Neural Excursions from Low-Dimensional Manifold Structure Explain Intersubject Variation in Human Motor Learning

Corson N. Areshenkoff,^{1,2,*} Dan J. Gale,¹ Joseph Y. Nashed,¹
Dominic Standage,³ J Randall Flanagan,^{1,2} and Jason P. Gallivan^{1,2,4}

¹*Centre for Neuroscience Studies, Queens University, Kingston ON, Canada*

²*Department of Psychology, Queens University, Kingston ON, Canada*

³*School of Psychology, Centre for Computational Neuroscience
and Cognitive Robotics, University of Birmingham*

⁴*Department of Biomedical and Molecular Sciences,
Queens University, Kingston ON, Canada*

Abstract Humans vary greatly in their motor learning abilities, yet little is known about the neural mechanisms that underlie this variability. Recent neuroimaging and electrophysiological studies demonstrate that large-scale neural dynamics inhabit a low-dimensional subspace or manifold, and that learning is constrained by this intrinsic manifold architecture. Here we asked, using functional MRI, whether subject-level differences in neural excursion from manifold structure can explain differences in learning across participants. We had subjects perform a sensorimotor adaptation task in the MRI scanner on two consecutive days, allowing us to assess their learning performance across days, as well as continuously measure brain activity. We find that the overall neural excursion from manifold activity in both cognitive and sensorimotor brain networks is associated with differences in subjects' patterns of learning and relearning across days. These findings suggest that off-manifold activity provides an index of the relative engagement of different neural systems during learning, and that intersubject differences in patterns of learning and relearning across days are related to reconfiguration processes in cognitive and sensorimotor networks during learning.

* Correspondence email address: areshenk@protonmail.com

Effective motor behaviour relies on the central nervous system’s ability to adapt to the environment by learning new mappings between motor commands and sensory outcomes. For example, an inexperienced youth hockey player, when walking from the change-room onto the ice surface, has to learn to adapt their stride in response to the drastic change in surface properties. This sensorimotor adaptation is known to rely partly on implicit learning processes, driven by sensory prediction errors reflecting the difference between the expected and sensed consequences of movement (Wolpert and Ghahramani 2000, Wolpert et al. 1995). In addition, there is evidence that explicit (declarative) learning processes, such as cognitive strategies, also help guide behavioural adjustments, and that these operate synergistically with implicit processes (see Krakauer and Mazzoni 2011, McDougle et al. 2016, Miyamoto et al. 2020, Wolpert et al. 2011). Recent behavioural and computational work (McDougle et al. 2015) has suggested that these two processes operate on different timescales, with implicit and explicit learning leading to slow and fast adaptation and de-adaptation, respectively.

Although the neural systems supporting these two learning processes are not perfectly understood, there is considerable evidence that implicit adaptation relies, in part, on sensory prediction errors computed in the cerebellum (Izawa et al. 2012, Taylor et al. 2014), which are used to drive adaptation in frontal and parietal sensorimotor regions in the cortex (McDougle et al. 2015). Explicit learning, by contrast, has been linked to higher-level working memory and inhibitory control processes (Anguera et al. 2010, Christou et al. 2016, Drummond et al. 2015, Keisler and Shadmehr 2010), and thus has been associated with a wide array of higher-order association areas across the neocortex (Anguera et al. 2007, Dayan and Cohen 2011, Doyon et al. 2009, Hardwick et al. 2013). There is also evidence that explicit learning is responsible for “savings”, or faster relearning upon reexposure to the task (Haith et al. 2015, Morehead et al. 2015). Moreover, emerging behavioural evidence indicates that individual differences in rates of adaptation and savings is largely driven by subject-level differences in the contribution of explicit processes to learning (de Brouwer et al. 2018, Fernandez-Ruiz et al. 2011). While the neural bases that underlie these subject-level differences in learning performance are likewise poorly understood, one hypothesis that follows from this prior work is that fast and slow learners are likely to differ in the relative degree to which they recruit brain networks associated with these distinct learning systems. These differences in the engagement of sensorimotor networks on the one hand, and higher-order cognitive regions on the other, may then reflect the differential contribution of explicit cognitive processes to their learning performance.

At the neural level, learning requires that the brain generate new patterns of activity. These

activity patterns are presumably constrained by the underlying neural circuitry and shaped by the information processing functions of single neurons or individual brain regions. This architecture results in correlations between the activity of neurons (and between brain regions), resulting in patterns of activation that primarily occupy a low-dimensional subspace or manifold (Gallego et al. 2017, Sadtler et al. 2014, Shenoy et al. 2013). In neurophysiological studies, this manifold structure has been observed both at the level of whole-brain functional networks (via functional MRI; Shine et al. 2019), and at the level of local neural populations (Gallego et al. 2017, Sadtler et al. 2014). Recent work has suggested that some types of learning are associated with a reconfiguration of this low-dimensional structure, allowing the brain to generate new patterns of off-manifold activity (Oby et al. 2019, Shine et al. 2019). Here we wondered whether individual differences in learning performance can be explained by subject-level differences in the generation of off-manifold activity during learning; and, if so, whether these neural excursions are present in both cognitive and sensorimotor brain networks.

We used fMRI to study changes in intrinsic manifold structure during sensorimotor adaptation on two consecutive days. First, we characterize network structure in two collections of regions: One comprising sensorimotor regions (which we refer to as the *sensorimotor network*), and another comprising higher-order association regions (the *cognitive network*). We studied both of these networks prior to (baseline), and during learning (early and late adaptation), and found that changes in cognitive network structure mirror subjects' task performance. Second, we examined excursions from intrinsic manifold structure during learning (i.e., off-manifold activity), and find that the degree of excursion across cognitive- and sensorimotor-related brain networks is related to differences in subjects' patterns of learning and relearning across days. Furthermore, we identified a subset of slow learners who achieved comparable performance to fast learners by the end of the first day, and relearned rapidly on the second day. These subjects, on the first day, exhibited patterns of neural excursion highly similar to fast learners, suggesting that they may have developed an explicit cognitive strategy, but simply implemented the strategy too slowly or too poorly to increase their rate of adaptation. Moreover, subjects who failed to reduce their errors early on day 1 showed a marked increase in cognitive network excursion upon readaptation on day 2. Together, these findings suggest that excursion from intrinsic manifold structure provides an index of the relative engagement of distributed brain networks during learning, and that fast versus slow learning is associated with different reconfiguration processes of cognitive and sensorimotor systems during learning.

I. RESULTS

To study adaptation learning, subjects ($N = 32$) performed a visuomotor rotation (VMR) task on two consecutive days, requiring them to launch a cursor to intercept a stationary target by applying a brief, isometric directional force pulse on an MRI-compatible force sensor. The task structure on each day was identical: After a series of baseline trials (Baseline epoch), the mapping between force and cursor motion was rotated clockwise by 45° , such that subjects had to learn to counteract the rotation in order to successfully hit the target. Following this Learning epoch, the rotation was removed (Washout epoch) to allow subjects to de-adapt before relearning on the second day (see Figure 1a for trial structure).

Here we take two approaches to study neural changes associated with learning. First, we examined changes in functional network structure during learning by estimating covariance networks from cognitive and sensorimotor regions during the baseline, early and late learning epochs. We then derive a low-dimensional embedding of these networks, which allows us to directly characterize subject differences in network structure across each of these epochs of the task, as well as across days (Figure 1c). Second, we specifically study learning-induced changes in network structure (i.e. following the onset of the visuomotor rotation) using an approach inspired by the work of Sadtler et al. (2014), Shine et al. (2019). This approach allows us to study the degree to which network activity during learning deviates from intrinsic manifold structure, assessed during a prior resting-state scan — a measure which we will hereby refer to as *excursion* (Figure 1d). We then relate the degree of excursion in cognitive and sensorimotor brain networks to subject learning performance within and across days.

