Towards a Cognitive Neuroscience of Non-Human Primates

Cognitive Neuroscience in Non-human Primates

Visuo-spatial memory; categorical judgments; strategic decision making, others including neuro-economic; list learning, including SCP; transitive inference; positive transfer; temporal discrimination; reward-preference; hierarchical behavior; imitation, emulation and social learning generally; aspects of numerosity; statistical learning; various operant and saccade-based paradigms

Functional Organization

- Primary rostro-caudal axis
  - 'temporal abstraction' (Koechlin)
  - (though see Badre & D'Esposito, 2009; Sakai & Passingham, 2006
- Lateral<-->Medial coupling
  - Kouneiher et al., 2009
  - Extent to which decisions must recruit temporally-extended behaviors, or consider temporally-removed goals / outcomes
- Laterally, Dorsal<-->Ventral coupling
  - How / What division (O'Reilly, 2010)
  - And see OFC (what?) and ACC (how?) medially
- Graziano, Nachev and others on PM, M1, SMC functional organization
  - Self organizing maps
- And see comparative anatomic work

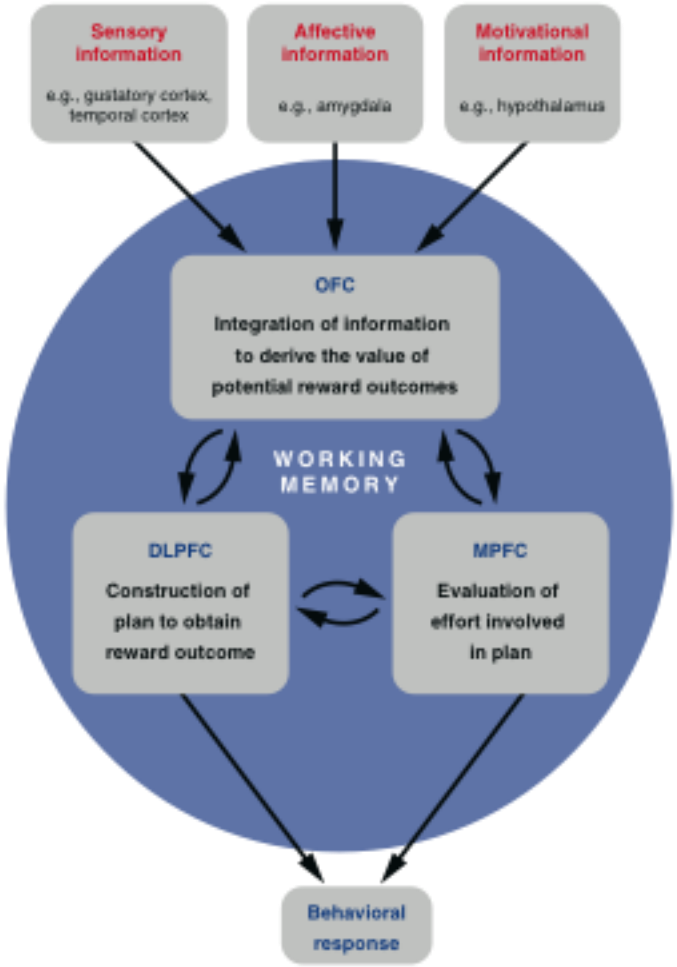

Figure 4

Model of the neuronal mechanisms underlying decision-making in PFC.

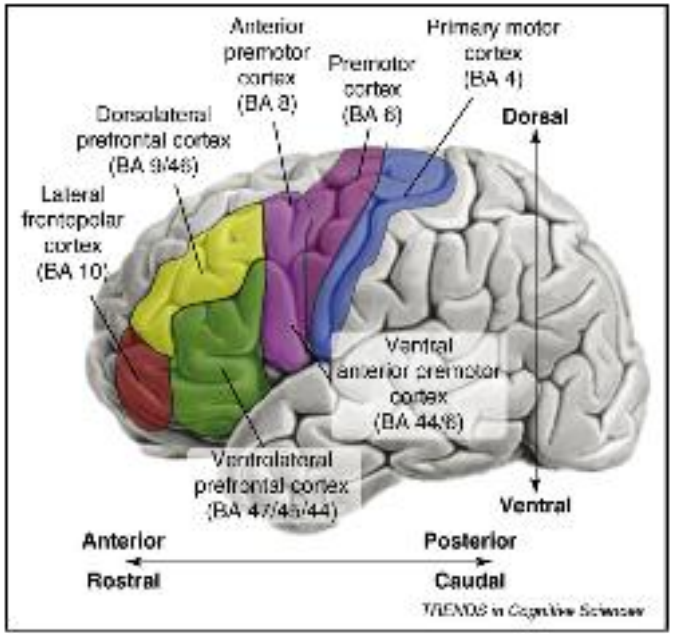

Figure 1. Schematic of major anatomical sub-divisions in the frontal lobes. Boundaries and Brodmann areas (BA) are only approximate. Arrows indicate anatomical directions of anterior/rostral (front) versus posterior/caudal (back) and dorsal (up) versus ventral (down). From caudal to rostral, labeled areas include motor cortex, dorsal (PMd) and ventral premotor cortex, dorsal (pre-PMd) and ventral aspects of anterior premotor cortex, ventro- (VLPFC) and dorsolateral PFC (DLPFC), and lateral frontal polar cortex (FPC).

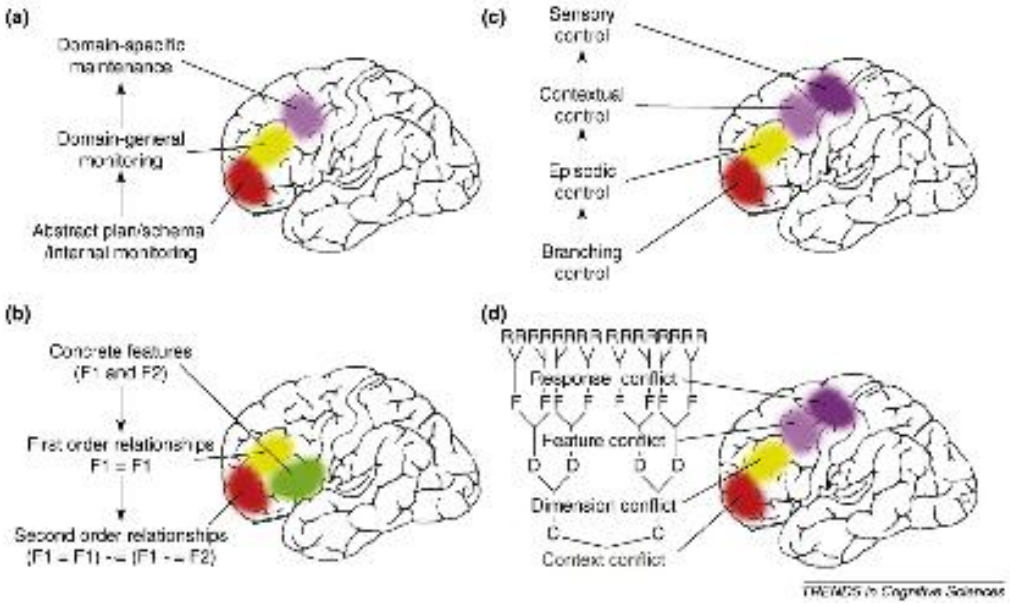
(a)
Domain-specific maintenance
Domain-general monitoring
Abstract plan/schema /internal monitoring
(b)
Concrete features (F1 and F2)
First order relationships F1 = F1
Second order relationships (F1 = F1) = (F1 = F2)
(c)
Sensory control
Contextual control
Episodic control
Branching control
(d)
Response conflict
Feature conflict
Dimension conflict
Context conflict
TRENDS in Cognitive Sciences

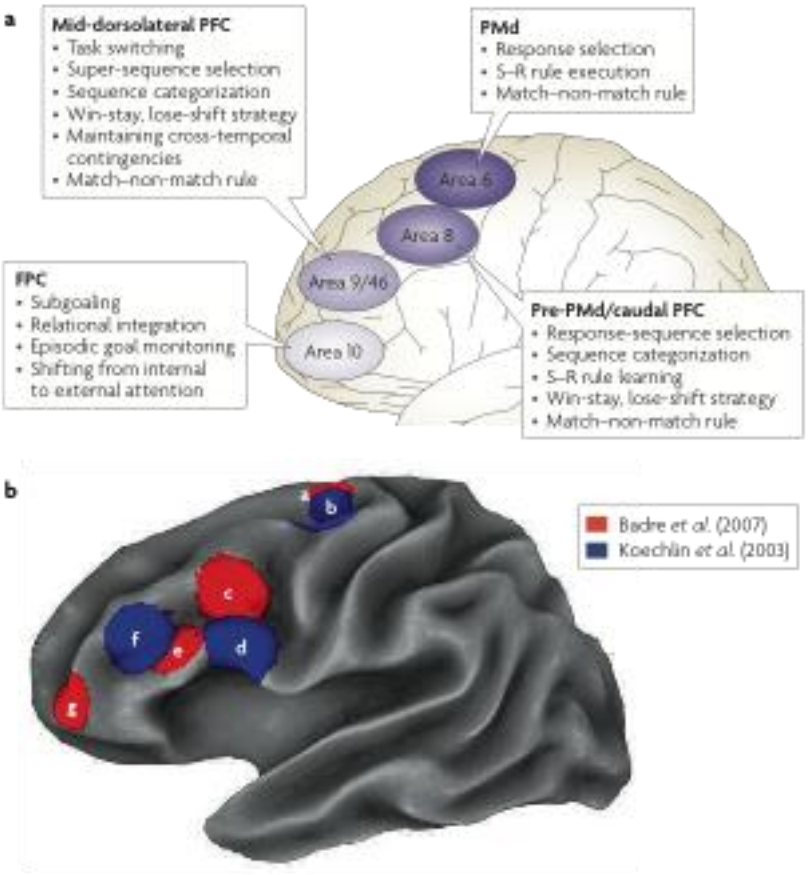
a
Mid-dorsolateral PFC
• Task switching
• Super-sequence selection
• Sequence categorization
• Win-stay, lose-shift strategy
• Maintaining cross-temporal contingencies
• Match–non-match rule
PMd
• Response selection
• S–R rule execution
• Match–non-match rule
Area 6
Area 8
Area 9/46
Area 10
FPC
• Subgoaling
• Relational integration
• Episodic goal monitoring
• Shifting from internal to external attention
Pre-PMd/caudal PFC
• Response-sequence selection
• Sequence categorization
• S–R rule learning
• Win-stay, lose-shift strategy
• Match–non-match rule
b
a
b
c
d
e
f
g
Badre et al. (2007)
Koechlin et al. (2003)

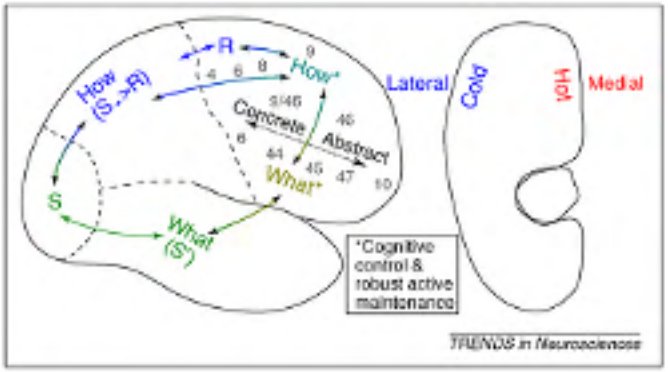

Figure 1. The complete set of broad functional organizations discussed, along each of the three major axes: dorsal [*How* = perception for action = Stimulus (S) to Response (R) mappings] versus ventral [*What* = Stimulus-driven semantic representations (S')]; within PFC rostral (abstract) versus caudal (concrete); medial (*Hot* value representations) versus lateral (*Cold* cognitive calculations). *How** (DLPFC) indicates a control system for the *How* posterior pathway, as *What** (VLPFC) does for the *What* pathway. Grey numbers indicate Brodmann areas on the lateral surface.

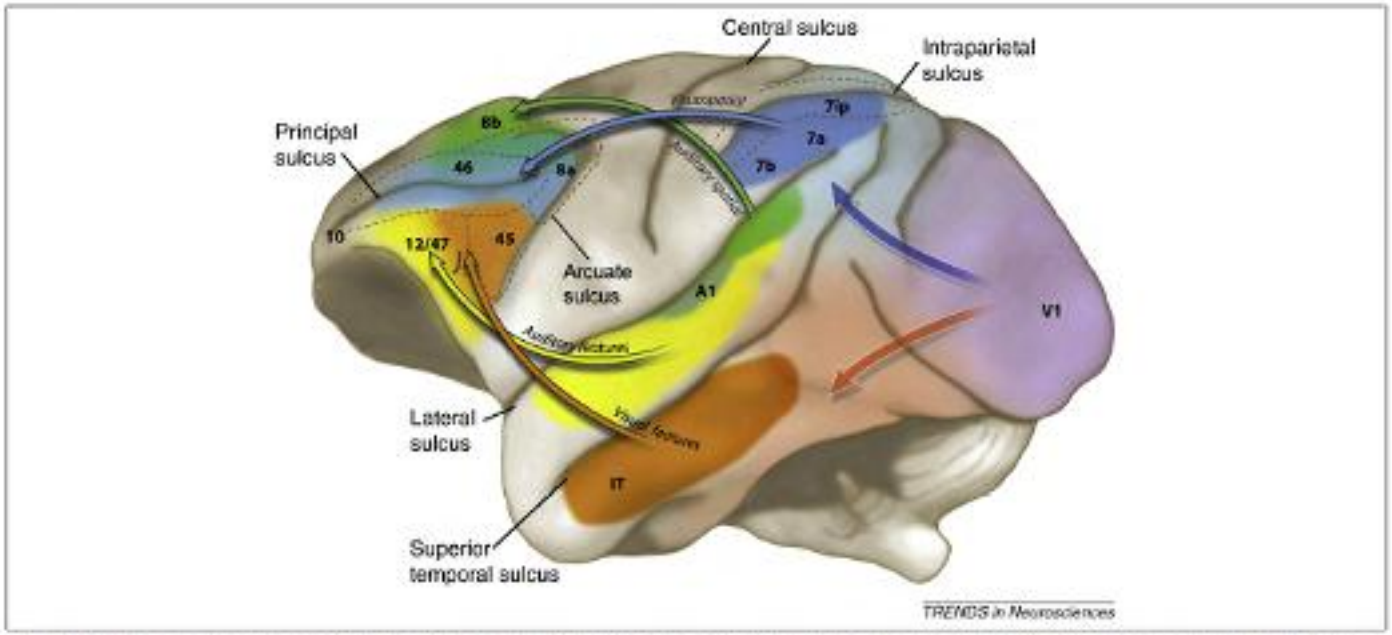

Figure 2. Summary of anatomical connectivity [in the macaque, where most of the anatomical work has been done] suggesting that dorsal versus ventral distinctions in posterior cortex should influence prefrontal cortex, due to dominant dorsal–dorsal and ventral–ventral connectivity. Reproduced, with permission, from Ref. [9].

Affordance Competition Hypothesis

- dPM (Cisek et al.) and PRR (Klaus et al.) code reach directions in parallel
  - Wait for cue to make decision
- M1 (Lu & Ashe), LPFC (Averbeck et a.), LIP (Cui & Andersen) also ‘load’ plans in parallel
  - Retrieve from memory a sequence of actions (Shima et al.)
- How does functional connectivity change as learning develops?
  - See Cromer et al., Pasupathy & Miller…
- Relate to Hikosaka, Graybiel, Schneider and others on Automatic Processing and ‘chunking’
  - Consolidation / optimization processes

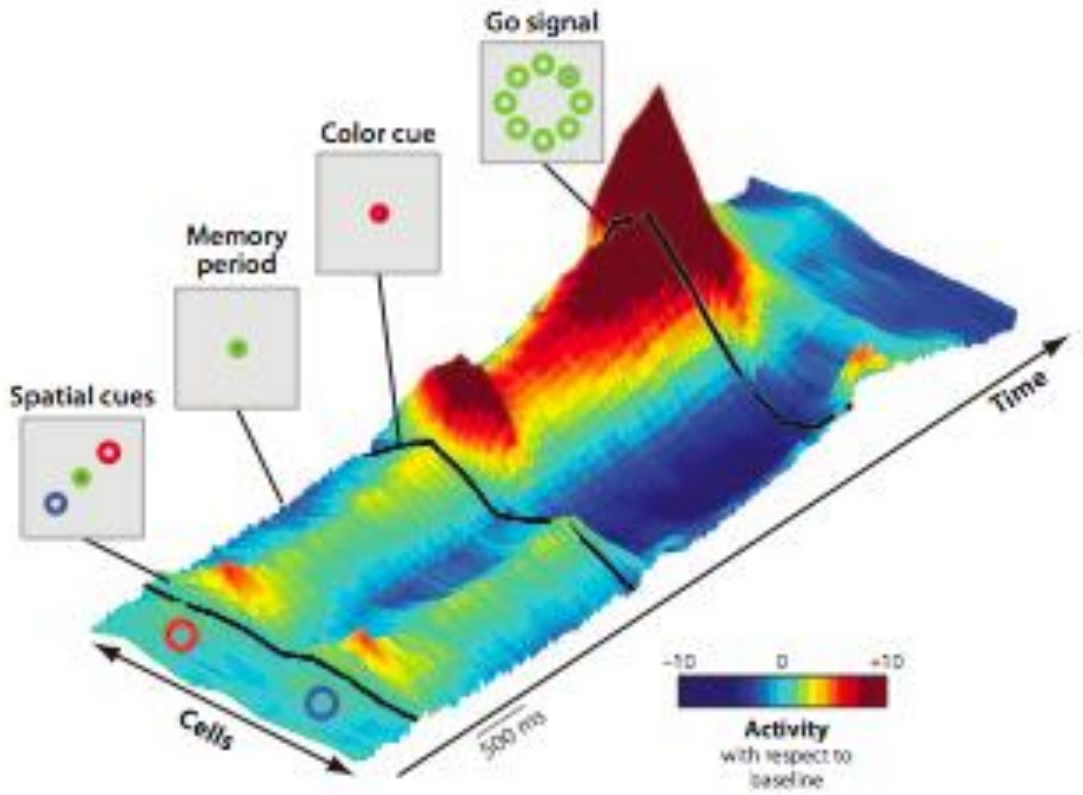

Figure 2

Population activity in the dorsal premotor cortex during a reach-selection task. The 3D colored surface depicts neural activity with respect to baseline, with cells sorted by their preferred direction along the bottom edge. Diagrams on the left show the stimuli presented to the monkey at different points during the trial (cross indicates the cursor). Note that during the period of ambiguity, even after stimuli vanished, the population encodes two potential directions. Data from Cisek & Kalaska (2005).

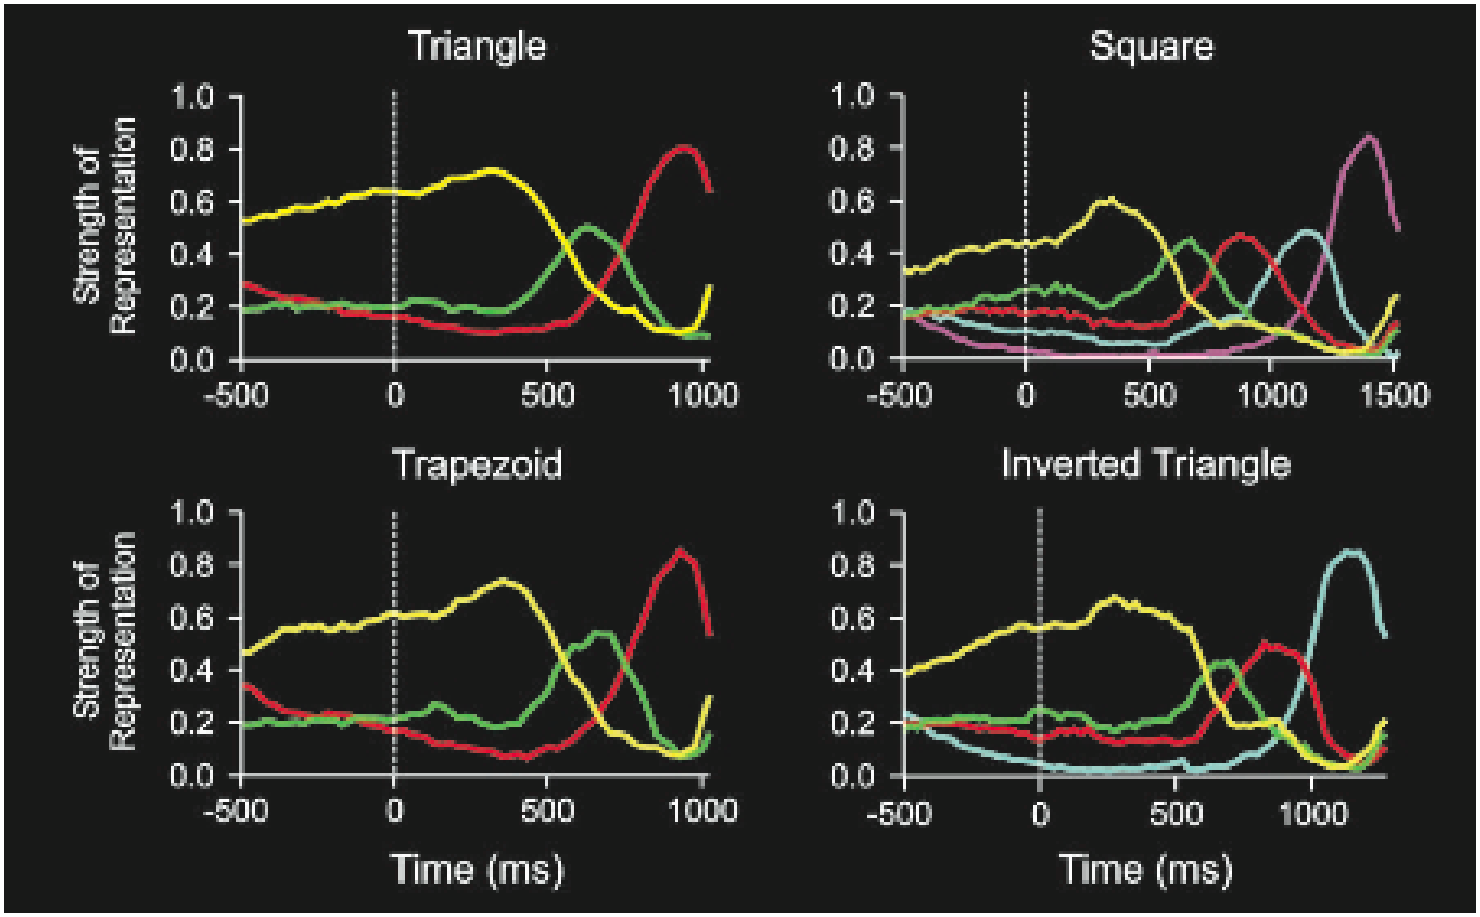

Fig. 2. Plots for all four shapes of strength of representation vs. time. Each plot shows the strength of the representation of each segment for each time bin of the task. Time 0 indicates the onset of the template. Time bins during hold period and RT are 25 ms. Length of segments were normalized to permit averaging across trials. Plots show parallel representation of segments before initiation of copying. Further, rank order of strength of representation before coping corresponds to the serial position of the segment in the series. The rank order evolves during the drawing to maintain the serial position code. Line color corresponds to segments as follows: yellow, segment 1; green, segment 2; red, segment 3; cyan, segment 4; magenta, segment 5. Not all lines are defined for all shapes.

