



Human midcingulate cortex encodes distributed representations of task progress

Clay B. Holroyd^{a,1,2}, José J. F. Ribas-Fernandes^{a,1}, Danesh Shahnazian^a, Massimo Silvetti^{b,c}, and Tom Verguts^{b,c}

^aDepartment of Psychology, University of Victoria, Victoria, BC V8W 2Y2, Canada; ^bDepartment of Experimental Psychology, Ghent University, B-9000 Ghent, Belgium; and ^cGhent Institute for Functional and Metabolic Imaging (GIFMI), Ghent University Hospital, 9000 Ghent, Belgium

Edited by Peter L. Strick, University of Pittsburgh, Pittsburgh, PA, and approved May 9, 2018 (received for review February 28, 2018)

The function of midcingulate cortex (MCC) remains elusive despite decades of investigation and debate. Complicating matters, individual MCC neurons respond to highly diverse task-related events, and MCC activation is reported in most human neuroimaging studies employing a wide variety of task manipulations. Here we investigate this issue by applying a model-based cognitive neuroscience approach involving neural network simulations, functional magnetic resonance imaging, and representational similarity analysis. We demonstrate that human MCC encodes distributed, dynamically evolving representations of extended, goal-directed action sequences. These representations are uniquely sensitive to the stage and identity of each sequence, indicating that MCC sustains contextual information necessary for discriminating between task states. These results suggest that standard univariate approaches for analyzing MCC function overlook the major portion of task-related information encoded by this brain area and point to promising new avenues for investigation.

midcingulate cortex | recurrent neural network | representational similarity analysis | fMRI | sequence execution

Anterior midcingulate cortex (MCC) (1) contributes to cognitive control and decision making, but its specific function is hotly debated (2–4). Attempts to build a general theory about it have been stymied by the fact that MCC neurons respond to a wide spectrum of task-related events (5–7) and because the region is activated in most functional neuroimaging studies (8). However, recent investigations have yielded two key insights into its function. First, application of multivariate statistical techniques (9) to the ensemble activity of MCC neurons in nonhuman animals (5, 10–14), and to the functional magnetic resonance imaging (fMRI) blood-oxygen-level dependent (BOLD) response in humans (15–18), has revealed distributed representations in MCC that are invisible to measures that pool across the activity of single neurons. As with distributed patterns observed elsewhere in cortex (19, 20), these representations suggest that the idiosyncratic sensitivities of MCC neurons to actions, stimuli, rewards, and other task-related events reflect their collective purpose (5; see ref. 21 for review). Second, recent evidence indicates that MCC evaluates and sustains the execution of extended, sequential behaviors (11, 22–26). MCC therefore appears to guide the agent through a task space by encoding goal-directed action sequences as dynamically evolving representations that are distributed across neural ensembles (13, 14).

Despite these advances, a formal framework that relates MCC function to a distributed encoding scheme remains to be established. As a step in this direction, we recently proposed that recurrent neural networks (RNNs) can be used to predict the signature features of MCC representations during sequence execution (Fig. 14) (21). RNN connectivity enables maintaining memories of past events to disambiguate otherwise identical states (27). When trained to predict action sequences, the activation patterns across a layer of so-called context units in RNNs can provide insight into the structure of the internal representations underlying the sequences (28). These patterns are broadly compatible with the ensemble activity of MCC neurons as nonhuman animals carry out goal-directed action sequences (21). On this

view, rather than encoding univariate signals such as conflicts or prediction errors (2–4), MCC represents the execution of extended tasks using a distributed, multivariate code that discriminates between the different steps of each task sequence.

However, testing this hypothesis in humans faces a key challenge, as intracranial studies of MCC neurons using multielectrode arrays are not readily amenable to human experiments. Moreover, although fMRI provides a noninvasive alternative, its spatial and temporal resolution is not comparable with the resolution of multielectrode recordings. Here we address this challenge with a combined model-based cognitive neuroscience (29) and multivariate statistical approach that characterizes the similarities between representations across representational systems (30). In particular, the approach identifies second-order isomorphisms (31) between model and neural data, thereby obviating the need to establish a precise (first-order) correspondence between neural firing, BOLD, and model neurons. In line with the predictions of our existing computational MCC model, the results reveal that human MCC encodes the execution of goal-directed action sequences as distributed representations that evolve dynamically as the task progresses.

Results

MCC and RNNs Encode Task Sequences Similarly. We adapted an “activities of daily living” (ADL) task to an fMRI design. Although ADL tasks like brushing one’s teeth can be carried out by healthy

Significance

Midcingulate cortex (MCC) contributes to cognitive control and decision making but its specific role in this process is highly mysterious. Here we use functional magnetic resonance imaging to test the predictions of a neural network model of MCC function. Confirming the model predictions, we find that human MCC encodes distributed, dynamically evolving representations of goal-directed action sequences. These results suggest that standard approaches for analyzing MCC function overlook the major portion of information encoded by this brain area. Rather, they indicate that MCC encodes the distances between representations of task events in task space, revealing how the MCC sustains the execution of extended behaviors.

