

Published in final edited form as:

Neuron. 2013 December 18; 80(6): 1384–1391. doi:10.1016/j.neuron.2013.09.028.

Causal evidence of performance monitoring by neurons in posterior cingulate cortex during learning

Sarah R. Heilbronner¹ and Michael L. Platt^{1,2}

¹Duke University Department of Neurobiology, Center for Cognitive Neuroscience

²Departments of Anthropology and Psychology and Neuroscience

Summary

The posterior cingulate cortex (CGp) is a major hub of the *default mode network* (DMN), a set of cortical areas with high resting activity that declines during task performance. This relationship suggests DMN activity contributes to mental processes that are antagonistic to performance. Alternatively, DMN may detect conditions under which performance is poor and marshal cognitive resources for improvement. To test this idea, we recorded activity of CGp neurons in monkeys performing a learning task while varying reward size and novelty. We found that CGp neurons responded to errors, and this activity was magnified by small rewards and novel stimuli. Inactivating CGp with muscimol impaired new learning when rewards were small but had no effect when rewards were large; inactivation did not affect performance on well-learned associations. Thus, CGp, and by extension the DMN, may support learning, and possibly other cognitive processes, by monitoring performance and motivating exploration.

Keywords

default mode network; learning; decision-making; performance; cingulate

Introduction

The default mode network (DMN) is a poorly understood set of brain regions with high baseline activity, which is suppressed during task engagement (Gusnard et al., 2001; Raichle et al., 2001; Shulman et al., 1997). In neuroimaging experiments, its component nodes—chiefly the posterior cingulate cortex (CGp) and the ventromedial prefrontal cortex (vmPFC)—show high levels of correlated spontaneous activity at rest (Vincent et al., 2007) and correlated patterns of activation and deactivation. Surprisingly, heightened DMN activity predicts poor performance in a variety of laboratory tasks (Hayden et al., 2009b; Weissman et al., 2006). One explanation for this relationship is that the DMN mediates cognitive processes, such as mind-wandering (Mason et al., 2007), introspection, and temporal projection of the self (Buckner and Carroll, 2007; Spreng and Grady, 2010), that are antagonistic to task performance. While this “antagonistic to cognition” interpretation is intuitive, an alternative is that DMN activation marshals neural resources to improve poor

© 2013 Elsevier Inc. All rights reserved.

Address for correspondence: Sarah R. Heilbronner, Department of Pharmacology and Physiology, University of Rochester Medical Center, 601 Elmwood Ave, Box 711, Rochester, NY 14642, sarah.heilbronner@gmail.com.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

task performance. Such an interpretation is consistent with the idea that the entire cingulate cortex detects the need for enhanced cognitive engagement and recruits additional resources (Botvinick et al., 2004; Pearson et al., 2009). This model predicts high DMN activity when performance is poor – but as an outcome, rather than as a cause, of poor performance.

There is some evidence supporting this possibility. CGp is often particularly active in situations that require high levels of cognitive engagement. Such situations include uncertain or conflict-rich environments where learning, adjustment, and executive control are advantageous. For example, CGp firing rates increase with reward uncertainty (McCoy and Platt, 2005) and with cognitively demanding exploration of alternatives (Pearson et al., 2009). Likewise, in a multi-attribute decision-making study, firing rates tracked decision salience (difference from standard), when the need for engagement was high (Heilbronner et al., 2011). Furthermore, microstimulation of CGp neurons causes monkeys to switch to an alternative option (Hayden et al., 2008), demonstrating a causal role in exploration and behavioral switching.

Together, these findings endorse the idea that CGp surveys the internal and external environment for evidence that the current model of the world is no longer working, thus motivating adaptive changes in cognition and behavior (Pearson et al., 2011). CGp, as part of the DMN, may marshal cognitive resources in difficult situations, when performance would otherwise be even worse. To test these ideas, we recorded the activity of single CGp neurons in two monkeys performing a conditional association task (cf. Halsband and Passingham, 1982; Petrides, 1982; Chen and Wise, 1995; Wirth et al., 2003) in which visual scenes were randomly paired with a saccade target. Effective performance required exploration and adaptation, an ideal probe for CGp function. Monkeys not only had to learn new scenes, but, on randomly interleaved trials, had to recall the actions associated with reference scenes used every day, thus probing recall and motor performance. Scenes were associated with either large or small rewards.

We found that neurons responded phasically to errors, and this signal was enhanced when scenes were new or rewards were small—situations where performance was poor. In addition, higher firing rates were associated with worse learning and performance. Importantly, reversible inactivation of CGp with the GABA_A agonist muscimol impaired learning in the most challenging condition, when monkeys were presented with novel, low-value scenes. This finding is inconsistent with the idea that CGp activity antagonizes task performance. Instead, increased activity in CGp, and perhaps the DMN more generally, may signal poor performance and help marshal the strategic recruitment of cognitive resources. In this task, such recruitment improves learning; in different scenarios (such as the reference scenes in this task), it may contribute to other cognitive processes.

Results

Learning varied with stimulus novelty and value

We used a conditional association task to elicit stimulus-action-outcome learning over the course of dozens of trials. On each trial, a photograph (scene) and two targets appeared on a computer monitor (Figure 1A). The scene indicated that a gaze shift to one target would be rewarded. This association could only be learned through trial-and-error. Incorrect choices were followed by no reward. Three types of scenes were interleaved each day: *reference*, *novel*, and *addition*; within each type, there were *high* and *low* value (juice reward) scenes (see Figure 1B). Reference scenes were held constant from session to session. Four novel scenes were introduced at the start of each session. For a subset of sessions, four new scenes, hereafter referred to as additions, were introduced after several hundred trials.

Monkeys performed well above chance (Figure 1C). For novel and addition scenes, performance was at chance on the first ten presentations (Novel: $M=50.9\%$ correct, $t(427)=8$, $p=.43$; Additions: $M=52.8\%$; $t(270)=1.83$, $p=.07$), but above it by the last ten presentations (Novel: $M=66.7\%$, $t(427)=12.12$, $p<.000001$; Additions: $M=61.8\%$, $t(279)=7.17$, $p<.000001$). Performance on each new scene presentation correlated with presentation number; thus, performance improved with time (Novel: $r=.37$, $p<.0001$; Additions: $r=.54$, $p<.0001$). Performance on reference scenes was already above chance during the first ten presentations each day ($t(425)=57.1$, $p<.000001$). Performance was higher on high-value scenes than on low-value scenes across all scene types ($p<0.0001$ in all 3 cases, Table 1). For new scenes, this difference emerged over the session: it was not present at the start (Novel: $t(211)=.43$, $p=.67$; Additions: $t(136)=.52$, $p=.60$), but was present at the end (Novel: $t(211)=5.8$; Additions: $p<.000001$; $t(136)=2.9$, $p=.004$) of the session. Thus, monkeys successfully learned the associations between stimuli, actions, and rewards, and subsequently recalled these associations with high fidelity.

