in our response to query 4a above, it is important to note that the timeseries data for each brain region was z-scored prior to analysis, with the aim of removing any mean changes in activity levels (note that this is a standard preprocessing step when performing functional connectivity analysis, given that mean signal changes are not the focus of interest in functional connectivity analyses).

To further emphasize these points, we have taken our z-scored timeseries data and calculated the mean signal for each region within each task epoch (Baseline, Early and Late learning, see panel A in figure below). The point of showing this data (where each z-score map looks near identical across the top, middle and bottom plots) is to demonstrate just how miniscule the mean signal changes are in the z-scored timeseries data. This point can also be observed when plotting the mean z-score signal across regions for each epoch (see panel B in figure below). Here we find that Baseline and Early learning have a near identical mean activation


level across regions (albeit with slightly different variability across subjects), whereas there is a slight increase during late learning — though it should be noted that our y-axis, which measures in the thousandths, really magnifies this effect.

To more directly address the reviewers' comments, using the z-score signal per region we have also performed the same statistical pairwise comparisons (Early > Baseline and Late>Early) as we performed in the main manuscript Fig. 4 (see panel C in Author response image 9 below). In this plot, areas in red denote an increase in activity from Baseline to Early learning (top plot) and from Early to Late learning (bottom plot), whereas areas in blue denote a decrease for those same comparisons. The important thing to emphasize here is that the spatial maps resulting from this analysis are generally quite different from the maps of eccentricity that we report in Fig. 4 in our paper. For instance, in the figure below, we see significant changes in the activity of visual cortex between epochs but this is not found in our eccentricity results (compare with Fig. 4). Likewise, in our eccentricity results (Fig. 4), we find significant changes in the manifold positioning of areas in medial prefrontal cortex (MPFC), but this is not observed in the activation levels of these regions (panel C below). Again, we are hesitant to make too much of these results, as the activation differences denoted as significant in the figure below are likely to be an effect on the order of thousandths of a z-score (e.g., $0.002 > 0.001$), but this hopefully assuages reviewers' concerns that our manifold results are solely attributable to changes in overall activity levels.

We are hesitant to include the results below in our paper as we feel that they don't add much to the interpretation (as the purpose of z-scoring was to remove large activation differences). However, if the reviewers strongly believe otherwise, we would consider including them in the supplement.

Author response image 9.

Examination of overall changes in activity across regions. (A) Mean z-score maps across subjects for the Baseline (top), Early Learning (middle) and Late learning (bottom) epochs. (B) Mean z-score across brain regions for each epoch. Error bars represent ± 1 SEM. (C) Pairwise contrasts of the z-score signal between task epochs. Positive (red) and negative (blue) values show significant increases and decreases in z-score signal, respectively, following FDR correction for region-wise paired t-tests (at $q < 0.05$).



<https://doi.org/10.7554/eLife.91928.2.sa4>