Author contributions: C.B.H. conceived the study; C.B.H., J.J.F.R.-F., D.S., M.S., and T.V. designed research; C.B.H., J.J.F.R.-F., D.S., and M.S. performed research; C.B.H., J.J.F.R.-F., and D.S. analyzed data; and C.B.H., J.J.F.R.-F., D.S., M.S., and T.V. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

Data deposition: fMRI data and computational modelling code have been deposited in the Open Science Framework repository (<https://osf.io/UHJCF/>; doi: 10.17605/OSF.IO/WXHTA).

¹C.B.H. and J.J.F.R.-F. contributed equally to this work.

²To whom correspondence should be addressed. Email: holroyd@uvic.ca.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1803650115/-DCSupplemental.

Published online June 4, 2018.




Fig. 1. Computational model. (A) The recurrent neural network is composed of 7 input units, 15 hidden units, and 8 output units with fully connected feed-forward projections. A fourth context layer with 15 units receives an exact copy of the hidden unit activation values on each sequence step and returns those values to the hidden layer on the subsequent step. The model is trained to predict each action in a sequence, given the previous action as input. For example, the model predicts a “stir” action (spoon unit) given an “add water” action (teapot unit) as input. Arrows indicate connections; squares represent individual units. Midcingulate cortex representations are derived from the activation values of the context layer units (“task representation”) (*SI Appendix*). (B) Representational dissimilarity matrix derived from the context unit activation values. The 24 states along each axis correspond to the six actions associated with each of the four task sequences given in Fig. 2. Blue vs. yellow squares indicate greater vs. lesser similarity between representations, respectively. (C) Multidimensional scaling of the context unit activation values (28, 41) across the 24 states for a representative simulation, revealing dimensions that are sensitive to sequence progression (“start” vs. “end”) and action type (“non-utensil” vs. “utensil”).

individuals relatively automatically, they consist of extended sequences subject to action slips that signal a need for supervisory control (32). Here we utilized a task that entails executing four different context-dependent action sequences, two that simulate coffee making and two that simulate tea making, according to six rules given to participants at the start of the experiment (Fig. 2) (*SI Appendix*). For example, subjects were told that cream and water can only be added to coffee, and sugar and water can only be added to tea. Thus, for states 2 and 4 (Fig. 2), the subjects were required to choose either cream or water when making coffee, or sugar and water when making tea (*SI Appendix*). In contrast to previous MCC studies in humans (2–4), this task requires subjects to remember their choice behavior—e.g., whether they have previously added water or sugar to their tea—to execute the sequences.

We used the RNN architecture, which we previously applied to account for distributed neuronal activity in nonhuman animals (21), to simulate how MCC encodes the sequences in the coffee–tea-making task (Fig. 1A). We expected that the context units of the model trained on this task would incorporate the overall task structure, revealing sensitivity to the contextual information needed to execute the action sequences correctly (21). In particular, we predicted that the context layer of the model and the MCC of human participants performing the task would represent the task states similarly. Further, because we expected that other brain regions might also exhibit activity partially similar to the model’s, we predicted that the correspondence with the model representations would be stronger for MCC than for other brain areas.

The model was trained to predict each state of the task based on the previous state as input (Figs. 1A and 3) (21) (*SI Appendix*). For example, following the first step of a coffee-making sequence, the model was trained to predict whether the next step would entail adding water, cream, or sugar. The training set consisted of four six-step action sequences (Fig. 2), for a total of 24 states. During a subsequent testing phase, the activation values of the 15 context units were read out for each of the 24 states. We then created

a representational dissimilarity matrix (RDM) for the model representations by calculating the Spearman’s rank correlation between the context unit activations for each pairwise combination of the 24 states, which yielded a 24×24 matrix (Figs. 1B and 3). This RDM enabled us to assess our key prediction that the MCC and the model would encode the task similarly, as a distributed representation tracking the progression of each sequence (30).

Eighteen participants performed the coffee–tea-making task in a 3T MRI scanner (Fig. 3) (*SI Appendix*). Errors occurred on 11.3% (SD = 5.0%) of trials (additional behavioral statistics appear in *SI Appendix*). A general linear model (GLM) to the MRI data (*SI Appendix*) yielded a separate activation pattern across the entire brain for each of the 24 task states, separately for each of the 18 participants. We then applied representational similarity analysis (RSA) (30) using a searchlight approach (33) by computing the similarity between the model RDM and same-size RDMs derived from each 27 -voxel vector associated with the $3 \times 3 \times 3$ voxel cube surrounding each voxel in the volume (*SI Appendix*). As we aimed to test the strong hypothesis that the model RDM corresponds to MCC representations better than to the representations of other brain regions (as opposed to the weak hypothesis that the model RDM simply correlates with MCC representations), the correlations were z transformed within subjects. Larger values therefore indicate brain areas that are more correlated than other brain areas with the model RDM. At the group level, these values were t tested against 0 across subjects, separately for each voxel; t values exceeding a threshold were then grouped into clusters. Protection against type-I errors was ensured with permutation testing (34, 35) (*SI Appendix*).