A. Clustering of behavior reveals qualitatively distinct groups of learners across days

Mean error curves across subjects for each day are shown in Figure 2a (top), confirming that subjects successfully reduced their error over the course of the learning epoch. When we examined individual subject performance, we observed subjects whose performance improved gradually on both days, subjects whose performance improved rapidly (within a few trials) on both days, and other subjects whose performance improved quite gradually on day 1 but rapidly on day 2 (Figure 2a, bottom). This variation cannot be well captured by standard measures of learning, such as early error or savings. For example, savings does not differentiate between the those subjects who reduce their errors very rapidly on both days and those who reduce their errors gradually on both

FIG. 1: Adaptation task and analysis approach for the neural data - a) VMR trial structure. Subjects used a force sensitive joystick to move a cursor to a cued target location. On rotation trials, the mapping between aim and cursor direction was rotated clockwise by 45° . **b)** Subjects adapted and de-adapted to the rotation during two MRI testing sessions on subsequent days. We studied functional network structure during baseline, early, and late learning on each day by estimating covariance matrices from the mean BOLD signals extracted from cognitive and motor ROIs. **c)** We isolated learning related changes in functional network structure by centering subjects with respect to their resting state network structure. We then jointly embedded networks estimated during each learning epoch into a common space in order to characterize changes in network structure over the timecourse of learning. **d)** We studied the degree to which learning is associated with the generation of novel patterns of neural activity by estimating resting-state manifolds for each subject, and then calculating the magnitude of the off-manifold component of the BOLD signal at each imaging volume during task. This measure, which we call *excursion*, provides a moment-by-moment index of functional reconfiguration of a brain network over the course of learning.

days—since the difference across days is minimal in both cases. Likewise, early error on day 1 does not distinguish between subjects who reduce errors gradually on day 1 but then reduce errors rapidly on day 2 (i.e., subjects who exhibit savings) and those who reduce errors gradually on both days. In order to profile these differences in performance across days, we clustered individual subject performance by early error on days 1 and 2 (defined as the mean error over the first four eight-trial blocks of the rotation phase), along with savings (the difference in early error across days). This clustering revealed three distinct subgroups of learners (Figure 2b), corresponding to subjects exhibiting fast learning on both days (FF, $N = 16$), slow learning on both days (SS, $N = 9$), and slow learning on day 1 but fast learning on day 2 (SF, $N = 7$; Figure 2c). We evaluated this cluster solution using the procedure described by Kimes et al. (2017). For numbers of clusters between two and eight, we computed the ratio of within/between cluster sums-of-squares and compared this to a null model of a multivariate normal distribution with covariance equal to the observed covariance between variables. Based on these simulations, we judged that a three cluster solution best explained the data (Figure 2b, bottom).

The rapid adaptation and de-adaptation shown by the FF group on both days, and by the SF group on day 2, are considered signatures of explicit learning (Benson et al. 2011, de Brouwer et al. 2018). In the context of the VMR task, this includes the use of a re-aiming strategy, in which a subject aims so as to counteract the rotation applied to the cursor. One possible interpretation of the SF group's performance is that they adopt an explicit strategy during day 2, after learning largely implicitly on day 1. Another alternative is that this group learns explicitly, even on day 1, but employs an explicit strategy either poorly or too late to reduce their early error. This second possibility is supported by similarity between the FF and SF groups later on day 1 (Figure 2d): both groups show comparable performance during late learning and washout (Figure 2c). The FF and SF group also display relatively long reaction times compared to the SS group (Figure 2d), which have been associated with explicit strategy use (e.g. Fernandez-Ruiz et al. 2011, Haith et al. 2015).

As explicit learning is believed to be highly cognitive, we may expect that the rapid learning exhibited by the FF group would be associated with the greater recruitment of cognitive brain regions outside of traditional sensorimotor networks, commonly associated with adaptation to sensorimotor errors. Moreover, if the late performance of the SF group on day 1 is indeed reflective of explicit strategy use, we might expect to see similar patterns of neural recruitment between the FF and SF groups.

FIG. 2: Clustering reveals three distinct groups of learners across days - a) (top) Mean angle error during baseline, rotation, and washout on both days. Grey bands denote early and late rotation periods, and early washout. Colored bands denote standard errors. (bottom) Early error on each day (mean error during the early learning window), as well as savings. **b)** (top) Variables used for clustering; (bottom, from left) Dendrogram obtained by hierarchical complete linkage clustering on the ℓ_1 distance for the above variables for each subject; Distance matrix displaying the distance between individual subjects' clustering variables; The cluster solution was evaluated using the procedure described by Kimes et al. (2017). The cluster solution was compared to a null model in which the data were sampled from a single multivariate normal distribution with identical covariance to the observed data. The observed data were compared with 1000 bootstrap samples for a number of clusters ranging from two to eight. The three cluster solution was chosen to form the basis for our analysis. **c)** Mean error curves for each behavioral group. **d)** Mean error and reaction time during early learning, late learning, and early washout on day 1.

B. FF and SF groups show similar changes in cognitive network structure during learning

We defined sensorimotor and cognitive networks using the whole-brain parcellation proposed by Seitzman et al. (2020), comprising 300 regions of interest (ROIs) across the cortex, basal ganglia, and cerebellum, and assigned functional classifications according to the parcellation of Power et al. (2011). Given that nearly every part of the brain can be said to be involved in "cognition" in some capacity, we simplified our division of motor versus cognitive networks such that it distinguished between higher-order association areas implicated in attention and cognitive control (frontoparietal, dorsal attention, and ventral attention networks, FP, DA and VA, respectively; Figure 3a; 62 ROIs), and regions primarily involved in the control and generation of movement (somatomotor dorsal and somatomotor lateral networks, SMD and SML, respectively; Figure 4a; 51 ROIs). We note that

our cognitive network comprises a large part of what has been characterized as the “multiple-demand system” (Duncan 2010) in frontal and parietal cortex, which has been associated with diverse cognitive demands and the implementation of strategies. The DA and VA networks in particular have been implicated in motor inhibition (Hsu et al. 2020) and in the re-orienting of spatial attention (Corbetta and Shulman 2002, Simpson et al. 2011), both processes of which are involved in explicit strategy use during visuomotor rotation learning (de Brouwer et al. (2018)).

For each of these networks, we estimated covariance matrices from the subject’s first resting state scan (177 TRs), and equivalent length periods on each day during the beginning of baseline, immediate following the onset of rotation (early learning), and at the end of the rotation block (late learning; see Figure 1b for an illustration).

1. *Covariance centering reveals task-related changes in network structure*

Previous work (e.g. Gratton et al. 2018) has found that subjects exhibit strong individual differences in functional connectivity, which could obscure learning- or task-related changes in network structure. To remove these subject-specific differences prior to analysis, we centered each subjects’ task-related covariance matrices with respect to their resting state covariance matrices. To illustrate the effect of this centering procedure—and why it is critical for investigating any learning-related effects here—we used uniform manifold approximation (UMAP; McInnes et al. 2018) to project the individual covariance matrices into two dimensions, both before and after centering. As can be seen in Figures 3b and 4b, prior to the centering, the network structure is completely dominated by subject-level clustering, with no hint of any task-related structure (i.e., separation of baseline, early and late learning). However, following the centering procedure, a clear task-related structure is now apparent, particularly between the baseline epoch and the two learning epochs (early and late). Notably, this task-related structure is present in both the cognitive and sensorimotor networks.

Next, we represented these centered covariance matrices as tangent vectors describing differences in functional connectivity (for each task epoch and day) with respect to the mean resting state scan (see Varoquaux et al. 2010). This approach is beneficial, as it has the effect removing static features of a subject’s network structure which do not change across task epochs (and thus are not of interest in the context of learning). We then embedded the ROIs within the cognitive and sensorimotor networks separately into two dimensions using the model of Wang et al. (2019), in order to identify components describing patterns of changes in functional connectivity relative to

rest. In this embedded space, two ROIs are located close together if they exhibit similar changes in functional connectivity with respect to rest.

Given these components $\mathbf{h}_k$, every tangent vector $\mathbf{T}_i$ is then approximated by $\mathbf{T}_i = \sum_k \mathbf{h}_k^\top \mathbf{h}_k \lambda_{ik}$. The matrices $\mathbf{h}_k^\top \mathbf{h}_k$ can be interpreted as patterns of changes in functional connectivity relative to rest, while the score λ_{ik} is the degree to which that pattern is expressed in the i 'th observation. By examining changes in the component scores across group, day, and task epoch, we are able to characterize learning-related changes in cognitive and sensorimotor network structure.

2. *Changes in cognitive and sensorimotor network structure during learning*

Figure 3 shows the results of the embedding of the cognitive networks. Figure 3c shows the grand mean (across all subjects) resting state covariance matrix to which subject task data were centered, revealing a clear clustering of the FP, DA, and VA networks, consistent with their designation as distinct functional networks. Figure 3d shows the joint embedding of all the cognitive network tangent vectors (across subjects, epochs and days) into two-dimensions, resulting in components describing patterns of changes in connectivity relative to rest. The first component appears to be a gradient separating the FP and DA networks, and the second appears to be a general cortical gradient spanning large portions of FP and DA networks. Notably, the VA network loads relatively weakly on both of these components, suggesting that connectivity both within and between this network may not be strongly modulated by task.

For each component, we computed the mean component score for each learning group and task epoch. These scores are shown in Figure 3e, where the arrows denotes the passage of time (from baseline, to early, to late learning), revealing, quite strikingly, that the FF and SF groups show almost identical trajectories on both days, distinct from the SS group. These components are examined separately in Figure 3f, where the leftmost panel shows each component in the brain, with each ROI color-coded according to its loading (as in Figure 3d).