CGp neurons dynamically encode errors

We recorded 140 CGp neurons (106 in Monkey L, 34 in Monkey B) as monkeys performed the conditional association task. Because learning in this task is based on trial-and-error, we hypothesized that CGp would track outcomes. Thus, we calculated the neuronal firing rates immediately following correct and incorrect choices. Figure 2B shows the average response of an example neuron aligned to the start of the outcome period (700 msec epoch beginning 600 msec after saccade offset, chosen to occur after peri-saccadic activity). This neuron showed a phasic response after errors (relative to correct trials) that lasted until nearly the start of the next trial ($t(1154)=16.3$, $p<.0001$, independent samples t-test on correct vs incorrect trial firing rates). For this neuron, errors elicited an average of 11.4 sp/s (46%) greater firing rate than correct responses. This pattern matched the population response (Figure 2A; $t(139)=4.08$, $p<.0001$: paired-samples t-test on average firing rate for correct vs error trials), which showed an average of 1.3 sp/s (17%) higher firing rate following errors compared with correct trials. 82 of the 140 cells we studied (58.5%) had significant error modulation (outcome period). Of these, 63 (76.8%) fired at higher rates following incorrect choices; the other 19 fired at higher rates following correct choices.

By start of the next trial, this modulation was no longer statistically significant (0.2 sp/sec), but still trended, $t(139)=1.75$, $p=.08$ (paired-samples t-test for correct vs error trial firing rates, first 300 msec of fixation period on subsequent trial). By the time the image appeared on the next trial, modulation by a previous error was gone, $t(139)=-.66$, $p=.51$ (paired-samples t-test for correct vs error trial firing rates, 500 msec following scene onset on subsequent trial).

If activation of CGp neurons competes with brain regions responsible for enhancing performance on a trial-by-trial basis, this error signal would be expected to be prepotent, beginning before feedback or overt choice. However, if neuronal activity in CGp facilitates performance, it should track feedback for monitoring and control purposes. We thus examined whether activity earlier in the trial (scene epoch—500 msec following scene onset) predicted performance. The error signal was not prepotent ($t(139)=.52$, $p=.6$, paired-samples t-test on average firing rate for correct vs error trials).

In this task, we varied scene novelty and value, parameters that should influence control signals. The error signal in CGp varied with stimulus novelty, and thus the need for learning. Specifically, errors on new scenes elicited higher firing rates than errors on reference scenes ($t(139)=2.19$, $p=.03$, paired-samples t-test on novel + addition scene vs reference scene errors, outcome period; Figure 2C). This was not true following correct choices ($p=.74$). The

novelty effect on the error signal was quite small (.45 sp/s). There was no effect of novelty on firing rates during scene presentation ($t(139)=-.60$, $p=.55$).

Reward size also modulated error signals (Figure 2D, $t(139)=-1.98$, $p<.05$, paired-samples t -test on high vs low value scenes, outcome period). This effect, too, was subtle, just .28 sp/s on average. The amount of juice received on correct trials did not affect firing rates ($p>.9$). Furthermore, the (missed) reward effect was limited to new (novel+addition) scenes ($t(137)=2.58$, $p=.01$). Reference scenes showed no such modulation ($t(77)=.57$, $p=.57$). There was no effect of reward size on firing rates during scene presentation ($t(139)=-.74$, $p=.46$).

Because of the small size of these effects, we performed additional single-cell analyses. Without regard to significance, 83/140 cells had higher firing rates following incorrect low value trials than incorrect high value trials (outcome period). This bias is significant (binomial test, $p=.034$). Moreover, of the 19 cells (13.5% of 140) with significant modulation, 16 had higher firing rates following incorrect low value than incorrect high value trials. This proportion is also significantly biased (binomial test, $p=.004$). Without regard to significance, 83/140 cells had higher firing rates following incorrect new (novel+addition) trials than following incorrect reference trials. This is also a significant bias (binomial test, $p=.034$). Only 11 cells had significant novelty modulation (7.9%), and this population had no bias in either direction ($p>.9$).

On the whole, CGp neurons thus appear to signal that the current level of task engagement is insufficient to perform optimally, due to scene novelty, low reward-rates, and failures to execute the appropriate action.

CGp activity predicts poor performance

Prior work shows that ACC drives rapid trial-to-trial adjustments in behavior (Hayden et al., 2011; Hayden et al., 2009a). By contrast, we hypothesize (Pearson et al., 2011) that CGp, as part of the DMN, modulates learning on a slower time scale, predicting performance across tens of trials. We therefore asked whether firing rates at the start of the session were predictive of final performance. Indeed, CGp activity strongly predicted future performance. We computed a normalized measure of neuronal activity by dividing outcome period firing rates by a cell's baseline value, defined as mean firing rate activity during a 500 ms pre-saccade epoch on all completed trials. This was done to allow us to compare activity across a heterogeneous population of cells. Note that this is not equivalent to z -score normalization. Higher firing rates at the start of the session (first 10 trials) predicted poorer future performance (last 10 trials, correlation between starting firing rates during the outcome period and end performance, $r=-.14$, $p<.0001$). To ensure this effect was not due to the difference in firing rate on correct vs. error trials, we separately examined the first 10 error and first 10 correct trials. In these cases, early high firing rates also predicted reduced future performance (correct: $r=-.10$, $p=.0002$; error: $r=-.09$, $p=.007$).

To determine whether these results could be explained by the subtle value and novelty effects, we performed a multiple regression analysis with starting firing rate, novelty level, and scene value for each scene analyzed as a function of end performance. Our findings remained unchanged (all trials: $\beta=-.02$, $p=.0006$; error trials: $\beta=-.02$, $p=.008$; correct trials: $\beta=-.02$, $p=.002$). Reference scene performance was not associated with variations in early-session firing rate ($p>.16$ in all 3 cases), presumably because learning was not necessary for this scene type, and thus they were fairly easy.

