Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex

Christopher C. Lapish[†‡], Daniel Durstewitz[‡§¶], L. Judson Chandler[‖], and Jeremy K. Seamans[†¶]

[†]Department of Psychiatry and Brain Research Centre, University of British Columbia, 2211 Wesbrook Mall, Vancouver, BC, V6T 2B5; [§]Centre for Theoretical and Computational Neuroscience, University of Plymouth, Portland Square, Plymouth, PL4 8AA, United Kingdom; and [‖]Department of Neuroscience, Medical University of South Carolina, 67 President Street, Charleston, SC 29425

Edited by Charles F. Stevens, The Salk Institute for Biological Studies, La Jolla, CA, and approved June 22, 2008 (received for review April 29, 2008)

Successful decision making requires an ability to monitor contexts, actions, and outcomes. The anterior cingulate cortex (ACC) is thought to be critical for these functions, monitoring and guiding decisions especially in challenging situations involving conflict and errors. A number of different single-unit correlates have been observed in the ACC that reflect the diverse cognitive components involved. Yet how ACC neurons function as an integrated network is poorly understood. Here we show, using advanced population analysis of multiple single-unit recordings from the rat ACC during performance of an ecologically valid decision-making task, that ensembles of neurons move through different coherent and dissociable states as the cognitive requirements of the task change. This organization into distinct network patterns with respect to both firing-rate changes and correlations among units broke down during trials with numerous behavioral errors, especially at choice points of the task. These results point to an underlying functional organization into cell assemblies in the ACC that may monitor choices, outcomes, and task contexts, thus tracking the animal's progression through "task space."

cell assemblies | decision making | multiple single-unit recordings | neural coding | population analysis

There is general agreement that processing of information in the neocortex is done by networks of cells operating in a coordinated fashion rather than working independently, as in a purely feed-forward-type architecture. The most popular concept to describe these processes dates back to Hebb's (1) proposal on cell assemblies, which has been investigated from a variety of experimental perspectives (2–8) and has formed the foundation of a number of computational approaches (9–12). There is now accumulating evidence that the transient organization of neurons into dynamic ensembles and the sequential transitions among them may form the basis for cortical information processing (2, 3, 5, 7, 13–15).

Although such processes have been investigated in some depth for perceptual and spatial domains, much less is known regarding the network dynamics that govern higher-order cognitive processes, such as action and outcome monitoring. The anterior cingulate cortex (ACC) has been the subject of increased interest as a region that plays a role in action monitoring, a supervisory cognitive function that is especially important for optimal decision making in challenging and novel situations (16–19). A number of diverse single-unit correlates accompanying these processes have been observed in the ACC (20–27), yet it is unknown how single neurons organize into functional networks that could serve these functions.

In the present study, we investigated functional ensemble dynamics within the ACC during the performance of the delayed win-shift radial-arm maze, a task with distinct cognitive phases and high ecological validity (28). The locations of rewards changed between and within trials, such that the animal had to continually monitor its own actions and track the changing environmental contingencies as it progressed through the task. Using sophisticated statistical methods, we found that the recorded population appeared to track each aspect of the task by entering distinct and separable network states that broke down on trials with numerous errors.

Results

Population Activity Organizes into Specific Patterns During Distinct Task Epochs. Multiple single-unit recordings from the rat ACC [dorsal medial agranular prefrontal cortex; see supporting information (SI) Fig. S1*A*] were obtained from 10 rats over a total of 27 trials of the delayed-win-shift radial arm-maze task. The task consisted of the following behaviorally dissociable epochs illustrated in Fig. 1*A*: the periods surrounding correct arm choices during the training phase (TrC) and test phase (TsC), respectively; periods surrounding incorrect choices during the test phase (TsI; there were very few incorrect choices during the training phase); the period surrounding the point when the animal reached a food cup (TrR, during the training phase, and TsR, during the test phase); the entire delay phase (Dl); and all of the remaining periods intervening between arm choices and reward epochs during the training (Tr) and test phases (Ts).

Consistent with previous studies (24, 29–33), we observed a number of single-unit correlates that were associated with arm entries, reward processing, specific movements, and behavioral errors (Fig. 1*E*). However, the focus of the present study was the functional patterns of network activity associated with the distinct task epochs rather than single-unit correlates.

To analyze population activity, spike trains of all units were first convolved with Gaussian functions and binned at 200 ms, yielding smoothed instantaneous firing rates (iFR) for each unit as a function of time. For each time bin, the iFR of the N simultaneously recorded units were combined into population vectors embedded within an N-dimensional state space, here termed multiple-unit activity (MUA) space. For the purpose of visualization, we used multidimensional scaling to obtain 3D projections of these N-dimensional MUA spaces as shown in Fig. 1 *B* and *C*. In these graphs, each dot represents the entire state of the recorded network within one 200-ms bin of the task, and all population vectors (dots) belonging to different 200-ms bins

Author contributions: C.C.L., D.D., L.J.C., and J.K.S. designed research; C.C.L. performed research; D.D. designed and implemented statistical analysis; C.C.L., D.D., and J.K.S. analyzed data; and C.C.L., D.D., L.J.C., and J.K.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

[‡]C.C.L. and D.D. contributed equally to this work.