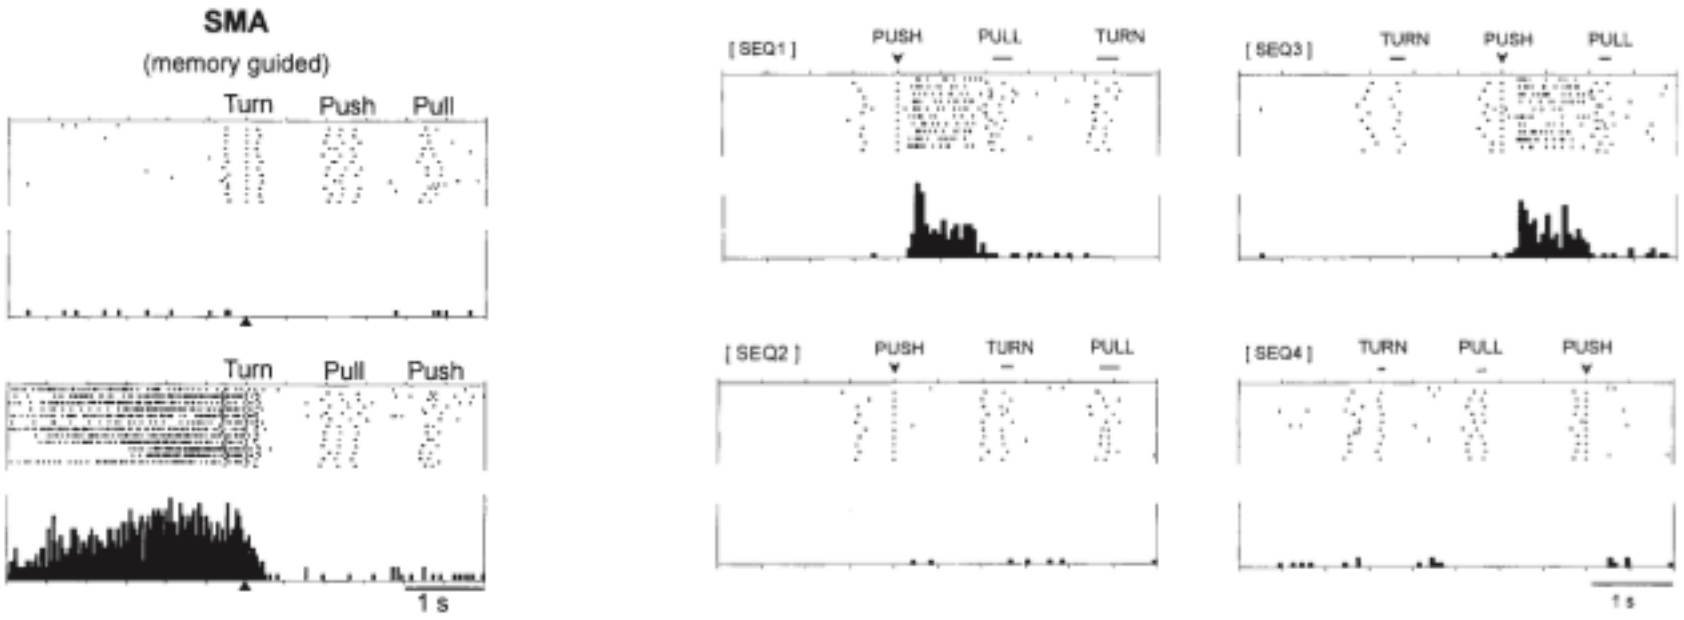
SMA
(memory guided)
Turn
Push
Pull
Turn
Pull
Push
1 s
[SEQ1]
PUSH
PULL
TURN
[SEQ3]
TURN
PUSH
PULL
[SEQ2]
PUSH
TURN
PULL
[SEQ4]
TURN
PULL
PUSH
1 s

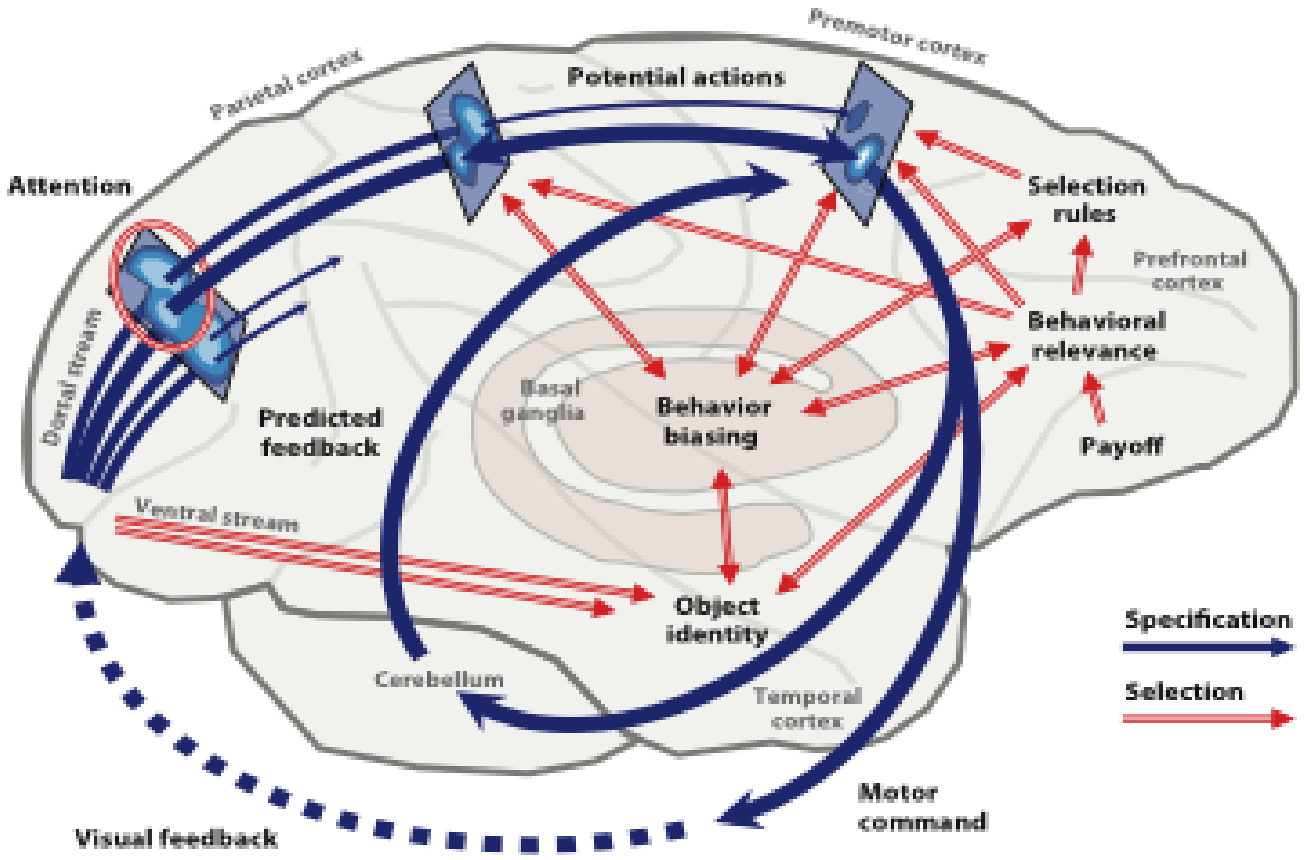
Premotor cortex
Parietal cortex
Potential actions
Attention
Selection rules
Prefrontal cortex
Behavioral relevance
Dorsal stream
Predicted feedback
Basal ganglia
Behavior biasing
Payoff
Ventral stream
Object identity
Specification
Selection
Cerebellum
Temporal cortex
Motor command
Visual feedback

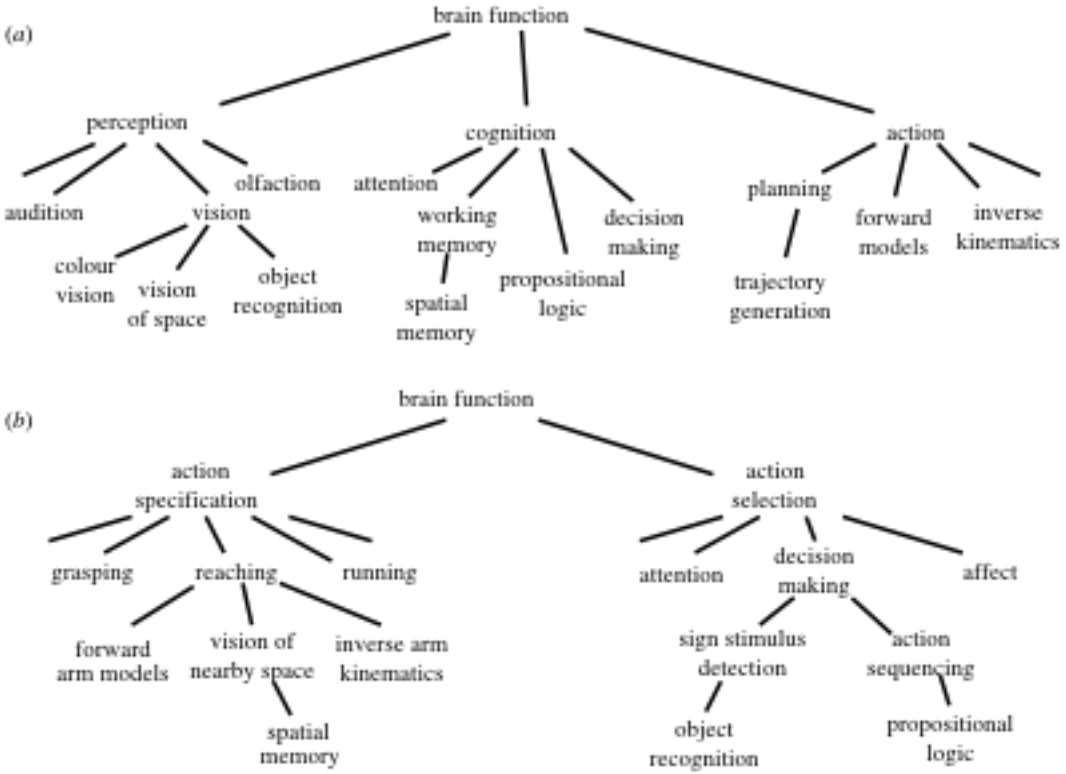
(a)
brain function
perception
cognition
action
audition
vision
olfaction
attention
working memory
decision making
planning
forward models
inverse kinematics
colour vision
vision of space
object recognition
spatial memory
propositional logic
trajectory generation
(b)
brain function
action specification
action selection
grasping
reaching
running
attention
decision making
affect
forward arm models
vision of nearby space
inverse arm kinematics
sign stimulus detection
action sequencing
spatial memory
object recognition
propositional logic

Encoding Strategies

- Reactive, event encoding responses versus prospective, state encoding / interval encoding responses
  - Phasic, tonic and anticipatory
- Campos et al. showed distributions of functional types in LIP and SEF
- Mita et al. on ‘relative interval’ encoding
- Saga et al on ‘multi-dimensional’ responses
- Genovesio et al. on ‘collapsing’ responses
- Numerosity encoding throughout Fronto-parietal network
- Note BOPs

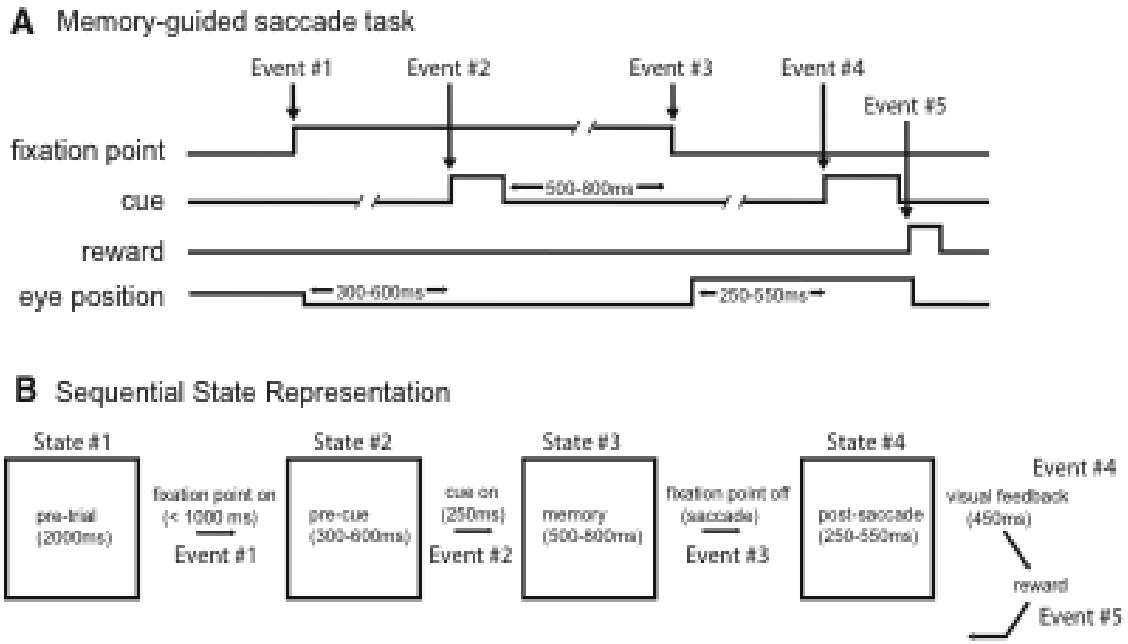

FIG. 1. Sequence of events and states in the oculomotor task. *A*: timeline of the task highlighting the 5 events that governed the task progression. *B*: representation of the task as a sequence of states with external events serving as transitions between states (see METHODS).

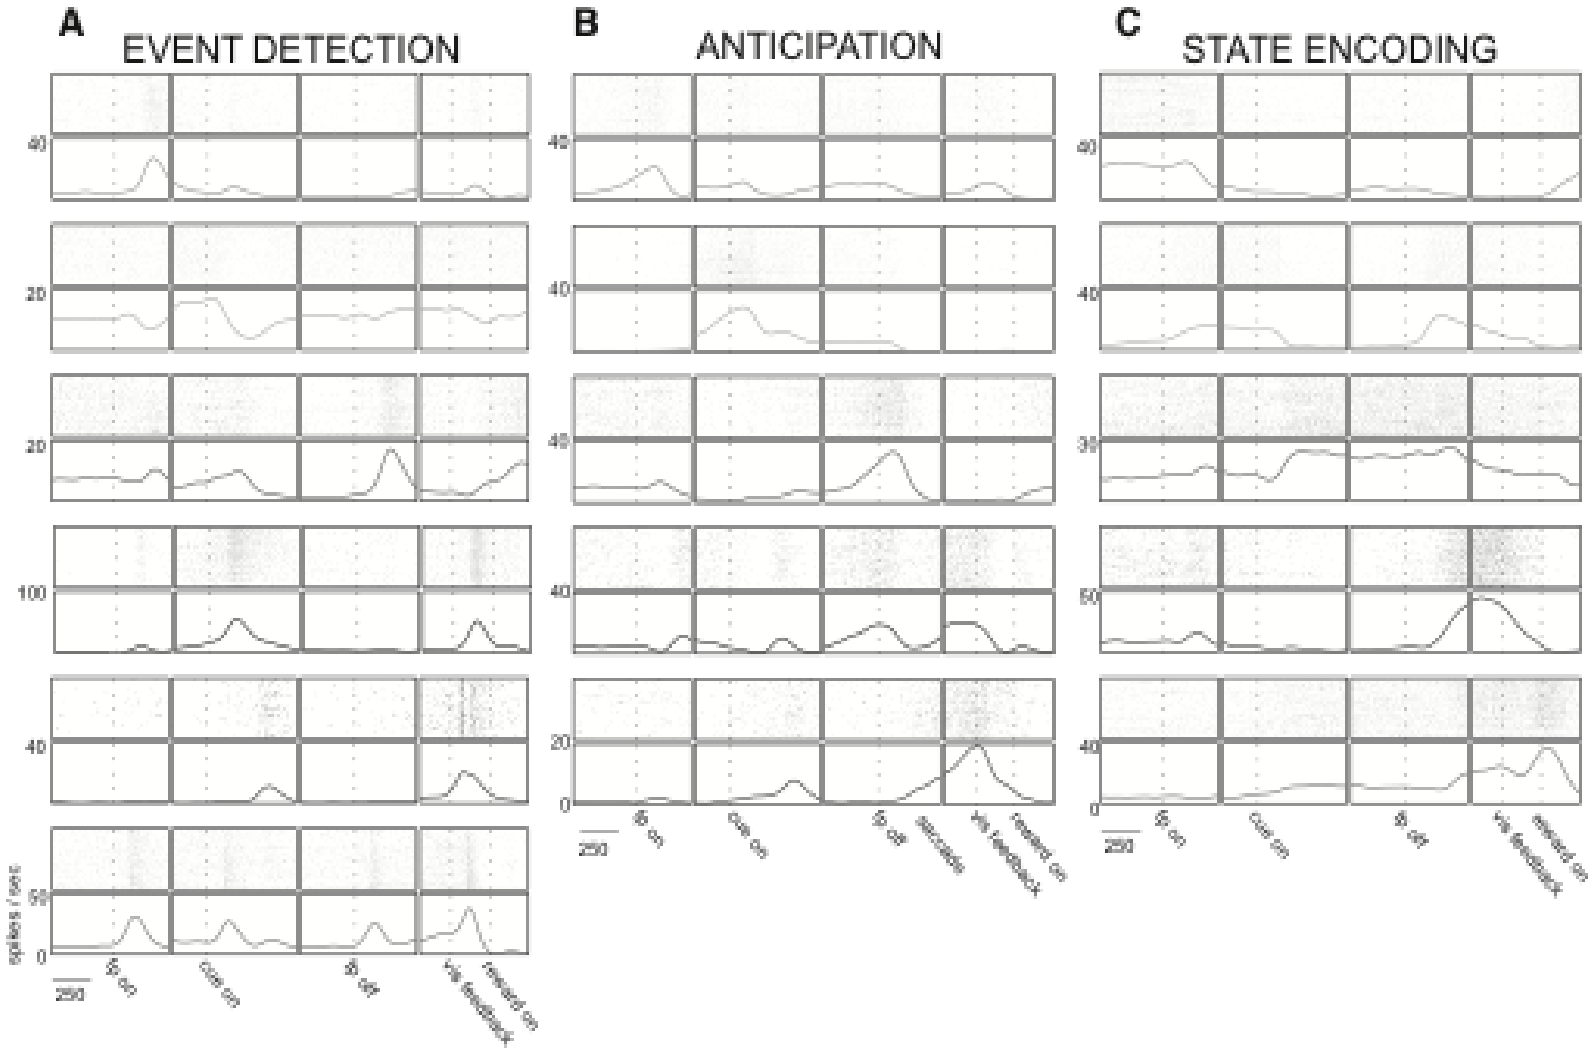

FIG. 3. Examples of activity types that support a sequential state representation. Each row consists of spike train rasters and smoothed (Gaussian kernel, SD = 50 ms) average firing rates for 1 example neuron in each of the 3 categories. All neurons are from the supplementary eye field (SEF). Numbers on the ordinate are the scale for the firing rate for each neuron. The columns include neurons with conspicuous event detection (*A*), anticipation (*B*), and state encoding (*C*) activities. See RESULTS for further description.

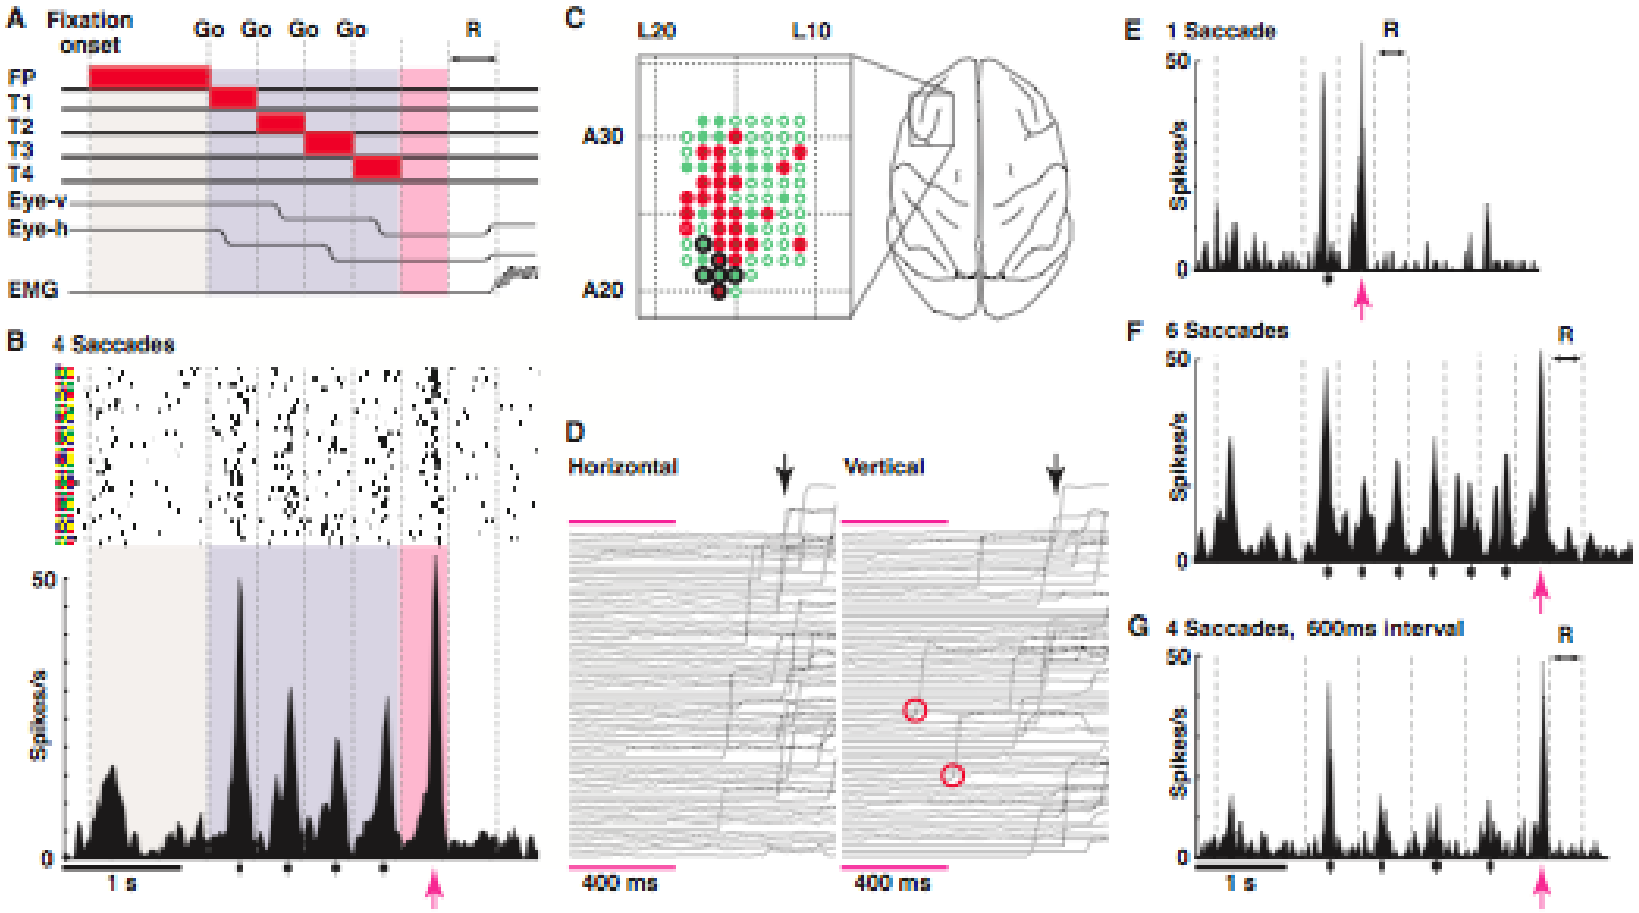

Fig. 1. Prefrontal neurons exhibit enhanced firing at the start and end of sequential saccade performance. (**A**) Standard four-saccade task (*13*), with example of vertical and horizontal (Eye-v, Eye-h) eye position traces and digastric electromyographic (EMG) activity accompanying licking. R, 400-ms reward delivery window. (**B**) Dot raster and peristimulus time histogram (PSTH) of activity of a single prefrontal neuron during 35 trials of this task [events aligned with those in (**A**)]; saccade directions indicated to left (yellow, down; red, up; green, left; blue, right). Pink denotes 400-ms extra-peak window; pink arrow indicates extra peak. Mean saccade onset times (vertical tics) and standard deviations (black boxes) are shown below *x* axis. (**C**) Recording and stimulation map of monkey M7 (A, anterior; L, lateral). Black circles are sites at which microstimulation elicited saccades. Colored circles show tracks with extra-peak activity (solid red), with task-related but not extra-peak activity (solid green), or lacking task-related activity (open green). (**D**) Rectified eye position traces recorded during 60 consecutive trials of the four-saccade task [35 trials in (**B**)]. Pink bars mark 400-ms extra-peak window, during which saccades occurred only in two trials (red circles). Black arrow, 811 ms after last target off. (**E** to **G**) Responses of the same neuron shown in (**B**) for blocks of one (**E**) and six (**F**) saccade sequences, and for four-saccade task with 600-ms target intervals (**G**).