In order to make the interpretation of the components more clear, it is useful to examine the matrix $\mathbf{h}^t \mathbf{o} \mathbf{p} \mathbf{h}$ for each component $\mathbf{h}$, which can be interpreted as a pattern of change in covariance during each task epoch, relative to rest (Figure 3c). Figure 3g shows the mean expression of each of these components during each day and task epoch. The first component can be seen to encode a decrease in functional connectivity within the DA network, along with an increase in functional connectivity between the DA and the FP networks. This effect appears to be strongest during the early learning epoch on both days, and to fade during the late learning epoch. In order

to examine learning-related changes in the expression of this component, we fit a mixed-effects ANOVA, with epoch and day as within-subjects factors, to the component scores. Due to the large number of possible comparisons and concerns for power given the small group sizes, we examined only the main effects of group, day, and epoch, and did not perform post-hoc testing. F-tests were conducted for each main effect, corrected for multiple comparisons (Benjamini and Hochberg 1995), and revealed a significant effect of epoch ($F_{2,29} = 34.65$, $p < .001$, $q < .001$). The second component encodes an increase in connectivity within and between the DA and FP networks, and also shows a significant main effect of epoch ($F_{2,29} = 9.45$, $p < .001$, $q < .001$), as well as group ($F_{2,29} = 4.38$, $p = .022$, $q = .032$). The expression of this component increases from baseline, to early, to late learning, and is expressed more strongly in the FF and SF groups than in the SS group.

The embedding of the sensorimotor networks is shown in Figure 4, in identical format to Figure 3. The mean resting state covariance matrix is shown in Figure 4c, showing strong connectivity within the somatomotor dorsal network. The embedding of the ROIs in Figure 4d reveals a component with positive loadings across almost all of the cortical sensorimotor network ROIs, and a second component separating the dorsal and lateral cortical somatomotor ROIs. Unlike in the cognitive networks, these components do not appear to clearly distinguish the groups, although they do appear to show consistent changes across task epochs (Figure 4e).

In Figure 4f, we see that the first component encodes the overall change in connectivity with the dorsal somatomotor regions, with baseline and early learning epochs showing a global decrease relative to resting state, which either disappears or reverses during late learning. This trend was present for all groups on both days, and a mixed-effects ANOVA showed a significant main effect of epoch on the expression of this component ($F_{2,29} = 21.71$, $p < .001$, $q < .001$).

The second sensorimotor component encodes an increase in connectivity within each of the dorsal and lateral somatomotor networks, and a decrease in connectivity between them. Although this component showed a significant main effect of epoch ($F_{2,29} = 20.19$, $p < .001$, $q < .001$), we could not discern any clear pattern to the effect, other than what appears to be a decrease during early learning as compared to baseline, which may suggest a reduced functional segregation between the dorsal and lateral regions during this epoch.

Taken together, these results suggest that, despite similarities in early learning performance between the SF and SS subjects, the SF subjects appear to engage their cognitive systems in a different fashion from the SS subjects over the course of learning. In particular, consistent with the idea that the FF and SF recruit their cognitive systems in a similar fashion, we find that

FIG. 3: Joint embedding of cognitive networks - **a)** ROIs comprising the cognitive networks, colored by the functional network assignment of (Power et al. 2011). **b)** Visualization of covariance matrices by uniform manifold approximation (UMAP; McInnes et al. 2018). Each dot represents a single covariance matrix. Covariance matrices are shown both before and after centering, and are colored by subject (top) and scan (bottom). Note the strong subject level clustering in each network before centering, which masks differences in task structure. **c)** Grand mean resting state covariance (across all subjects), to which subject task data were centered. Matrix was ordered using single-linkage hierarchical clustering. **d)** (top) Components extracted by the embedding. ROIs are colored by functional network assignment (left) and anatomical location (right). (bottom) Embedding components displayed on the brain. ROIs are colored by their loading on each component. **e)** Mean component score in each group for each task epoch and day. Arrow denotes the direction of time (baseline to early learning to late learning). **f)** Components displayed as tangent vectors, reflecting changes in functional connectivity relative to rest. Each matrix represents the overall mean effect in each task epoch and day. All matrices are ordered identically to the grand mean resting state covariance, for ease of comparison. **g)** Components scores for each group during each task epoch and day. Error bars denote standard errors.

FIG. 4: Joint embedding of sensorimotor networks - **a)** ROIs comprising the sensorimotor networks, colored by the functional network assignment of (Power et al. 2011). **b)** Visualization of covariance matrices by uniform manifold approximation (UMAP; McInnes et al. 2018). Each dot represents a single covariance matrix. Covariance matrices are shown both before and after centering, and are colored by subject (top) and scan (bottom). Note the strong subject level clustering in each network before centering, which masks differences in task structure. **c)** Grand mean resting state covariance (across all subjects), to which subject task data were centered. Matrix was ordered using single-linkage hierarchical clustering. **d)** (top) Components extracted by the embedding. ROIs are colored by functional network assignment (left) and anatomical location (right). (bottom) Embedding components displayed on the brain. ROIs are colored by their loading on each component. **e)** Mean component score in each group for each task epoch and day. Arrow denotes the direction of time (baseline to early learning to late learning). **f)** Components displayed as tangent vectors, reflecting changes in functional connectivity relative to rest. Each matrix represents the overall mean effect in each task epoch and day. All matrices are ordered identically to the grand mean resting state covariance, for ease of comparison. **g)** Components scores for each group during each task epoch and day. Error bars denote standard errors.

both groups show greater functional connectivity within and between the dorsal attention and frontoparietal networks (cognitive component 2).

C. FF and SF groups show greater excursion from intrinsic manifold structure at the onset of learning

The results of the previous section suggest that behavioral groups differ in their recruitment of cognitive brain networks, with FF and SF groups showing greater functional connectivity within and between frontoparietal and dorsal attention networks during task performance. We now turn our attention to the dynamics of cognitive and sensorimotor network structure during learning, and, in particular, to the reconfiguration of these networks at the onset of rotation. Our approach is motivated by recent neural recording work (Sadler et al. 2014) investigating whether deviations from low-dimensional manifold structure are related to the capacity for learning, as well as fMRI work investigating how these deviations relate to task performance (Shine et al. 2019). We used principal component analysis to estimate a low dimensional subspace capturing the dominant patterns of BOLD signal covariance during each subject’s resting-state scan. We then examined excursions from this intrinsic structure during the task, defined as the proportion of the BOLD signal activation vector which lies off of the subject’s intrinsic manifold. This measure provides an index of the degree to which activation within the cognitive and sensorimotor networks deviate from the structure present during rest. For each subject, we computed this neural excursion at each time-point (imaging volume) in the cognitive and sensorimotor networks on each day. We then baseline-centered these excursion curves in order to isolate changes in excursion during the learning phase (see Methods for a more detailed description).

Figure 5a shows the mean excursion across subjects in the cognitive and sensorimotor networks on each day. Notably, we observed a sharp peak in excursion in the cognitive network immediately after the onset of rotation, which is most pronounced on day 1, but also prominent on day 2. We defined *early excursion* to be the mean excursion during the first four 8 trials blocks – equivalent to our behavioural definition of early error – and compared excursion across networks and days using a 2x2 factorial ANOVA with day and network as within-subjects effects. After correcting for multiple comparisons (Benjamini and Hochberg 1995), we observed significant main effects of day ($F_{1,29} = 10.197$, $p = .002$, $q = .003$) and of network ($F_{1,29} = 25.766$, $p < .001$, $q < .001$), but no day by network interaction ($F_{1,29} = 3.691$, $p = .06$, $q = .06$). Figure 5b shows the post-rotation peak broken down by behavioral group, revealing what appears to be greater peak magnitude in

the FF and SF groups (compared to SS) in the cognitive network on day 1, and a greater peak in the SF group on day 2, where these subjects succeed in reducing their early error in a similar fashion to FF subjects.

Due to the complexity of these data – which comprise excursion at each timepoint, in each network, across two days – we sought an interpretable, low-dimensional representation of the patterns of excursion across days and across networks in order to reduce the total number of possible comparisons. To this end, we used functional principal component analysis (mfPCA) to derive a low-dimensional summary of the main patterns of excursion across individual subjects during learning. Just as standard PCA finds a set of components capturing the dominant patterns of variability of data lying in ordinary Euclidean space $\mathbb{R}^n$, functional PCA attempts to find a basis of component functions capturing the dominant patterns of variability in functional data (see Figure 6a). As each subject contributes four excursion curves – one for each network and each day – we used *multivariate* fPCA (Happ and Greven 2018) to summarize the dominant patterns of excursion across days and networks. Based on variance explained by each component, we retained three components for analysis, which together describe 77% of the variance across all excursion curves. Each of these components describes a characteristic pattern of excursion, represented as a deviation from the overall mean. Note that the identification of these components was done over all subjects, and thus is agnostic to group membership.

The first of these components is shown in Figure 6b. In the left panel, the black lines denote the mean excursion across all subjects, while the colored bands denote the modes of variation encoded by the component, with red (resp. blue) denoting the effects of positive (resp. negative) scores. The right panel shows mean scores on this component for each of the three groups. In this case, the first component appears to describe overall excursion on day 1, with positive scores denoting greater excursion than average in both networks. Positive scores on the second component appear to encode greater excursion in the cognitive network on Day 2. Finally, component three appears to encode cognitive excursion on day 1, with a slight decrease in motor excursion on both days.