[¶]To whom correspondence may be addressed. E-mail: daniel.durstewitz@plymouth.ac.uk or seamans@interchange.ubc.ca.

This article contains supporting information online at www.pnas.org/cgi/content/full/0804045105/DCSupplemental.

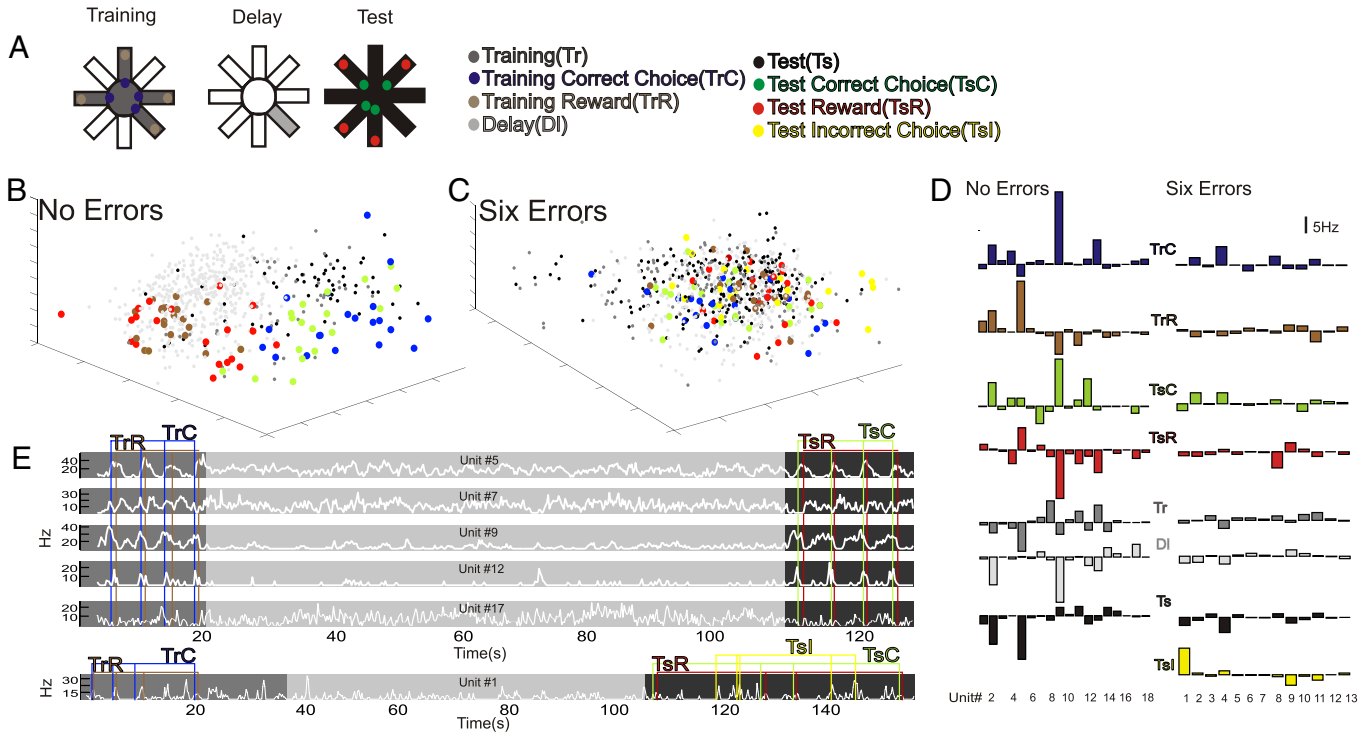

Fig. 1. Population analysis of ACC neuronal firing properties with respect to task epochs on a trial when an animal performed with no errors and a trial when the same animal made six errors. (*A*) Illustration of the delayed win-shift radial arm maze with the separate task epochs chosen. (*B*, *C*) MUA space representation. Each color-coded dot represents the state of the network (i.e., the iFR population vector of all recorded neurons) in a 200-ms time bin at some time point during the task. The blue and green dots are from 1-s periods centered around the points of arm entries (threshold crossings) for correct arm choices during the training phase (TrC) and test phase (TsC), respectively. The yellow dots are from 1-s time periods centered around the points of arm entries for incorrect choices during the test phase (TsI). The brown and red dots are from 1-s periods starting 200 ms before the points where the animal's nose reached a food cup during the training phase (TrR) and test phase (TsR), respectively. The smaller light gray, dark gray, and black dots represent delay (Dl), training (Tr), and test phase (Ts) epochs, respectively, excluding those periods defined by the choice and reward events. The axes of this 3D projection correspond to the simultaneous activation and mutual suppression of several units. (*D*) iFR prototypes showing for each task epoch the average deviation of the iFR of each neuron from the grand average across all task epochs from the same trials as in *B* and *C* with no errors (*D*, *Left*) and six errors (*D*, *Right*). (*E*) White lines represent the Gaussian smoothed iFRs of 5 units from the 0-error trial (*Upper*) and 1 unit from the six-error trial (*Lower*) as the trial progressed from the training phase (dark gray) through the delay (light gray) to the test phase (black). Units were selected based on the prototypes in *D*. Blue and green vertical lines represent the points of arm entries, whereas brown and red lines indicate points of reward during the training and test phase, respectively.