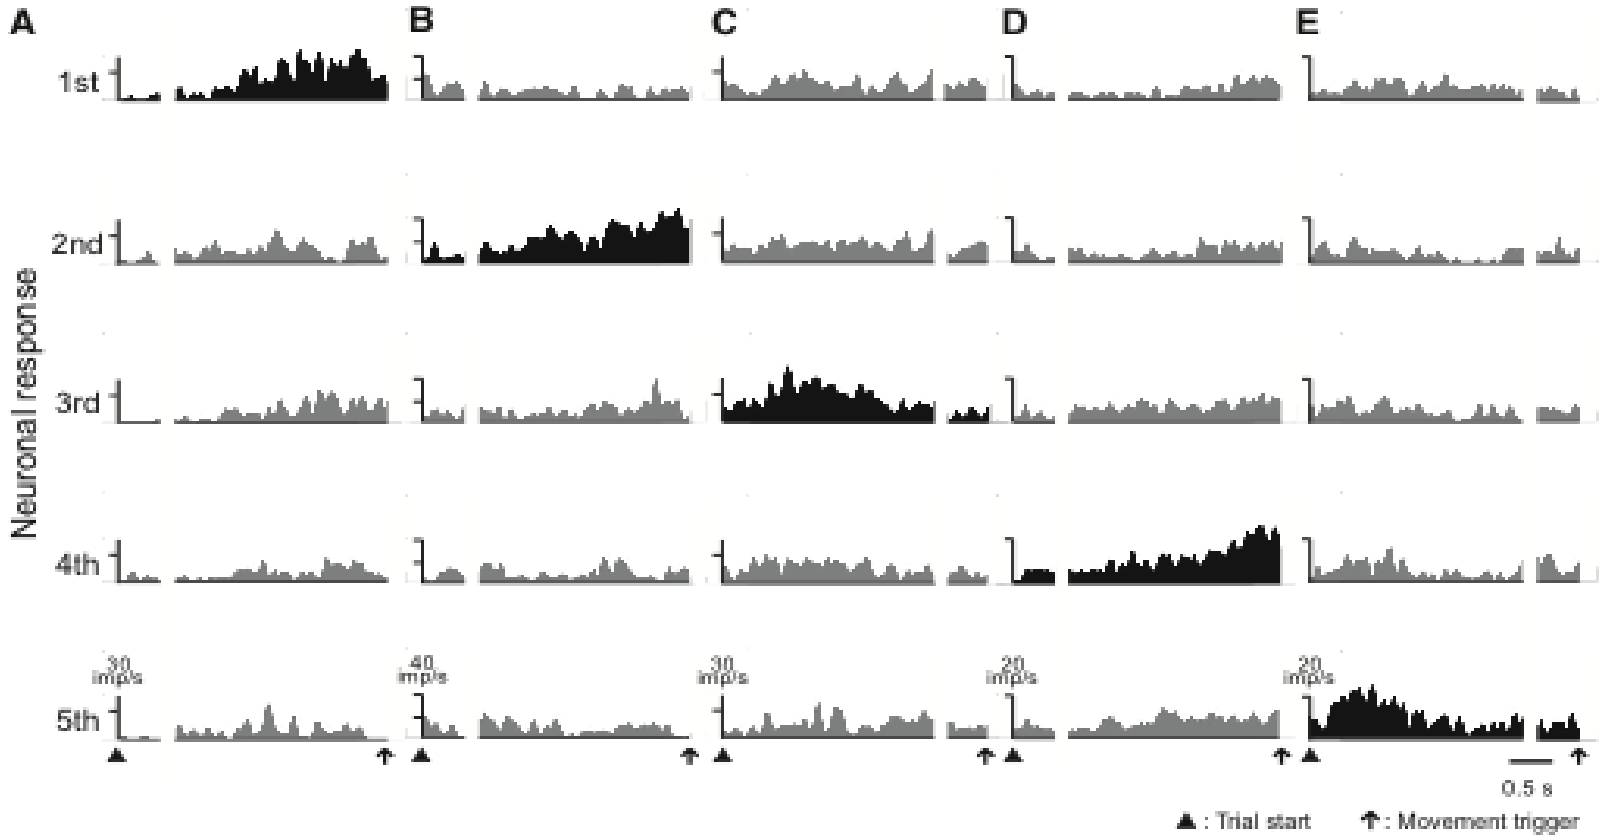
A
B
C
D
E
1st
2nd
3rd
4th
5th
Neuronal response
30 imp/s
40 imp/s
30 imp/s
20 imp/s
20 imp/s
0.5 s
▲ : Trial start
↑ : Movement trigger

Context / Structural Responses

- LPFC, PM, SMC, PPC, SEF…
- Non-reward, non-time-discounted state and event encoders
- Categorical neurons; strategy-responsive (and ‘retrospective’) neurons (Tsujimoto et al.); rule responsive neurons (Miller et al.); ‘rank’ responsive; ‘set’ encoding (Mansouri et al., 2006); response suppression (Sakagami et al.; Hasegawa et al.);
- Model-based learning and decision-making (Pan et al., 2008; Fan et al., 2011)
  - Hybrid RL-models combing model-free and model-based learning (Daw et al., 2005)

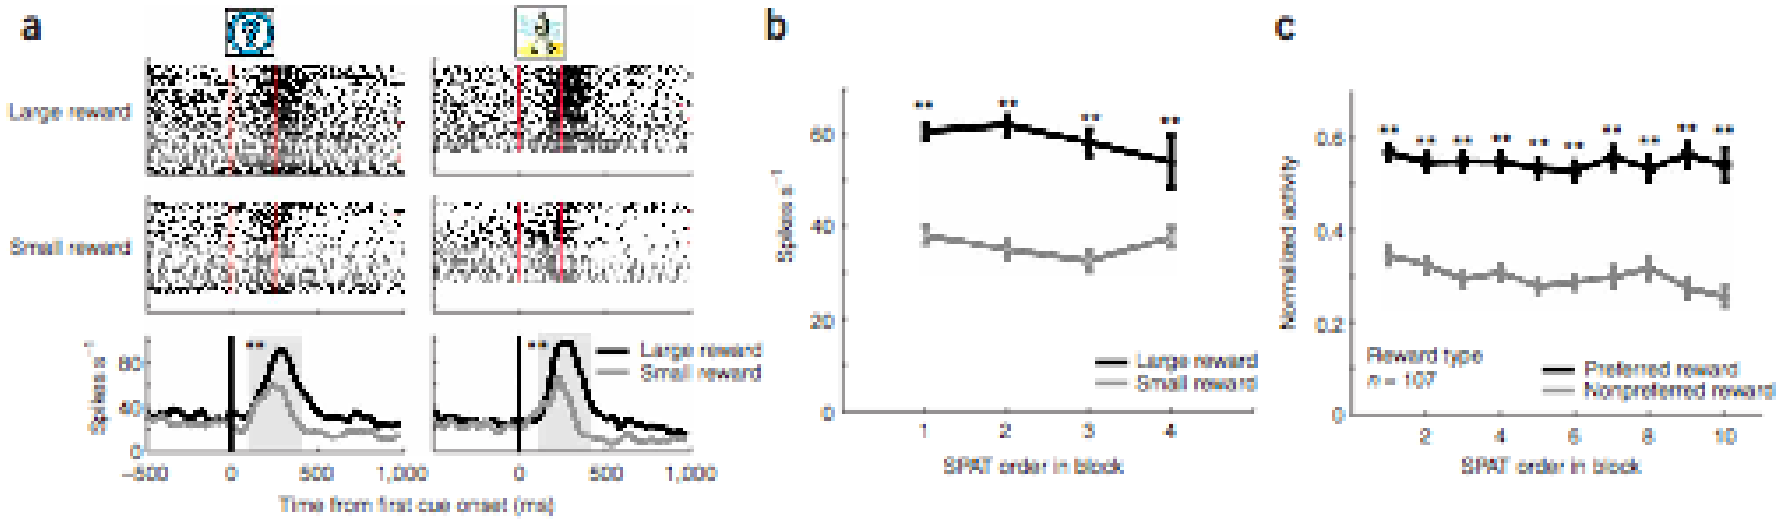

Figure 2 An example of a reward-type cell and population activity. (**a**) Single-unit activity of a reward-type neuron. All trials were sorted by four conditions—the first cue stimuli (A1 versus A2) and the two reward conditions (larger reward versus small reward)—and aligned on the first cue onset. In the raster-grams, red ticks mark the onset and offset of the first cue. In the histograms, the black curve represents data from large-reward trials and the gray curve represents data from small-reward trials. The gray area indicates the first cue period that was used for analysis of neuronal activities. Neuronal activity was also compared between two reward conditions for each stimulus separately in this time window (two tailed *t*-test, ** $P < 0.01$). (**b**) Averaged activity of this neuron to the first cue is plotted as a function of SPAT order in blocks. The activity was sorted by large reward (black curve) and small reward (gray curve). (**c**) Population-averaged activity of reward-type neurons in the first cue period as a function of SPAT order in blocks. Activities were separated by two conditions: preferred reward condition (black curve) and nonpreferred reward condition (gray curve). Statistical significance was determined by two tailed *t*-test (** $P < 0.01$). Error bars indicate s.e.m.

Compare population activity in LPFC and striatum

1. Reward neurons to old stimuli

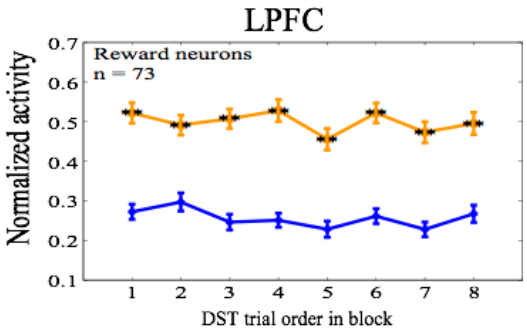

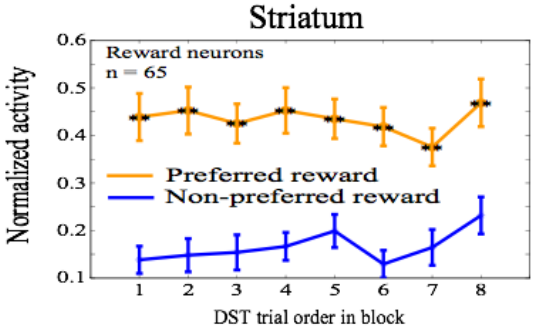

2. Reward neurons to new stimuli

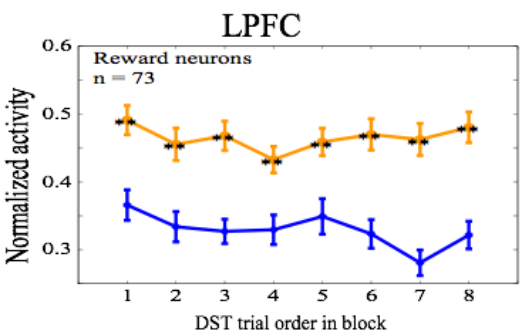

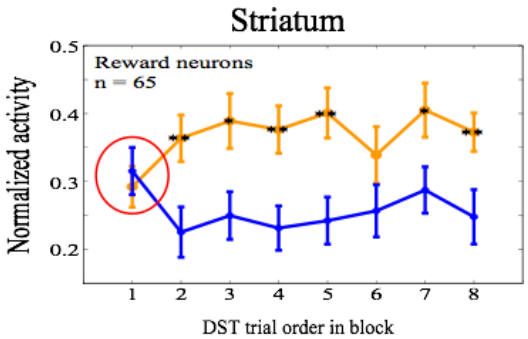

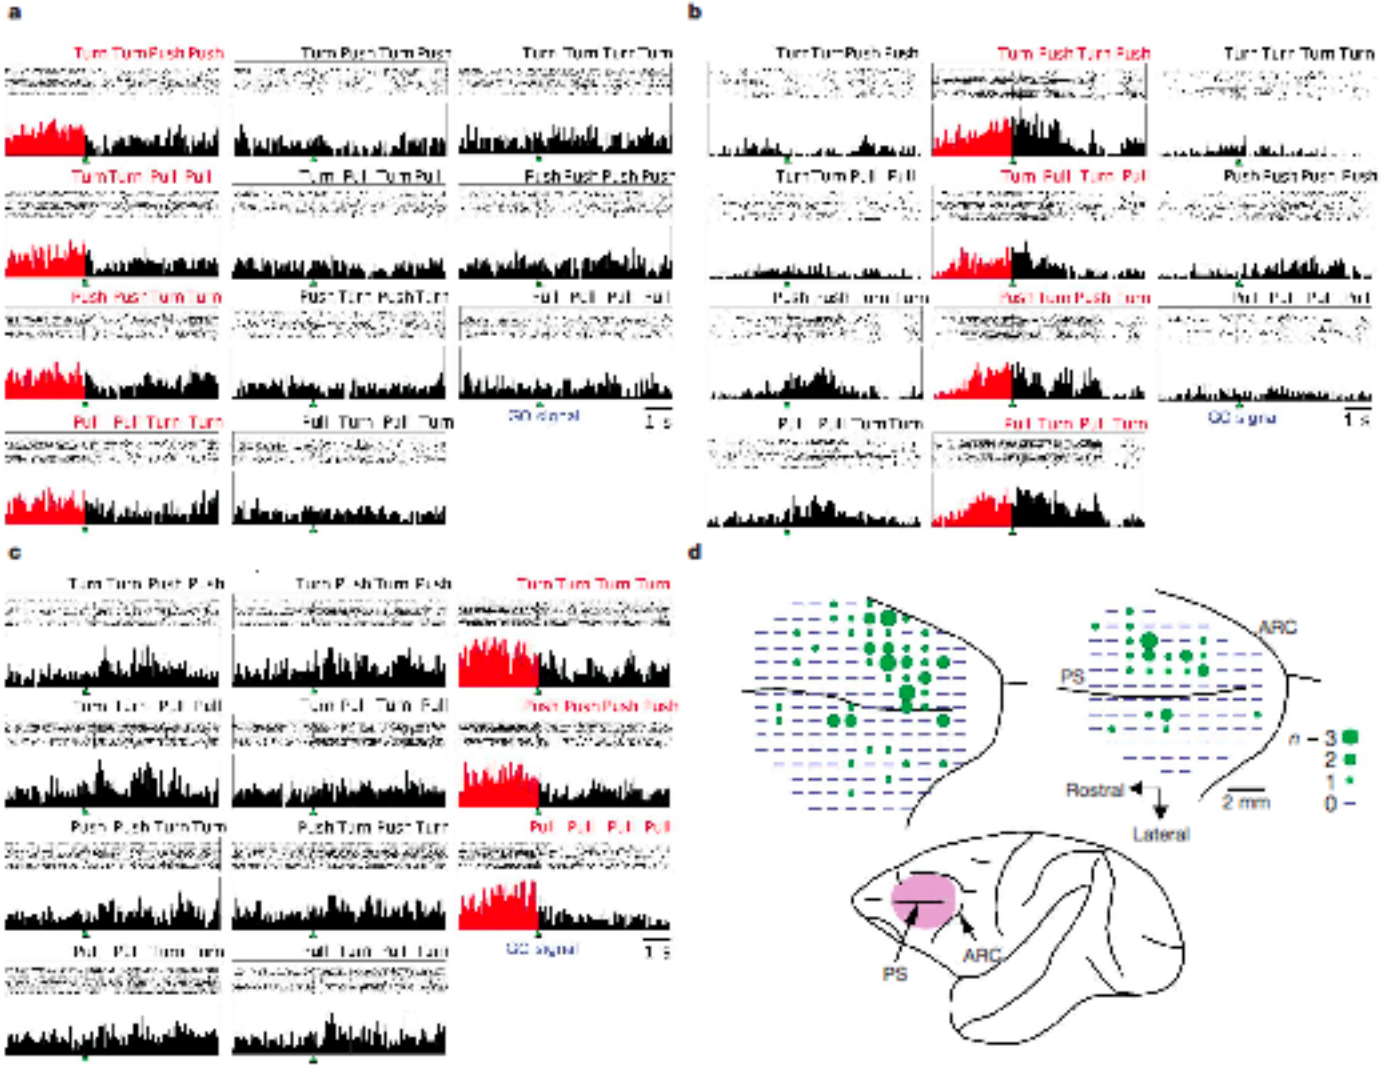

Figure 1 | Activity of PF cells selective for a category of sequences during planning. **a**, Raster displays and peri-event histograms illustrating the cellular activity selective for the 'paired' category. The displays are aligned on the appearance of the GO signal (green-filled triangle) for the first of the memorized movements. **b**, Activity selective for the 'alternate' category. **c**, Activity selective for the 'four-repeat' category. **d**, Top: recording sites of category-selective cells. The size of the circle is proportional to the number of selective cells at each site. Bottom: a cortical surface map showing the surveyed area. PS, principal sulcus; ARC, arcuate sulcus.

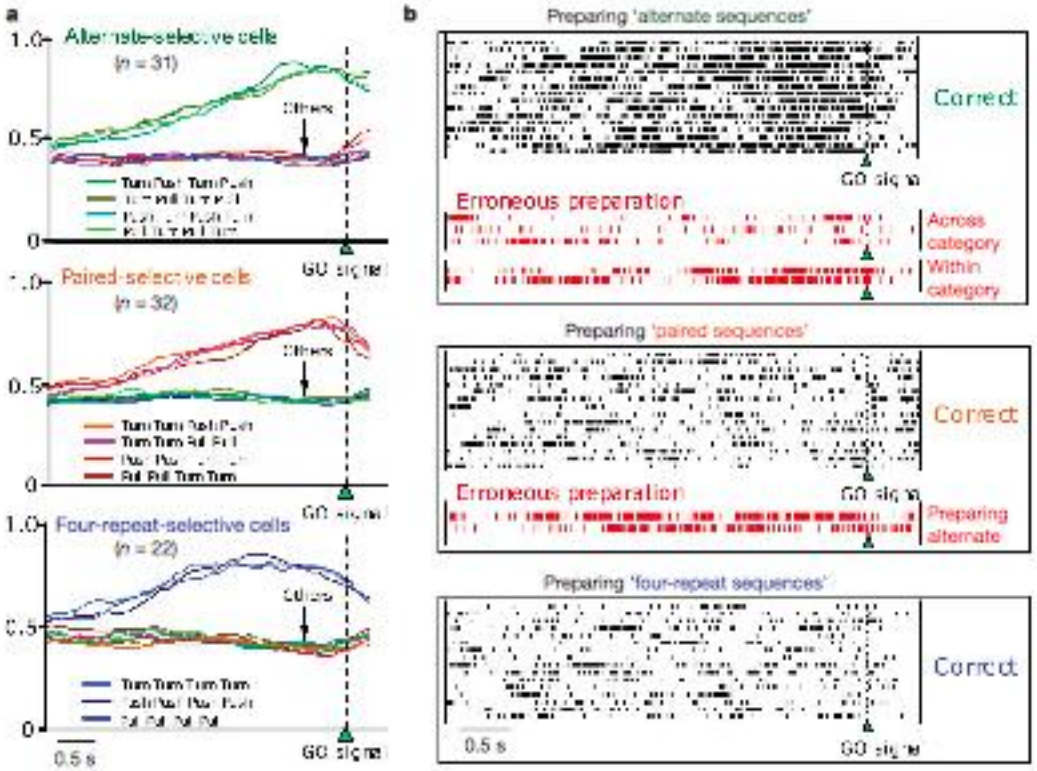

Figure 2 | Time courses of the activity of cell populations showing category selectivity and a comparison of cellular activity before the monkey executed the correct or an incorrect sequence. **a**, Spike density function for populations of PF cells that were selectively active as the monkey prepared behavioural sequences from one of the three categories. The traces compare the activity during preparation for alternate (top), paired (middle) and four-repeat (bottom) sequences with the activity during preparation for sequences of other categories. **b**, Activity of a PF cell when the monkey was preparing to perform the correct or an incorrect behavioural sequence.

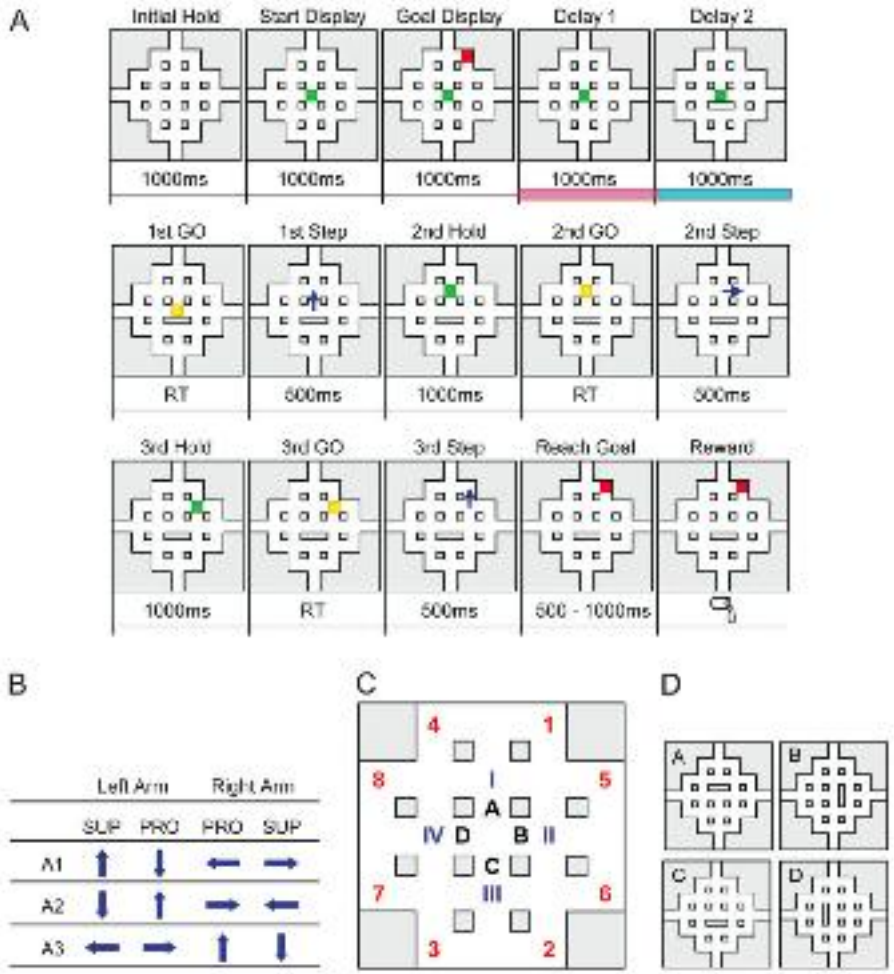

Figure 1. The path-planning behavioral task. (*A*) Temporal sequence of events in the task used in the present study. The behavioral sequence is depicted from the top left to the bottom right. Each panel represents a maze that was displayed on a monitor. Green squares denote current cursor positions and red squares indicate the position of the final goal. Yellow squares represent movement initiation (GO) signals. Blue arrows delineate cursor movements. Light red and blue bars at the bottom of the upper panels indicate the task periods that were defined as delay 1 and delay 2, respectively. RT, response time (<1 s). (*B*) Directions of cursor movements assigned to supination (SUP) or pronation (PRO) of either arm. Arm movements were assigned to each cursor movement in three different ways. A1–A3 represent arm–cursor assignments 1–3. (*C*) Goal and path-block positions within the maze. The final goal in each trial corresponded to one of the eight positions indicated by the red numbers (1–8). During the recordings of neuronal activity, we used either positions 1–4 or 5–8 as a set of goals. Four positions (blue, I–IV) were defined as immediate goals. One of four paths adjacent to the central starting position (black, A–D) was randomly removed (referred to as a path-block), as illustrated in (*D*). (*D*) Each panel shows the location of a path-block that obscured a path at the top, right, left, or bottom of the start position.

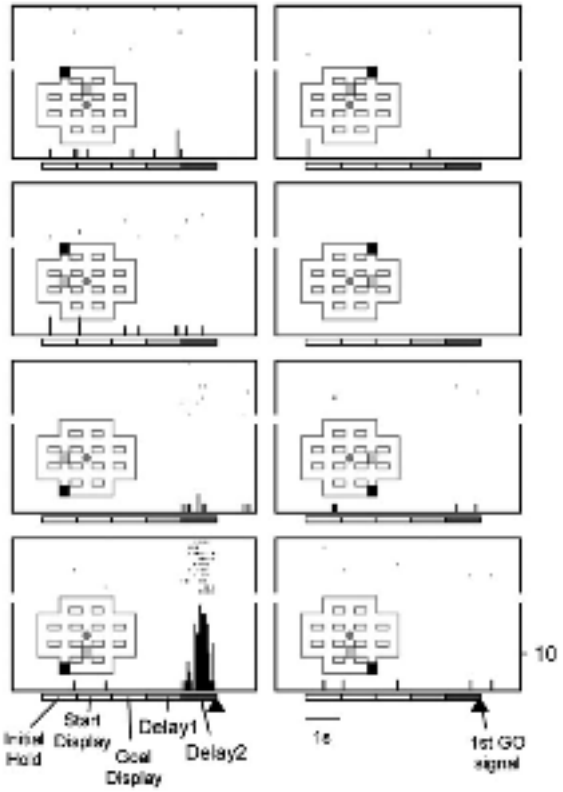

Figure 6. Example of final and immediate goal-selective neural activity. This PFC neuron discharged during delay 2 only when the animal intended to move the cursor to immediate goal II to reach final goal 3. Conventions are as in Figure 4.

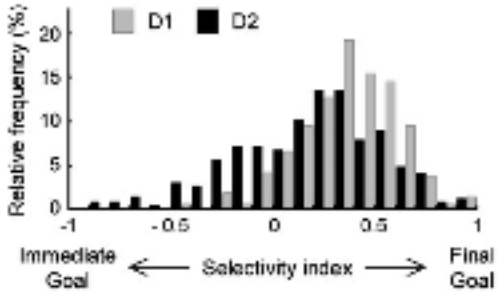

Figure 7. Comparison of behavioral goal selectivity during the delay periods. The frequency of occurrence of neurons that were either final or immediate goal-selective is depicted. The selectivity index (SI) was calculated as described in Materials and Methods. Positive SI values indicate final goal selectivity, and negative values indicate immediate goal selectivity. Each histogram depicts the relative frequency of neurons with SI values grouped in bins of 0.1. Data are for neurons that were active during either delay 1 (D1; 213 neurons, median for SI = 0.38) or delay 2 (D2; 268 neurons, median for SI = 0.24).

Outcome-Related Processing

- ACC, OFC, MFC, FPC…
- Vicarious reinforcement
- Processing outcomes of both self and other
- For updating strategies, sets
- Retrieve choices at time of feedback (for evaluation)
- Important for reversal learning
  - OFC, ACC, set-shifting

Box 1. Task used to study the FPC in monkeys

Figure I illustrates the task used to study the monkey FPC [4]. In the main task, a visual cue (Figure Ib) instructed rhesus monkeys either to 'stay' with the goal they had chosen on the previous trial or to 'shift' to a different goal. Success on this task required both a memory of the previous goal and the application of the correct strategy (stay or shift). After the cue and a delay period, the monkeys chose a goal to the left or right of a central fixation point, made a saccade to their choice, and continued fixating it until feedback arrived. The trial depicted in Figure Ia, for example, followed one in which the monkey had chosen the left goal. The purple square cued the shift strategy, which led to choosing the right goal. If the monkeys chose correctly, feedback consisted of one or two drops of fruit juice as a reward. When the monkeys chose incorrectly, red squares appeared over both potential goals instead. Two variants of the task served as controls. In one, either one or two drops of juice were delivered during the cue period, instead of a visual cue (Figure Ib). One drop meant 'stay'; two half-drops meant 'shift'. Accordingly, the monkeys received a reward twice on each correct trial: the first time as a cue, the second as feedback. In the other control task, a spatial cue instructed the goal on each trial, and the monkeys simply obeyed these instructions without reference to the previous goal.