The largest cluster identified by the searchlight (348 voxels) was in medial–frontal cortex (MFC). As predicted, the cluster was centered on MCC (348 voxels), with a center of mass and peak voxel in the right paracingulate gyrus in Brodmann area (BA) 32 (Fig. 4A and *SI Appendix*, Fig. S1 and Table S1). This cluster extended rostrally through anterior MCC (including BA 9 and 32)




Fig. 2. The coffee–tea making task. Each trial consisted of a sequence of six decisions followed by feedback (indicated by digits on *Left*). Each decision comprised three images of different items corresponding to three possible actions (state 1: coffee, pepper, tea; state 2: cream, sugar, water; state 3: fork, spoon, knife; state 4: sugar, water, cream; state 5: spoon, knife, fork; state 6: prepared coffee, prepared tea, orange juice). For each decision, subjects selected one of the three actions by pressing a spatially corresponding button. Each decision always entailed selecting between the same three actions, but the positions of the images within each state varied across and within trials. Application of the task rules resulted in execution of four different task sequences. Arrows indicate allowable actions; the arrows diverge from states in which the action must be remembered to complete the sequence correctly (e.g., state 1 to state 2), and converge onto states when that information is no longer needed (e.g., state 4 to state 5). State 7 indicated correct (happy face) or incorrect (sad face) performance.

and caudally through posterior MCC and the presupplementary motor area (including BA 8 and BA 6) (36, 37). This cluster was much larger than the next biggest cluster, which comprised 106 voxels and was located in right frontal opercular cortex (*SI Appendix, Table S1*). Only the MFC cluster survived the permutation test correction. These findings indicate that MCC and the model represent the coffee–tea task sequences similarly, and that this correspondence is stronger for MCC than for other brain areas. To confirm that the model RDM also corresponded to MCC at the cluster level, we computed an RDM for the 24-action state space derived from the entire 348-voxel MFC cluster (Fig. 4*B*). As expected, the model RDM (Fig. 1*B*) was also positively correlated across subjects with this MFC-cluster RDM [mean $\rho = 0.30$, $SD = 0.15$, $t(17) = 8.86$, $P < 0.001$].

MCC Representations Are Distributed. One might wonder whether the RSA results reflect truly distributed representations of BOLD signal across the MFC cluster, or univariate signals reflecting the average activity of the cluster (38, 39). For example, whereas the model predicts similar activation patterns during state 1 across the four different sequences (indicated in blue in Fig. 1*B*), and different activation patterns between state 1 and state 6 for each pair of sequences (indicated in yellow in Fig. 1*B*), a univariate hypothesis might predict greater overall MCC activation during

state 6 compared with state 1. However, because correlations are insensitive to condition means, and because the searchlight RDMs were based on across-voxel correlations of the fMRI BOLD signal, the univariate hypothesis predicts 0 correlation between all pairs of conditions along the off-diagonal elements of the RDM. This is not what we observed (Fig. 4*B*).

For further validation, we assessed univariate MCC activity by averaging the GLM beta values across voxels within the MFC cluster separately for each condition, which yielded a single (i.e., scalar) value for each of the 24 conditions. We then computed the MFC-cluster RDM based on the Euclidean distances between each pairwise combination of these 24 condition means, reflecting absolute distances between univariate (scalar) rather than multivariate representations. Thus, this RDM described the degree to which each pair of conditions differentially activates the MFC. When the multivariate MFC-cluster RDM was regressed onto this univariate MFC-cluster RDM for each subject separately, the MFC-cluster RDM derived from the residuals—hereafter the “residual MFC-cluster RDM”—remained positively correlated with the model-RDM across subjects when tested against 0 [mean $\rho = 0.28$, $SD = 0.14$, $t(17) = 8.73$, $P < 0.001$]. This confirms that the model predictions are sensitive to information uniquely encoded by distributed representations in MCC, above and beyond the information encoded as univariate signals in MCC.

MCC Representations Do Not Depend on Time-On-Task and Related Theoretical Constructs. Theoretical accounts of MCC related to response conflict, cognitive effort, and time on task predict a positive association between MCC activity and response time (RT) (2–4, 40). To investigate whether the distributed representations in MCC are also sensitive to RT, we computed an RDM for each subject based on the absolute value of the difference between each pairwise combination of the average RTs for the 24 states (Fig. 5, response times), and then repeated the searchlight analysis. Fig. 4*C* indicates that the overlap between the MFC cluster and a midline cluster identified by this analysis is mainly confined to posterior medial cortex near the presupplementary




Fig. 3. Data collection and analysis pipeline. We applied representational similarity analysis (RSA) using a searchlight approach that computed the similarity between the representations predicted by a computational model and the representations encoded across local volumes of BOLD signal activation. See the text for details (*MCC and RNNs Encode Task Sequences Similarly*).