of the same task epoch are shown in the same color. As Fig. 1*B* illustrates (see also Movies S1 and S2 and Fig. S2), different task epochs segregate in MUA space, i.e., MUA vectors belonging to the same task epoch tend to cluster within similar regions of MUA space, whereas MUA vectors belonging to different task epochs populate different regions, implying that the population as a whole differentiates between individual task epochs. Averaging across all points in MUA space belonging to the same task epoch yields an *N*-dimensional "prototype vector" representing the mean iFRs of all units within this epoch, illustrating the differential patterns of activation as a function of task epoch (Fig. 1*D*).

This organization into distinct patterns appeared to be functionally relevant to the successful completion of a trial. Fig. 1*B* illustrates population activity for an animal that performed perfectly on a trial. However, when the population activity for the same animal was reexamined during a trial in which six errors were committed, the segregation in MUA space broke down (Fig. 1*C*; see also Movies S1 and S2 and Fig. S2). This collapse of segregation in MUA space was also reflected in the lack of differentiation between task-epoch iFR prototypes (Fig. 1*D*). On this trial, only one unit showed clear differential activity, and it appeared to be most tightly linked to the signaling of behavioral errors (Fig. 1*E*). Therefore, successful task performance was characterized by the neural population attaining distinct task-epoch-specific patterns that changed dynamically with the cognitive requirements of the task. A failure to exhibit such distinct population patterns was correlated with errors in choosing the correct arms.

To quantify these phenomena statistically, activity patterns were analyzed across all animals used in this study by computing a separation error for each trial and every pair of task epochs. This was done by fitting a hyperplane (linear discriminant function) that optimally separated the clouds of points associated with each pair of task epochs using discriminant analysis and by determining for all of these pairs the relative number of population vectors (i.e., dots in MUA space) that were assigned to the wrong task epoch according to this linear classifier. Because the separability of task-epoch points is also affected by other factors such as the dimensionality of the MUA space and the total number of points/epoch, conservative surrogate data, obeying the same temporal continuity constraints as the original data, were constructed and used to test the significance of separation between any pair of task epochs (see *Methods* for details). Fig. 2*A* shows the separation errors for each of the original pairwise comparisons and for the respective surrogates constructed for each pair of task epochs averaged across all trials and animals. Each of the pairwise comparisons reached significance except for the comparison between test-phase incorrect (TsI) and correct (TsC) choices (paired *t* tests using corrected α^*-levels according to the Holm–Bonferroni method, (34), $\alpha = 0.05$; see *SI Text* and Fig. S3 for additional separation measures and