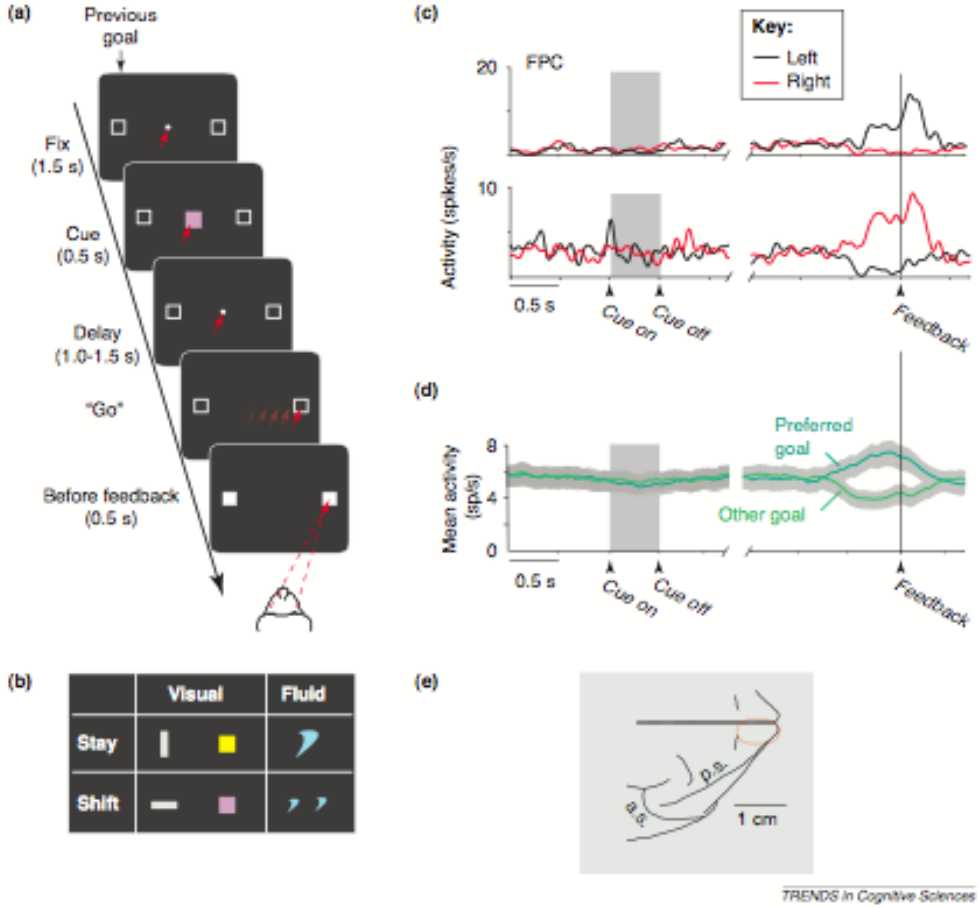

Figure I. Task design and activity in the FPC (**a**) Sequence of task events, from top to bottom. Each large rectangle corresponds to the video monitor at a given time. The white squares served as possible goals for a saccadic eye movement on each trial. Note that the term goal, in this sense, is not meant to distinguish spatial goals from movement targets or the parameters of movement. The red arrows show the target of the monkey's visual fixation. At first, the monkey had to maintain fixation (Fix) on the white circle. After 1.5 s, a cue appeared for 0.5 s, followed by a variable delay period of 1.0–1.5 s. Then the white circle disappeared as a 'go' signal for a saccade to the left or right goal. The red arrows and dashed lines depict a saccade to the right. Feedback usually arrived 0.5 s later. (**b**) Strategy cues. Visual cues were gray rectangles or colored squares. For fluid cues, either one drop of juice (large blue shape) or two half-drops (smaller blue shapes) cued a strategy. (**c**) Activity of two FPC cells. Red lines show the activity rate when the monkey correctly selected the right goal; black lines are for left goals. (**d**) Population activity averaged over all FPC neurons with goal selectivity. Activity for the preferred goal (blue) and the alternative goal (green). Shading: S.E.M. (**e**) Rostral part of the right cerebral cortex of a macaque monkey, with the FPC region studied bounded by the orange dashed line. Abbreviations: a.s., arcuate sulcus; p.s., principal sulcus. Rostral is to the right; dorsal view.

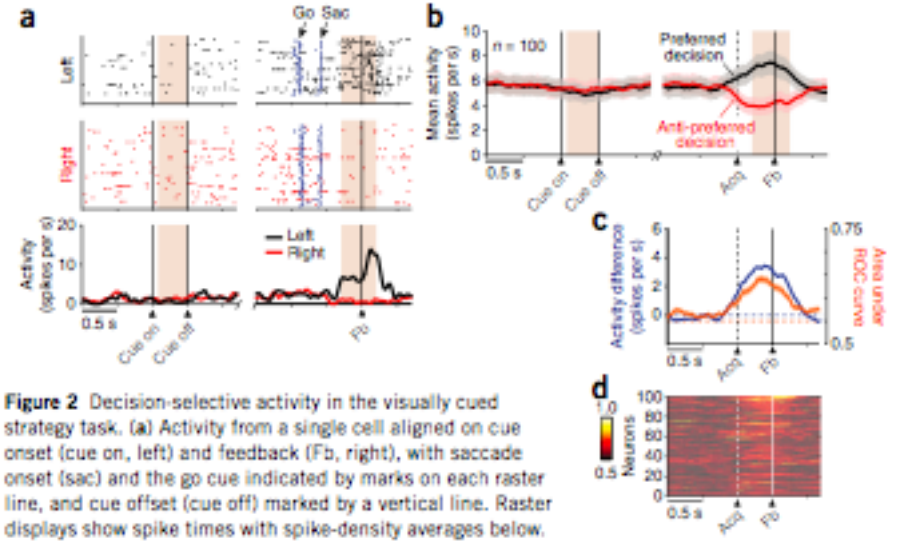

Figure 2 Decision-selective activity in the visually cued strategy task. (**a**) Activity from a single cell aligned on cue onset (cue on, left) and feedback (Fb, right), with saccade onset (sac) and the go cue indicated by marks on each raster line, and cue offset (cue off) marked by a vertical line. Raster displays show spike times with spike-density averages below. Background shading indicates the analysis periods. Feedback-period activity for left decisions (8.2 ± 5.3 spikes per s, mean ± s.d.) significantly exceeded that for right decisions (0.8 ± 1.7 spikes per s, two-way ANOVA, $F_{1, 87} = 85.0$, $P << 0.001$). (**b**) Population activity for decision-selective FPC neurons, computed separately for each neuron's preferred (black) and anti-preferred (red) decision. Shading indicates s.e.m. Bin width was 20 ms, three-bin moving average. The dashed vertical line indicates target acquisition (acq). (**c**) The activity difference between preferred and anti-preferred decisions (blue) from **b** and the mean ROC value from **d** (orange, shading indicates s.e.m.). The blue dashed line is at 0 spikes per s and the orange dashed line is at the mean of shuffled ROC values. (**d**) ROC plots for decision-selective FPC neurons, with the area under the ROC curve color coded for each cell (scale at left), ranked according to values during the feedback period.

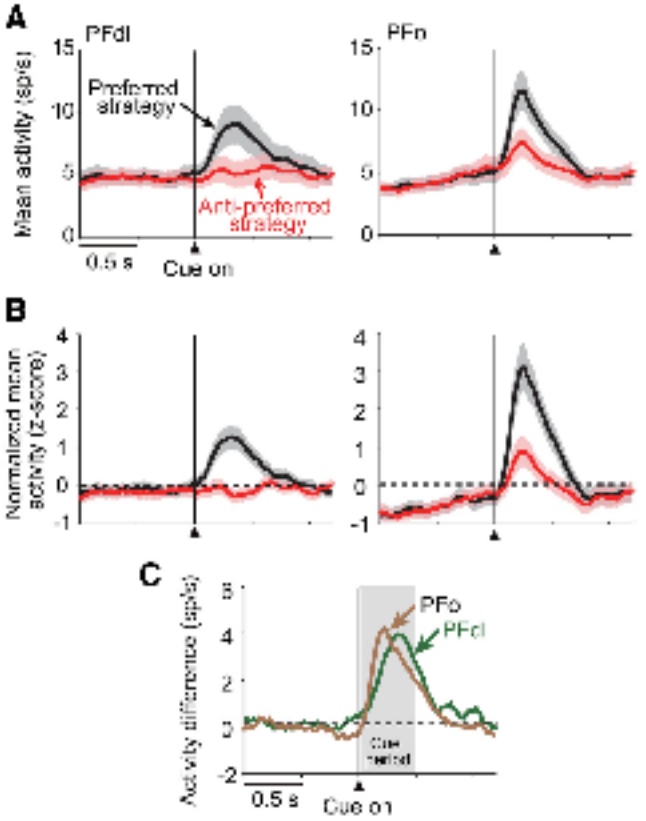

Figure 5. Strategy-selective population means. ***A***, Mean population activity (shaded area, SEM) for PFdl (left) and PFo (right) neurons with significant strategy-selective activity during the cue period ($n = 48$ and 49, respectively), computed separately for each neuron's preferred (black) and anti-preferred (red) strategy. Vertical line shows the time of cue onset. ***B***, Normalized (*z*-score) population averages for the same data as shown in ***A***. ***C***, Difference in activity between the preferred and anti-preferred strategies.

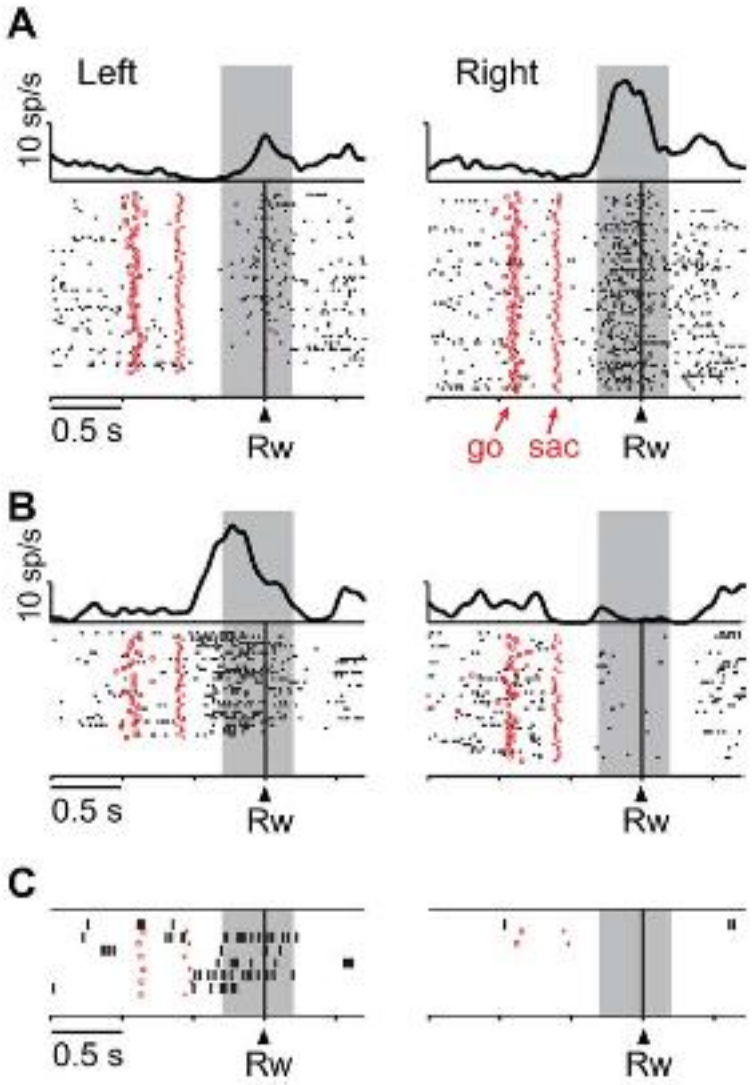

Figure 2. Two response-selective PFo neurons. ***A***, Neuron from monkey 2. Activity aligned on reward (Rw), sorted chronologically from top to bottom, with saccade onset (sac) and the go cue indicated by marks on each raster line. Raster displays show spike times with spike-density averages above each display. Left and right responses shown separately. Correct trials only are shown. Background shading: feedback period. ***B***, Neuron from monkey 1, in format of ***A***. ***C***, From the cell in ***B***, for error trials, in format of ***A***. Note that on error trials, feedback was visual, in contrast to reward feedback on correct trials. Trials are sorted by responses made (not by correct responses).

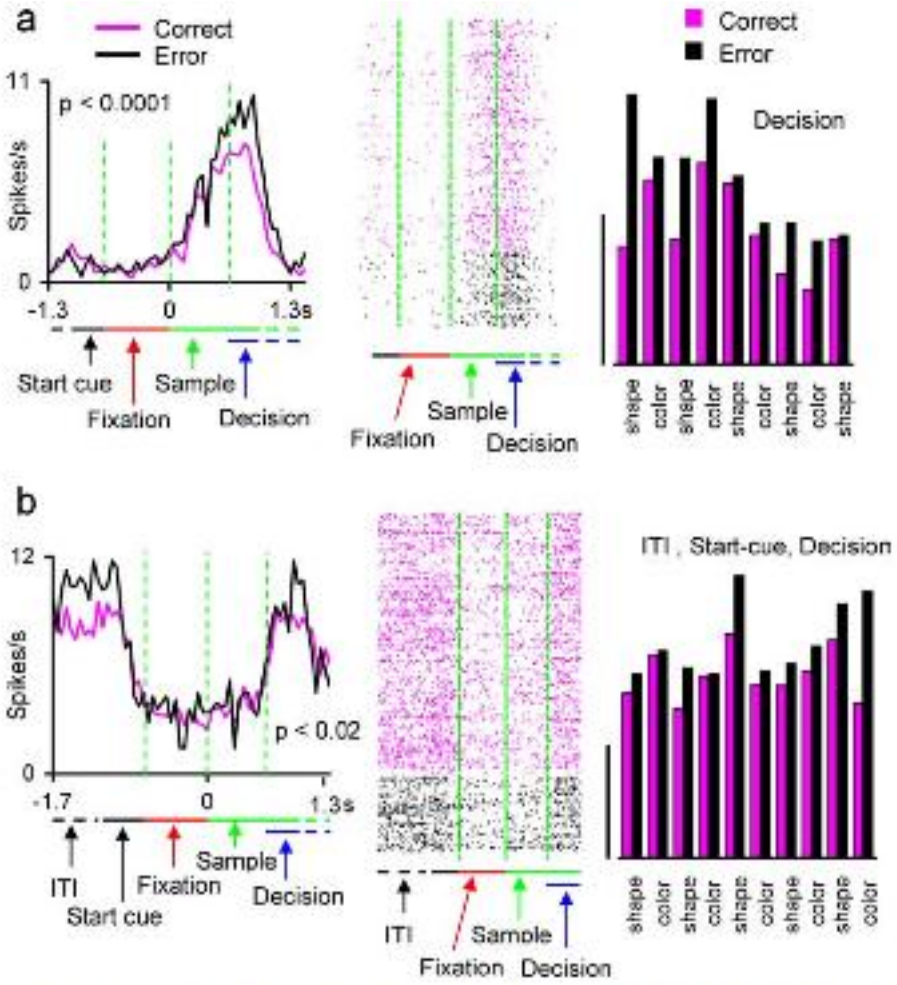

Figure 5. Activity differences between correct and error trials in two PFC cells. ***a***, ***b***, The line graphs at the left show the averaged firing rates in correct and error trials, aligned at sample onset. The *p* values indicate the significance of effects of response type in a three-way ANOVA (rule × response type × epoch). The bin size is 50 ms. The raster grams in the middle show activity in individual correct and error trials of the first four blocks, also aligned at sample onset. The bar graph at the right represents the mean firing rate in the epoch noted above individual graphs for correct and error trials in consecutive blocks. The activity differences between correct and error trials were consistently seen in the same direction in consecutive color and shape blocks. Pink and black colors of lines, bars, and dots indicate correct and error trials, respectively. Vertical calibration bar, 5 spikes/s.

Reward / Value / Motivational Modulations

- LPFC, OFC, SMC, FEF, BG, LIP, PM…
- Reward size, reward probability, reward preference, reward expectancies
- Also, effort, information quality, time-discounting, satiation
  - Note that social information inherently rewarding
  - Need ‘multi-dimensional’ view of motivation
- Bernacchia et al 2011 on temporally-extended modulations widely broadcast
- Can increase or decrease firing rate
- Need integrative framework for combining ‘plans’ / ‘context’ (lateral) and ‘values’ / ‘motivations’ (medial)
  - Watanabe & Sakagami, 2007
  - Executability X Desirability

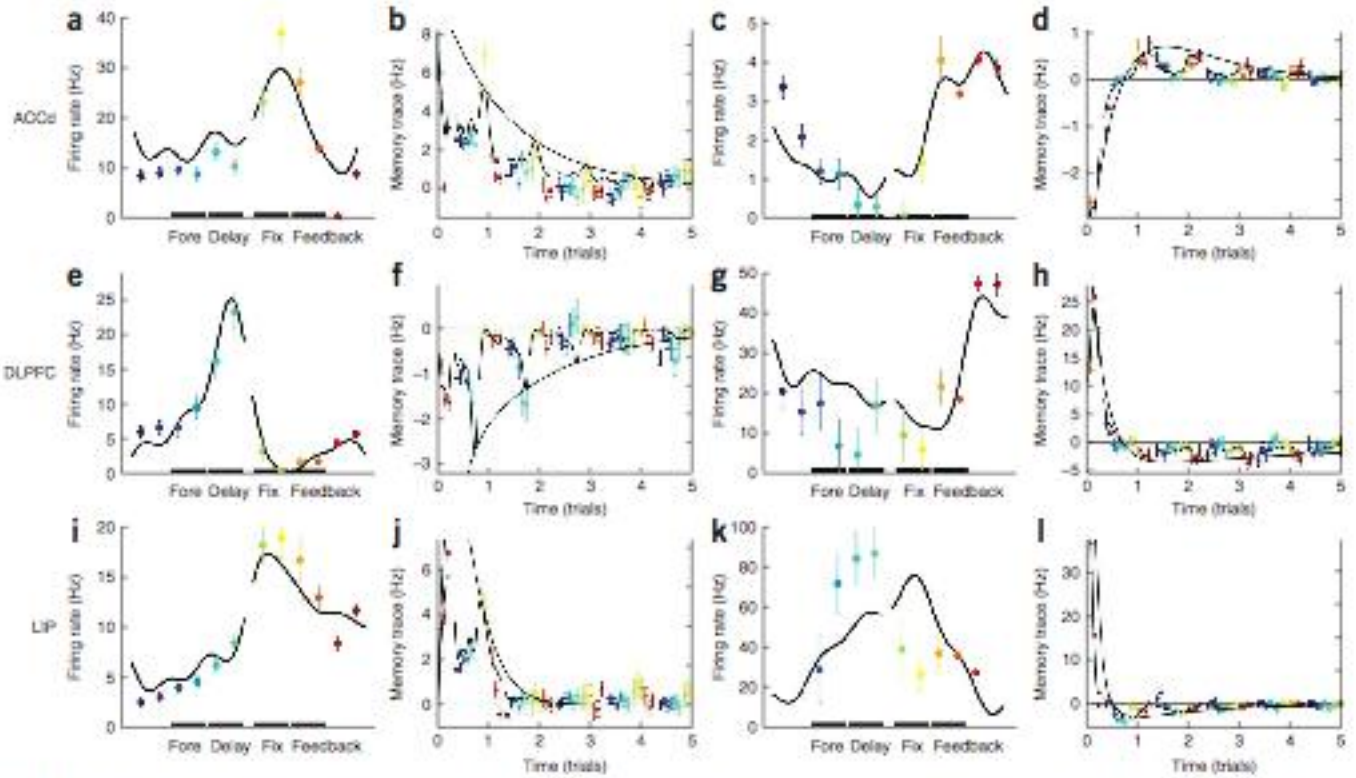

Figure 3 Firing rates and memory traces for six neurons, two for each of the three recorded areas. For each of the six neurons, epoch codes (first and third column) and memory traces (second and fourth column) are shown, presented as in **Figure 2a**,**b**. The second column shows monotonic decay of the memory trace and the fourth column shows biphasic memory traces (double exponential). Different neurons had different firing rates, both in magnitude and time course, and different types of memory decay, but they were all consistent with an exponential (single or double) decay of the memory modulated by the epoch code. The factorization indexes for those neurons are 0.98 (**a**,**b**), 0.91 (**c**,**d**), 0.98 (**e**,**f**), 0.84 (**g**,**h**), 0.97 (**i**,**j**) and 0.61 (**k**,**l**).

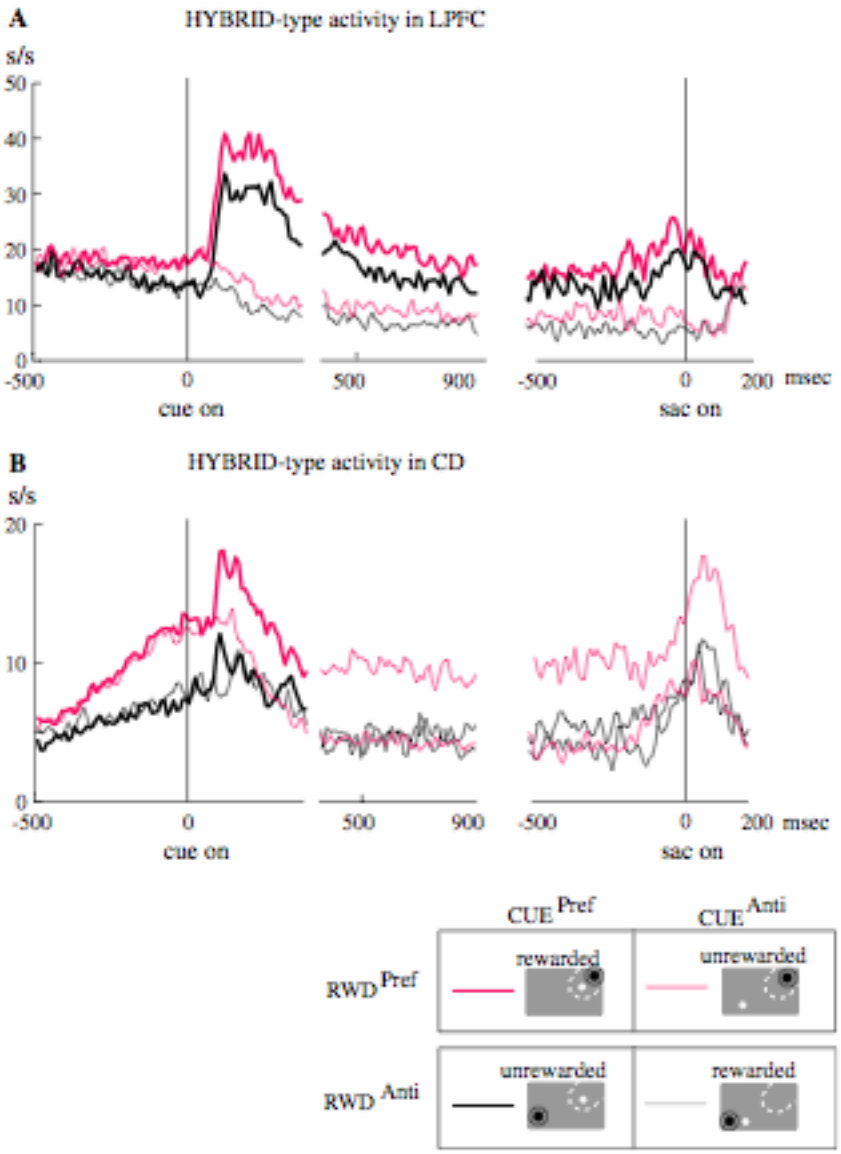

Fig. 8 Population histograms of HYBRID-type cells. Averaged discharge rates of HYBRID-type cells, which had a preferred direction for cue and reward in the same direction, are shown for the LPFC population (**a**) and CD population (**b**). A *white dashed line* in an inset schema indicates the preferred direction. The cue was presented either in the preferred direction (CUE^{Pref}, white circle inside the *white dashed line* in the schema; thick line in the histograms) or anti-preferred direction (CUE^{Anti}, white circle out of the *white dashed line* in the schema; *thin line* in the histograms). Immediate reward was associated either with the preferred direction (RWD^{Pref}, bull's eye mark inside the *white dashed line* in the schema; *red line* in the histograms) or non-preferred direction (RWD^{Anti}, bull's eye mark out of the *white dashed line*; *black line* in the histograms). Immediate reward was available in the CUE^{Pref}–RWD^{Pref} condition and the CUE^{Anti}–RWD^{Anti} condition, whereas immediate reward was not available in the CUE^{Pref}–RWD^{Anti} condition and the CUE^{Anti}–RWD^{Pref} condition. The *left vertical line* indicates the time of cue onset and the *right vertical line* indicates the time of saccade onset

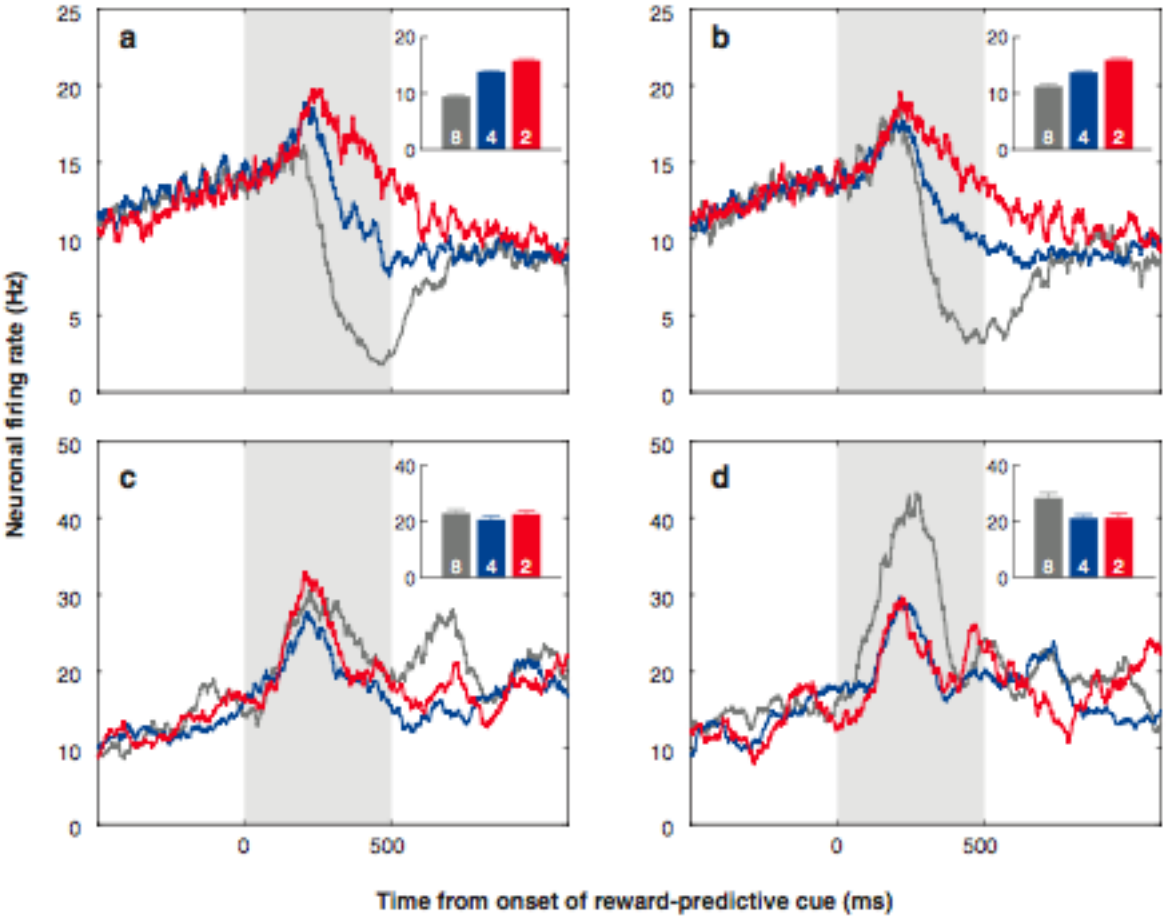

Figure 3

Spike histograms from two single neurons encoding the expected reward and/or the monkey's response (a left or right saccade). Inset bar graphs indicate the mean neuronal firing rate (± standard error) during the presentation of the reward-predictive cue (the first 500 ms). Gray indicates that the cue predicted the delivery of eight drops of juice, blue four drops, and red two drops. (*a*, *b*) OFC neuron encoding the predicted reward in a parametric fashion irrespective of saccade direction. This neuron showed a depression in its firing rate that was greatest for eight drops of juice, less for four drops, and least of all for two drops. Its firing rate, however, was the same irrespective of whether the monkey would make a left or right saccade to earn the reward. Significantly more OFC neurons (28%) showed this pattern of selectivity compared with DLPFC neurons (13%, chi-squared = 9.8, P < 0.005). (*c*, *d*) A DLPFC neuron that showed a complex pattern of selectivity that encoded a combination of the reward and the upcoming saccade. During the cue epoch, the neuron discriminated between the different expected reward amounts only when the monkey would make a rightward saccade (showing a high firing rate when eight drops of juice were expected). In contrast, during the subsequent period the same neuron was reward-selective only when the monkey would make a leftward saccade. Significantly more DLPFC neurons (43%) encoded a combination of the reward and response compared with OFC neurons (19%, chi-squared = 19, P < 0.00,005).