Reversible inactivation of CGp impairs learning of low-value scenes

Because these results are correlative, they leave open two conflicting possibilities: (1) DMN signals cause poor performance and (2) they marshal resources to mitigate poor performance. To adjudicate between these two hypotheses, we injected the GABA_A agonist muscimol into CGp. As a control, on different days, we injected vehicle (saline). We predicted deleterious effects of muscimol on learning of new scenes, but not on performance on familiar reference scenes.

For reference scenes, a three-way ANOVA with treatment (muscimol/saline), reward size (low/high value scene), and target location as factors for performance on each scene revealed a main effect of motivational state on performance (high value scenes > low value scenes, $F(1, 56)=36.41$, $p<0.000001$), but no other significant main effects or interactions (Figure 3). Thus, for reference scenes, CGp inactivation had no effect.

For new scenes, we restricted our analysis to the first set (novel scenes), as additional scenes were too rare to analyze. There was a main effect of scene value ($F(1, 56)=8.56$, $p=.005$) and a significant interaction between treatment and scene value ($F(1, 56)=4.25$, $p=.04$). Post-hoc tests revealed an effect of muscimol on performance on low-reward scenes (paired-samples t -test, $t(15)=3.56$, $p=.003$), but not on high-value scenes ($t(15)=0.5$, $p=.6$).

Each monkey performed poorly on novel, low-reward scenes during muscimol. Monkey L's data were significant alone, $t(9)=-2.51$, $p=.03$ (5 session pairs, 16.5% difference in performance). Monkey B's data trended toward significance, $t(5)=-2.39$, $p=.062$ (3 session pairs, 21.6% difference in performance) (Figure 3C). We therefore examined 3 measures of motivation: time to fixation, reaction time, and fixation breaks. Muscimol did not affect times to fixation or reaction times (Supplementary Materials). For fixation breaks, we excluded trials in which breaks occurred prior to scene onset. For reference scenes, fixation break frequencies were higher during muscimol sessions ($F(1, 56)=4.43$, $p<.04$.) These effects were mainly driven by low-value scenes ($t(15)=-3.14$, $p=0.007$); effects on high-value scenes were not significant ($p=.26$). Although novel scenes were not significant ($p=.12$), the other muscimol results led us to examine, posthoc, the effects on low value, novel scenes exclusively. These trended toward more fixation breaks ($p=.06$).

These results suggest that DMN activity is particularly important when conditions are more challenging—learning novel scenes for small rewards. If CGp activity merely antagonizes task performance, CGp inactivation should have improved learning; instead, we found the opposite.

Discussion

CGp activity is associated with reduced performance in a variety of laboratory tasks. One possibility is that this activity reflects cognitive processes that compete with task competence (Gusnard et al., 2001; Sonuga-Barke and Castellanos, 2007; Weissman et al., 2006). Alternatively, it may reflect control processes that compensate for conditions that reduce performance. We tested these ideas by recording and altering neuronal activity in CGp during an associative learning and memory task.

We report three main findings. First, CGp neurons carry a prominent error signal; second, the size of this signal is correlated with poor performance on a learning task; and third, inactivation of CGp retards performance in low stakes conditions. These results are consistent with the idea that CGp detects conditions leading to poor performance and acts to ameliorate suboptimal behavior, but are not concordant with the alternative 'antagonistic' hypothesis. To our knowledge, this is the first study to reversibly inactivate primate CGp

(which has no rodent homologue), although our results are broadly consistent with those from studies of permanent lesions in medial prefrontal DMN regions, which show cognitive control and decision-making deficits (e.g., Gläscher et al., 2012).

DMN activity is negatively correlated with task performance (Eichele et al., 2008; Hayden et al., 2009b). Error signals were not prepotent in our experiment, meaning firing rates were not higher on error trials prior to feedback. Thus, we infer that CGp activity *does not* directly lead to moment-to-moment failures of attention or other lapses that may cause errors. Instead, CGp is responsive to task difficulty and priorities—in this case, inherent scene difficulty, novelty and value. For certain tasks, this will likely correlate with single-trial performance, and error modulation will be prepotent (Hayden et al., 2009b). For the present task, however, variations in performance may be more closely linked with inherent difficulty than with attention per se.

Our data present an apparent contradiction: high firing rates early in learning predict poor future performance, but suppressing firing rates with muscimol also impairs performance. One hypothesis is that firing rates increase early for scenes with poor learning because factors associated with poor performance—low reward rates, scene difficulty, lack of focus—are detected and subsequently, but only partially, remedied by CGp. By way of analogy, firemen are often found near a fire, but it would be a mistake to conclude that they cause the fire. On the contrary, preventing the firemen from showing up (muscimol in this analogy) would only make the fire worse. Together, our results thus suggest that higher CGp activity does not *cause* poor task performance, but instead detects and possibly endeavors to improve it.

These data indicate that, when rewards are small, the normal processes that motivate performance are insufficient, and control areas are recruited to learn the new scene. Inactivating CGp impairs this process, effectively blocking new learning in the low reward condition. This idea is somewhat different from more conventional theories of motivation and control, such as those associated with ACC and vmPFC. Interestingly, portions of these medial prefrontal areas are also part of the DMN (Buckner et al., 2008), and CGp is anatomically and functionally connected to both of them (Vogt et al., 1987; Greicius et al., 2003). Like CGp, ACC also increases its firing following errors (Niki and Watanabe, 1979). ACC and vmPFC are associated with executive control and decisions, but tend to have antagonistic relationships (Kolling et al., 2012; Boorman et al., 2013). Variations in functional connectivity between CGp and ACC/vmPFC may bias decision-making and learning. While these different components of the DMN have distinct functions, we suggest that none of them acts in opposition to task performance.

While the present results are limited to learning contexts, CGp may have a role in regulating other cognitive processes (Weissman et al., 2006; Vogt et al., 1992; Small et al., 2003; Castellanos et al., 2008). Moreover, in conjunction with a rich literature showing that ACC has dynamic outcome signals (Quilodran et al., 2008), our results indicate that monitoring performance and marshalling cognitive resources is a more general property of cingulate cortex (Botvinick et al., 2004; Carter et al., 1998). We find that muscimol increases the proportion of fixation breaks for low-reward scenes, hinting that one specific mechanism by which CGp may improve performance is to increase motivation and/or attention. We previously developed a model suggesting one function of CGp is to detect slow changes in environmental statistics important for learning (Pearson et al., 2011). The present findings endorse this idea.