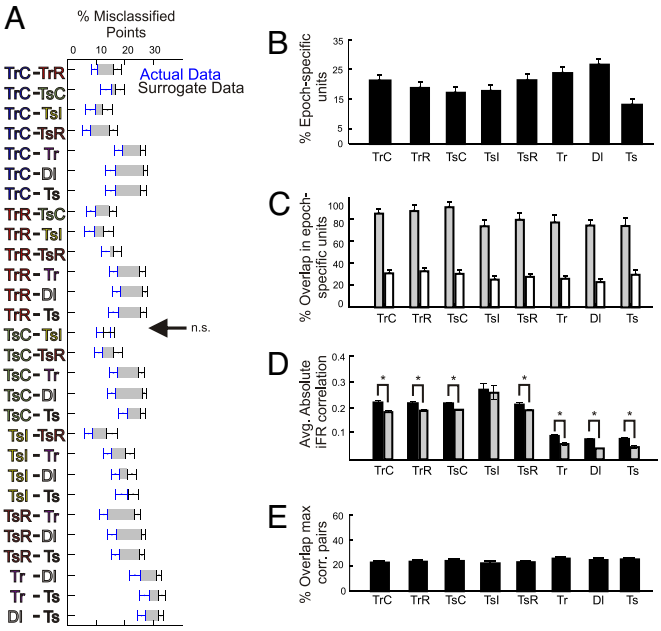

Fig. 2. Summary statistics revealing distinct network states in terms of both firing rate properties and correlations for each task epoch. (*A*) Separability between any two epochs in MUA space was quantified as the percentage of incorrectly classified points in the original data set (blue points and error bars) relative to those in the surrogates (black points and error bars) averaged across all datasets. Error bars = SEM. The gray shadings between the error bars surrounding the means of the original and surrogate data highlight the degree of difference between them. All pairs of task epochs were statistically separable except for the comparison TsC − TsI as indicated by the arrow. (*B*) The percentage of cells that have a significant iFR selectivity index for each task epoch. (*C*) For each task epoch, the percentage of the epoch-selective cells shown in *B* that exhibited a significant modulation of their firing rate in at least one other epoch (gray bars), and the percentage of selective cells shared on average with any one of the other epochs (white bars). The light-gray bars provide an indication of how many cells are involved in multiple task epochs, whereas conversely the white bars indicate how well the subpopulation of cells with significant selectivity can dissociate between any two task epochs. (*D*) For each task epoch, averaged (across all cell pairs and datasets) absolute iFR correlations ± SEM for the original data (black bars) and for surrogates composed of shuffled iFR bins (gray bars). For all task epochs, the differences between correlations among the original iFR time series and those within the shuffled surrogates were significant as denoted by an asterisk, except for the incorrect choice epoch. (*E*) Relative proportion of the 20% highest-correlated pairs a given task epoch shares on average with any other task epoch. The relatively low values indicate that the cell pairs exhibiting the strongest functional couplings change from one task epoch to the next.

analysis). Moreover, in >40% of all trials (averaged across all pairwise comparisons), a significant separation could be achieved on a single trial basis (with significant contributions coming from all animals and trials; Fig. S3*A*). Hence, the statistical analysis confirmed what visual inspection of MUA spaces in Fig. 1 suggested: in many more single trials than expected by chance (5% according to a $P < 0.05$ criterion), and overall on average, all pairs of task epochs significantly differed in terms of population behavior, except for one comparison involving incorrect choices (which is related to the general breakdown of MUA space separation especially at choice points in high-error trials as shown below).

Individual Neurons Are Generally not Selective for Single-Task Epochs. As described above, ensemble activity significantly differentiated between almost all pairs of task epochs. To address the question to what extent this differentiation may be rooted in highly task-epoch selective responses of single ACC neurons, a selectivity index for each unit and task epoch was computed that compared the average activity of a unit during a given epoch with the average activity across all other epochs (see *Methods* for details), and the significance of this index was again tested using surrogates. A significant selectivity index of a unit for a particular epoch implies that its average activity during that epoch significantly deviated (positively or negatively) from the average activity across all remaining epochs. For each of the eight task epochs, between 13% and 27% of all units were found to be task-epoch-selective according to this definition (Fig. 2*B*). However, cells may significantly modulate their average iFR compared with the grand mean during more than one task epoch (e.g., units 5 and 9 in Fig. 1*E*). Indeed, ≈80% of task-epoch-selective units exhibited a significant modulation in at least one other epoch (Fig. 2*C*, gray bars). Conversely, this implies that <4% of all units recorded were highly selective for just one epoch. On the other hand, if any two of the epochs were compared with each other, 72% of the epoch-selective units were active in only one of the two instances and could therefore account for their separation (Fig. 2*C*, white bars). Thus, although each unit was generally modulated during more than one epoch, all pairs of task epochs differed in their constitution of selective units.

Different Task Epochs Are also Associated with Unique Coalitions Among Neurons. Our results thus far show that different cognitive phases of the task are associated with specific patterns of iFR activity across all recorded units. To examine whether these iFR patterns are also accompanied by epoch-specific correlations among units, the absolute (zero time lag), iFR Pearson correlation coefficients averaged across all pairs of units as a function of task epoch were computed (using 200-ms bins as before; Fig. 2*D*, dark bars). For each task epoch, except for the incorrect choices, iFR comodulations were significantly above chance (*t* tests, all $P < 0.005$) based on comparisons with data sets where we shuffled the iFR bins (Fig. 2*D*, light gray bars). To assess whether these correlated activity changes were indeed task-epoch-specific or simply a general feature of the neuronal pairs recorded, we extracted the pairs with the 20% highest correlations from each epoch and determined the overlap between task epochs with regard to the most highly correlated pairs. A given task epoch shared only ≈20–30% on average of its most highly correlated pairs with any other task period (Fig. 2*E*; Fig. S4), suggesting that coalitions among neurons formed and dispersed with each task epoch. As a further confirmation of this observation, we constructed surrogates by recombining iFR bins from different task epochs while maintaining the temporal relation among units (i.e., no shuffling of iFR bins). In all cases except for the TsR epochs (*t* test, $P < 0.1$), these across-task-epoch iFR correlations were significantly lower than the within-task-epoch correlations (*t* tests, all $P < 0.005$), even after correction by the shuffle predictor (Fig. S5). Hence, different task epochs are not only differentiated by unique patterns of changes in firing rate but also through task epoch specific coalitions among units.