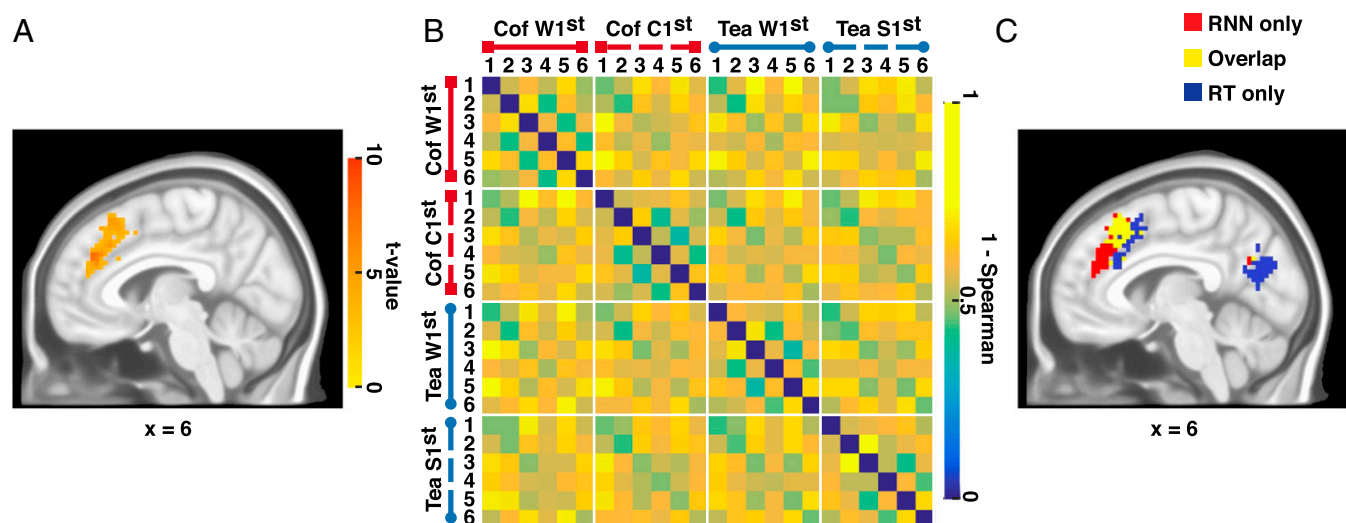


Fig. 4. Representational similarity analysis (RSA) results. (A) Medial-frontal cortex (MFC) cluster identified by the searchlight RSA based on the recurrent neural network model representational dissimilarity matrix (RDM). The color scale indicates the across-subject t test statistics of the within-subject normalized Spearman correlation values between the model RDM and local searchlight RDMs. Only the cluster that survived the permutation analysis is shown. (B) RDM derived from the MFC cluster. Labels are as given in Fig. 1B. Cof, coffee. (C) Overlap between the RSA results based on the model RDM and the individual subject RDMs derived from the average response times (RTs) for each condition. Red, regions associated purely with the model RDM; blue, regions associated purely with the individual subject RT RDMs; yellow, regions of overlap. Coordinates are indicated in MNI space and the underlay is International Consortium for Brain Mapping MNI 152 sixth generation.

motor area [center of mass of the RT cluster, i.e., the union of the yellow and blue clusters in MFC in Fig. 4C: 3, 22, 47, Montreal Neurological Institute (MNI) space] (36, 37). By contrast, the anterior MCC correlates with the model predictions irrespective of RT. As a check, when the MFC-cluster RDM was regressed onto the individual subject RT RDMs, the MFC-cluster RDM derived from the residuals remained positively correlated with the model RDM across subjects when tested against 0 [$M = 0.062$, $SD = 0.087$, $t(17) = 3.14$, $P < 0.005$]. This indicates that the observed association between model and MCC representations is not accounted for by RT differences across conditions.

MCC Representations Confirm Specific Predictions of the RNN Model.