Our findings bear on a recent debate concerning the source of the error-related negativity (ERN), the distinctive change in brain activity following error commission observed with

electroencephalography (Gehring et al., 1993). Most studies localize the source of the ERN to the ACC (Holroyd and Coles, 2002). A recent study, however, localized the source of the ERN to CGp (Agam et al., 2011). Here, we confirm that CGp does increase its firing rate following error commission.

Because these results were found in a learning context, they may have implications for learning and memory impairments. Patients with Alzheimer's Disease (AD) have hypometabolism in CGp (Minoshima et al., 1997), often before symptoms become severe (Valla et al., 2001). Metabolic abnormalities in CGp predict AD symptom severity (Yoshiura et al., 2002), including deficits in learning (Hirono et al., 1998). Our results suggest that in its early stages, AD involves aberrant control of learning because of abnormalities in CGp. We therefore hypothesize that tests of cognitive regulation of learning may offer particularly sensitive instruments for diagnosis of AD.

Experimental Procedures

Surgical and behavioral procedures

Procedures were approved by the Duke University Institutional Animal Care and Use Committee and were conducted in compliance with the NIH *Guide for the Care and Use of Animals*. Two adult male rhesus macaques (*Macaca mulatta*) were implanted with a head-holding prosthesis and plastic recording chamber (Crist Instruments), and habituated to laboratory conditions. Animals received analgesics and antibiotics after surgeries. The chamber was kept sterile with regular antibiotic washes and sealed with sterile caps.

Eye positions were sampled at 1000 Hz by an infrared camera system (Eyelink 1000). Stimuli were controlled by a computer running Matlab with the Psychophysics Toolbox (Brainard, 1997) and the Eyelink Toolbox (Cornelissen et al., 2002). A solenoid valve controlled the duration of juice delivery. Reward amounts were 0 msec of solenoid open time (incorrect, 0 mL juice), 150 msec (correct, low value, 0.2 mL juice), and 250 msec (correct, high value, 0.33 mL juice).

Each trial began when a single yellow square appeared at the center of the screen, which the monkey fixated ($\pm 0.5^\circ$). The fixation epoch was 800 msec (monkey L) or 300 msec (monkey B). Following the fixation epoch, an arbitrary photograph ("scene") was displayed on the center of the monitor (512×384 pixels), along with 2 gray targets located in the upper left and lower right portions of the screen, outside the scene. The monkey maintained fixation on the center square during scene presentation (500 msec). When the scene was extinguished, the monkey maintained fixation for another 400 (monkey L) or 700 (monkey B) msec (delay period). At the end of the delay period, the central fixation point was turned off, and the monkey was permitted to shift his gaze eccentrically to one of the 2 available gray targets. The chosen target became red for 250 msec. Then, all targets were extinguished, and reward was delivered. The monkey was not required to maintain fixation through the reward period.

Sessions included reference, novel, and (sometimes) addition scenes. Four reference scenes were present every day (high/low value, left/ right saccade), and their saccade/reward contingencies did not change. Four novel scenes were used each session, then never seen again. Addition scenes were 4 new scenes introduced after trial 400 (or 500, increased as the performance improved). Scenes were selected randomly on each trial; thus, any scene could follow any other scene, including itself. Proportions of novel, addition, and reference scenes were adjusted between sessions to boost performance. Within a type, all scenes were presented in equal proportions. Scenes gave no cues about value or direction independent of learning.

Monkeys learned to associate each scene with upper-left or lower-right gaze shifts (Figure 1A–B). Half of scenes were high-value, meaning a correct response resulted in a large juice reward; the others were low-value. An incorrect response resulted in reward omission.

Scenes did not repeat following incorrect trials, but, during certain sessions, scenes repeated if the monkey broke fixation early. Monkey L was always required to do this; Monkey B was required to do this on 14 ‘remedial’ sessions to improve performance.

Recording techniques

Single electrodes (Frederick Haer) were lowered with a hydraulic microdrive (Kopf Instruments) until the waveform(s) of 1–3 neurons were isolated with sorting software (SortClient, Plexon, Inc). Neurons were chosen for recording based solely on isolation and not on any task-related properties. We approached CGp through a standard recording grid (Crist Instruments). CGp was identified either by magnetic resonance images taken before the experiment (Monkey L) or by ultrasound and stereotactic map (Monkey B). We confirmed that electrodes were in CGp using stereotactic measurements, and by listening for characteristic sounds of white and gray matter during recording. Recordings were made 1–2 mm lateral from the midline, between the positions of 3 mm posterior and 3 mm anterior to the interaural plane, with most occurring 1 mm posterior to 2 mm anterior.

Neuronal analyses

Our analyses (performed with MATLAB software) focused on the outcome epoch, a 700 msec period beginning 350 msec following the possible onset of reward. The scene epoch was the 500 msec while the scene is on the computer monitor. The fixation epoch was the 300 msec beginning when fixation was achieved at trial start. Where normalization occurred, pre-saccade firing rates were used (500 msec prior to saccade offset).

Reversible inactivation

In each session, we injected either 4 µg (.8 µl volume, 5 µg/µl in saline vehicle) of muscimol (Sigma/Aldrich), a GABA_A agonist, or .8 µl of saline vehicle, per site, with 1–4 sites. We used a 10 µl volume Hamilton syringe with a 28-gauge needle. We compared muscimol sessions to vehicle (saline) sessions, controlling for the presence of the needle, and isolating the effects of muscimol. The task did not begin until all of the sites had been injected for the day. Each muscimol session was paired with a vehicle session pre-hoc, to allow for paired statistical comparisons. Vehicle and saline sessions occurred on different days, and each day contained only one session. Because muscimol effects may extend up to 24 hours, we waited at least 40 hours between sessions.

The order of muscimol (M)/vehicle (V) injections was: Monkey L—VMMVMVMM; Monkey B—MVVMMV. Vehicle sessions were equally likely to occur before and after paired muscimol sessions. For all but 2 saline-muscimol pairs, there were 4 injection sites (locations). For one session, we injected at two positions. For another, we only injected at one site. We injected slowly (.2 ul/2 min) to avoid damaging the brain. Inactivation experiments were performed on the same monkeys used in the recording experiments, but began after the recording phase was complete. Within animal, injections were performed unilaterally (Monkey L: left; Monkey B: right). Paired muscimol and vehicle injections always matched on number of sites.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by NIH grants RO1 EY013496 (MLP) and F31 DA028133 (SRH). We thank Benjamin Hayden, John Pearson, and 3 anonymous reviewers for useful comments.