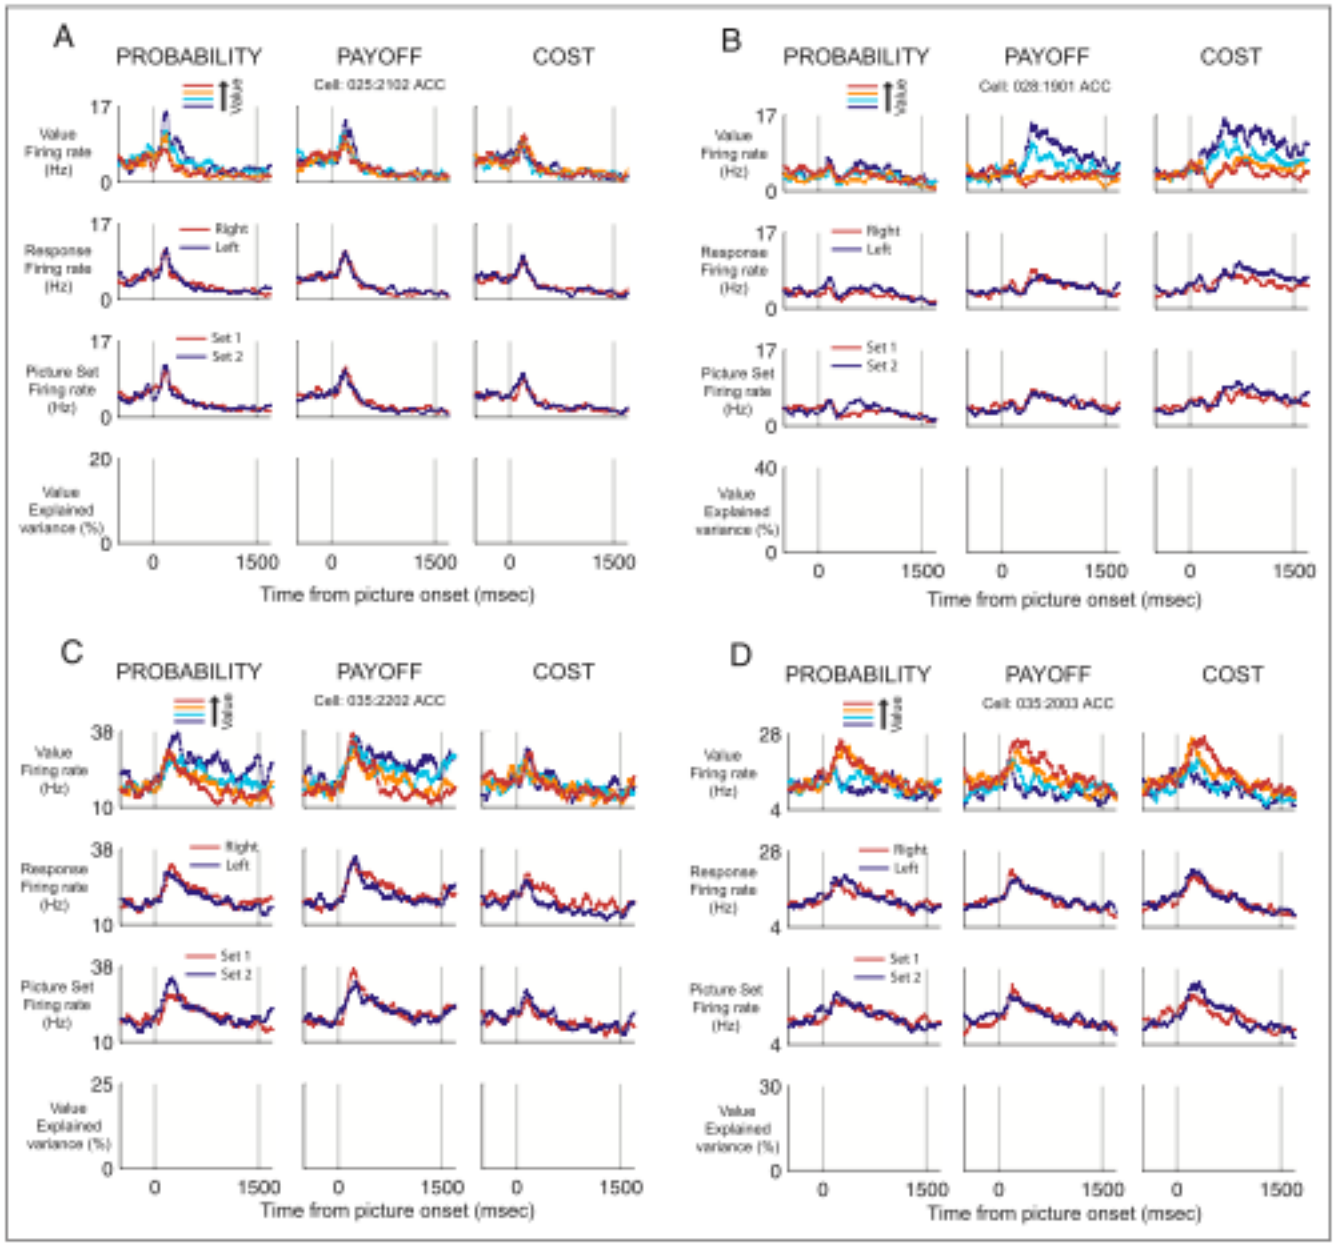

Figure 3. (A) An ACC neuron that increases its firing rate as choice values decrease for probability decisions only. The top three rows of panels display spike density histograms illustrating the mean firing rate of the neuron on trials sorted according to choice value, direction of the behavioral response, and picture set, respectively. Each column illustrates the neuron's response to each of the decision variables. The lowest row of panels quantifies the strength of encoding of the choice's value by calculating the percentage of explained variance based on the results of a sliding linear regression analysis (see Methods). Significant time bins are colored red. (B) An ACC neuron that encodes the choice values for payoff and cost decisions. (C) An ACC neuron that encodes the choice values for probability and payoff. (D) An ACC neuron that encodes the value of the choices for all three decision variables.

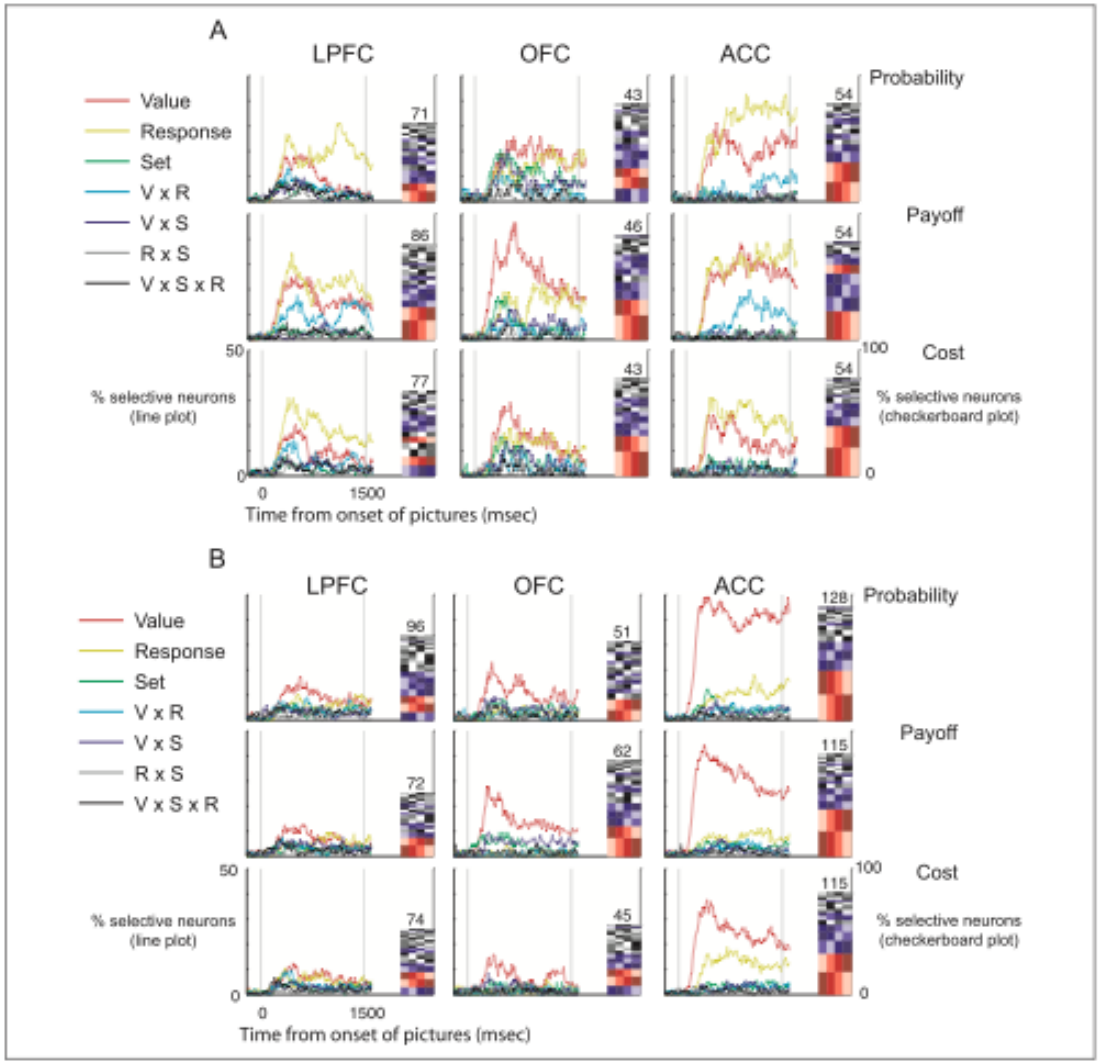

Figure 8. The line plots indicate the percentage of all neurons according to the three-way ANOVA that encoded a given factor or interaction of factors at each time point across the course of a trial for (A) Subject A and (B) Subject B. The checkerboard indicates the proportion of different value-encoding schemes. Each row indicates a specific ordering of values, with the values arranged from left to right according to how strongly the neuron fired. The intensity of shading indicates a specific value (with lowest to highest value shaded light to dark). The height of each row indicates the percentage of neurons that encoded value with that particular ordering. The different orderings are arranged vertically in decreasing prevalence. As an example, consider the encoding of value for probability decisions in ACC (A, top right). The most common ordering (bottom row) is those neurons that fired least to the highest value (leftmost darkest shading) and then showed a progressive increase in firing as value progressively decreased, with their highest firing rate occurring to the lowest value (rightmost lightest shading). The next most prevalent ordering (row second from bottom) is those neurons that fired least to the lowest value and then showed a progressive increase in firing as value progressively increased. We have color-coded the different orderings to highlight those of particular interest. The red shading indicates the two orderings that correspond to monotonic encoding schemes, whereas the blue shading indicates the six orderings that correspond to single transpositions from monotonic encoding. The gray-scale shadings reflect the remaining 16 orderings that were neither monotonic nor a single transposition from monotonic. The overall height of the checkerboard indicates the proportion of the neuronal population that encoded value, whereas the absolute number of neurons that did so is printed atop the checkerboard.

Decision Processes

- Combine ‘contextual’ with ‘motivational’ information
- Evidence accumulates to drive a population over threshold
- Must drive downstream processes towards a decision

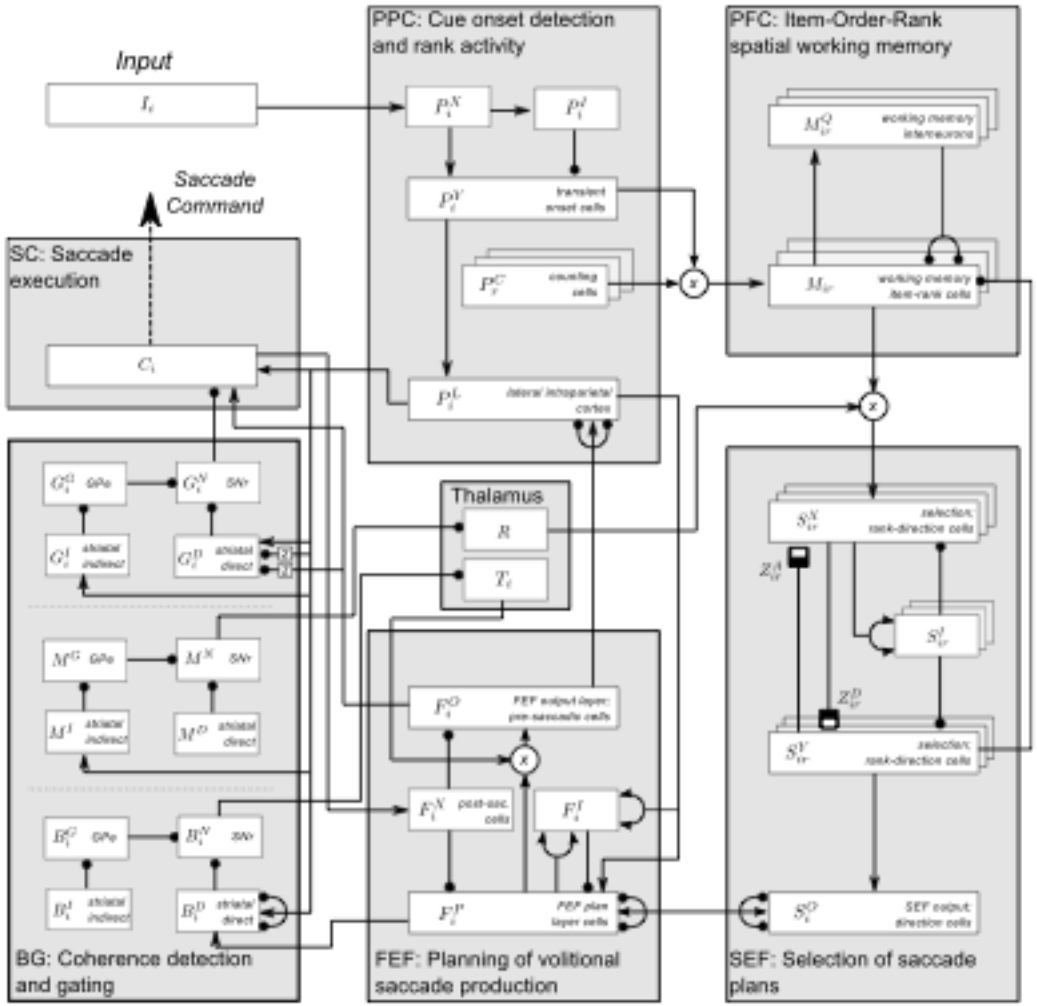

Fig. 1. An Item-Order-Rank spatial working memory and performance model. Each gray box represents a brain region within which fields of cells, represented by white inset boxes, share similar functional roles. Arrowheads denote excitatory connections between cells, and filled circles represent inhibitory connections. Curved branches at the ends of connections represent one-to-many fan-out connections that impact all other cells in the field. Half-filled boxes at the ends of connections represent habituative gates which exhibit activity-dependent changes in synaptic efficacy. White circles containing a multiplication sign ($\times$) represent multiplicative interaction between two signals. Boxes containing a sigma (Σ) represent the sum of outputs from all cells in the field that gave rise to the projection. Stacked field representations denote populations of rank-sensitive cells.

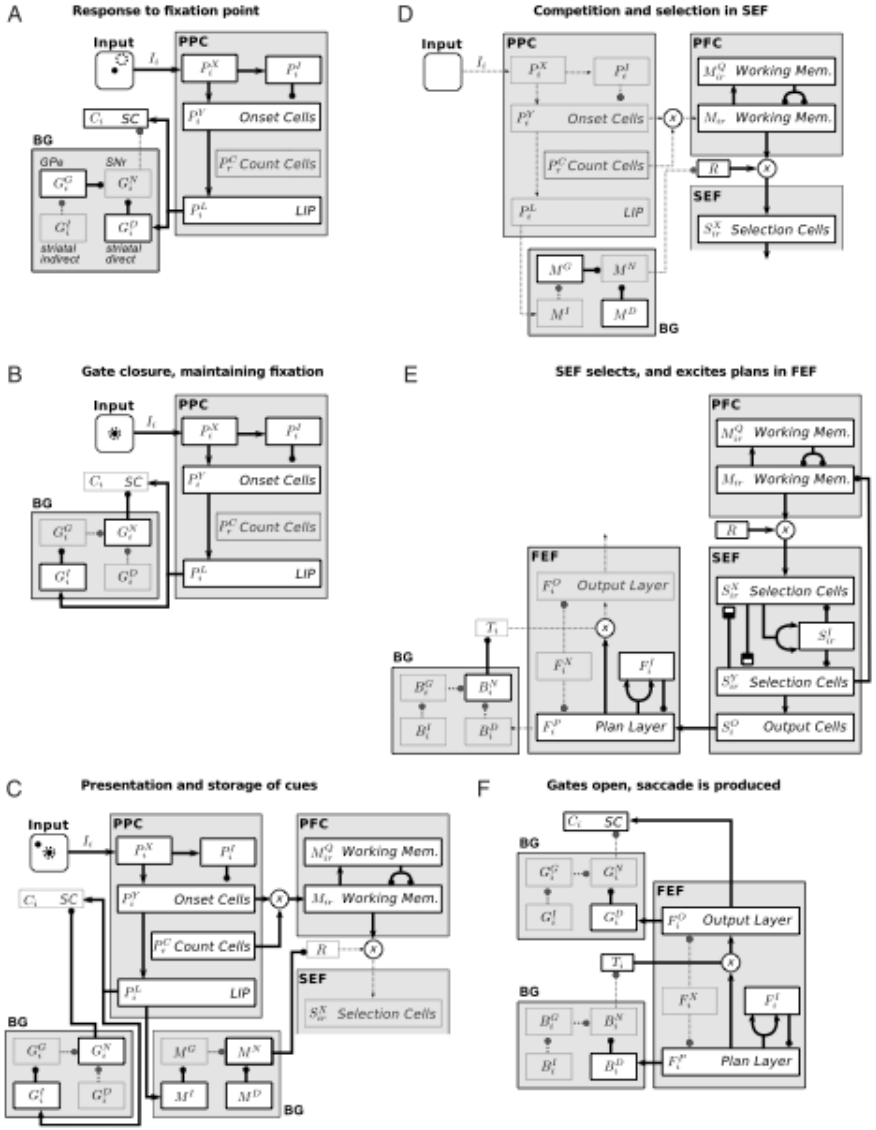

Fig. 10. Model stages during the ISR task. Each panel depicts portions of the model relevant for a particular process as the model solves the ISR task. Gray fields are those that are effectively removed from the system by either inhibition or lack of excitation. White, bold stroke fields are active, and participate in the process being explained in each panel. The dotted circle in the input space of panels A–C represents the direction of gaze. (A) The model saccades to the fixation point. Anatomical labels shown in this panel for BG apply to all panels. (B) Fixation is maintained. (C) A sequence of spatial cues is presented and stored. Neither saccades nor selection can occur because gates are held closed by fixation-related LIP activity. (D) Fixation point removal opens the rehearsal gate R via the BG working memory loop (Fig. 4[A]), and allows selection to begin in SEF (cf. Grossberg & Pearson, 2008). (E) Selected saccade plans are excited in FEF and deleted from working memory. (F) FEF and collicular gates open, allowing the saccade plan to flow through FEF and to SC, which generates a saccade. Operations in panels D–F repeat until no representations remain in the working memory.

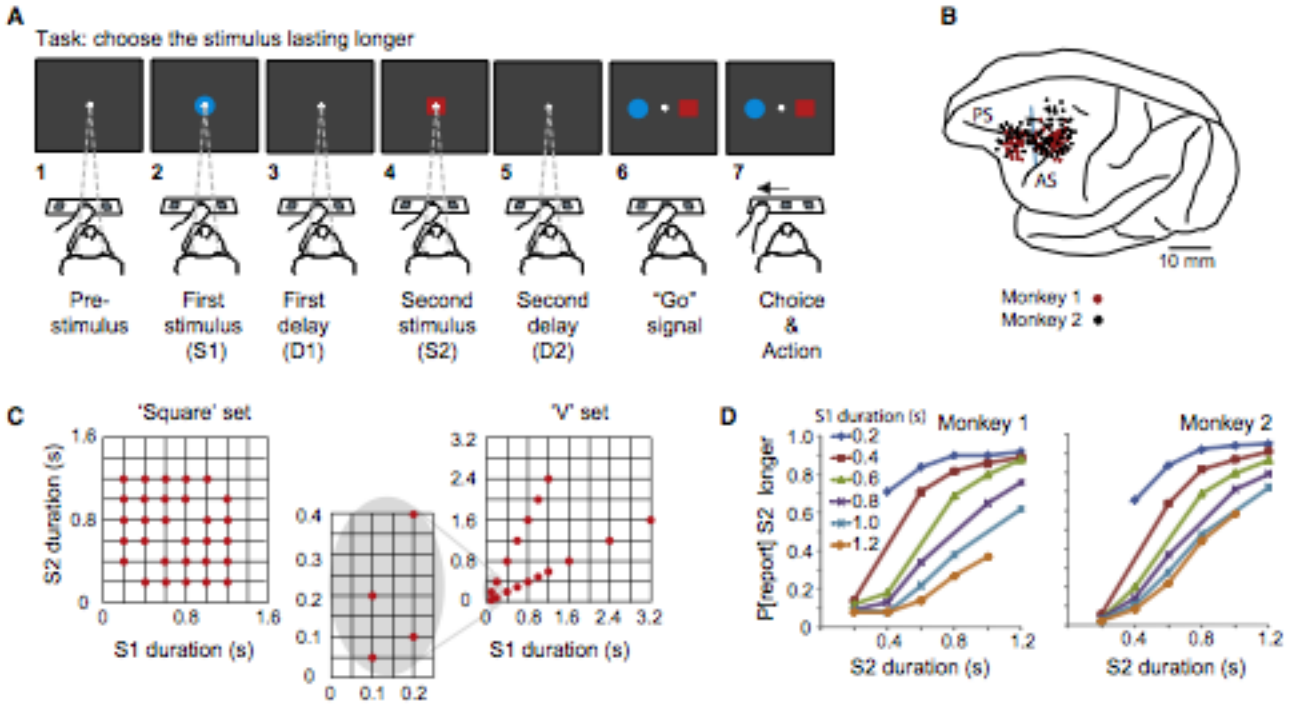

Figure 1. Task Design, Recording Locations, and Psychometric Curves

(A) Sequence of task events. Each gray rectangle represents the video screen.