We next asked whether the MCC representations reflect specific features of the model. Toward this end, multidimensional scaling (MDS) (41) of the context unit activations across the 24 states was conducted, revealing that the representations were characterized by three salient features (Fig. 1C) (SI Appendix). Each of these features was statistically robust across 100 runs of the model (SI Appendix). First, one dimension of the MDS (y axis of Fig. 1C) is sensitive to the progression of each sequence from start to end (state 1 to state 6). To look for similar representations in the MFC cluster, we created an RDM indicative of sequence position only (coding nearby states in each sequence as similar and faraway states as dissimilar; Fig. 5, sequence position); the correlation between this sequence position RDM and the MFC-cluster RDM for each subject was positive across subjects [mean $\rho = 0.15$, $SD = 0.16$, $t(17) = 4.76$, $P < 0.001$]. Moreover, when this sequence position RDM was correlated with the residual MFC-cluster RDM (which controlled for the average cluster activity, see above), the correlation remained positive across subjects [mean $\rho = 0.14$, $SD = 0.15$, $t(17) = 3.98$, $P < 0.001$]. This result indicates that MCC encodes sequence position as a distributed, across-voxel representation that evolves in time as the sequence progresses, comparable to how the model context layer encodes the sequence.

The second dimension of the MDS (x axis of Fig. 1C) revealed that the model strongly discriminates between the “utensil” states (states 3 and 5, which always required the same response, namely, selecting a spoon to “stir”), and all of the other states (states 1, 2,

4, and 6, which required varying responses), irrespective of the sequence. To look for this dimension in the MFC cluster, we created an RDM that distinguished between the utensil states and all other states (Fig. 5, utensil state). The mean of the (Spearman) within-subject correlation between the sequence identity RDM and the MFC-cluster RDM was significantly different from 0 [mean $\rho = 0.31$, $SD = 0.18$, $t(17) = 7.5$, $P < 0.001$]. Further, when this utensil RDM was correlated with the residual MFC-cluster RDM, the correlation remained positive across subjects [mean $\rho = 0.30$, $SD = 0.20$, $t(17) = 6.59$, $P < 0.001$]. These results indicate that, as predicted by the model, the MCC also strongly discriminates between the stir actions and the other actions in the sequence.

A third notable feature of the MDS is that, except for state 1 (which is identical across the four sequences) and state 2 (which discriminates between but not within the coffee and tea sequences), each state is represented differently by the network across the four sequences (Fig. 1C). These noncollinear trajectories reflect the task requirement to maintain memories of past events to disambiguate the different sequences (28). To investigate this property in the MFC cluster, we created an RDM that discriminates all of the actions within each sequence from all of the actions in the other sequences (Fig. 5, sequence identity). The (Spearman) correlation between this sequence identity RDM with the MFC-cluster RDM was positive [mean $\rho = 0.07$, $SD = 0.09$, $t(17) = 3.40$, $P < 0.005$], indicating that, as predicted by the model, MCC differentiates between the different action sequences. Further, when this sequence identity RDM was correlated with the residual MFC-cluster RDM, the correlation remained positive across subjects [mean $\rho = 0.07$, $SD = 0.08$, $t(17) = 4.1$, $P < 0.001$]. These results indicate that, as predicted by the model, MCC encodes contextual information that discriminates between the action sequences.

Finally, we performed a within-subject general linear model that simultaneously regressed all three predictors (the sequence identity RDM, the sequence position RDM, and the utensil RDM) onto the MFC-cluster RDM; each predictor remained statistically significant [sequence identity, $t(17) = 4.24$; sequence position, $t(17) = 4.10$; utensil, $t(17) = 7.26$; all $P < 0.001$]. The same was true when we regressed these predictors onto the residual MCC-cluster RDM [sequence identity, $t(17) = 6.51$, $P < 0.001$; sequence




Fig. 5. Representational dissimilarity matrices (RDMs) associated with follow-up analyses. Response times: RDM derived from the average response times of individual participants, averaged across subjects; lower values correspond to pairs of conditions with relatively similar average response times. Sequence position: sequence position RDM coding temporally distant states as dissimilar. Utensil state: utensil RDM discriminating utensil states from all other states. Sequence identity: sequence identity RDM discriminating within-sequence actions from across-sequence actions. The 24 states along each axis correspond to the six actions associated with each of the four task sequences [coffee (cof), water first; coffee, cream first; tea, water first; tea, sugar first; see Fig. 1B]. Blue indicates high similarity; yellow indicates low similarity. Units are arbitrary.

position, $t(17) = 3.33$, $P < 0.005$; utensil, $t(17) = 4.24$, $P < 0.001$]. Thus, like the model, MCC differentially encodes sequence identity and sequence position while also discriminating between stir versus other actions. Note that the model RDM and the MCC cluster RDM may also share other similarities that are not readily apparent with MDS. Additional exploratory analyses are reported in *SI Appendix*.

Discussion

We applied a computational model of MCC (21) to show the existence of distributed representations in human MCC that track the execution of action sequences. These patterns evolved in time, displaying unique sensitivity to each step of each sequence while discriminating between the different sequences. Moreover, the representations were encoded more robustly in MCC than in other brain areas, implicating MCC specifically in this process (21). These results suggest that the ubiquitous MCC activity observed across neuroimaging studies (8) reflects the pooled contribution of fine-grained patterns of activation that support task execution. Further, they suggest that the diverse array of neuronal responses associated with task-related events (5–7) reflects the contributions of individual MCC neurons to these global patterns.