References

- Agam Y, Hamalainen MS, Lee AK, Dyckman KA, Friedman JS, Isom M, Makris N, Manoach DS. Multimodal neuroimaging dissociates hemodynamic and electrophysiological correlates of error processing. *Proc Natl Acad Sci U S A*. 2011; 108:17556–17561. [PubMed: 21969565]
- Boorman ED, Rushworth MF, Behrens TE. Ventromedial prefrontal and anterior cingulate cortex adopt choice and default reference frames during sequential multi-alternative choice. *J Neurosci*. 2013; 33:2242–2253. [PubMed: 23392656]
- Botvinick MM, Cohen JD, Carter CS. Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn Sci*. 2004; 8:539–546. [PubMed: 15556023]
- Brainard DH. The Psychophysics Toolbox. *Spat Vis*. 1997; 10:433–436. [PubMed: 9176952]
- Buckner RL, Andrews-Hanna JR, Schacter DL. The brain's default network: anatomy, function, and relevance to disease. *Ann N Y Acad Sci*. 2008; 1124:1–38. [PubMed: 18400922]
- Buckner RL, Carroll DC. Self-projection and the brain. *Trends Cogn Sci*. 2007; 11:49–57. [PubMed: 17188554]
- Carter CS, Braver TS, Barch DM, Botvinick MM, Noll D, Cohen JD. Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*. 1998; 280:747–749. [PubMed: 9563953]
- Castellanos FX, Margulies DS, Kelly C, Uddin LQ, Ghaffari M, Kirsch A, Shaw D, Shehzad Z, Di Martino A, Biswal B, et al. Cingulate-precuneus interactions: a new locus of dysfunction in adult attention-deficit/hyperactivity disorder. *Biol Psychiatry*. 2008; 63:332–337. [PubMed: 17888409]
- Chen LL, Wise SP. Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. *J Neurophysiol*. 1995; 73:1101–1121. [PubMed: 7608758]
- Cornelissen FW, Peters EM, Palmer J. The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behav Res Methods Instrum Comput*. 2002; 34:613–617. [PubMed: 12564564]
- Eichele T, Debener S, Calhoun VD, Specht K, Engel AK, Hugdahl K, von Cramon DY, Ullsperger M. Prediction of human errors by maladaptive changes in event-related brain networks. *Proc Natl Acad Sci U S A*. 2008; 105:6173–6178. [PubMed: 18427123]
- Gehring WJ, Goss B, Coles MG, Meyer DE, Donchin E. A neural system for error detection and compensation. *Psychol Sci*. 1993; 4:385–390.
- Gläscher äJ, Adolphs R, Damasio H, Bechara A, Rudrauf D, Calamia M, Paul LK, Tranel D. Lesion mapping of cognitive control and value-based decision making in the prefrontal cortex. *Proc Natl Acad Sci U S A*. 2012; 109:14681–14686.
- Greicius MD, Krasnow B, Reiss AL, Menon V. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proc Natl Acad Sci U S A*. 2003; 100:253–258. [PubMed: 12506194]
- Gusnard DA, Akbudak E, Shulman GL, Raichle ME. Medial prefrontal cortex and self-referential mental activity: relation to a default mode of brain function. *Proc Natl Acad Sci U S A*. 2001; 98:4259–4264. [PubMed: 11259662]
- Halsband U, Passingham R. The role of premotor and parietal cortex in the direction of action. *Brain Res*. 1982; 240:368–372. [PubMed: 7104700]
- Hayden BY, Heilbronner SR, Pearson JM, Platt ML. Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. *J Neurosci*. 2011; 31:4178–4187. [PubMed: 21411658]
- Hayden BY, Nair AC, McCoy AN, Platt ML. Posterior cingulate cortex mediates outcome-contingent allocation of behavior. *Neuron*. 2008; 60:19–25. [PubMed: 18940585]
- Hayden BY, Pearson JM, Platt ML. Fictive reward signals in the anterior cingulate cortex. *Science*. 2009a; 324:948–950. [PubMed: 19443783]