Notably, we observed significant group differences in the expression of the first (Kruskal-Wallis, $\chi^2_2 = 6.03$, $p = .04$) and the second components ($\chi^2_2 = 10.16$, $p = .006$). For each of these components, we conducted pairwise Conover-Iman post-hoc tests (Conover and Iman 1979) between groups, correcting for multiple comparisons (Benjamini and Hochberg 1995). For component 1, scores were significantly higher in the FF group than in the SS group ($z = 2.62$, $p = .02$), as well as higher in the SF group than in the SS group ($z = 2.25$, $p = .02$). For component 2, scores were significantly higher in the SF group than in the FF group ($z = 3.29$, $p = .004$), as well as higher

FIG. 5: **Excursion from resting manifold structure - a)** (left) Mean excursion in cognitive and sensorimotor networks on each day. In order to isolate the effect of rotation, subject excursion curves were standardized with respect to baseline. Vertical dashed line denotes the onset of rotation. (right) Early excursion, defined as the mean excursion during the first 4 eight trial blocks (shaded region in left panel), analogous to our definition of early error. **b)** Mean excursion for each behavioral group. Error bars represent standard errors.

in the SS group than in the FF group ($z = 2.35$, $p = .02$). Finally, for component 3, we found no significant effect of group.

In summary, interpreting both the results of the mfPCA and the mean excursion shown in Figure 5a, both the FF and SF groups show a large degree of excursion in the cognitive network on the first day, with the SF group likewise showing a spike in excursion on day 2. These results are consistent with the idea that the rapid learning exhibited by the FF group involves the recruitment of cognitive brain regions outside the sensorimotor cortices. The fact that the SF group expresses a similar excursion profile on day 1 lends evidence to the interpretation that the SF does indeed engage cognitive learning processes on day 1, as reflected in their late performance, but merely fails to learn a strategy early enough to bring down their early error.

II. DISCUSSION

Prior work has shown that some neural activity patterns are more difficult to generate than others, and that this constrains learning (Sadler et al. (2014)). Here, using human fMRI and by investigating changes in cognitive and sensorimotor network structure during sensorimotor adaptation, we asked whether similar principles might also account for individual differences in learning. Clustering of subject learning data revealed distinct profiles of learning, including a group which adapted rapidly on both days, achieving near complete reduction of error during late learning, as well as rapid de-adaptation (FF); and a group which adapted and de-adapted slowly on both

FIG. 6: Functional principal component analysis of network excursion data - a) Given a set of curves (top) – for example, excursion curves for each scan – functional principal component analysis (fPCA) extracts components describing the dominant patterns of variability in functional data (second row). In this example, the first component encodes whether the curve is higher or lower than average, while the second encodes a sinusoidal pattern in which the curve begins higher or lower than average, and then reverses direction. The score associated with a component then describes the degree to which that component is expressed in a particular observation (third row). These scores provide a low dimensional representation of the overall shape of the curve. For example, in the bottom row, we see three example curves, and their corresponding scores on the two components. **b)** Functional principal components. Each panel corresponds to a single component, representing a pattern of variability in excursion across each network and day. The left panel illustrates each component as a deviation from the mean excursion (solid black line). Red (resp. blue) bands denote the effect of positive (resp. negative) scores. The right panel shows the mean score in each group. Error bars denote standard errors, while stars denote significant differences for pairwise Conover-Iman tests with Benjamini-Hochberg correction for multiple comparisons within each component.

days (SS). Based on prior work, these learning profiles suggest that the FF group was employing cognitive strategies during learning. Notably, we also observed a third group of subjects (SF), who displayed large early errors and slow adaptation during the first day, but also showed comparable rates of de-adaptation to the FF group, and rapid readaptation on the second day. We interpreted this as evidence that the SF group may have been attempting the use of cognitive re-aiming strategies during initial learning, but that it took them longer to figure out the appropriate strategy. Consistent with this idea, the embedding of subject cognitive networks revealed that network structure in the FF and SF groups were highly similar during learning across both days (Figure 3e). In particular, both groups showed high expression of a component encoding an increase in connectivity within and between the frontoparietal and dorsal attention networks. Our excursion analysis further revealed that these groups' cognitive networks exhibited greater excursion from intrinsic network structure during the learning phase, after onset of the visuomotor rotation. Taken

together, these results suggest that both the FF and SF groups, who display behavioral characteristics consistent with the use of explicit, cognitive strategy use, engage their cognitive brain networks in a manner that is distinct from gradual learners, and to a greater degree.

A. Cognitive and sensorimotor network structure during learning

Our cognitive network embedding revealed a pair of components encoding, predominantly, connectivity within and between frontoparietal (FP) and dorsal attention (DA) networks. The first of these components appeared to encode a decrease in connectivity within the DA network but an increase in connectivity between the FP and DA networks. This component showed a decrease in expression from early to late learning in all three behavioral groups, suggesting that it may reflect a general feature of learning or task performance (e.g., decreasing errors). The second component, encoding an increase in connectivity both within and between the FP and DA networks, was more strongly expressed in both the FF and SF groups than in the SS group. The DA network in particular is implicated in the top-down orientation of spatial attention ([Corbetta and Shulman 2002](#), [Simpson et al. 2011](#)), and is known to be modulated by the FP network during goal-directed cognition ([Spreng et al. 2013](#)). The increase in connectivity within and between these networks in the FF and SF groups may then reflect the deliberate reorienting of attention during task performance, which would be consistent with the finding of [de Brouwer et al. \(2018\)](#) that explicit learners alter their gaze behavior to reflect their intended aim direction in response to a visuomotor rotation. The DA network has also been implicated in motor inhibition ([Hsu et al. 2020](#)), which in the VMR task is thought to be essential for the initial suppression of a motor response to the actual target prior to implementation of a re-aiming strategy ([Drummond et al. 2015](#)). The reduced connectivity in the DA network in the SS compared to the FF/SF groups might then reflect weaker engagement of this system. This would also be consistent with the shorter response times in SS, as this initial inhibition, and the successful deployment of a re-aiming strategy, are believed to require longer preparation ([Fernandez-Ruiz et al. 2011](#)).

It is notable that subcortical and cerebellar regions loaded relatively weakly on the components extracted in both the cognitive and sensorimotor networks. This, however, should not be taken to suggest that these regions are not functionally engaged during learning, as the cerebellum in particular is known to play a key role in sensorimotor adaptation. Rather, because our components reflect *changes* in functional connectivity with respect to rest, it may suggest that the ways in which these regions are situated in the cognitive and sensorimotor networks are similar across task and

rest, or that changes in functional connectivity are relatively larger within cortical networks. In our cognitive network, both components reflected changes in connectivity within and between cortical FP and DA regions, which may suggest that these regions drive the majority of functional reconfiguration during our task.

B. Neural excursions from manifold structure

Low-dimensional structure in the activity of local neural populations is generally believed to result from the underlying network organization, with structural connections imposing strong constraints on the activity patterns expressible by the network. Several authors, such as [Sadler et al. \(2014\)](#), have suggested that this structure may, in turn, impose constraints on learning, such that behaviors requiring the generation of on-manifold activity are learned more quickly than behaviors requiring novel, off-manifold activation. Although whole-brain fMRI BOLD activity is likewise known to display low-dimensional structure (e.g. [Gao et al. 2020](#)), its origins are likely to involve not only the underlying structural connectivity, but also shared function, shared sensory input, and correlations introduced by shared activation during cognitive tasks. In the context of large-scale brain networks, then, excursion from resting manifold structure is likely to reflect, at least partly, a change in the functional organization of these networks.

We observed a large spike in excursion after rotation onset, located predominantly in the cognitive networks, and which was largest on the first day. This spike was most prominent in the FF and SF groups, which show behavioral traits consistent with explicit learning. The excursion does not appear to be a response to error, as the SS group showed less cognitive excursion than both the FF and SF groups, despite greater errors on both days. Nor does it appear to reflect the pure recall of a strategy itself, as the spike was dampened in the FF group on the second day, while being maintained in the SF group. This may suggest that the excursion reflects cognitive network engagement involved in the learning or refinement of an explicit strategy, and not in the recall or implementation of the strategy itself.

C. Methodological considerations

Our characterization of brain network structure during visuomotor rotation learning comes with a few caveats. First, we consider cognitive and sensorimotor networks separately, while ignoring possible functional interactions between them. In truth, sensorimotor regions receive substantial

top-down modulation from higher order cognitive regions, and these connections are likely essential for successful learning and performance (Mars et al. 2011). Our decision to treat these networks separately was motivated largely by interpretability: Whole-brain networks spanning multiple functional domains are more difficult to analyze and summarize, and because our hypotheses concerned the role of higher-order cognitive regions specifically, their separation from sensorimotor regions allowed us to more directly probe the structure of these networks during learning. It is worth noting that our fPCA analysis (Figure 6b) of subject excursion revealed higher excursion in sensorimotor as well as cognitive networks in FF and SF learners, which suggests a functional reconfiguration of sensorimotor network structure in explicit learners. This may itself be partly a consequence of a change in cognitive networks, and a resulting change in the functional connections between cognitive and sensorimotor regions.