As shown above, the correspondence with the model RDM could not result from mean (univariate) activation differences across conditions. This dissociation is true for any univariate signal, including those predicted by prominent theories of MCC function (2–4). Further, the model RDM correlated with the MCC patterns at both local-searchlight and global-cluster levels. Although overlap of the searchlight with the boundaries of

circumscribed areas of univariate, task-related activity could induce spurious correlations with the model RDM, such correlations are unlikely to occur at the cluster level. Conversely, the mean activation of the entire cluster could average across local differences in univariate activity. However, because such local responses would be anticorrelated, they would be less likely to correlate with the model RDM at the searchlight level.

Three observations suggest that this mechanism is hierarchically organized along the rostral–caudal extent of the frontal midline. First, although the task representations were distributed across a broad region of MFC, the presupplementary motor area was relatively more sensitive to differences in RT (Fig. 4C). This region modulates response speed (42) and contributes to sequence execution (43) and thus may be more concerned with the physical implementation of the sequences. Second, as we predicted (21), the peak t value was located in the anterior region of MCC (1) (*SI Appendix*, Fig. S1), which is consistent with an extensive literature implicating this area in monitoring and sustaining the execution of goal-directed action sequences (10, 11, 13–15, 22, 25, 44–48). Third, an exploratory analysis revealed that the frontal pole robustly encoded sequence identity (*SI Appendix*, Fig. S2 and Table S2). Rostral parts of the frontal midline are said to be involved in the hierarchical control over sequential behaviors (49, 50), as well as in representing hidden states (51) and in facilitating switches between task strategies (18, 24, 52), all of which feature in the coffee–tea task. Taken together, these results suggest that hierarchical action sequences are processed along a rostral–caudal cortical gradient (3, 24), from rostral prefrontal cortex, which selects the sequences, to anterior MCC, which monitors the progress of the sequences, to posterior MCC, which implements the sequences. These midline processes are likely supported by recurrent loops through the basal ganglia (53–56), which are uniquely positioned to create, monitor, and regulate the expression of hierarchically organized action sequences (57), and by interactions with the hippocampus (58), which may differentiate between ambiguous states for the purpose of future episodic retrieval (59–61).

The distributed mechanism described here provides a way to understand the role of MCC in sustaining extended behaviors. First, because RNNs are trained to predict each event in a sequence, the networks yield a large prediction error whenever a predicted event fails to match the actual event. We have shown previously that these prediction errors behave like the univariate conflict- and error-related signals commonly observed in MCC (21), which suggests that the size of these signals reflects the distance in task space between the predicted and obtained network states. Second, points of near convergence between the state–space trajectories of different action sequences can permit action slips across the sequences (28). Increasing the separation between the trajectories should therefore prevent against such errors. In this regard, it is notable that tonic dopamine levels in frontal cortex, which are associated with both the rate of reward delivery and the deployment of cognitive effort, have been proposed to regulate a gain parameter that can control the overall distances between sequence representations (24, 62). Consistent with this, amphetamine—a dopamine agonist—exhibits an inverted-U dose-dependent effect on neural state–space separation between task trajectories (12), and reward incentives improve the discriminability of brain states associated with task performance (63). We suggest that cognitive effort expenditure associated with dopaminergic modulation of MCC (24, 62) promotes separation between different task trajectories, thereby preventing against error commission when persisting at extended behaviors.

Methods

The experimental protocol was approved by the Ethical Committee of Ghent University Hospital and by the University of Victoria Human Research Ethics Board. Participants signed an informed consent and the experiment was

conducted in accordance with the 1964 Declaration of Helsinki. Eighteen volunteers participated in this experiment for pay. Structural and functional images were acquired at Ghent University Hospital using a Siemens 3T Magnetom Trio MRI scanner. We utilized an existing RNN model of MCC function (21) to simulate MCC representations during execution of the coffee-tea-making task. Model predictions were tested by applying RSA to the fMRI data. Protection against type-I errors was ensured with permutation testing. Please see *SI Appendix* for complete experimental details.