- Hayden BY, Smith DV, Platt ML. Electrophysiological correlates of default-mode processing in macaque posterior cingulate cortex. *Proc Natl Acad Sci U S A*. 2009b; 106:5948–5953. [PubMed: 19293382]
- Heilbronner SR, Hayden BY, Platt ML. Decision salience signals in posterior cingulate cortex. *Front Neurosci*. 2011; 5:55. [PubMed: 21541308]
- Hirono N, Mori E, Ishii K, Ikejiri Y, Imamura T, Shimomura T, Hashimoto M, Yamashita H, Sasaki M. Hypofunction in the posterior cingulate gyrus correlates with disorientation for time and place in Alzheimer's disease. *J Neurol Neurosurg Psychiatry*. 1998; 64:552–554. [PubMed: 9576555]
- Holroyd CB, Coles MG. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol Rev*. 2002; 109:679–709. [PubMed: 12374324]
- Kolling N, Behrens TE, Mars RB, Rushworth MF. Neural mechanisms of foraging. *Science*. 2012; 336:95–98. [PubMed: 22491854]
- Mason MF, Norton MI, Van Horn JD, Wegner DM, Grafton ST, Macrae CN. Wandering minds: the default network and stimulus-independent thought. *Science*. 2007; 315:393–395. [PubMed: 17234951]
- McCoy AN, Platt ML. Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat Neurosci*. 2005; 8:1220–1227. [PubMed: 16116449]
- Minoshima S, Giordani B, Berent S, Frey KA, Foster NL, Kuhl DE. Metabolic reduction in the posterior cingulate cortex in very early Alzheimer's disease. *Ann Neurol*. 1997; 42:85–94. [PubMed: 9225689]
- Niki H, Watanabe M. Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Res*. 1979; 171:213–224. [PubMed: 111772]
- Pearson JM, Hayden BY, Raghavachari S, Platt ML. Neurons in posterior cingulate cortex signal exploratory decisions in a dynamic multioption choice task. *Curr Biol*. 2009; 19:1532–1537. [PubMed: 19733074]
- Pearson JM, Heilbronner SR, Barack DL, Hayden BY, Platt ML. Posterior cingulate cortex: adapting behavior to a changing world. *Trends Cogn Sci*. 2011; 15:143–151. [PubMed: 21420893]
- Petrides M. Motor conditional associative-learning after selective prefrontal lesions in the monkey. *Behav Brain Res*. 1982; 5:407–413. [PubMed: 7126320]
- Quilodran R, Rothe M, Procyk E. Behavioral shifts and action valuation in the anterior cingulate cortex. *Neuron*. 2008; 57:314–325. [PubMed: 18215627]
- Raichle ME, MacLeod AM, Snyder AZ, Powers WJ, Gusnard DA, Shulman GL. A default mode of brain function. *Proc Natl Acad Sci U S A*. 2001; 98:676–682. [PubMed: 11209064]
- Shulman GL, Fiez JA, Corbetta M, Buckner RL, Miezin FM, Raichle ME, Petersen SE. Common blood flow changes across visual tasks: II. Decreases in Cerebral Cortex. *Journal of Cognitive Neuroscience*. 1997; 9:648–663.
- Petrides M. Motor conditional associative-learning after selective prefrontal lesions in the monkey. *Behav Brain Res*. 1982; 5:407–413. [PubMed: 7126320]
- Small DM, Gitelman DR, Gregory MD, Nobre AC, Parrish TB, Mesulam MM. The posterior cingulate and medial prefrontal cortex mediate the anticipatory allocation of spatial attention. *Neuroimage*. 2003; 18:633–641. [PubMed: 12667840]
- Sonuga-Barke EJ, Castellanos FX. Spontaneous attentional fluctuations in impaired states and pathological conditions: a neurobiological hypothesis. *Neurosci Biobehav Rev*. 2007; 31:977–986. [PubMed: 17445893]
- Spreng RN, Grady CL. Patterns of brain activity supporting autobiographical memory, prospection, and theory of mind, and their relationship to the default mode network. *J Cogn Neurosci*. 2010; 22:1112–1123. [PubMed: 19580387]
- Valla J, Berndt JD, Gonzalez-Lima F. Energy hypometabolism in posterior cingulate cortex of Alzheimer's patients: superficial laminar cytochrome oxidase associated with disease duration. *J Neurosci*. 2001; 21:4923–4930. [PubMed: 11425920]
- Vincent JL, Patel GH, Fox MD, Snyder AZ, Baker JT, Van Essen DC, Zempel JM, Snyder LH, Corbetta M, Raichle ME. Intrinsic functional architecture in the anaesthetized monkey brain. *Nature*. 2007; 447:83–86. [PubMed: 17476267]

- Vogt BA, Finch DM, Olson CR. Functional heterogeneity in cingulate cortex: the anterior executive and posterior evaluative regions. *Cereb Cortex*. 1992; 2:435–443. [PubMed: 1477524]
- Vogt BA, Pandya DN. Cingulate cortex of the rhesus monkey: II. Cortical afferents. *J Comp Neurol*. 1987; 262:271–289. [PubMed: 3624555]
- Weissman DH, Roberts KC, Visscher KM, Woldorff MG. The neural bases of momentary lapses in attention. *Nat Neurosci*. 2006; 9:971–978. [PubMed: 16767087]
- Wirth S, Yanike M, Frank LM, Smith AC, Brown EN, Suzuki WA. Single neurons in the monkey hippocampus and learning of new associations. *Science*. 2003; 300:1578–1581. [PubMed: 12791995]
- Yoshiura T, Mihara F, Ogomori K, Tanaka A, Kaneko K, Masuda K. Diffusion tensor in posterior cingulate gyrus: correlation with cognitive decline in Alzheimer's disease. *Neuroreport*. 2002; 13:2299–2302. [PubMed: 12488815]

Highlights

CGp neurons responded phasically to errors and predicted future poor performance.

Reversible inactivation of posterior cingulate blocked new learning.

Default mode network activity is necessary for cognitively demanding performance.