Second, while our joint embedding and our excursion measure probe the static and dynamic structure of cognitive and sensorimotor networks during learning, respectively, they do so in slightly different ways, and their respective conclusions are difficult to directly compare. Specifically, our excursion measure isolates the off-manifold component of the BOLD signal, and so quantifies the degree to which a brain network exhibits covariance patterns qualitatively different from rest; though it doesn't provide a characterization of the nature of the difference (though this is also true in the neural manifold work of e.g. Sadtler et al. 2014). The embedding, by contrast, is conducted on the centered tangent vectors encoding differences in covariance relative to rest, which include both on-manifold (a change in expression of components already present during rest) and off-manifold components. It is not necessarily the case, then, that the components extracted by our embedding describe the changes in network structure underlying the increase in excursion during learning. Future work will be required to establish this relationship. Moreover, excursion allows for a moment-by-moment characterization of changes in network structure, whose timecourse may not be fully captured by the windows we have used for our embedding.

Finally, although we interpret several features of subject performance; such as rapid learning, savings, and long reaction times, as reflecting explicit, cognitive learning, our study lacked a direct assay of subjects' explicit strategy. This is often assayed through the incorporation of report trials (Taylor et al. 2014), in which subjects are periodically asked to report their aim direction prior to performing a movement. The difference between the target location and reported aim direction then provides a direct measure of a subject's explicit knowledge, which might have allowed stronger claims about the relationship between explicit strategy use and the structure of cognitive brain networks. However, there is considerable evidence (e.g. de Brouwer et al. 2018, Leow et al. 2017)

that these trials can alter subject performance on the task, and magnify the use of these strategies. Moreover, their inclusion may have contaminated the neuroimaging data, which would then reflect not only the performance of the main task, but also the secondary reporting task. For these reasons, report trials were not included. Nevertheless, subject reports strongly correlate both with our measures of error, and other assays of explicit learning (de Brouwer et al. 2018, Leow et al. 2017), and measures such as savings and learning rate are linked to explicit learning processes by a large literature, both behavioral and computational (e.g. McDougle et al. 2015, Morehead et al. 2015).

III. METHODS

A. Subjects

Forty right-handed individuals between the ages of 18 and 35 ($M = 22.5$, $SD = 4.51$; 13 males) participated in the study and received financial compensation for their time. Data from 8 participants were excluded due to either head motion in the MRI scanner ($N = 4$; motion greater than 2 mm or 2° rotation within a single scan) or their inability to learn the rotation task ($N = 4$), leaving 32 participants in the final analysis. Right-handedness was assessed with the Edinburgh handedness questionnaire (Oldfield et al. 1971). Participants' written, informed consent was obtained before commencement of the experimental protocol. The Queen's University Research Ethics Board approved the study and it was conducted in accordance with 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).

B. Visuomotor rotation task

1. Apparatus

Participants performed hand movements directed towards a target by applying a directional force onto an MRI-compatible force sensor (ATI technologies) using their right index finger and thumb. The target and cursor 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 fixed to the MRI coil directly above participants' eyes, thus preventing participants from being able to see the hand. The force

sensor and associated cursor movement were sampled at 500 Hz.

2. Procedure

This experiment used a well-established visuomotor rotation task (Krakauer 2009) to probe sensorimotor adaptation. To start each trial, the cursor (20-pixel radius) appeared in a central start position (25-pixel radius). A white target circle (30-pixel radius) appeared at one of eight locations (0, 45, 90, 135, 180, 225, 270, 315°) on an invisible ring around the central position (300-pixel radius) and filled in (white) following a 200 ms delay. Once launched, the cursor would travel the 300-pixel distance to the ring over a 750 ms period (with a bell-shaped velocity profile) before becoming stationary at the ring to provide participants with end-point error feedback. If the cursor overlapped with the target to any extent, the target would become green, signifying a target "hit". Each trial was separated by 4 s and within this period, participants had 2.6 s from target presentation to complete the trial (including the 200 ms target delay, participants' own reaction time, and the 750 ms cursor movement; any remaining time was allotted to providing the end-point error feedback). At 2.6 s the trial was ended, the screen was blanked, the data saved, and participants would briefly wait for the next trial to begin. Reaction times were not stressed in this experimental procedure. On trials in which the reaction time exceeded 2.6 s, the trial would end, and the participant would wait for the next trial to begin. These discarded trials were rare (0.56% across all trials, all participants) and were excluded from behavioural analyses, but were kept in the neuroimaging analysis due to the continuous nature of the fMRI task and our focus on functional connectivity analyses.

During each testing session, 120 baseline trials (15 sets of 8 trials) were completed without a rotation of the cursor. Following these trials, 320 learning trials (40 sets of 8 trials) were completed, wherein a 45° clockwise rotation of the cursor was applied. The baseline and learning trials were completed during one continuous fMRI scan. Following this scan, conditions were restored to baseline (i.e. no rotation of cursor) in a separate scan and participants performed 120 washout trials. These washout trials allowed us to probe participants' rate of relearning 24 hours later (and thus, their savings). In addition to these VMR-related task components, we also interspersed three 6-minute resting state fMRI scans prior to, between, and following VMR Learning and Washout. During these resting-state scans, participants were instructed to rest with their eyes open, while fixating a central cross location presented on the screen. The total testing time was 75 minutes on each testing day.

3. *Behavioral data processing*

On each trial we measured the angle error between the target and the final cursor position. Trials with reactions times greater than 2 s or less than 100 ms were discarded (the latter value was chosen as a conservative threshold on prepotent or anticipatory responses). Additionally, because subjects occasionally had difficulty utilizing the force sensor, leading to trials in with extreme cursor deviations (e.g. in the opposite direction of the target), we performed outlier removal by first detrending each block by subtracting a smoothed error curve estimated using a cubic spline basis with one knot per 40 trials, and then excluding those trials more than 3 standard deviations from the fitted curve. Excluded trials were then interpolated using the same cubic spline fit. In total, 4.3% of trials were removed and interpolated across all subjects and all task blocks.

4. *Behavioral summary measures*

For each subject, we computed the mean early error during the rotation block on each day, defined as the circular mean error during the first four blocks of eight trials. We defined savings as the difference between early error on days 1 and 2 (Haith et al. 2015, Morehead et al. 2015). These measures (day 1 early error, day 2 early error, and savings) were clustered using hierarchical complete linkage clustering on the Manhattan distance between the standardized variables. We extracted three clusters, corresponding roughly to subjects which learned the rotation rapidly on both days (FF learners), subjects which learned the rotation slowly on both days (SS), and subjects which learned slowly on day 1 and quickly on day 2 (SF).

We assessed the significance of this cluster solution through a parametric bootstrap procedure inspired by Kimes et al. (2017). We drew 5000 bootstrap samples from a null model in which the data were derived from a single multivariate normal distribution with covariance equal to that of the observed data. For each of these bootstrap samples, we performed hierarchical clustering and evaluated the ratio of within cluster to total sums of squares for a number of clusters between two and eight (Figure 2b, right). Only solutions with three or four clusters produced within-cluster sums of squares lower than what would be expected under the null model (Figure 2c), and as the three cluster solution was both slightly superior and more parsimonious, this solution formed the basis of our analysis.

C. fmri preprocessing

Preprocessing was performed using *fMRIPrep* 1.4.0 (Esteban et al. (2018b); Esteban et al. (2018a); RRID:SCR_016216), which is based on *Nipype* 1.2.0 (Gorgolewski et al. (2011); Gorgolewski et al. (2018); RRID:SCR_002502).

1. Anatomical data preprocessing

T1-weighted (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 2 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 two standard spaces (MNI152NLin6Asym, MNI152NLin2009cAsym) was performed through nonlinear registration with `antsRegistration` (ANTs 2.2.0), using brain-extracted versions of both T1w reference and the T1w template. The following templates were selected for spatial normalization: *FSL's MNI ICBM 152 non-linear 6th Generation Asymmetric Average Brain Stereotaxic Registration Model* [Evans et al. (2012), RRID:SCR_002823; TemplateFlow ID: MNI152NLin6Asym], *ICBM 152 Nonlinear Asymmetrical template version 2009c* [Fonov et al. (2009), RRID:SCR_008796; TemplateFlow ID: MNI152NLin2009cAsym].