- Vogt BA (2016) Midcingulate cortex: Structure, connections, homologies, functions and diseases. *J Chem Neuroanat* 74:28–46.
- Ebitz RB, Hayden BY (2016) Dorsal anterior cingulate: A Rorschach test for cognitive neuroscience. *Nat Neurosci* 19:1278–1279.
- Silvetti M, Alexander W, Verguts T, Brown JW (2014) From conflict management to reward-based decision making: Actors and critics in primate medial frontal cortex. *Neurosci Biobehav Rev* 46:44–57.
- Vassena E, Holroyd CB, Alexander WH (2017) Computational models of anterior cingulate cortex: At the crossroads between prediction and effort. *Front Neurosci* 11:316.
- Enel P, Procyk E, Quilodran R, Dominey PF (2016) Reservoir computing properties of neural dynamics in prefrontal cortex. *PLoS Comput Biol* 12:e1004967.
- Kennerly SW, Dahmubed AF, Lara AH, Wallis JD (2009) Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci* 21:1162–1178.
- Wang C, Ulbert J, Schomer DL, Marinkovic K, Halgren E (2005) Responses of human anterior cingulate cortex microdomains to error detection, conflict monitoring, stimulus-response mapping, familiarity, and orienting. *J Neurosci* 25:604–613.
- Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD (2011) Large-scale automated synthesis of human functional neuroimaging data. *Nat Methods* 8:665–670.
- Cunningham JP, Yu BM (2014) Dimensionality reduction for large-scale neural recordings. *Nat Neurosci* 17:1500–1509.
- Balaguer-Ballester E, Lapiush CC, Seamans JK, Durstewitz D (2011) Attracting dynamics of frontal cortex ensembles during memory-guided decision-making. *PLoS Comput Biol* 7:e1002057.
- Blanchard TC, Strait CE, Hayden BY (2015) Ramping ensemble activity in dorsal anterior cingulate neurons during persistent commitment to a decision. *J Neurophysiol* 114:2439–2449.
- Lapiush CC, Balaguer-Ballester E, Seamans JK, Phillips AG, Durstewitz D (2015) Amphetamine exerts dose-dependent changes in prefrontal cortex attractor dynamics during working memory. *J Neurosci* 35:10172–10187.
- Lapiush CC, Durstewitz D, Chandler LJ, Seamans JK (2008) Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. *Proc Natl Acad Sci USA* 105:11963–11968.
- Ma L, Hyman JM, Phillips AG, Seamans JK (2014) Tracking progress toward a goal in corticostriatal ensembles. *J Neurosci* 34:2244–2253.
- Balaguer J, Spiers H, Hassabis D, Summerfield C (2016) Neural mechanisms of hierarchical planning in a virtual subway network. *Neuron* 90:893–903.
- Kahnt T, Grueschow M, Speck O, Haynes J-D (2011) Perceptual learning and decision-making in human medial frontal cortex. *Neuron* 70:549–559.
- Schapiro AC, Rogers TT, Cordova NI, Turk-Browne NB, Botvinick MM (2013) Neural representations of events arise from temporal community structure. *Nat Neurosci* 16:486–492.
- Schuck NW, et al. (2015) Medial prefrontal cortex predicts internally driven strategy shifts. *Neuron* 86:331–340.
- Rigotti M, et al. (2013) The importance of mixed selectivity in complex cognitive tasks. *Nature* 497:585–590.
- Yuste R (2015) From the neuron doctrine to neural networks. *Nat Rev Neurosci* 16:487–497.
- Shahbazian D, Holroyd CB (2018) Distributed representations of action sequences in anterior cingulate cortex: A recurrent neural network approach. *Psychon Bull Rev* 25:302–321.
- Cowen SL, Davis GA, Nitz DA (2012) Anterior cingulate neurons in the rat map anticipated effort and reward to their associated action sequences. *J Neurophysiol* 107:2393–2407.
- Holroyd CB, Yeung N (2012) Motivation of extended behaviors by anterior cingulate cortex. *Trends Cogn Sci* 16:122–128.
- Holroyd CB, McClure SM (2015) Hierarchical control over effortful behavior by rodent medial frontal cortex: A computational model. *Psychol Rev* 122:54–83.
- Hoshi E, Sawamura H, Tanji J (2005) Neurons in the rostral cingulate motor area monitor multiple phases of visuomotor behavior with modest parametric selectivity. *J Neurophysiol* 94:640–656.
- Kolling N, et al. (2016) Value, search, persistence and model updating in anterior cingulate cortex. *Nat Neurosci* 19:1280–1285.
- Elman JL (1990) Finding structure in time. *Cogn Sci* 14:179–211.
- Botvinick M, Plaut DC (2004) Doing without schema hierarchies: A recurrent connectionist approach to normal and impaired routine sequential action. *Psychol Rev* 111:395–429.
- Forstmann BU, Wagenmakers E-J (2015) *An Introduction to Model-Based Cognitive Neuroscience* (Springer, Heidelberg).
- Kriegeskorte N, Mur M, Bandettini P (2008) Representational similarity analysis—Connecting the branches of systems neuroscience. *Front Syst Neurosci* 2:4.
- Edelman S (1998) Representation is representation of similarities. *Behav Brain Sci* 21:449–467, discussion 467–498.