Ensemble Organization Is Diminished on Trials with Numerous Behavioral Errors. To address the functional importance of the distinct network patterns, changes in population activity as a function of behavioral errors were examined. For the entire dataset, trials were divided according to a median split based on the number of incorrect choices. The resulting groups agreed with previously defined criteria of asymptotic performance on this task (0–1 vs. 2 + incorrect choices; ref. 35). For the combined delay and test phases, separability was significantly worse for trials with many incorrect choices versus trials with 0–1 incorrect choices (Fig. 3*A*; *t* test, $P < 0.01$). Furthermore, Fig. 3*A* shows that the breakdown in MUA space affected mostly comparisons involving test-phase choice and reward epochs (TsC and TsR), and less

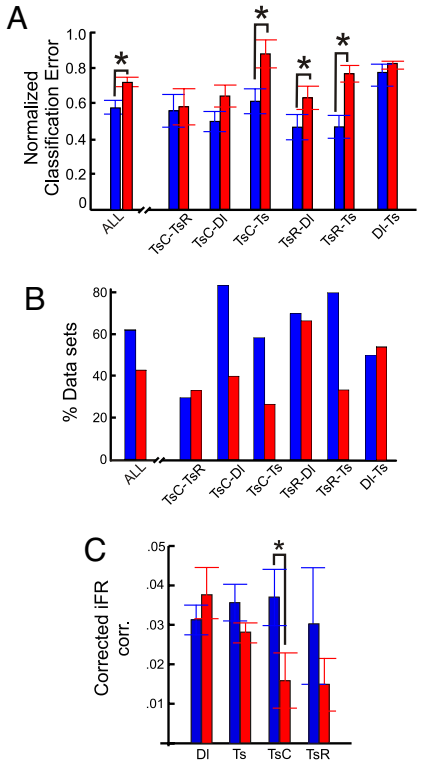

Fig. 3. Breakdown of separability in MUA space and decrease of correlations are associated with a rise in behavioral errors. (*A*) Separability of task epochs in MUA space as given by the normalized (to the surrogates) classification errors for the low (0–1; blue bars) and high (≥2; red bars) behavioral error trials. The graph shows overall separability averaged across all pairs of delay and test phase comparisons and separately for all epoch-pair comparisons from the delay + test phases. Asterisks denote a significant ($P < 0.05$) difference in separability for high and low behavioral error trials. (*B*) Percentage of individual datasets for which a significant separation of task epochs was reached for the low (blue bars) and high (red bars) behavioral error groups, averaged across all delay + test phase comparisons, and separately for all delay + test phase comparisons. (*C*) Absolute iFR correlations averaged across all cell pairs in the low (blue bars) vs. high (red bars) behavioral error groups as a function of task epoch, corrected by subtracting the average iFR correlations obtained for the respective iFR-shuffled surrogates.

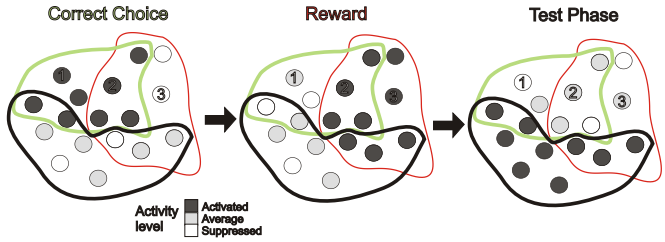

Fig. 4. Schema of a functional architecture that could explain the present results. Shown are three cell assemblies corresponding to correct choices (green), reward epochs (red), and basic test epochs (black). Neurons enclosed by the same colored outline are embedded within the same assembly, with neurons shared among assemblies being located within the intersections of the colored outlines. The firing rate of a neuron is indicated by its gray level. As the task progresses, first choice assemblies are activated (*Left*), followed by reward assemblies (*Center*), followed by a general test phase assembly (*Right*). Neurons 1 and 2 give an example of a pair that would be correlated during activation of the choice-assembly, as they form part of the same activated assembly but may be uncorrelated during the reward and test epochs. Neuron 3 is an example for a cell that is suppressed during correct choices due to mutual inhibition among assemblies (not illustrated for clarity), activated during the reward epoch, and fires at its average baseline activity in the test epoch. Neuron 2, in contrast, would be activated in both the choice and reward epochs. Recording from such a neuron would therefore allow one to differentiate between choice or reward epochs and the test phase, but not between choice and reward epochs. Hence, depending on a neuron's participation in various assemblies, numerous patterns of differential activity may be observed for single units, whereas at the same time, the cell assembly gives rise to a unique pattern at the network level.

so basal test vs. delay phase epochs (Ts vs. Dl). There was not only an overall significant decrease in separation but also a lower number of individual trials that yielded significant separation for comparisons involving TsC and TsR epochs (Fig. 3*B*). In addition, task-epoch specific iFR correlations among units also tended to decrease as a function of behavioral errors, but only with TsC did this reduction reach statistical significance (Fig. 3C; *t* test, $P < 0.05$). In conclusion, on trials where animals made two or more choice errors, network dynamics were severely compromised as evidenced by a failure to organize into distinct task-epoch-specific patterns both in terms of firing rate activity and correlations among units. This was especially true at choice points during the test phase, suggesting that errors are associated with a failure of the ACC to enter into unique and coordinated network states particularly at those times where decisions about arm entries are made.