2. Functional data preprocessing

For each scan, the following preprocessing was performed. First, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. 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 nine

degrees of freedom to account for distortions remaining in the BOLD reference. 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 time-series, were resampled to surfaces on the following spaces: *fsaverage*. The BOLD time-series (including slice-timing correction when applied) were resampled onto their original, native space by applying a single, composite transform to correct for head-motion and susceptibility distortions. These resampled BOLD time-series will be referred to as *preprocessed BOLD in original space*, or just *preprocessed BOLD*. The BOLD time-series were resampled into several standard spaces, correspondingly generating the following *spatially-normalized, preprocessed BOLD runs*: MNI152NLin6Asym, MNI152NLin2009cAsym. First, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. Automatic removal of motion artifacts using independent component analysis (ICA-AROMA, [Pruim et al. 2015](#)) was performed on the *preprocessed BOLD on MNI space* time-series after removal of non-steady state volumes and spatial smoothing with an isotropic, Gaussian kernel of 6mm FWHM (full-width half-maximum). Corresponding "non-aggressively" denoised runs were produced after such smoothing. Additionally, the "aggressive" noise-regressors were collected and placed in the corresponding confounds file. Several confounding time-series were calculated based on the *preprocessed BOLD*: framewise displacement (FD), DVARS and three region-wise global signals. FD and DVARS are calculated for each functional run, both using their implementations in *Nipype* (following the definitions by [Power et al. 2014](#)). The three global signals are extracted within the CSF, the WM, and the whole-brain masks. Additionally, a set of physiological regressors were extracted to allow for component-based noise correction (*CompCor*, [Behzadi et al. 2007](#)). Principal components are estimated after high-pass filtering the *preprocessed BOLD* time-series (using a discrete cosine filter with 128s cut-off) for the two *CompCor* variants: temporal (tCompCor) and anatomical (aCompCor). tCompCor components are then calculated from the top 5% variable voxels within a mask covering the subcortical regions. This subcortical mask is obtained by heavily eroding the brain mask, which ensures it does not include cortical GM regions. For aCompCor, components are calculated within the intersection of the aforementioned mask and the union of CSF and WM masks calculated in T1w space, after their projection to the native space of each functional run (using the inverse BOLD-to-T1w transformation). Components are also calculated separately within the WM and CSF masks. For each CompCor decomposition, the k components with the largest singular values are retained, such

that the retained components' time series are sufficient to explain 50 percent of variance across the nuisance mask (CSF, WM, combined, or temporal). The remaining components are dropped from consideration. The head-motion estimates calculated in the correction step were also placed within the corresponding confounds file. The confound time series derived from head motion estimates and global signals were expanded with the inclusion of temporal derivatives and quadratic terms for each (Satterthwaite et al. 2013). Frames that exceeded a threshold of 0.5 mm FD or 1.5 standardised DVARS were annotated as motion outliers. All resamplings can be performed with *a single interpolation step* by composing all the pertinent transformations (i.e. head-motion transform matrices, susceptibility distortion correction when available, and co-registrations to anatomical and output spaces). Gridded (volumetric) resamplings were performed using `antsApplyTransforms` (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels (Lanczos 1964). Non-gridded (surface) resamplings were performed using `mri_vol2surf` (FreeSurfer).

Many internal operations of *fMRIPrep* use *Nilearn* 0.5.2 (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](#).

D. Network embedding

We defined motor and cognitive networks using the whole-brain parcellation proposed by Seitzman et al. (2020), comprising 300 ROIs across the cortex, subcortex, and cerebellum. Our **Motor** network was comprised of those ROIs functionally assigned to the SomatomotorDorsal and SomatomotorLateral networks, while our **Cognitive** network comprised the FrontoParietal, DorsalAttention, and VentralAttention networks. The later was chosen to include a wide array of higher order association regions in the frontal and parietal cortices. Mean time courses were extracted for each ROI by intersecting an 8mm sphere with a grey matter mask, excluding those voxels lying in the intersection of multiple spheres. See Figures 3a and 4a for visualizations of each network.

1. Covariance estimation, centering, and embedding

For each subject, we used the estimator proposed by Ledoit and Wolf (2004) to derive covariance matrices for the day 1 rest scan, as well as periods of equivalent length (177 TRs) at the end of baseline, and the beginning and end of rotation (Figure 1b). Note that we use only the first resting scan in order to avoid any effects due to learning related changes in resting state network structure.

We chose periods of equal length to the resting state scan in order to avoid introducing bias into our centering procedure, as the regularization applied by our covariance estimator is a function of the sample size. We then centered each subjects covariance matrices so that their resting scan aligned with the grand mean covariance during rest. For each subject i , we performed centering separately on the cognitive and sensorimotor networks as follows:

Let $\bar{R}$ be the mean resting state covariance matrix (see [Moakher 2005](#)), R_i the subject's observed resting state covariance matrix, and S_{ijd} one of the task covariance matrices, where $j \in \{baseline, early, late\}$ and $d \in \{Day1, Day2\}$. We then projected the observed covariance S_{ijd} onto the tangent space around R_i in order to obtain a tangent vector T_{ijd} (a symmetric matrix) as follows:

$$T_{ijd} = R_i^{1/2} \log(R_i^{-1/2} S_{ijd} R_i^{-1/2}) R_i^{1/2}, \quad (1)$$

where $\log()$ denotes the matrix logarithm. As each covariance matrix X is symmetric and positive-definite, it can be written $X = V\Lambda V^\top$, where V is a matrix of eigenvectors and Λ a diagonal matrix of eigenvalues. The matrix square root $X^{1/2}$ can then be computed as $X^{1/2} = V\Lambda^{1/2}V^\top$, and similarly for $X^{-1/2}$.

We then translated each tangent vector to the grand mean resting state scan $\bar{R}$ using the transport proposed by [Zhao et al. \(2018\)](#), obtaining a centered tangent vector

$$T_{ijd}^c = G T_{ijd} G^\top \quad (2)$$

where $G = \bar{R}^{1/2} R^{-1/2}$. This tangent vector can then be projected back onto the space of covariance matrices by

$$S_{ijd}^c = \bar{R}^{1/2} \exp(\bar{R}^{-1/2} T_{ijd}^c \bar{R}^{-1/2}) \bar{R}^{1/2}, \quad (3)$$

in order to obtain a centered covariance matrix. Rather than work directly with the covariance, we focus our analysis on the centered tangent vectors themselves, which can be interpreted as differences in covariance relative to resting state ([Varoquaux et al. 2010](#)).

In order to visualize the effects of centering, we derived a two-dimensional embedding of the

covariance matrices before and after centering using uniform manifold approximation (UMAP; McInnes et al. 2018) on the matrix of pairwise geodesic distances between covariance matrices (Figures 3a and 4a). This embedding reveals that centering was successful in eliminating static differences in covariance, and in isolating the effect of task condition.

We then constructed weighted graphs from these tangent vectors, and derived a joint two-dimensional embedding of the graphs in each network using the model proposed by Wang et al. (2019). Intuitively, the embedding seeks a two-dimensional representation of the ROIs in each network, with each individual observation (group, scan, day) receiving a score indicating the degree to which it expresses each component of the embedding. These components are shown in Figures 3d and Figures 4d, while scores for each observation are given in Figures 3g and Figures 4g.

E. Excursion

F. Network excursion

For each subject, we estimated resting state manifolds for each of the motor and cognitive networks by applying PCA to the ROI timecourses from first resting state scan on Day 1, retaining the number of components required to explain 75% of the variance. Note that we experimented with several procedures for selecting the number of components, including explained variance from 50% to 90%, and the use of k-fold cross-validation. As all methods returned similar results, we report the results with a threshold of 75% as a median compromise.

The retained components form a basis for a subspace containing the majority of variability in the observed resting BOLD activation. Assembling these components in the columns of a matrix $\mathbf{V}$, we define the on-manifold activation at time t to be the projection of the BOLD activation $\mathbf{x}_t$ onto the the subspace spanned by $\mathbf{V}$, given by the projection matrix

$$\mathbf{P}^{\text{on}} = \mathbf{V}^{\top}(\mathbf{V}^{\top}\mathbf{V})^{-1}\mathbf{V}$$

The off-manifold activation is then the projection onto the orthogonal compliment of this subspace, given by the projection matrix

$$\mathbf{P}^{\text{off}} = \mathbf{I} - \mathbf{P}^{\text{on}}$$

Letting $\mathbf{x}_t^{\text{on}} = \mathbf{P}^{\text{on}}\mathbf{x}_t$ and $\mathbf{x}_t^{\text{off}} = \mathbf{P}^{\text{off}}\mathbf{x}_t$ be the on and off manifold activation, respectively, we

define the off-manifold activation at time t to be $e_t = \|\mathbf{x}_t^{\text{off}}\|$. As this activation is correlated with the overall BOLD activation $\|\mathbf{x}_t\|$, we define the excursion $r_t = e_t/\|\mathbf{x}_t\|$ to be the ratio of the off-manifold activation to the overall BOLD activation. Intuitively, this measure roughly describes the proportion of the overall BOLD activation which is *off-manifold* (see (Figure 1d))

After estimating resting-state manifolds for each subject, we computed the excursion within the task scan, comprising both the pre- and post-rotation periods. We then standardized the excursion trajectories using the mean and standard deviation during the pre-rotation trials, to account for baseline differences in excursion.