- Schwartz MF, et al. (1998) Naturalistic action impairment in closed head injury. *Neuropsychology* 12:13–28.
- Kriegeskorte N, Goebel R, Bandettini P (2006) Information-based functional brain mapping. *Proc Natl Acad Sci USA* 103:3863–3868.
- Groppe DM, Urbach TP, Kutas M (2011) Mass univariate analysis of event-related brain potentials/fields I: A critical tutorial review. *Psychophysiology* 48:1711–1725.
- Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Hum Brain Mapp* 15:1–25.
- Picard N, Strick PL (2001) Imaging the premotor areas. *Curr Opin Neurobiol* 11:663–672.
- Kim JH, et al. (2010) Defining functional SMA and pre-SMA subregions in human MFC using resting state fMRI: Functional connectivity-based parcellation method. *Neuroimage* 49:2375–2386.
- Coutanche MN (2013) Distinguishing multi-voxel patterns and mean activation: Why, how, and what does it tell us? *Cogn Affect Behav Neurosci* 13:667–673.
- Davis T, et al. (2014) What do differences between multi-voxel and univariate analysis mean? How subject-, voxel-, and trial-level variance impact fMRI analysis. *Neuroimage* 97:271–283.
- Grimband J, et al. (2011) The dorsal medial frontal cortex is sensitive to time on task, not response conflict or error likelihood. *Neuroimage* 57:303–311.
- Shepard RN (1980) Multidimensional scaling, tree-fitting, and clustering. *Science* 210:390–398.
- Bogacz R, Wagenmakers E-J, Forstmann BU, Nieuwenhuis S (2010) The neural basis of the speed-accuracy tradeoff. *Trends Neurosci* 33:10–16.
- Cona G, Semenza C (2017) Supplementary motor area as key structure for domain-general sequence processing: A unified account. *Neurosci Biobehav Rev* 72:28–42.
- Procyk E, Tanaka YL, Joseph JP (2000) Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nat Neurosci* 3:502–508.
- Koechlin E, Danek A, Burnod Y, Grafman J (2002) Medial prefrontal and subcortical mechanisms underlying the acquisition of motor and cognitive action sequences in humans. *Neuron* 35:371–381.
- Mulder AB, Nordquist RE, Örgüt O, Pennartz CM (2003) Learning-related changes in response patterns of prefrontal neurons during instrumental conditioning. *Behav Brain Res* 146:77–88.
- Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* 11:823–833.
- Hayden BY, Pearson JM, Platt ML (2011) Neuronal basis of sequential foraging decisions in a patchy environment. *Nat Neurosci* 14:933–939.
- Desrochers TM, Chatham CH, Badre D (2015) The necessity of rostral lateral prefrontal cortex for higher-level sequential behavior. *Neuron* 87:1357–1368.
- Mansouri FA, Koechlin E, Rosa MGP, Buckley MJ (2017) Managing competing goals—A key role for the frontopolar cortex. *Nat Rev Neurosci* 18:645–657.
- Schuck NW, Cai MB, Wilson RC, Niv Y (2016) Human orbitofrontal cortex represents a cognitive map of state space. *Neuron* 91:1402–1412.
- Powell NJ, Redish AD (2016) Representational changes of latent strategies in rat medial prefrontal cortex precede changes in behaviour. *Nat Commun* 7:12830.
- Nakahara H, Doya K, Hikosaka O (2001) Parallel cortico-basal ganglia mechanisms for acquisition and execution of visuomotor sequences—A computational approach. *J Cogn Neurosci* 13:626–647.
- Haruno M, Kawato M (2006) Heterarchical reinforcement-learning model for integration of multiple cortico-striatal loops: fMRI examination in stimulus-action-reward association learning. *Neural Netw* 19:1242–1254.
- Collins AG, Frank MJ (2013) Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychol Rev* 120:190–229.
- Keramati M, Gutkin B (2013) Imbalanced decision hierarchy in addicts emerging from drug-hijacked dopamine spiraling circuit. *PLoS One* 8:e61489.
- Schiffer AM, Waszak F, Yeung N (2015) The role of prediction and outcomes in adaptive cognitive control. *J Physiol Paris* 109:38–52.
- Remondes M, Wilson MA (2013) Cingulate-hippocampus coherence and trajectory coding in a sequential choice task. *Neuron* 80:1277–1289.
- Brown TI, Ross RS, Keller JB, Hasselmo ME, Stern CE (2010) Which way was I going? Contextual retrieval supports the disambiguation of well learned overlapping navigational routes. *J Neurosci* 30:7414–7422.
- Ginther MR, Walsh DF, Ramus SJ (2011) Hippocampal neurons encode different episodes in an overlapping sequence of odors task. *J Neurosci* 31:2706–2711.
- Hsieh LT, Gruber MJ, Jenkins LJ, Ranganath C (2014) Hippocampal activity patterns carry information about objects in temporal context. *Neuron* 81:1165–1178.
- Westbrook A, Braver TS (2016) Dopamine does double duty in motivating cognitive effort. *Neuron* 89:695–710.
- Etzel JA, Cole MW, Zacks JM, Kay KN, Braver TS (2016) Reward motivation enhances task coding in frontoparietal cortex. *Cereb Cortex* 26:1647–1659.