Discussion

The present study explored population coding within the ACC while animals foraged for food in an ecologically valid radial arm maze task involving distinct cognitive phases. Many single-unit correlates (e.g., of choices or rewards) were observed. In most cases, however, single units did not limit their firing rate changes to one specific event type but were active across multiple cognitively defined task epochs. On the other hand, each task epoch was characterized by a unique pattern of firing-rate changes across units and correlations among units. The behavioral importance of this functional organization was supported by the fact that these population patterns tended to break down on trials with numerous errors.

Cell Assemblies and Dynamic Population Patterns. The cell assembly framework first introduced by Hebb (1) could provide a mechanistic basis for the observed ensemble patterns. A cell assembly is defined by a functional group of neurons entertaining relatively strong recurrent excitatory connections among each other, while potentially inhibiting pyramidal cells of other assemblies. ACC networks may consist of various partially overlapping cell assemblies encoding various cognitive events within the decision making process as illustrated in Fig. 4. As a result of such functional arrangements, at the single-neuron level, there would be numerous combinations of enhanced and depressed activity changes across different cognitive events, whereas each assembly as a whole would be associated with a unique pattern of firing rate changes and a unique pattern of within-assembly correlations (due to the recurrent excitatory connectivity within but not across assemblies), as observed here (Figs. 1 and 2).

Many other approaches, both similar to and different from ours, have been used to identify ensemble organization in multiple single-unit data. These include Hidden Markov Models, which also revealed distinct patterns and transitions among them *in vivo* (14, 36), and principal component or cluster analysis to unravel functional groupings (37), hierarchical organization of task-coding assemblies (7), or simply as a means for visualizing distinct population states and their connecting trajectories (15). At the level of precise spiking times, task-phase-specific alliances among neurons were demonstrated in primate prefrontal areas using spike train cross-correlograms (13) and repeating patterns of spikes aligned among multiple neurons (6, 38) were also taken as evidence for ensemble organization. Hence, a number of

studies employing different analysis methodologies provide converging evidence for the organization of neurons into functional ensembles at different temporal scales.

ACC and Functional Assemblies in Behavioral Monitoring. Transient lesions of the rat ACC produced significant impairments on the same task used here (35), causing numerous revisits to previously baited arms indicative of an inability to update responding as reward contingencies changed. Consistent with the putative function of the ACC in monitoring and updating response policies (39), single-unit activity in this area is most often related to specific movements, movement sequences, or response choices, as well as to rewards, specific action-reward pairs, or to the detection of response errors (21–27, 29–33). One important aspect of these single-unit correlates is their highly dynamical nature, with neural firing rates changing as a function of task context and reward magnitude even for the very same movements (24, 32, 40), or, vice versa, with the same neurons coding for very different movements in different tasks (29, 40). This high degree of flexibility and context dependence supports the view of the ACC as a device for monitoring or attending to actions (16, 17).

Within this framework, the functional organization into assemblies and the transitions among them may help to bind different cognitive attributes of a response strategy together, and to connect actions to subsequent outcomes (39). Performing a complex or novel task requires an animal to integrate various sensory, motor, reward, and memory aspects, and the transient formation of dynamic assemblies in the ACC as observed here may reflect this process. In this sense, the assemblies monitor the animal's progression in a task-dependent frame of reference or "task space." The sequential transitions among ACC assemblies and the transepoch activity of some of the units recorded may in addition help to connect the temporally separated components of action-outcome chains. Sequential transitions among active populations as well as persistent activity of single frontal neurons have been proposed previously as a means to link events through time (33, 41–43). Thus, we propose that the capacity of neural networks within the ACC to achieve and move between distinct dynamical patterns of activity may provide the neural bases for the putative role of the ACC in the selection and monitoring of action-outcome sequences (39).

Methods

Data Acquisition. Male Long-Evans rats (225–250 g, Harlan) were deeply anesthetized with 100 mg/kg ketamine and 10 mg/kg xylazine and placed into a sterotaxic device where multielectrode arrays were inserted through a cranial hole centered at A/P 2.2; M/L 0.8; D/V 2.5 relative to bregma and offset 10° from the vertical. After recovery, rats received 1 trial per day of the delayed win-shift radial arm maze described in detail elsewhere (35). During the training phase, four open arms were chosen randomly and baited, whereas the remaining four arms were blocked by a door. Upon visiting all four baited arms, the animal was locked into the last arm it visited, and the lights were turned out. After a 60- to 90-sec delay, the light was turned on, the door opened, and the animal began the testing phase. During the testing phase, all eight arms were open, and the rat had to visit the four arms that were blocked during the training phase. An error was scored as an entry into an arm that has been visited previously during either the training or test phase. All behavior was recorded with an online frame-capture COHU camera synchronized with timestamps created by the Neuralynx recording system (Custom Software). Off-line analysis of the video was used to collect the timestamps of behaviorally relevant events, from which event files were created with Event Session Splitter (Neuralynx). At the end of the experiments, each animal was deeply anesthetized with pentobarbital and transcardially perfused with 4% paraformaldehyde (P6148, Sigma). The brains were then collected, sectioned at 50 μm, and placements were observed with a dissecting microscope at 20× (Cambridge Instruments). See *SI Text* for in-depth description of arrays and unit recording and isolation parameters.