(B) Penetration sites. Composite from both monkeys, relative to sulcal landmarks. Vertical blue line: division between periarcuate (right) and dorsolateral prefrontal (left) areas. Abbreviations: AS, arcuate sulcus; PS, principal sulcus.

(C) Stimulus sets.

(D) Psychometric curves showing the probability of reporting (P[report]) that S2 lasted longer as a function of S2 duration for the square set of durations.

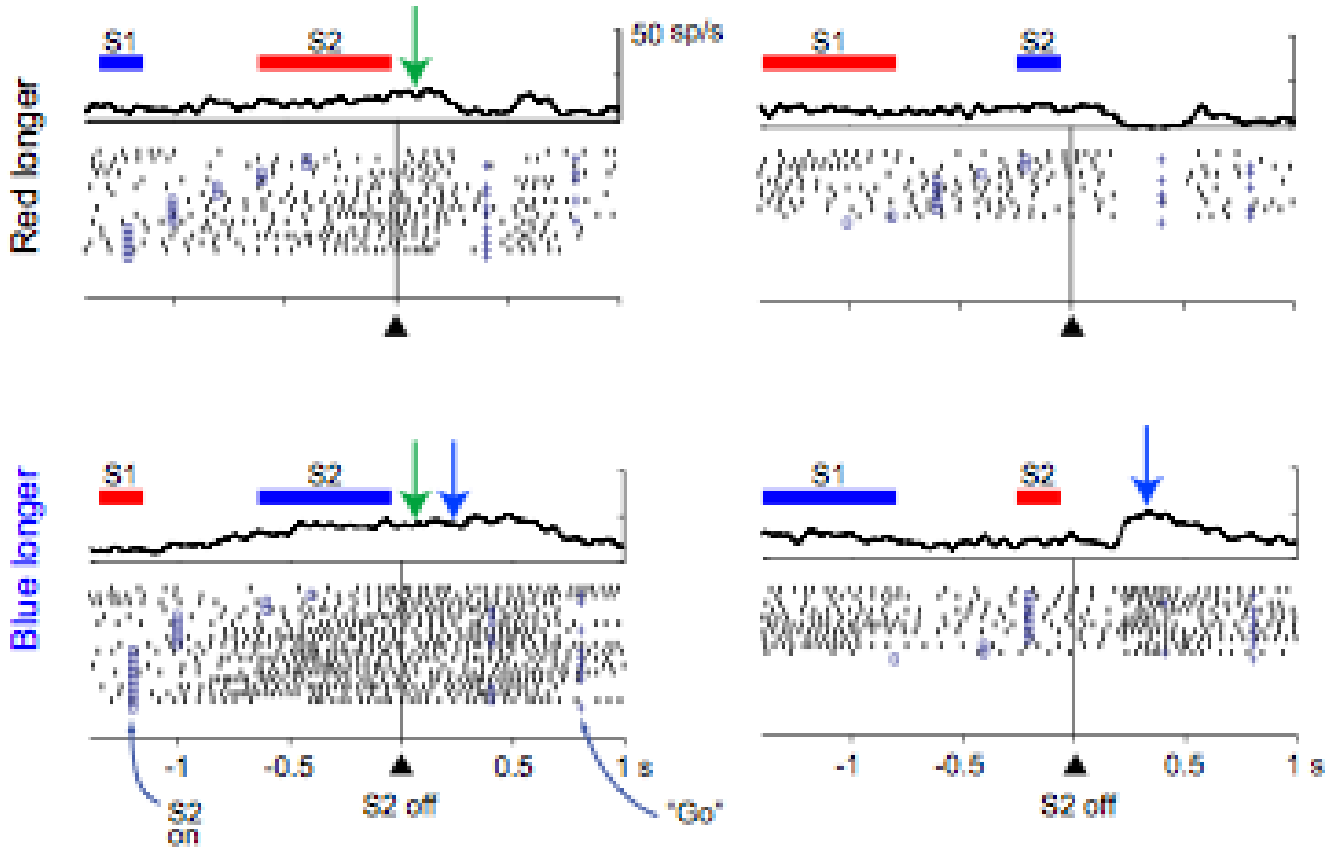

Figure 5. Neuron Encoding Relative Duration

This neuron encoded the relative duration on the basis of both the stimulus (i.e., whether the blue or red stimulus had lasted longer, $p < 0.001$) and the order of presentation (i.e., whether S2 or S1 lasted longer, $p < 0.001$). The time course of these two signals differed in the same way as the average ROC values shown in Figure 6F. The representation of relative duration based on stimulus order had already developed by the beginning of the D2 period (green arrows). The cell discharged more when S2 was longer (left column) than when S1 was longer (right column). The representation of relative duration based on the stimuli (red or blue) emerged later (blue arrows). By ~200 ms after S2 offset, the neuron was more active when the blue stimulus was longer (bottom row) than when the red one was longer (top row). The format is as in Figure 2.

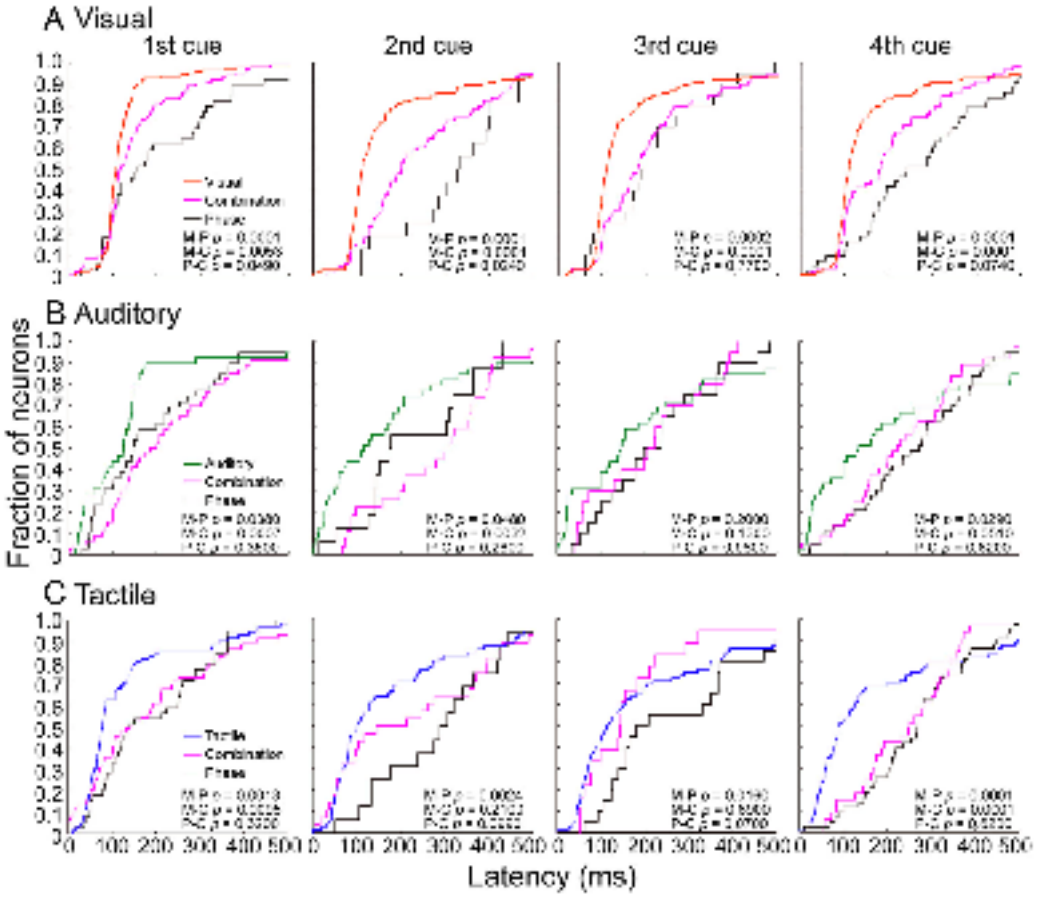

Figure 14. Latencies of individual neurons. ***A–C***, Cumulative distributions of latencies at which individual neurons began to respond to the appearance of the cue. ***A***, Visual trials. ***B***, Auditory trials. ***C***, Tactile trials. In ***A–C***, the four panels from left to right indicate data after the onset of the first, second, third, and fourth cue, respectively. "M" shows the onset latencies of the unimodal neurons with early cue responses and with a preference for the visual cue (***A***, the red line labeled "Visual"), the auditory cue (***B***, the green line labeled "Auditory"), and the tactile cue (***C***, the blue line labeled "Tactile"). "P" (black line labeled "Phase") shows the onset latencies of the single-phase selective neurons in the cue period. The latencies were calculated according to the preference for the phase position (first to fourth cue) and for the modality [visual trials (***A***), auditory trials (***B***), and tactile trials (***C***)]. "C" shows the onset latencies of neurons selective for the combination of phase and modality in the cue period (indicated with "Combination"). The latencies were calculated separately for neurons preferring a specific combination of phase and modality (Table 5). The *p* values of the statistical analysis (the two-sample Kolmogorov–Smirnov test) are displayed for each comparison of the onset latencies. M-P, Between unimodal neurons with early cue responses and neurons with single-phase selectivity; M-C, between unimodal neurons and combination neurons; P-C, between phase neurons and combination neurons.

Learning

- How does functional connectivity change as a result of learning?
- Model-based versus model-free learning
  - How do these systems functionally segregate?
  - How is the ‘model’ implemented?
- What is the time course of learning in different systems?
- Consolidation / automaticity
  - Chunking of programs versus task expertise
- Pay attention to lesion data…

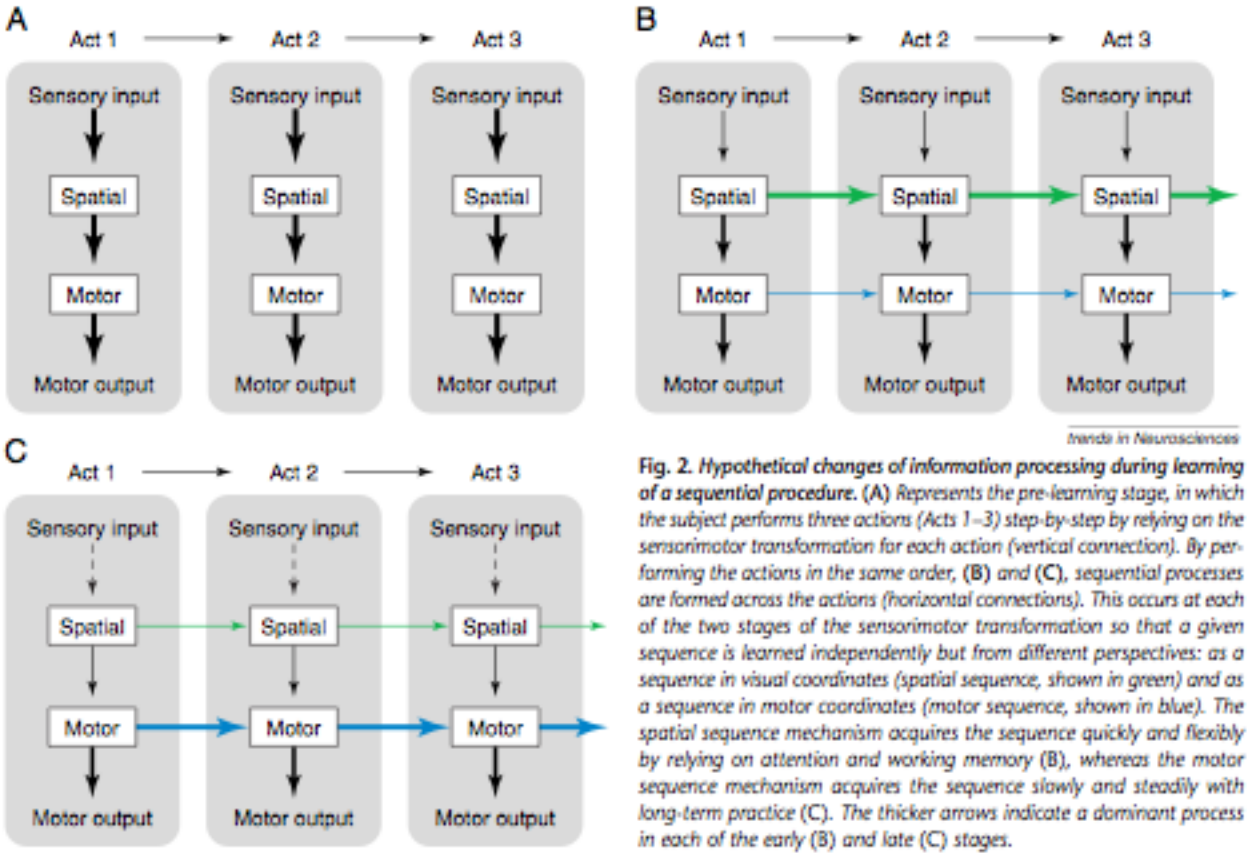

Trends in Neurosciences

Fig. 2. ***Hypothetical changes of information processing during learning of a sequential procedure.*** **(A)** *Represents the pre-learning stage, in which the subject performs three actions (Acts 1–3) step-by-step by relying on the sensorimotor transformation for each action (vertical connection). By performing the actions in the same order,* **(B)** *and* **(C)***, sequential processes are formed across the actions (horizontal connections). This occurs at each of the two stages of the sensorimotor transformation so that a given sequence is learned independently but from different perspectives: as a sequence in visual coordinates (spatial sequence, shown in green) and as a sequence in motor coordinates (motor sequence, shown in blue). The spatial sequence mechanism acquires the sequence quickly and flexibly by relying on attention and working memory* (B)*, whereas the motor sequence mechanism acquires the sequence slowly and steadily with long-term practice* (C)*. The thicker arrows indicate a dominant process in each of the early* (B) *and late* (C) *stages.*

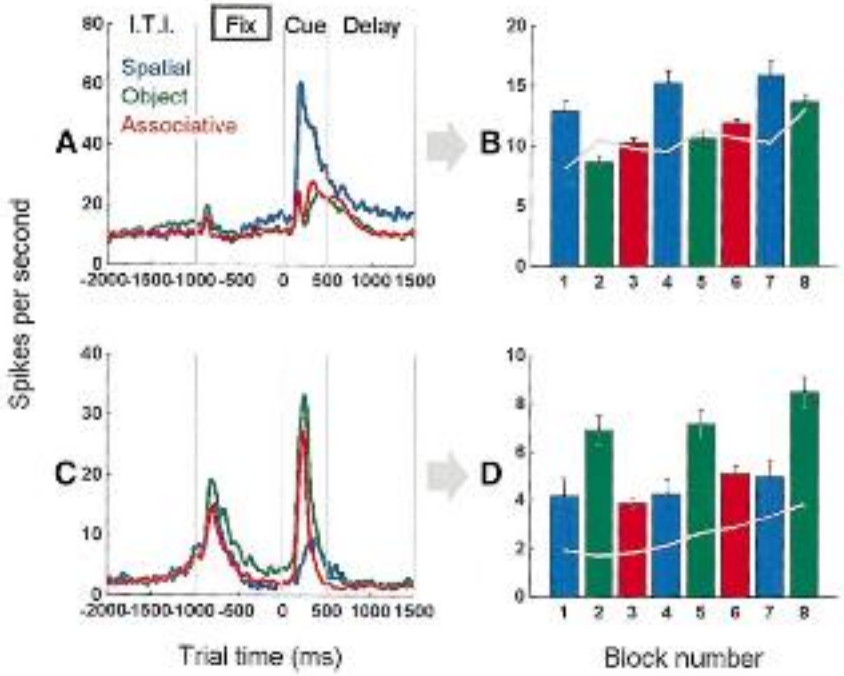
I.T.I.
Fix
Cue
Delay
Spatial
Object
Associative
A
B
C
D
Spikes per second
Trial time (ms)
Block number
-2000 -1500 -1000 -500 0 500 1000 1500
1 2 3 4 5 6 7 8

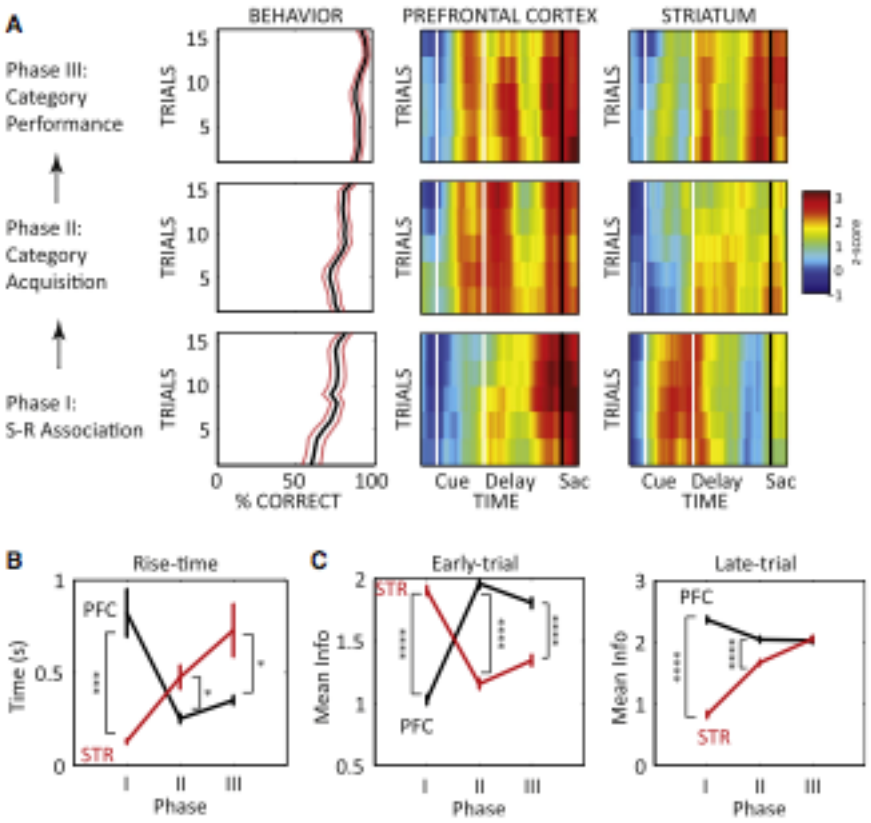
A
BEHAVIOR
PREFRONTAL CORTEX
STRIATUM
Phase III:
Category
Performance
Phase II:
Category
Acquisition
Phase I:
S-R Association
TRIALS
15
10
5
0
50
100
% CORRECT
Cue
Delay
Sac
TIME
z-score
B
Rise-time
Time (s)
PFC
STR
I
II
III
Phase
C
Early-trial
Late-trial
Mean Info

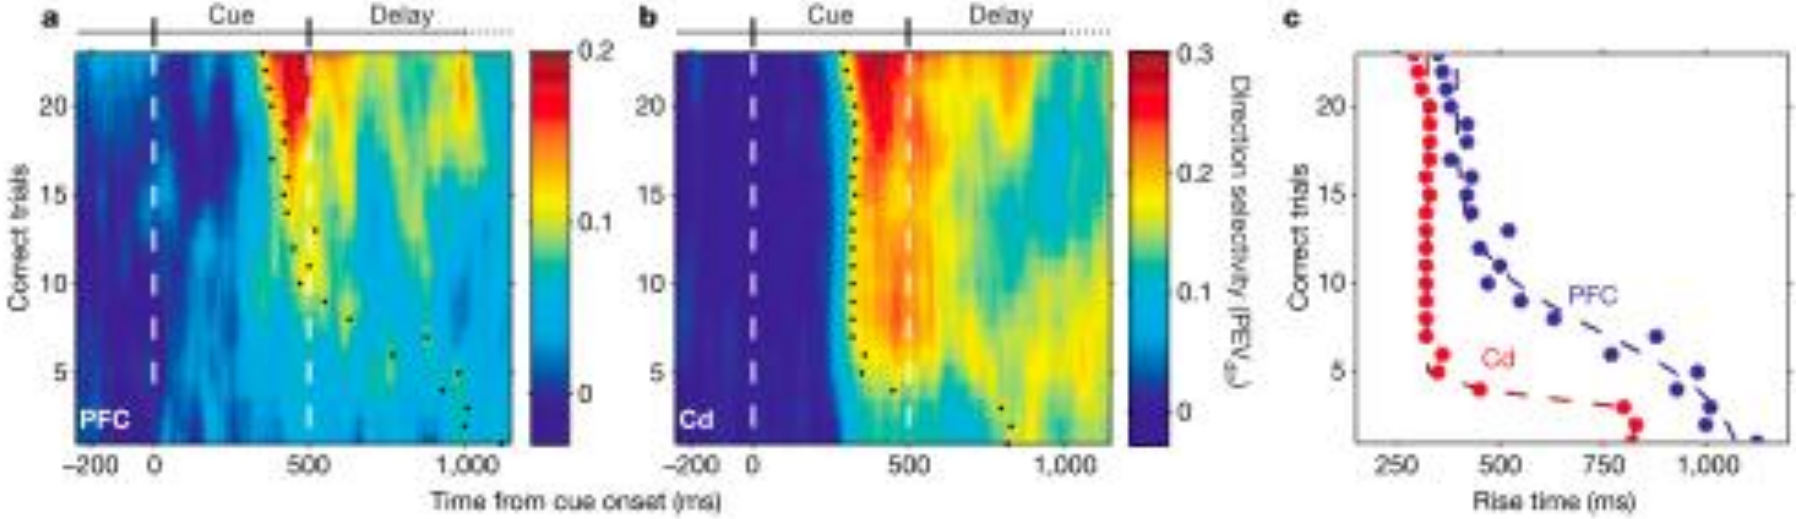

Figure 2 Change in peri-cue saccade direction selectivity in prefrontal cortex and caudate nucleus with learning. **a**,**b**, Population strength of direction selectivity (PEV_{dir}: proportion of explainable variance by direction factor) (colour scale) shown as a function of correct trials and time from cue onset for PFC (**a**) and caudate nucleus (Cd) (**b**) during cue (white lines) and delay periods. Black dots indicate 'rise time' (time to half-maximum selectivity). Selectivity strength increases and appears earlier in both areas as learning takes place. Changes appear earlier and reach an asymptote sooner in the caudate nucleus than the PFC. **c**, Rise times for PFC (blue) and Cd (red). Dotted lines show sigmoids of best fit. Data shown in **a**–**c** are based on correct trials collapsed across all blocks (reversals).

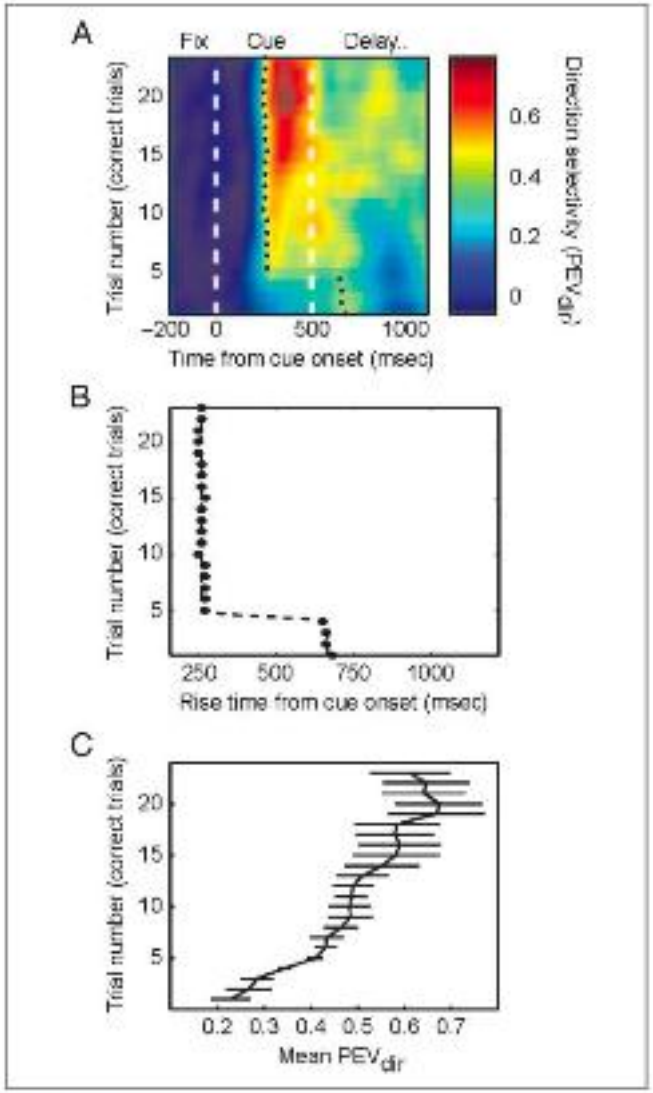

Figure 3. (A) Change in peri-cue direction selectivity during association learning without reversals. Population percentage of variance explained by direction (PEV_{dir}, color scale) shown as a function of correct trials and time from cue onset, averaged across blocks and cues. Black dots indicate the half-maximal PEV_{dir} or rise times. (B) Rise times replotted and fit with a sigmoid curve show bistable learning, with initial late activity (Trials 1–4) followed by an increase in early trial direction selectivity starting with Trial 5. (C) Mean PEV_{dir} from the second half of the cue period (250–500 msec after cue onset) also shows a jump by Trial 5 but continues to increase with learning. Error bars show standard deviation of the mean.