G. Excursion components

For each subject, we derived a set of four excursion trajectories, comprising the excursion at each TR for both days and both networks. We then used multivariate functional principal component analysis (mfPCA; [Happ and Greven 2018](#)) to identify the dominant patterns of excursion across networks and across days. Just as PCA estimates a basis for a low dimensional subspace capturing variability in a set of data lying in Euclidean space $\mathbb{R}^n$, so that each observation can be approximated by the sum of a small set of (ideally, interpretable) components, functional PCA likewise attempts to approximate a set of *functions* by a smaller set of basis functions. This functional approach is often more suitable in cases where observations represent continuous functions, as the basis expansion used in functional PCA allows features like smoothness or periodicity to be encoded into the components. As each subject contributes a set of four observations – excursion trajectories for each day and each network – we seek a *multivariate* decomposition to derive components describing patterns of excursion across days and networks.

We first represented each excursion trajectory as a continuous function using a cubic spline basis with knots at each observation and smoothing penalty selected by generalized cross validation ([Härdle 1990](#)). Multivariate functional PCA was then performed using the R package MFPCA ([Happ-Kurz 2020](#)) using a penalized spline basis. We extracted three components (Figure 5b), sufficient to explain the majority of the variance in subject excursion trajectories (Figure 5c). These components encoded, respectively, the total excursion in cognitive and sensorimotor networks on day 1; the excursion in the cognitive network on day 1; and the excursion in the cognitive network on day 2. For each component, we compared scores in each group using a Kruskal-Wallis one-way analysis of variance. Post-hoc comparisons were made using the Conover-Iman procedure ([Conover and Iman 1979](#)), corrected for multiple comparisons using the method of [Benjamini and Hochberg](#)

(1995).

ACKNOWLEDGMENTS

This work was supported by operating grants from the Canadian Institutes of Health Research (CIHR) awarded to J.P.G (MOP126158). C.N.A. was supported by a Natural Sciences and Engineering Research Council (NSERC) graduate award. J.P.G. was also supported by a NSERC Discovery Grant, as well as funding from the Canadian Foundation for Innovation.

-
- Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., and Varoquaux, G. (2014). Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics*, 8.
- Anguera, J. A., Reuter-Lorenz, P. A., Willingham, D. T., and Seidler, R. D. (2010). Contributions of spatial working memory to visuomotor learning. *Journal of cognitive neuroscience*, 22(9):1917–1930.
- Anguera, J. A., Russell, C. A., Noll, D. C., and Seidler, R. D. (2007). Neural correlates associated with intermanual transfer of sensorimotor adaptation. *Brain research*, 1185:136–151.
- Avants, B., Epstein, C., Grossman, M., and Gee, J. (2008). Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, 12(1):26–41.
- Behzadi, Y., Restom, K., Liau, J., and Liu, T. T. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fmri. *NeuroImage*, 37(1):90–101.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300.
- Benson, B. L., Anguera, J. A., and Seidler, R. D. (2011). A spatial explicit strategy reduces error but interferes with sensorimotor adaptation. *Journal of neurophysiology*, 105(6):2843–2851.
- Christou, A. I., Miall, R. C., McNab, F., and Galea, J. M. (2016). Individual differences in explicit and implicit visuomotor learning and working memory capacity. *Scientific reports*, 6(1):1–13.
- Conover, W. J. and Iman, R. L. (1979). On multiple-comparisons procedures. Technical report, Technical report, Los Alamos Scientific Laboratory.
- Corbetta, M. and Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature reviews neuroscience*, 3(3):201–215.
- Cox, R. W. and Hyde, J. S. (1997). Software tools for analysis and visualization of fmri data. *NMR in Biomedicine*, 10(4-5):171–178.
- Dale, A. M., Fischl, B., and Sereno, M. I. (1999). Cortical surface-based analysis: I. segmentation and surface reconstruction. *NeuroImage*, 9(2):179–194.
- Dayan, E. and Cohen, L. G. (2011). Neuroplasticity subserving motor skill learning. *Neuron*, 72(3):443–454.

- de Brouwer, A. J., Albaghdadi, M., Flanagan, J. R., and Gallivan, J. P. (2018). Using gaze behavior to parcellate the explicit and implicit contributions to visuomotor learning. *Journal of neurophysiology*, 120(4):1602–1615.
- Doyon, J., Bellec, P., Amsel, R., Penhune, V., Monchi, O., Carrier, J., Lehericy, S., and Benali, H. (2009). Contributions of the basal ganglia and functionally related brain structures to motor learning. *Behavioural brain research*, 199(1):61–75.
- Drummond, N. M., Cressman, E. K., and Carlsen, A. N. (2015). Inhibition of motor-related activation during a simple reaction time task requiring visuomotor mental rotation. *Behavioral neuroscience*, 129(2):160.
- Duncan, J. (2010). The multiple-demand (md) system of the primate brain: mental programs for intelligent behaviour. *Trends in cognitive sciences*, 14(4):172–179.
- Esteban, O., Blair, R., Markiewicz, C. J., Berleant, S. L., Moodie, C., Ma, F., Isik, A. I., Erramuzpe, A., Kent, James D. andGoncalves, M., DuPre, E., Sitek, K. R., Gomez, D. E. P., Lurie, D. J., Ye, Z., Poldrack, R. A., and Gorgolewski, K. J. (2018a). fmriprep. *Software*.
- Esteban, O., Markiewicz, C., Blair, R. W., Moodie, C., Isik, A. I., Erramuzpe Aliaga, A., Kent, J., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh, S., Wright, J., Durnez, J., Poldrack, R., and Gorgolewski, K. J. (2018b). fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nature Methods*.
- Evans, A., Janke, A., Collins, D., and Baillet, S. (2012). Brain templates and atlases. *NeuroImage*, 62(2):911–922.
- Fernandez-Ruiz, J., Wong, W., Armstrong, I. T., and Flanagan, J. R. (2011). Relation between reaction time and reach errors during visuomotor adaptation. *Behavioural brain research*, 219(1):8–14.
- Fonov, V., Evans, A., McKinstry, R., Almlie, C., and Collins, D. (2009). Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage*, 47, Supplement 1:S102.
- Gallego, J. A., Perich, M. G., Naufel, S. N., Ethier, C., Solla, S. A., and Miller, L. E. (2017). Multiple tasks viewed from the neural manifold: Stable control of varied behavior. *BioRxiv*, page 176081.
- Gao, S., Mishne, G., and Scheinost, D. (2020). Non-linear manifold learning in fmri uncovers a low-dimensional space of brain dynamics. *bioRxiv*.
- Gorgolewski, K., Burns, C. D., Madison, C., Clark, D., Halchenko, Y. O., Waskom, M. L., and Ghosh, S. (2011). Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python. *Frontiers in Neuroinformatics*, 5:13.
- Gorgolewski, K. J., Esteban, O., Markiewicz, C. J., Ziegler, E., Ellis, D. G., Notter, M. P., Jarecka, D., Johnson, H., Burns, C., Manhães-Savio, A., Hamalainen, C., Yvernault, B., Salo, T., Jordan, K., Goncalves,

M., Waskom, M., Clark, D., Wong, J., Loney, F., Modat, M., Dewey, B. E., Madison, C., Visconti di Oleggio Castello, M., Clark, M. G., Dayan, M., Clark, D., Keshavan, A., Pinsard, B., Gramfort, A., Berleant, S., Nielson, D. M., Bougacha, S., Varoquaux, G., Cipollini, B., Markello, R., Rokem, A., Moloney, B., Halchenko, Y. O., Wassermann, D., Hanke, M., Horea, C., Kaczmarzyk, J., de Hollander, G., DuPre, E., Gillman, A., Mordom, D., Buchanan, C., Tungaraza, R., Pauli, W. M., Iqbal, S., Sikka, S., Mancini, M., Schwartz, Y., Malone, I. B., Dubois, M., Frohlich, C., Welch, D., Forbes, J., Kent, J., Watanabe, A., Cumba, C., Huntenburg, J. M., Kastman, E., Nichols, B. N., Eshaghi, A., Ginsburg, D., Schaefer, A., Acland, B., Giavasis, S., Kleesiek, J., Erickson, D., Küttner, R., Haselgrove, C., Correa, C., Ghayoor, A., Liem, F., Millman, J., Haehn, D., Lai, J., Zhou, D., Blair, R., Glatard, T., Renfro, M., Liu, S., Kahn, A. E., Pérez-García, F., Triplett, W., Lampe, L., Stadler, J., Kong, X.-Z., Hallquist, M., Chetverikov, A., Salvatore, J., Park, A., Poldrack, R., Craddock, R. C., Inati, S., Hinds, O., Cooper, G., Perkins, L. N., Marina, A., Mattfeld, A., Noel, M., Snoek, L., Matsubara, K., Cheung, B., Rothmei, S., Urchs, S., Durnez, J., Mertz, F., Geisler, D., Floren, A., Gerhard, S., Sharp, P., Molina-Romero, M., Weinstein, A., Broderick, W., Saase, V., Andberg, S. K., Harms, R., Schlamp, K., Arias, J., Papadopoulos Orfanos, D., Tarbert, C., Tambini, A., De La Vega, A., Nickson, T., Brett, M., Falkiewicz, M., Podranski, K., Linkersdörfer, J., Flandin, G., Ort, E., Shachnev, D., McNamee, D., Davison, A., Varada, J., Schwabacher, I., Pellman, J., Perez-Guevara, M., Khanuja, R., Pannetier, N., McDermottroe, C., and Ghosh, S. (2018). Nipype. *Software*.