Data Analysis. All spike trains were first convolved with Gaussian functions to yield smoothed firing-rate functions and minimize the impact of random spike-time jitter at the borders between bins. Spike trains were then binned at 200 ms (approximately the inverse of the average firing rate of ≈4.8 Hz), and all simultaneously recorded neurons were combined into N-dimensional vectors **iFR**(t). To confirm the visually apparent separation among task epochs statistically, a linear classifier was constructed for each trial and pair of task epochs by determining an optimally separating $(N-1)$-dimensional hyperplane via discriminant analysis (e.g., ref. 44). The relative number of misclassified points was taken as an index of separability and was compared to surrogates constructed in the following way: Pairs of task epochs were first combined, and from these unions of points k contingent segments were randomly drawn, where k is the number of contingent segments for the original task epoch. To determine significance for each dataset individually, a nonparametric test was applied: significance at the 5% level was assigned by establishing whether the original classification error was among the 5% lowest within the set of 1 original and 99 surrogate classification errors.

To examine percentages of epoch-selective cells, for each cell n and task epoch p, a selectivity index was computed as

$$s_n(p) = \frac{|\langle iFR_n(t)\rangle_p - \langle iFR_n(t)\rangle_{q\neq p}|}{\langle iFR_n(t)\rangle_p + \langle iFR_n(t)\rangle_{q\neq p}},$$

where $\langle\cdot\rangle$ denotes the average across segments of epoch p, or across all of the combined other epochs q, respectively. Significance of each index $s_n(p)$ at the 5% level was again established by a nonparametric comparison with 99 surrogates, where in this case k contingent segments, with k being the number of segments making up epoch p, were drawn at random from the combination of all epochs in the task.

For the correlation analysis, absolute standard zero-lag Pearson correlation coefficients $|r_p(n,m)|$ were computed among scalar times series $iFR_{n,p}(t)$ and $iFR_{m,p}(t)$ for all pairs of simultaneously recorded neurons $n,m \in \{1..N\}$, $n \neq m$, and separately for all task epochs p. To test significance nonparametrically, 99 surrogates were constructed for each $r_p(n,m)$ by randomly shuffling the bins within series $iFR_{n,p}(t)$ and $iFR_{m,p}(t)$. For comparison of within-task-epoch to across-task-epoch correlations, equal time slices were first drawn from all task epochs to ensure that the surrogates drew from every task epoch with equal likelihood. Surrogates were then constructed by drawing for each comparison at random 20 bins from the union of time-equalized task epochs, original correlations were also recomputed for the time-equalized task epochs, and both original and across-epoch surrogate correlations were corrected by the mean from 99 shuffles.

For the comparison of the two behavioral error groups, all separation errors for each dataset were first normalized to the average separation error within the respective set of surrogates. However, the two behavioral error groups based on the median split did not differ with regard to either the average number of recorded neurons (ME ≈13.8 for the low and ME ≈14.2 for the high error group, two-sided t test: $P > 0.86$) or the standard deviation of this number (SD ≈5.2 for the low and SD ≈5.8 for the high error group), i.e., the dimensionality of the MUA spaces for those two groups was the same.

Further details of methods and methodological considerations can be found in *SI Text*.

ACKNOWLEDGMENTS. We thank Drs. A. G. Phillips, J. Cohen, J. Hyman, N. Vittoz, T. Wennekers, and D. Euston for helpful comments during the preparation of the manuscript. This work was funded by the National Institutes of Health (Grants MH065924, MH064569, and AA10983) and the Canadian Institutes of Health Research (Grant MOP-84319).

1. Hebb DO (1949) *The Organization of Behaviour* (Wiley, New York).
2. Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. *Annu Rev Neurosci* 18:555–586, review.
3. Sakurai Y (1999) How do cell assemblies encode information in the brain? *Neurosci Biobehav Rev* 23:785–796.
4. Stopfer M, Laurent G (1999) Short-term memory in olfactory network dynamics. *Nature* 402:664–668.
5. Pasupathy A, Connor CE (2002) Population coding of shape in area V4. *Nat Neurosci* 5:1332–1338.
6. Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsaki G (2003) Organization of cell assemblies in the hippocampus. *Nature* 424:552–556.
7. Lin L, Osan R, Tsien JZ (2006) Organizing principles of real-time memory encoding: Neural clique assemblies and universal neural codes. *Trends Neurosci* 29:48–57.
8. Reddy L, Kanwisher N (2006) Coding of visual objects in the ventral stream. *Curr Opin Neurobiol* 16:408–414.
9. Amit DJ, Brunel N, Tsodyks MV (1994) Correlations of cortical Hebbian reverberations: theory versus experiment. *J Neurosci* 14:6435–6445.