Non-PFC cognitive functions

- Categorization w/o PFC (Minamimoto et al., 2010)
  - Seger et al.
- Categorization in Temporal and Parietal
  - Striatum, too (Miller et al.)
- Striatum in reversal learning (Pasupathy & Miller, 2005)
- Premotor in abstract rule task (Wallis & Miller, 2003)
- PPC in WCST analog (Kamigaki et al., 2009)
  - And SMC and ACC (lesion; Buckley, et al., 2009)

LIP as ‘Behavioral Priority Map’

Multi-modal, spatially-tuned; bottom-up saliency; covert attention; intentional signals; reward considerations: reward likelihood, probability, preference, magnitude; quality of decision variables; attention; saccadic planning; categorical judgments; numerosity; task states; top-down inhibition; social gaze cues; or anything really (including color)!

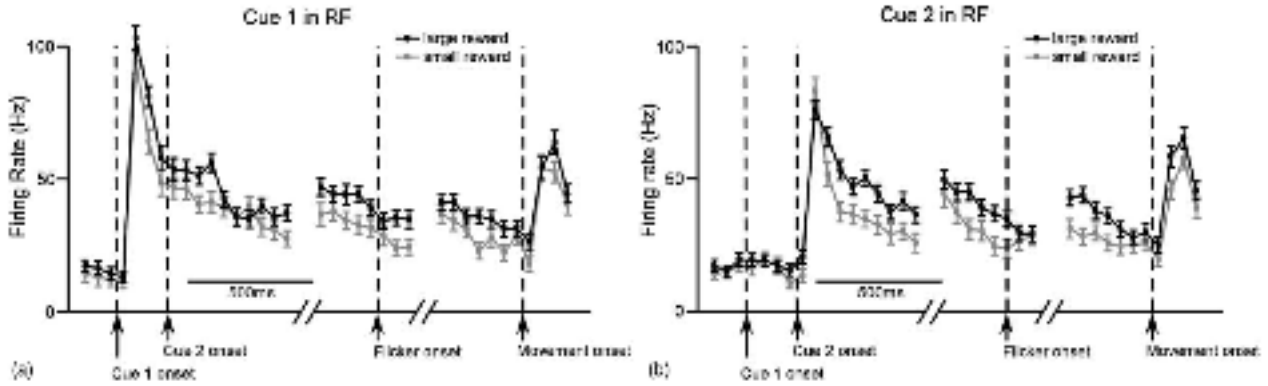

Fig. 3. Sustained visual responses in LIP are modulated by reward size. Firing rates for a single LIP neuron (±S.E.) plotted in 50 ms bins aligned on the cue onsets, flicker onset, and movement onset (left to right, indicated by arrows). Black curves: large reward blocks; grey curves: small reward blocks. (a) Cue 1 in response field. (b) Cue 2 in response field.

??

- Match vs. non-match responses (miller et al. 1996)
- ‘shift’ in activations (Cromer et al., 2011; Histed et al., 2009; Pasupathy & Miller, 2005)
  - See Hikosaka et al 1999
- Task set

??

- 2nd next neurons
- Kennerley & Wallis, 2009: VLPFC vs. DLPFC
  - VLPFC earlier, stronger and more sustained
  - And see VLPFC as task set related
- Buckley et al. shows different effects of lesions during WCST analog

Interactive Processes

- Premotor→Prefrontal
- Wallis & Miller, 2003: OFC and DLPFC show similar reward modulation
  - DLPFC shows conjunction of reward expectancy and response
  - OFC shows earlier reward response

- Use of new technologies in non-human primates
- EEG, ERP, PET in apes

Non-human Primate Social Behavior

Hunt in groups (with tools); share meat, fruit, plants, and tools; respond to fiat currency; respond to unequal pay; cooperate for resources, access to mates, etc.; compete for resources, mates, space, etc.; engage in deception; vocalize differentially to stimuli; vocalize differentially to cognitive states; communicate flexibly with gesture (including in sequence bouts); pantomime or use iconic gestures; monitor comprehension of others

Non-human Primate Social Behavior

Place high value on social information; attend to others' attention and follow gaze (especially of high-status individuals); infer or understand dominance relations; differentially process intentional versus accidental actions; understand actions not in one's own motor repertoire; attribute cognitive states to others; reciprocate, act prosocially or otherwise respond according to past behavior of others

Non-human Primate Social Behavior

show aggression (according to complex considerations); reconcile after conflict; form short- and long-term bonds; benefit positively from social bonds; maintain social bonds through grooming, or policing functions; sensitive to emotional state of others; process vicarious reinforcement; learn complex skills from others; have primitive cultures!

Non-human Primate Social Behavior

- But compared to humans:
  - Cooperate and act prosocially less
  - Communicate less flexibly
  - Fail false-belief tests with more frequency
  - Achieve lower scores on comparative social learning tests, or show different response profiles (over-imitation)
  - Lack ‘shared intentionality’
  - And recall ‘Ai’
- Also, distinguish ‘active’ and ‘passive’ social learning

Social Neuroscience

- Responses from DLPFC, OFC, SMC, ACC, LIP, caudate, along with F5, AIP, STS, etc. ‘mirror’ regions
  - Yoshida et al. ‘other’ neurons
  - Lee et al. DLPFC neurons monitoring history of choices
- Novel designs incorporating computer or monkey opponents
  - ‘pennies’; ‘paper, scissors, rock’; team-based; mixed-strategy
- ERP, EEG, PET along with DTI, MRI and other anatomic and cytoarchitectonic methods
  - Can use synthetic techniques to relate to more direct data
- Selective lesioning an effective method too
- Tsujimoto et al. (2009) FPC recordings

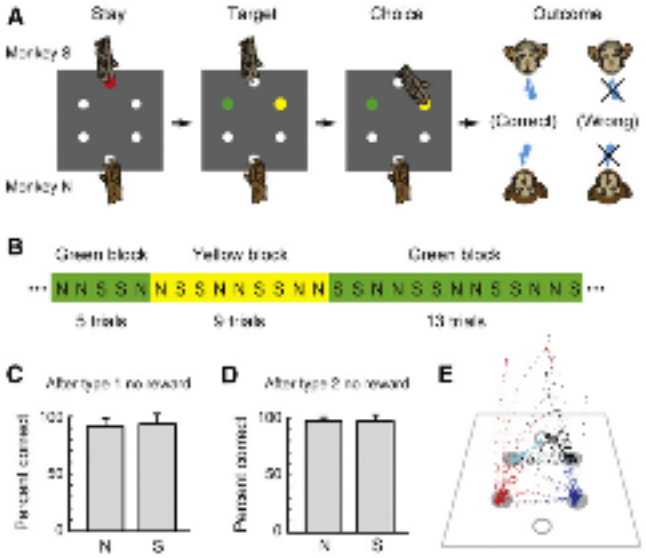

Figure 1. Behavioral Task and Animals' Choice Behavior

(A) Temporal sequence of events in role-reversal task. Shown is an example of a single trial in which monkey S was the actor and monkey N was the observer. The reward outcome differed depending on the color-reward contingency ("Outcome").

(B) In the task, two monkeys alternated in the role of the actor every two trials, and the color-reward contingency switched unpredictably every 5–17 trials (blocked design). "Green block" means that the green target was associated with a reward. N and S indicate the acting monkey.

(C) Percent correct choice in trials immediately after partner's type 1 no reward. Error bars indicate standard error of the mean (SEM). See Figure S1A for further explanation.

(D) Percent correct choice in trials immediately after partner's type 2 no reward. Error bars indicate SEM. See Figure S1B for further explanation.

(E) Gaze directions of monkey N (near side) during action period. The start buttons are depicted as open gray circles and the target buttons as filled gray circles. These color codes are for illustrative purposes only; during actual experiments, the target buttons were illuminated in green and yellow. Red dots indicate gaze directions when monkey N reached for his left target (as actor), whereas dark blue dots indicate gaze directions when monkey N reached for his right target. Light blue dots indicate gaze directions of monkey N (as observer) when his partner reached for the left target (viewed from monkey N), whereas black dots indicate gaze directions when his partner reached for the right target (viewed from monkey N).

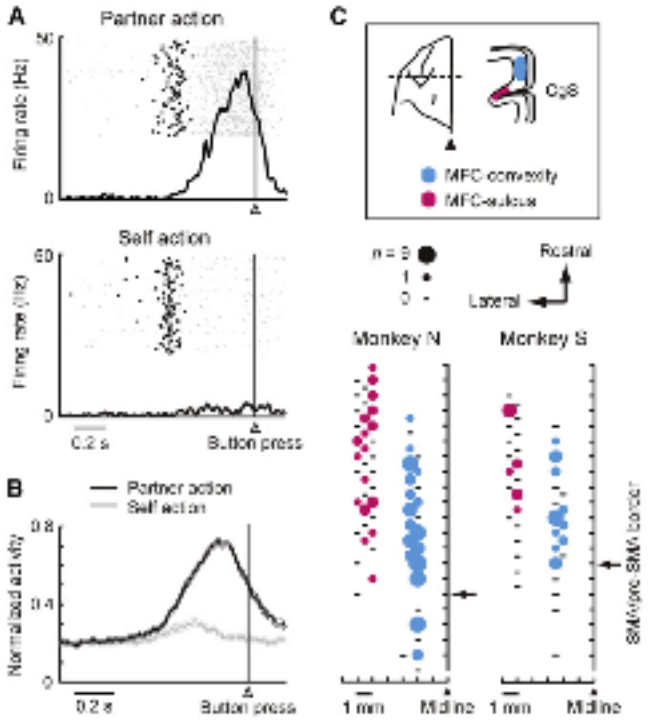

Figure 2. Firing Property of Partner-Type Neurons and Their Spatial Distribution in Medial Frontal Cortex

(A) Raster displays and spike density functions illustrating activity of a single partner-type neuron. Smaller dots represent the time of individual action potentials, and larger dots indicate the time of target onset. The displays are aligned on the onset of target button press (vertical lines and gray triangles).

(B) Spike density functions for population of partner-type neurons (mean ± SEM).

(C) Recording sites of partner-type neurons (left hemispheres). Neurons in the medial frontal cortex (MFC) convexity are depicted in blue, and those in the MFC sulcus are depicted in red (see inset). The size of the circles is proportional to the number of neurons at each site. Horizontal arrows indicate the physiologically defined border between the supplementary motor area (SMA) and pre-SMA. Black triangles represent the cortical midline. The inset above shows the top view of the left frontal lobe (left) and a coronal section (right) at the rostrocaudal level indicated by a broken line. CgS indicates cingulate sulcus.

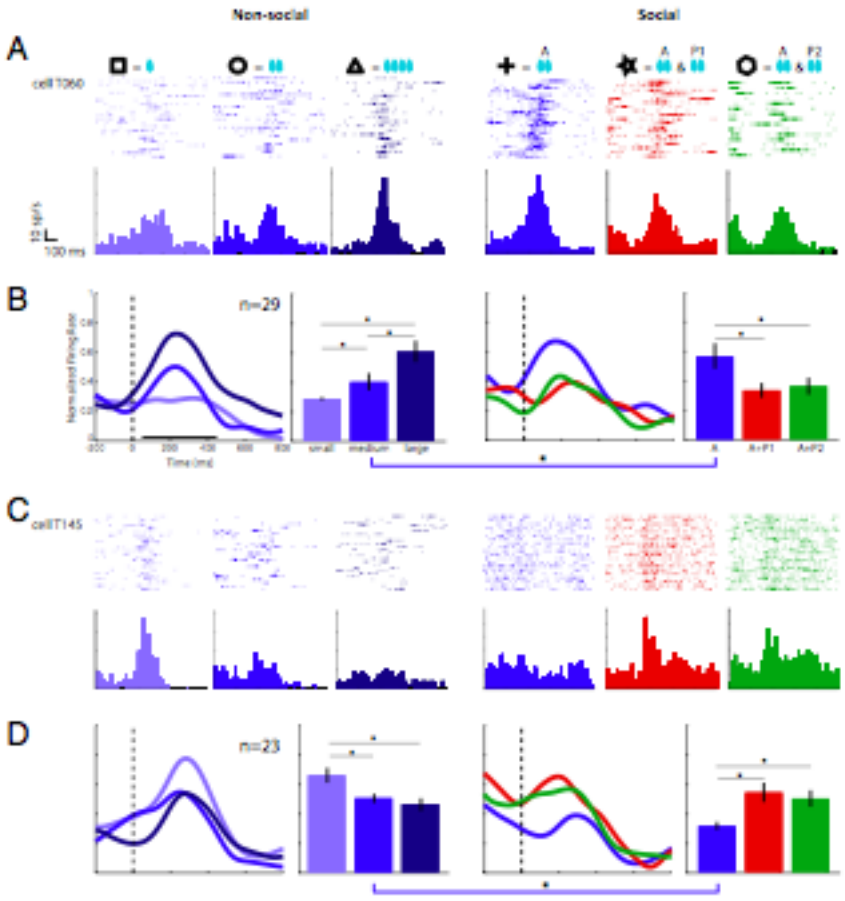

Fig. 3. Single-unit and population activity of OFC neurons in the *nonsocial* and *social* blocks. (*Left* three panels in *A* and *C*) Examples of single units with, respectively, up-modulated and down-modulated responses as a function of reward size; (*Right* three panels) Activity of the same two neurons in the social block. (*B* and *D*) Normalized spike density curves and mean discharge rate for the neuron population from which the corresponding single-unit examples are drawn. The thick horizontal bar below the spike density curves in *B* indicates the time window used for computing all statistical tests on mean population activity (*Right*). The asterisks and thin black or blue horizontal lines indicate significant pair-wise comparisons ($P < 0.01$).

Social Neuroscience

- Subiaul et al., 2004 claims to show ‘cognitive imitation’
  - We claim expertise on task greatly facilitates observational learning; not ‘imitation’
- Animals must link responses from others to responses in self, according to analysis of its outcome
- Fictive reward processing, vicarious reinforcement

Modeling

- Hierarchical Learning
  - Braun et al.. Propose Hierarchical Bayesian method
  - Botvinick, Barto and others propose HRL
  - Also, how to consolidate motor programs?
- Category formation
  - Hinaut & Dominey, 2011
  - Rougier et al., 2005

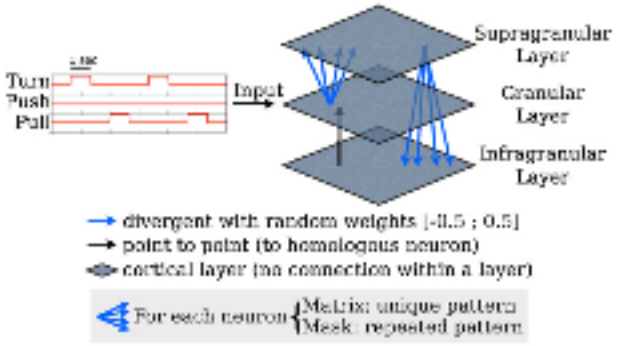

Fig. 3. Three-layered cortical model. Each layer corresponds to a 5 × 5 matrix of leaky integrator neurons, with the indicated connectivity. Divergent connections (in blue) can be defined by the matrix connectivity or by the mask connectivity (see Fig. 2). The three-dimensional input (sequence of movements) is presented to the granular layer. The activity produced by the inputs is propagated through the recurrent connections forming a pattern specific to each sequence.

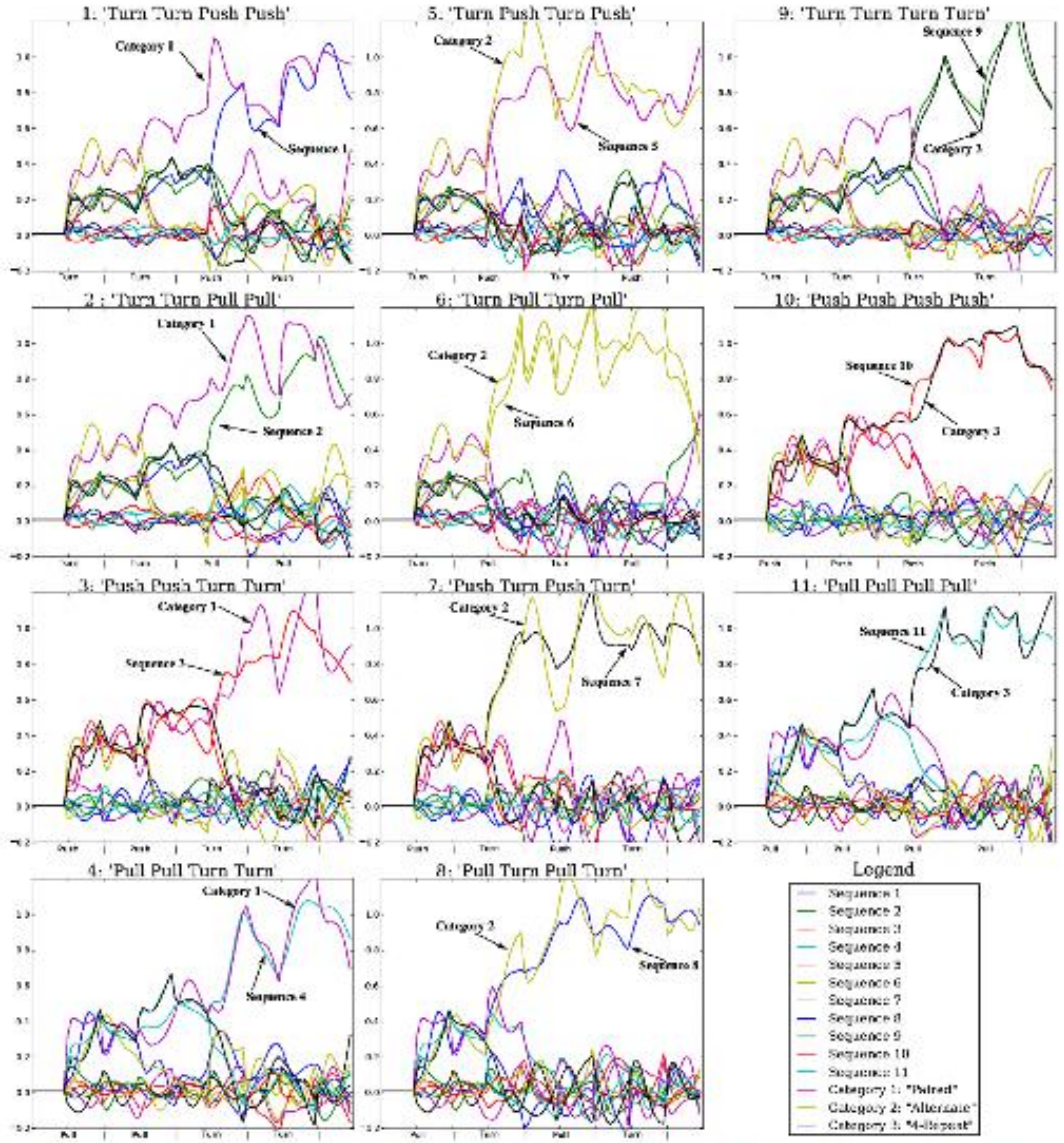

Fig. 9. Activity of the read-out layer (striatum model) after learning. The teacher output "since 1st stimulus" was used here to force the network to answer as quickly as it can. One can see the activity in the readout layer (corresponding to the striatum) for sequences 1–11. For each signal one can see a "pseudo" probability of which sequence and category the network predicts. We can see that as soon as the network has enough input information to discriminate between 11 sequences, it shows a high activity for the corresponding sequence and category. In the first column are the sequences of category "paired", "alternate" sequences are in the second column, and "4-repeat" sequences are in the third column.

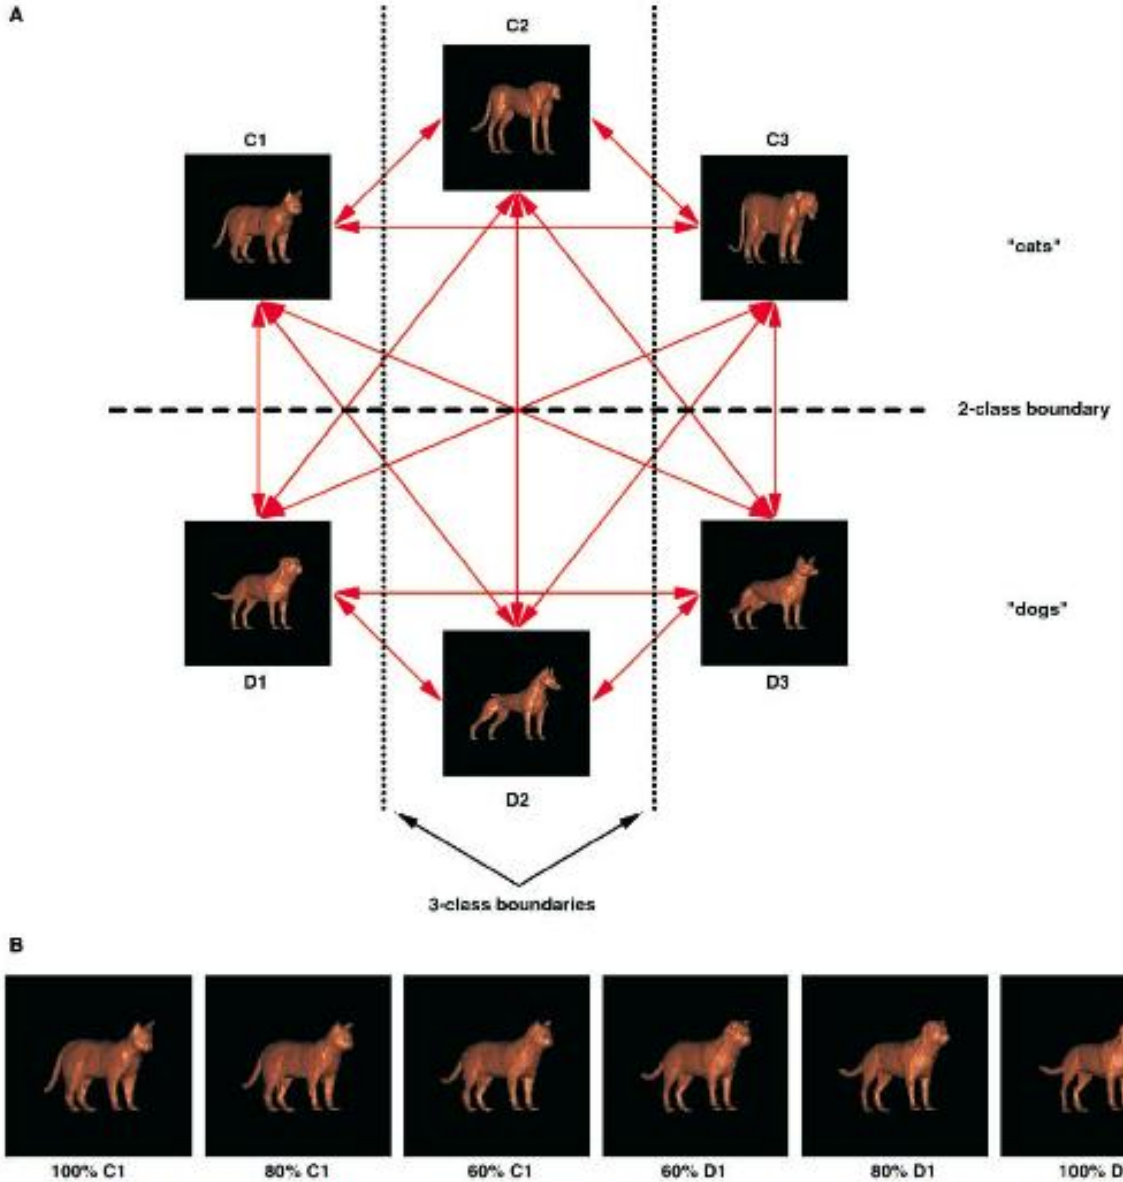

Fig. 1. The stimuli. (**A**) Monkeys learned to categorize randomly generated "morphs" from the vast number of possible blends of six prototypes. For neurophysiological recording, 54 sample stimuli were constructed along the 15 morph lines illustrated here. The placement of the prototypes in this diagram does not reflect their similarity. (**B**) Morphs along the C1-D1 line.

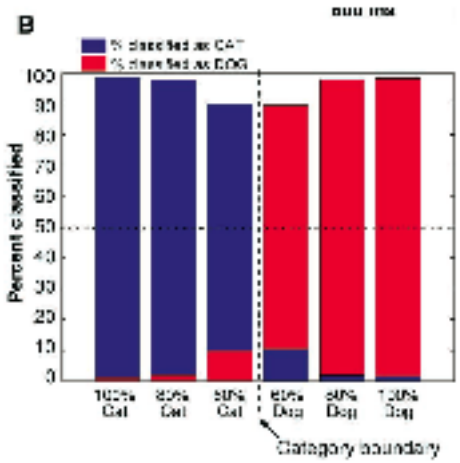

Fig. 3. Recording locations and single neuron example. (**A**) Recording locations in both monkeys. A, anterior; P, posterior; D, dorsal; V, ventral. There was no obvious topography to task-related neurons. (**B**) The average activity of a single neuron in response to stimuli at the six morph blends. The vertical lines correspond (from left to right) to sample onset, offset, and test stimulus onset. The inset shows the neuron's delay activity in response to stimuli along each of the nine between-class morph lines (see Fig. 1). The prototypes (C1, C2, C3, D1, D2, and D3) are represented in the outermost columns; each appears in three morph lines. A color scale indicates the activity level.