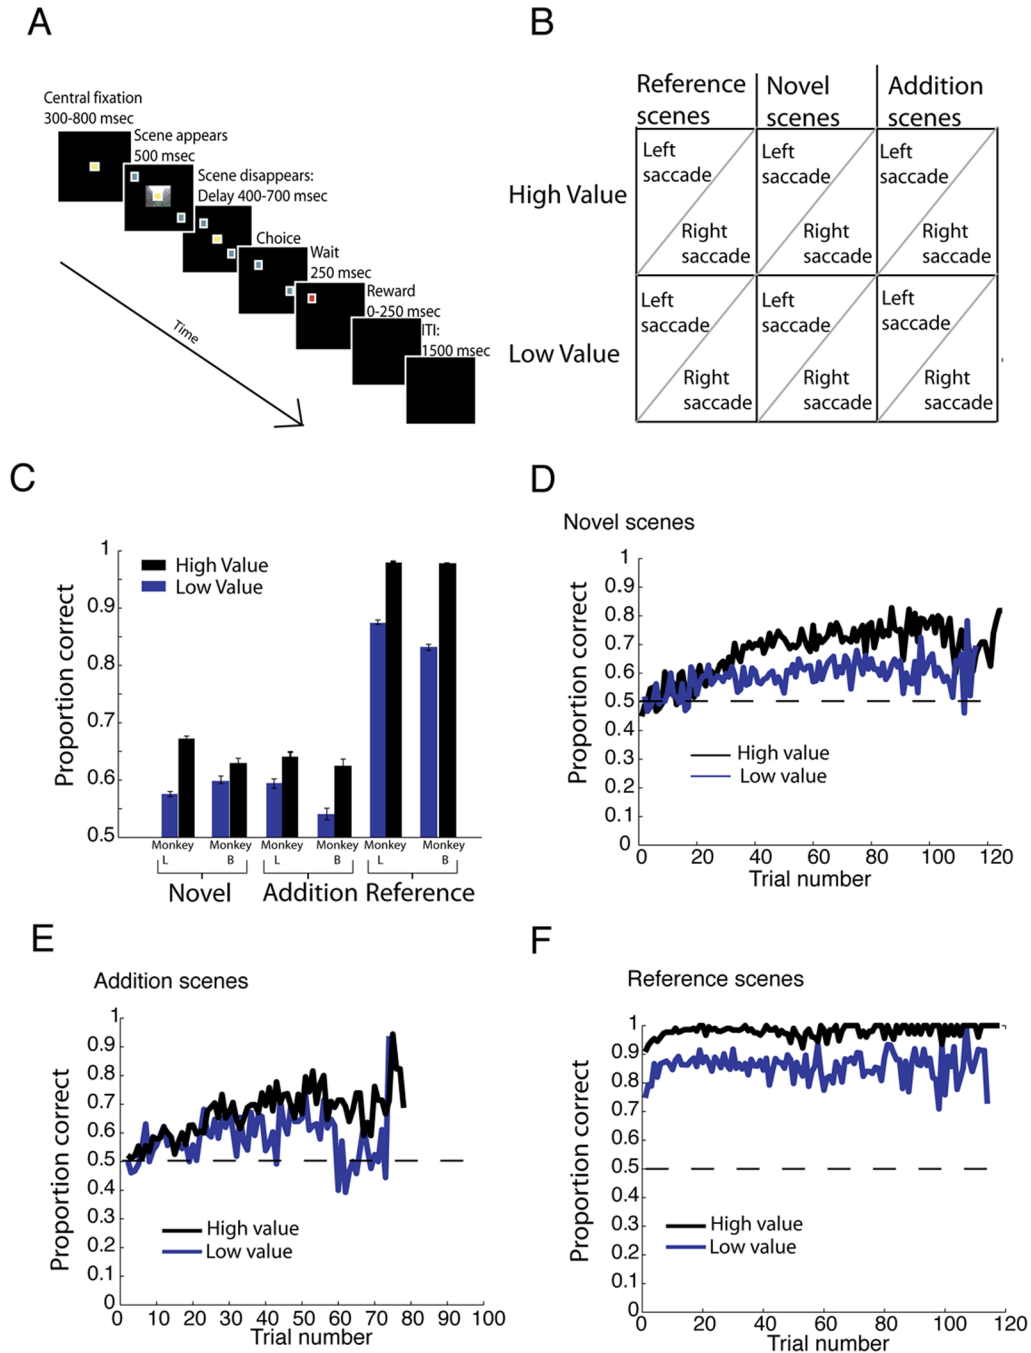


Figure 1. Novelty and reward size shape learning in monkeys

A. Trial structure. Trials began with a fixation period, followed by scene display, another fixation period, and then choice. **B.** Trial types: Scenes could be high value or low value, with correct eye movement assigned to left or right. Some scenes were reference (viewed every day), some were new scenes introduced at the start of the day (novel), and some were new scenes introduced midway through the session (addition). **C.** Behavioral performance. **D, E, F:** Averaged performance across all scenes within type, by trial number within session. Performance on novel (**D**), addition (**E**), reference (**F**) scenes.

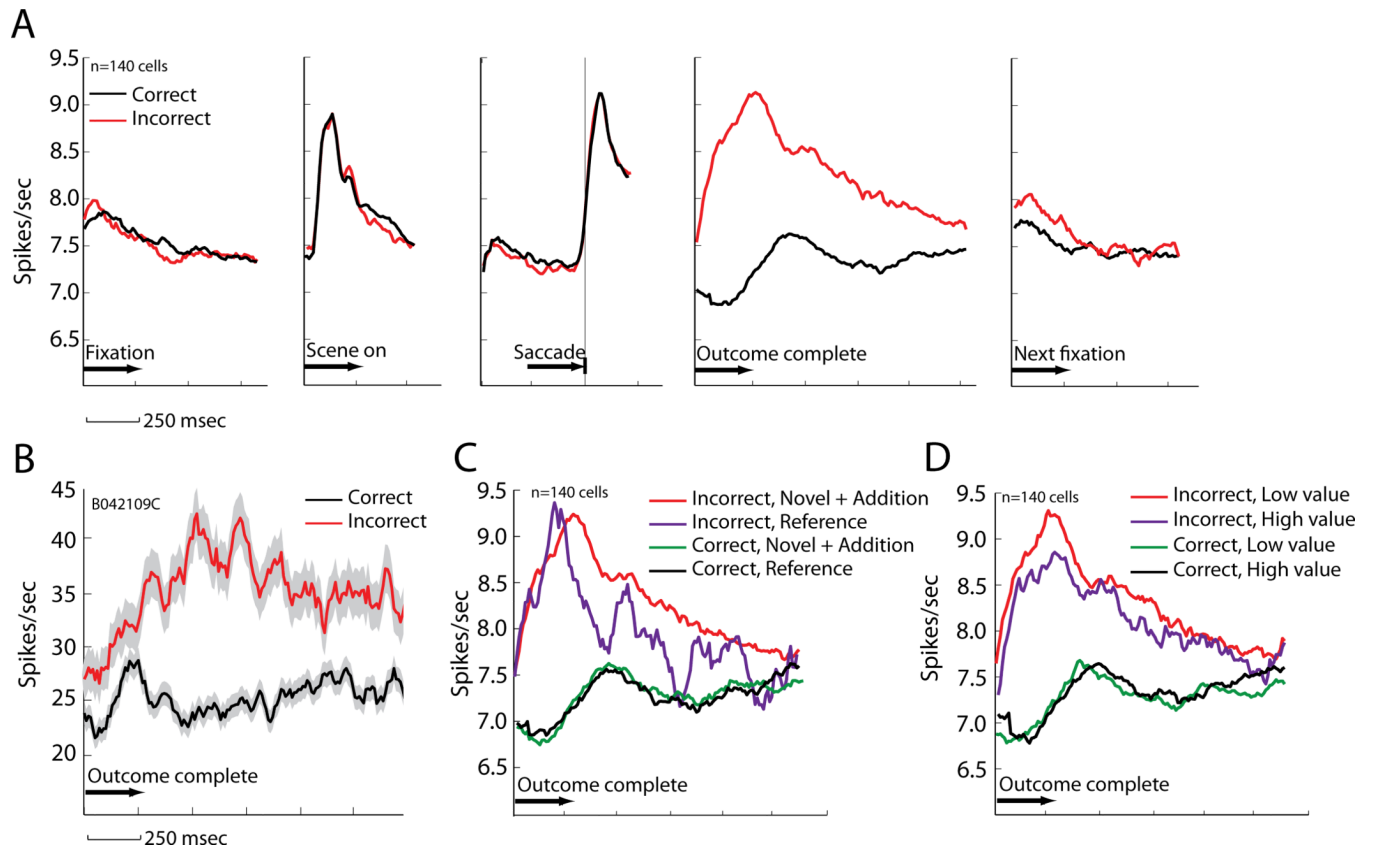


Figure 2. Error signals carried by CGp neurons reflect reward size and novelty

A. Error signal averaged across all recorded CGp cells during different trial epochs. Signal is aligned to key task events and re-plotted for consistency. The error signal was not present prior to outcome, and it disappeared at the start of the following trial. **B.** Error signal (+/- SEM) during post-outcome period for a single CGp cell. **C.** Population error signal by scene type. **D.** Population error signal by scene value.