10. O'Reilly RC (2006) Biologically based computational models of high-level cognition. *Science* 314:91–94.
11. Durstewitz D, Seamans JK, Sejnowski TJ (2000). Neurocomputational models of working memory. *Nat Neurosci* 3 Suppl:1184–1191.
12. Grinvald A, Arieli A, Tsodyks M, Kenet T (2003) Neuronal assemblies: single cortical neurons are obedient members of a huge orchestra. *Biopolymers* 68: 422–436.
13. Sakurai Y, Takahashi S (2006) Dynamic synchrony of firing in the monkey prefrontal cortex during working-memory tasks. *J Neurosci* 26:10141–10153.
14. Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E (1996) Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J Neurosci* 16:752–768.
15. Mazor O, Laurent G (2005) Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* 48:661–673.
16. Posner MI, Petersen SE (1990) The attention system of the human brain. *Annu Rev Neurosci* 13:25–42.
17. Passingham RE (1996) Attention to action. *Philos Trans R Soc London Ser B* 351:1473–1479.
18. van Veen V, Carter CS (2002) The timing of action-monitoring processes in the anterior cingulate cortex. *J Cognit Neurosci* 14:593–602.
19. Botvinick MM, Braver TS, Barch DM, Carter CS, Cohen JD (2001) Conflict monitoring and cognitive control. *Psychol Rev* 108:624–652.
20. 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.
21. Niki H, Watanabe M (1979) Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Res* 171:213–224.
22. Picard N, Strick PL (1996) Motor areas of the medial wall: a review of their location and functional activation. *Cereb Cortex* 6:342–353.
23. Shima K, Tanji J (1998) Role for cingulate motor area cells in voluntary movement selection based on reward. *Science* 282:1335–1338.
24. Pratt WE, Mizumori SJ (2001) Neurons in rat medial prefrontal cortex show anticipatory rate changes to predictable differential rewards in a spatial memory task. *Behav Brain Res* 123:165–183.
25. Matsumoto K, Suzuki W, Tanaka K (2003) Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science* 301:229–232.
26. Isomura Y, Ito Y, Akazawa T, Nambu A, Takada M (2003) Neural coding of ''attention for action'' and ''response selection'' in primate anterior cingulate cortex. *J Neurosci* 23:8002–8012.
27. Ito S, Stuphorn V, Brown JW, Schall JD (2003) Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science* 302:120–122.
28. Timberlake W (2002) Niche-related learning in laboratory paradigms: the case of maze behavior in Norway rats. *Behav Brain Res* 134:355–374.
29. Jung MW, Qin Y, McNaughton BL, Barnes CA (1998) Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cereb Cortex* 8:437–450.
30. Cowen SL, McNaughton BL (2007) Selective delay activity in the medial prefrontal cortex of the rat: contribution of sensorimotor information and contingency. *J Neurophysiol* 98:303–316.
31. Euston DR, McNaughton BL (2006) Apparent encoding of sequential context in rat medial prefrontal cortex is accounted for by behavioral variability. *J Neurosci* 26:13143–13155.
32. Kargo WJ, Szatmary B, Nitz DA (2007) Adaptation of prefrontal cortical firing patterns and their fidelity to changes in action-reward contingencies. *J Neurosci* 27:3548–3559.
33. Narayanan NS, Laubach M (2006) Top-down control of motor cortex ensembles by dorsomedial prefrontal cortex. *Neuron*. 52:921–931.
34. Holm S (1979) A simple sequentially rejective multiple test procedure. *Scan J Stat* 6:65–70.
35. Seamans JK, Floresco SB, Phillips AG (1995) Functional differences between the prelimbic and anterior cingulate regions of the rat prefrontal cortex. *Behav Neurosci* 109:1063–1073.
36. Jones LM, Fontanini A, Sadacca BF, Miller P, Katz DB (2007) Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci USA* 104:18772–18777.
37. Chapin JK, Nicolelis MA (1999) Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations. *J Neurosci Methods* 94:121–140.
38. Riehle A, Grun S, Diesmann M, Aertsen A (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278:1950–1953.
39. Rushworth MF, Walton ME, Kennerley SW, Bannerman DM (2004) Action sets and decisions in the medial frontal cortex. *Trends Cognit Sci* 8:410–417.
40. Isomura Y, Takada M (2004) Neural mechanisms of versatile functions in primate anterior cingulate cortex. *Rev Neurosci* 15:279–291.
41. Baeg EH, Kim YB, Huh K, Mook-Jung I, Kim HT, Jung MW (2003) Dynamics of population code for working memory in the prefrontal cortex. *Neuron* 40:177–188.
42. Chang JY, Chen L, Luo F, Shi LH, Woodward DJ (2002) Neuronal responses in the frontal cortico-basal ganglia system during delayed matching-to-sample task: Ensemble recording in freely moving rats. *Exp Brain Res* 142:67–80.
43. Fuster, J. M (1997) Network memory. *Trends Neurosci* 20:451–459.
44. Krzanowski WJ (2000) *Principles of Multivariate Analysis* (Oxford Univ Press, Oxford, UK), 2nd Ed.