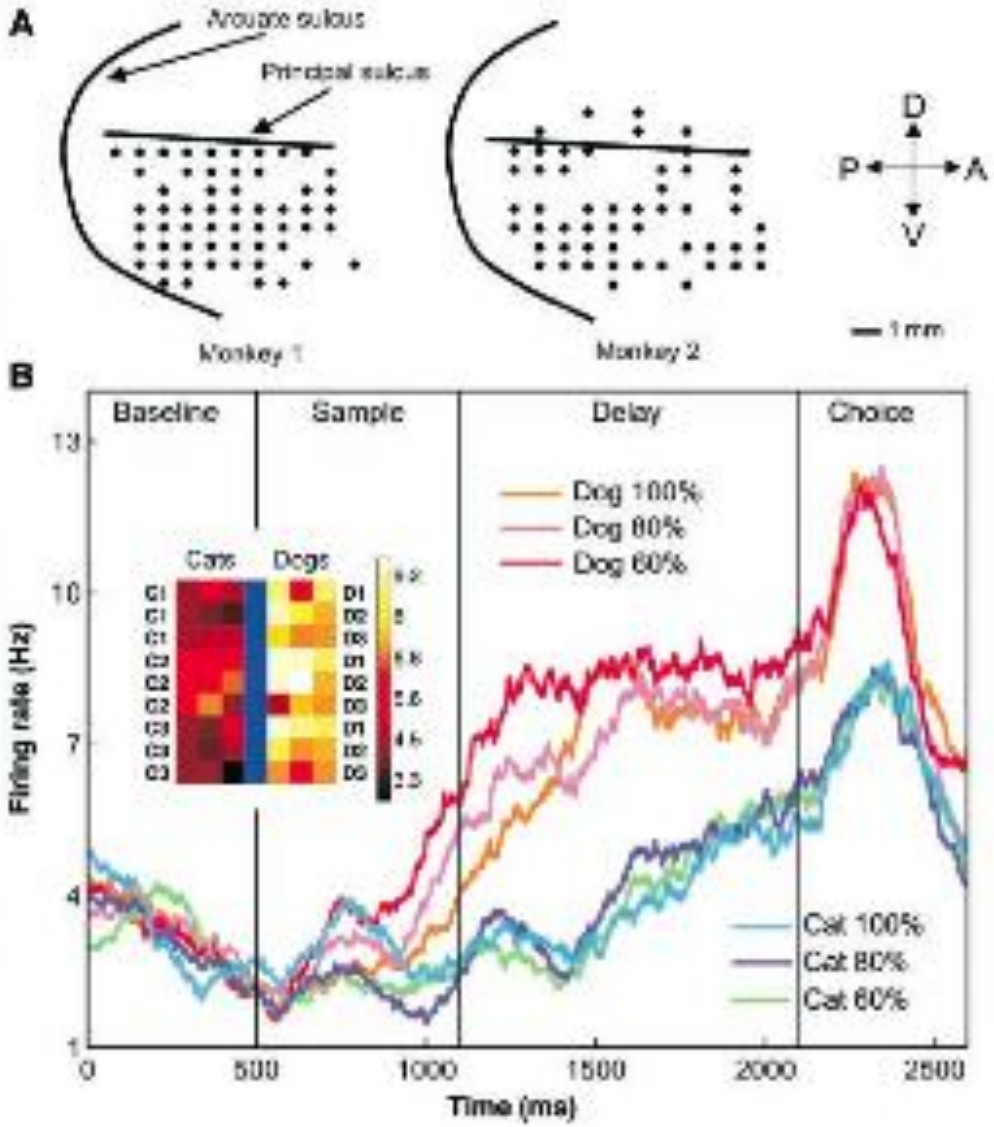

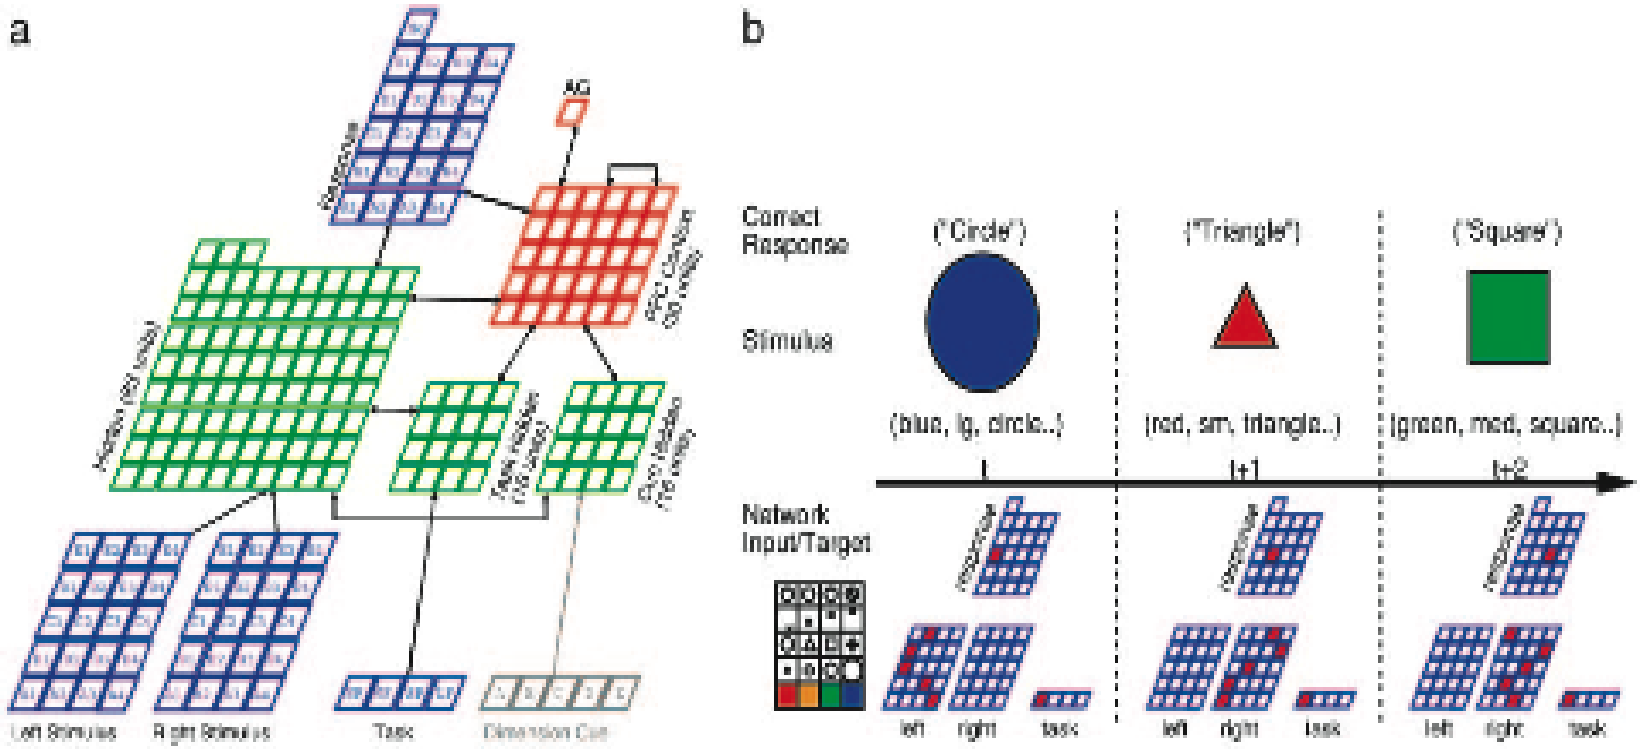

Fig. 1. Model and example stimuli. (*a*) The model with the complete PFC system. Stimuli are presented in two possible locations (left, right). Rows represent different stimulus dimensions (e.g., color, size, shape, etc., labeled A–E for simplicity), and columns represent different features (red, orange green, and blue; small, medium, etc., numbered 1–4). Other inputs include a task input indicating current task to perform (NF, name feature; MF, match feature; SF, smaller feature; LF, larger feature), and, for the "instructed" condition (used to control for lack of maintenance in non-PFC networks), a cue to the currently relevant dimension. Output responses are generated over the response layer, which has units for the different stimulus features, plus a "No" unit to signal nonmatch in the matching task. The hidden layers represent posterior cortical pathways associated with different types of inputs (e.g., visual and verbal). The AG unit is the adaptive gating unit, providing a temporal differences (TD) based dynamic gating signal to the PFC context layer. The weights into the AG unit learn via the TD mechanism, whereas all other weights learn using the Leabra algorithm that combines standard Hebbian and error-driven learning mechanisms, together with k-winners-take-all inhibitory competition within layers and point-neuron activation dynamics (26) (also see supporting information). (*b*) Example stimuli and correct responses for one of the tasks (NF) across three trials where the current rule is to focus on the Shape dimension (the same rule was blocked over 200 trials to allow networks plenty of time to adapt to each rule). The corresponding input and target patterns for the network are shown below each trial, with the unit meanings given by the legend in the lower left. The network must maintain the current dimension rule to perform correctly.

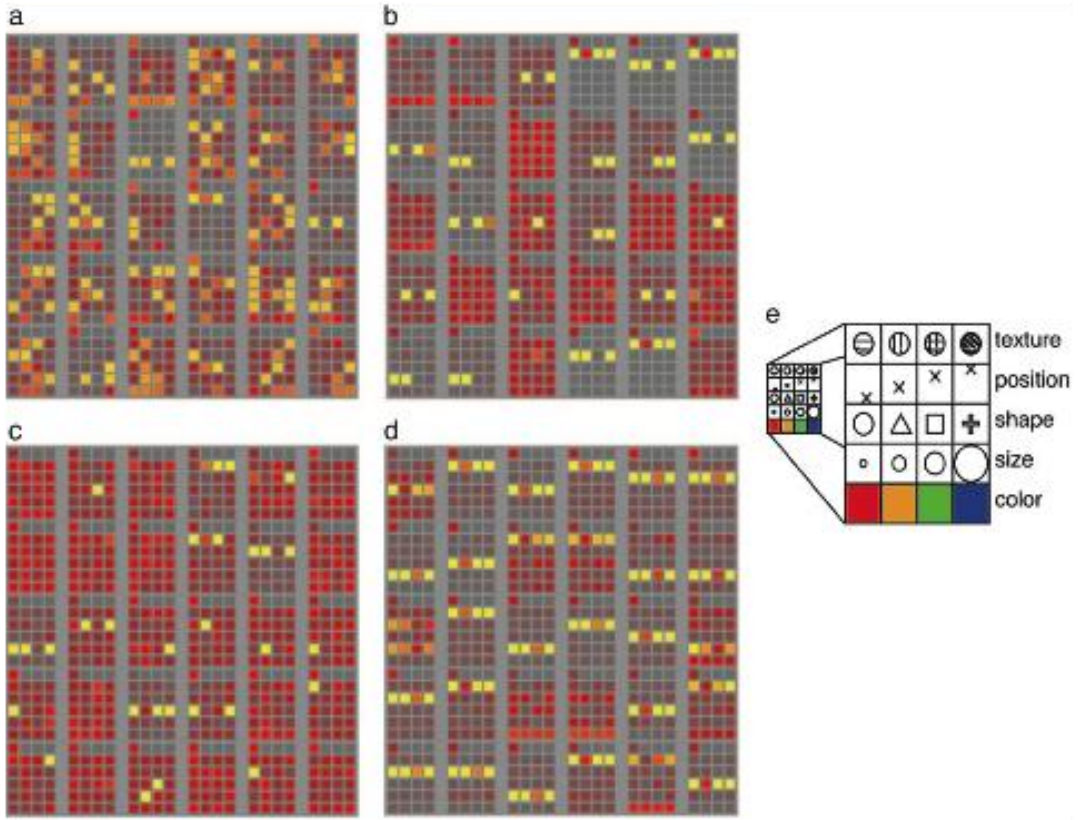

Fig. 2. Representations (synaptic weights) that developed in four different network configurations. (*a*) Posterior cortex only (no PFC) trained on all tasks. (*b*) PFC without the adaptive gating mechanism (all tasks). (*c*) Full PFC trained only on task pairs (name feature and match feature in this case). (*d*) Full PFC (all tasks). Each image shows the weights from the hidden units (*a*) or PFC (*b*–*d*) to the response layer. Larger squares correspond to units (all 30 in the PFC and a random and representative subset of 30 from the 145 hidden units in the posterior model), and the smaller squares within designate the strength of the connection (lighter = stronger) from that unit to each of the units in the response layer. Note that each row designates connections to response units representing features in the same stimulus dimension (as illustrated in *e* and Fig. 1). It is evident, therefore, that each of the PFC units in the full model (*d*) represents a single dimension and, conversely, that each dimension is represented by a distinct subset of PFC units. This pattern is less evident to almost entirely absent in the other network configurations (see text for additional analyses).

Modeling Challenges

- ‘Deep-time’ / trans-interval modeling
  - Recall Campos et al., Bernacchia et al., Saito et al.
- Reward modulations
  - Much deeper issue, see social brain modeling
  - Outcome processing, value updating, etc.
- Conjunctive coding
- Activation-based versus weight-based (O’Reilly, 1998; O’Reilly & Munataka, 2000)
  - Learning, reward expectancies, etc. encoded in rate
- Social learning (including theoretical issues)
  - E.g., Acerbi et al., 2011
- Social brain modeling
  - Perhaps carries different challenges

BOPs

- Change matters
- Conjunctive coding
  - Multi-plexing / collapsing representation
- Reward modulation
  - Across long time scales (Bernacchia et al., 2011)
- Adaptive Coding
  - LIP responds to color
- Consolidation / learning-related changes
- Interactive activation
- Prospective and reactive encoding
- Task sets / hierarchical organization of control

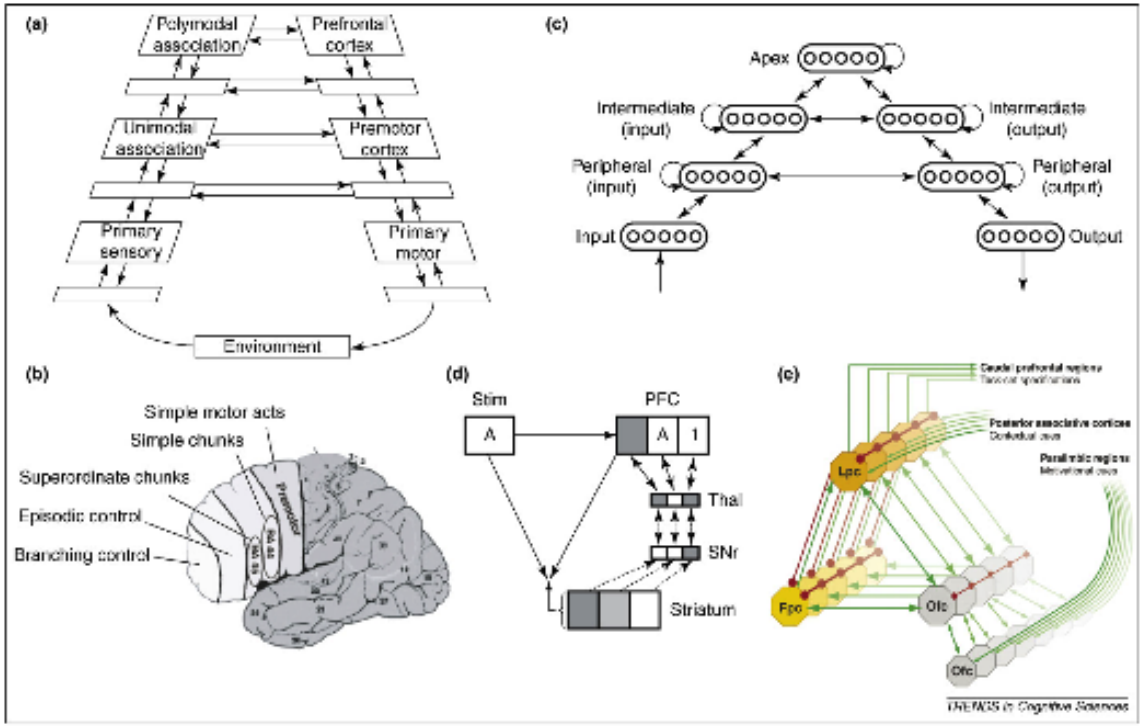

Figure 3. Hierarchical organization in frontal cortex. **(a)** The position of the DLPFC within a hierarchy of cortical areas, as described by Fuster [45]. **(b)** Levels of control represented in different sectors of frontal cortex, according to Koechlin [9]. Representations become progressively more abstract towards the rostrum. **(c)** The hierarchically structured network studied by Botvinick [35], showing only a subset of units in each layer. Arrows indicate all-to-all connections. When trained on a hierarchically structured task, units in the apical group spontaneously come to represent context information more strongly than do groups further down the hierarchy. **(d)** Schematic of the gating model proposed by O'Reilly and Frank [14], during performance of a task requiring maintenance of the stimuli '1' and 'A' in working memory. At the point shown, a '1' has already occurred and has been gated into a prefrontal (PFC) stripe via a pathway through the striatum, substantia nigra (SNr) and thalamus (thal). At the moment shown, an 'A' stimulus occurs (Stim) and is gated into another PFC stripe. Two levels of context are thus represented. **(e)** Koechlin's [37] model of FPC function. Orbitofrontal cortex (Ofc) encodes the incentive value of various tasks. When two tasks are both associated with a high incentive value, the one with the highest value is selected within lateral PFC (Lpc) for execution, while the runner-up is held in a pending state by the frontopolar cortex (Fpc). Part (a) reprinted, with permission, from Ref. [45] (http://www.sciencedirect.com/science/journal/08966273); part (b) reprinted, with permission, from Ref. [9] (www.oup.com); part (c) reproduced, with permission, from Ref. [35] (http://publishing.royalsociety.org/); part (d) adapted, with permission, from Ref. [14] (http://mitpress.mit.edu); part (e) reprinted, with permission, from Ref. [37].

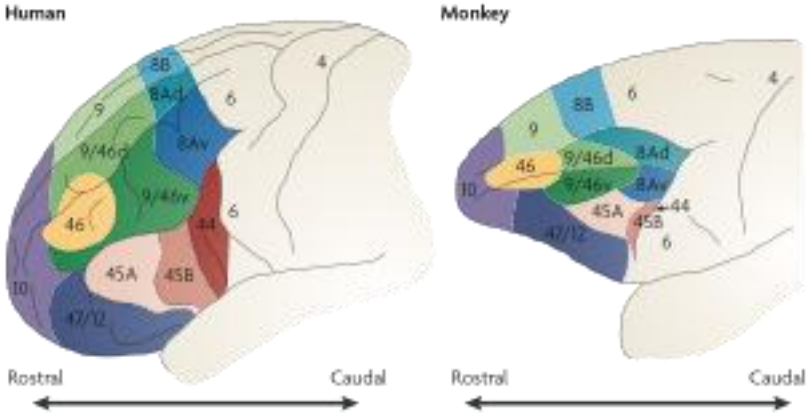
Human
8B
8Ad
8Av
9
9/46d
9/46v
46
44
6
4
45A
45B
10
47/12
Rostral
Caudal
Monkey
6
4
8B
9
46
9/46d
8Ad
9/46v
8Av
10
45A
44
45B
47/12
6
Rostral
Caudal

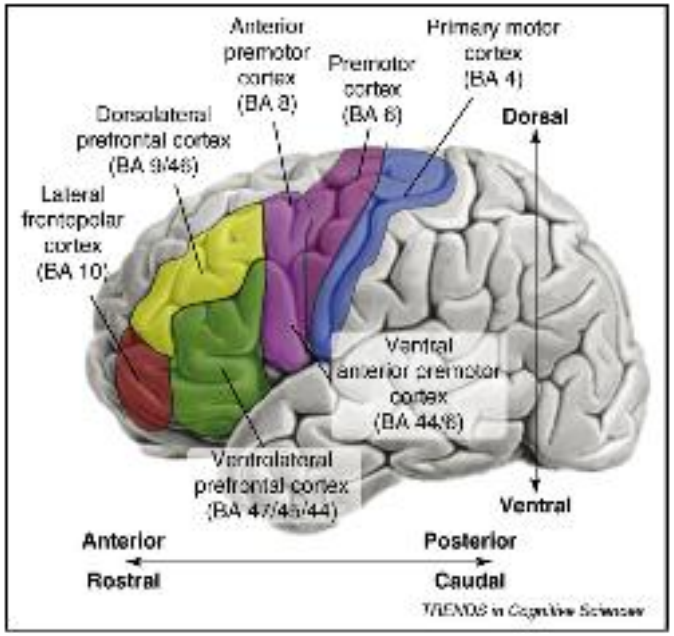

Figure 1. Schematic of major anatomical sub-divisions in the frontal lobes. Boundaries and Brodmann areas (BA) are only approximate. Arrows indicate anatomical directions of anterior/rostral (front) versus posterior/caudal (back) and dorsal (up) versus ventral (down). From caudal to rostral, labeled areas include motor cortex, dorsal (PMd) and ventral premotor cortex, dorsal (pre-PMd) and ventral aspects of anterior premotor cortex, ventro- (VLPFC) and dorsolateral PFC (DLPFC), and lateral frontal polar cortex (FPC).

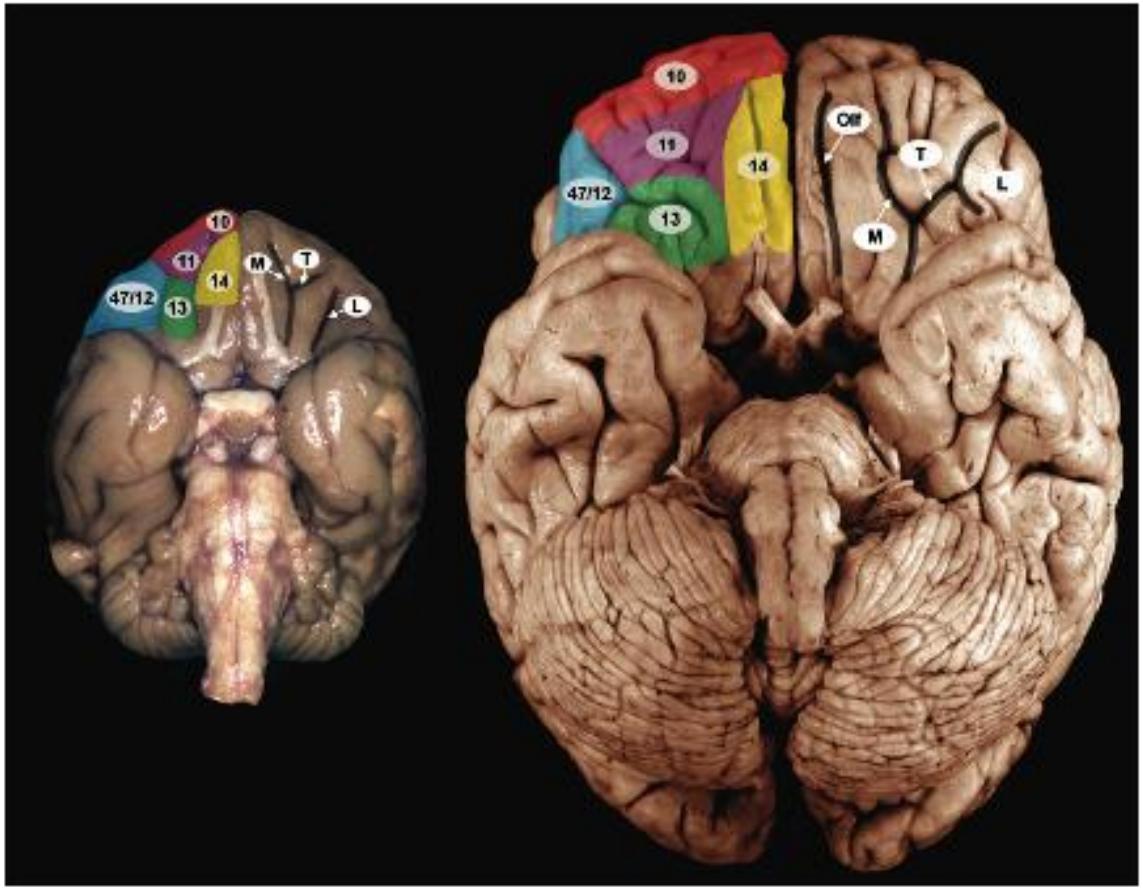
10
11
14
47/12
13
M
T
L
10
11
14
47/12
13
Orb
T
L
M

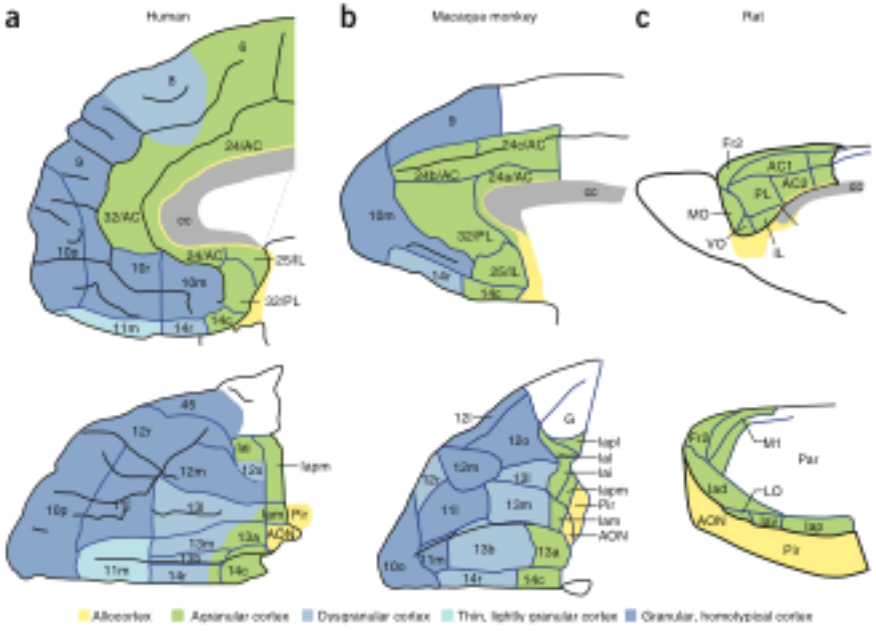

Allocortex Agranular cortex Dysgranular cortex Thin, lightly granular cortex Granular, homotypical cortex