Gratton, C., Laumann, T. O., Nielsen, A. N., Greene, D. J., Gordon, E. M., Gilmore, A. W., Nelson, S. M., Coalson, R. S., Snyder, A. Z., Schlaggar, B. L., et al. (2018). Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. *Neuron*, 98(2):439–452.

Greve, D. N. and Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, 48(1):63–72.

Haith, A. M., Huberdeau, D. M., and Krakauer, J. W. (2015). The influence of movement preparation time on the expression of visuomotor learning and savings. *Journal of neuroscience*, 35(13):5109–5117.

Happ, C. and Greven, S. (2018). Multivariate functional principal component analysis for data observed on different (dimensional) domains. *Journal of the American Statistical Association*, 113(522):649–659.

Happ-Kurz, C. (2020). *MFPCA: Multivariate Functional Principal Component Analysis for Data Observed on Different Dimensional Domains*. R package version 1.3-6.

Härdle, W. (1990). *Applied nonparametric regression*. Number 19. Cambridge university press.

Hardwick, R. M., Rottschy, C., Miall, R. C., and Eickhoff, S. B. (2013). A quantitative meta-analysis and review of motor learning in the human brain. *Neuroimage*, 67:283–297.

Hsu, H. M., Yao, Z.-F., Hwang, K., and Hsieh, S. (2020). Between-module functional connectivity of the salient ventral attention network and dorsal attention network is associated with motor inhibition. *PloS*

one, 15(12):e0242985.

Izawa, J., Criscimagna-Hemminger, S. E., and Shadmehr, R. (2012). Cerebellar contributions to reach adaptation and learning sensory consequences of action. *Journal of Neuroscience*, 32(12):4230–4239.

Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):825–841.

Keisler, A. and Shadmehr, R. (2010). A shared resource between declarative memory and motor memory. *Journal of Neuroscience*, 30(44):14817–14823.

Kimes, P. K., Liu, Y., Neil Hayes, D., and Marron, J. S. (2017). Statistical significance for hierarchical clustering. *Biometrics*, 73(3):811–821.

Klein, A., Ghosh, S. S., Bao, F. S., Giard, J., Häme, Y., Stavsky, E., Lee, N., Rossa, B., Reuter, M., Neto, E. C., and Keshavan, A. (2017). Mindboggling morphometry of human brains. *PLOS Computational Biology*, 13(2):e1005350.

Krakauer, J. W. (2009). Motor learning and consolidation: the case of visuomotor rotation. In *Progress in motor control*, pages 405–421. Springer.

Krakauer, J. W. and Mazzoni, P. (2011). Human sensorimotor learning: adaptation, skill, and beyond. *Current opinion in neurobiology*, 21(4):636–644.

Lanczos, C. (1964). Evaluation of noisy data. *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis*, 1(1):76–85.

Ledoit, O. and Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis*, 88(2):365–411.

Leow, L.-A., Gunn, R., Marinovic, W., and Carroll, T. J. (2017). Estimating the implicit component of visuomotor rotation learning by constraining movement preparation time. *Journal of neurophysiology*, 118(2):666–676.

Mars, R. B., Neubert, F.-X., and Rushworth, M. F. (2011). Top-down control over the motor cortex. *Neural basis of motivational and cognitive control*. MIT Press, Cambridge.

McDougle, S. D., Bond, K. M., and Taylor, J. A. (2015). Explicit and implicit processes constitute the fast and slow processes of sensorimotor learning. *Journal of Neuroscience*, 35(26):9568–9579.

McDougle, S. D., Ivry, R. B., and Taylor, J. A. (2016). Taking aim at the cognitive side of learning in sensorimotor adaptation tasks. *Trends in cognitive sciences*, 20(7):535–544.

- McInnes, L., Healy, J., and Melville, J. (2018). Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Miyamoto, Y. R., Wang, S., and Smith, M. A. (2020). Implicit adaptation compensates for erratic explicit strategy in human motor learning. *Nature neuroscience*, 23(3):443–455.
- Moakher, M. (2005). A differential geometric approach to the geometric mean of symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications*, 26(3):735–747.
- Morehead, J. R., Qasim, S. E., Crossley, M. J., and Ivry, R. (2015). Savings upon re-aiming in visuomotor adaptation. *Journal of neuroscience*, 35(42):14386–14396.
- Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Byron, M. Y., Chase, S. M., and Batista, A. P. (2019). New neural activity patterns emerge with long-term learning. *Proceedings of the National Academy of Sciences*, 116(30):15210–15215.
- Oldfield, R. C. et al. (1971). The assessment and analysis of handedness: the edinburgh inventory. *Neuropsychologia*, 9(1):97–113.
- Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., Vogel, A. C., Laumann, T. O., Miezin, F. M., Schlaggar, B. L., et al. (2011). Functional network organization of the human brain. *Neuron*, 72(4):665–678.
- Power, J. D., Mitra, A., Laumann, T. O., Snyder, A. Z., Schlaggar, B. L., and Petersen, S. E. (2014). Methods to detect, characterize, and remove motion artifact in resting state fmri. *NeuroImage*, 84(Supplement C):320–341.
- Pruim, R. H. R., Mennes, M., van Rooij, D., Llera, A., Buitelaar, J. K., and Beckmann, C. F. (2015). Ica-AROMA: A robust ICA-based strategy for removing motion artifacts from fmri data. *NeuroImage*, 112(Supplement C):267–277.
- Reuter, M., Rosas, H. D., and Fischl, B. (2010). Highly accurate inverse consistent registration: A robust approach. *NeuroImage*, 53(4):1181–1196.
- Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Byron, M. Y., and Batista, A. P. (2014). Neural constraints on learning. *Nature*, 512(7515):423–426.
- Satterthwaite, T. D., Elliott, M. A., Gerraty, R. T., Ruparel, K., Loughead, J., Calkins, M. E., Eickhoff, S. B., Hakonarson, H., Gur, R. C., Gur, R. E., and Wolf, D. H. (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(1):240–256.

- Seitzman, B. A., Gratton, C., Marek, S., Raut, R. V., Dosenbach, N. U., Schlaggar, B. L., Petersen, S. E., and Greene, D. J. (2020). A set of functionally-defined brain regions with improved representation of the subcortex and cerebellum. *Neuroimage*, 206:116290.
- Shenoy, K. V., Sahani, M., and Churchland, M. M. (2013). Cortical control of arm movements: a dynamical systems perspective. *Annual review of neuroscience*, 36:337–359.
- Shine, J. M., Hearne, L. J., Breakspear, M., Hwang, K., Müller, E. J., Sporns, O., Poldrack, R. A., Mattingley, J. B., and Cocchi, L. (2019). The low-dimensional neural architecture of cognitive complexity is related to activity in medial thalamic nuclei. *Neuron*, 104(5):849–855.
- Simpson, G. V., Weber, D. L., Dale, C. L., Pantazis, D., Bressler, S. L., Leahy, R. M., and Luks, T. L. (2011). Dynamic activation of frontal, parietal, and sensory regions underlying anticipatory visual spatial attention. *Journal of neuroscience*, 31(39):13880–13889.
- Spreng, R. N., Sepulcre, J., Turner, G. R., Stevens, W. D., and Schacter, D. L. (2013). Intrinsic architecture underlying the relations among the default, dorsal attention, and frontoparietal control networks of the human brain. *Journal of cognitive neuroscience*, 25(1):74–86.
- Taylor, J. A., Krakauer, J. W., and Ivry, R. B. (2014). Explicit and implicit contributions to learning in a sensorimotor adaptation task. *Journal of Neuroscience*, 34(8):3023–3032.
- Tustison, N. J., Avants, B. B., Cook, P. A., Zheng, Y., Egan, A., Yushkevich, P. A., and Gee, J. C. (2010). N4itk: Improved n3 bias correction. *IEEE Transactions on Medical Imaging*, 29(6):1310–1320.
- Varoquaux, G., Baronnet, F., Kleinschmidt, A., Fillard, P., and Thirion, B. (2010). Detection of brain functional-connectivity difference in post-stroke patients using group-level covariance modeling. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 200–208. Springer.
- Wang, S., Arroyo, J., Vogelstein, J. T., and Priebe, C. E. (2019). Joint embedding of graphs. *IEEE transactions on pattern analysis and machine intelligence*.
- Wolpert, D. M., Diedrichsen, J., and Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature Reviews Neuroscience*, 12(12):739–751.
- Wolpert, D. M. and Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature neuroscience*, 3(11):1212–1217.
- Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 269(5232):1880–1882.

Zhang, Y., Brady, M., and Smith, S. (2001). Segmentation of brain MR images through a hidden markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*, 20(1):45–57.

Zhao, Q., Kwon, D., and Pohl, K. M. (2018). A riemannian framework for longitudinal analysis of resting-state functional connectivity. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 145–153. Springer.