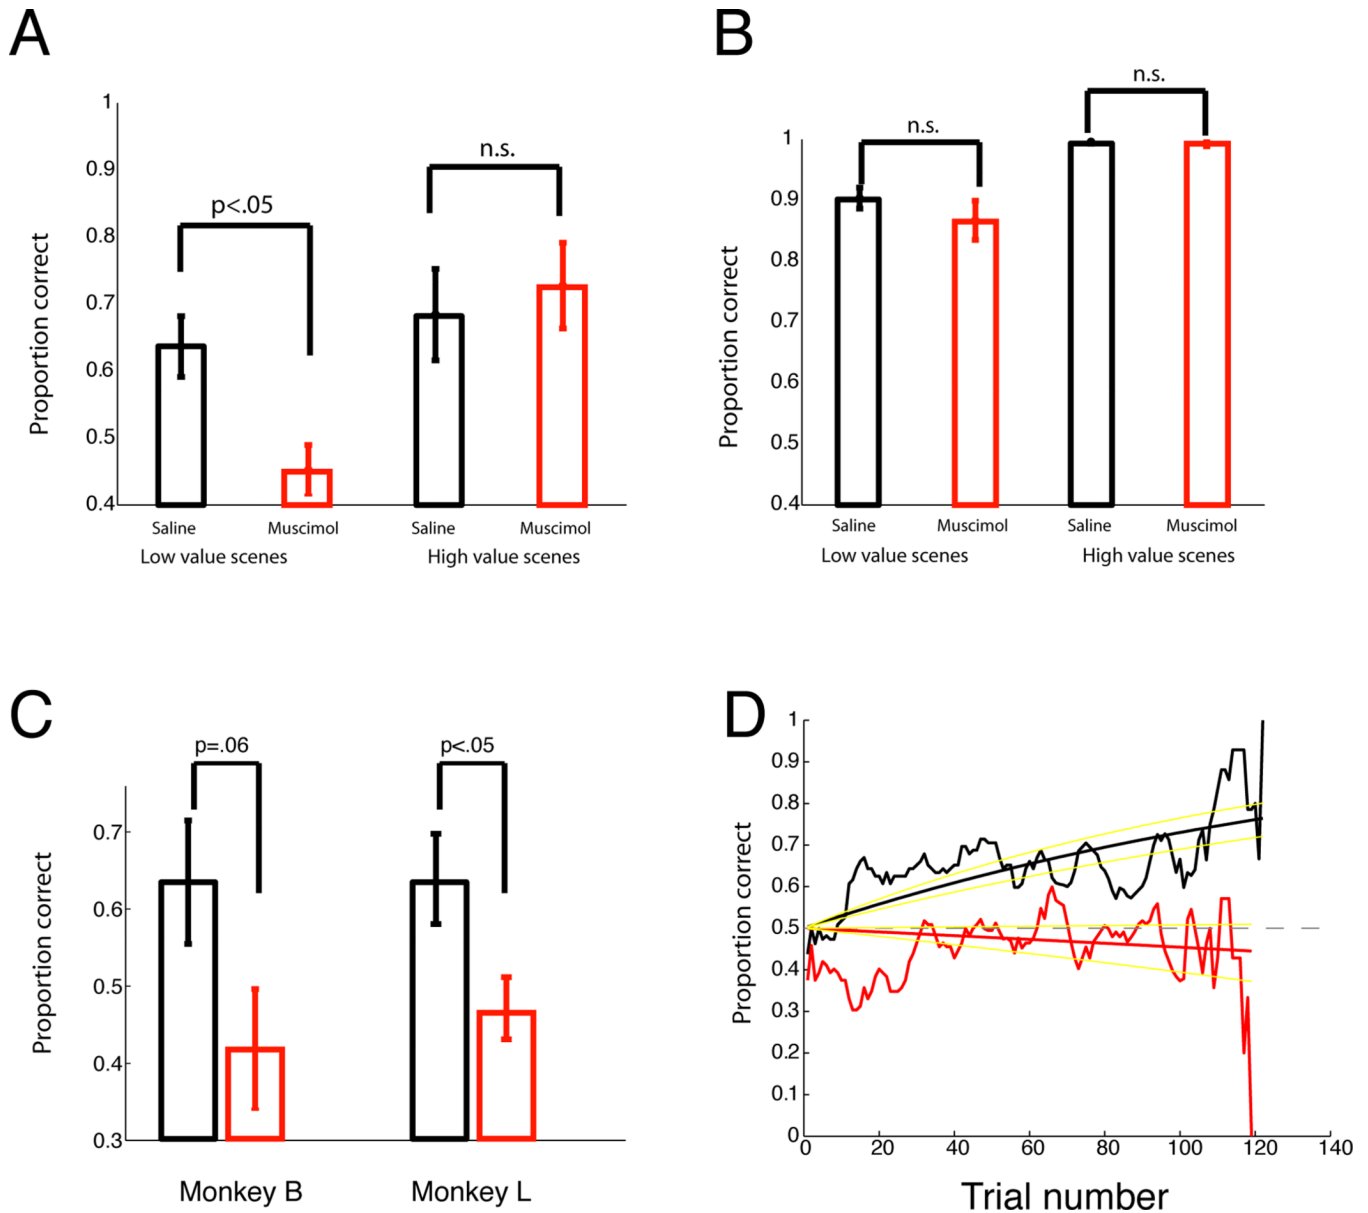


Figure 3. CGp inactivation abolishes new learning but not recall in low motivational states
A. Mean performance (\pm SEM) on novel scenes. **B.** Performance on reference scenes. **C.** Performance of each monkey on novel, low-reward scenes during saline (black) and muscimol (red) sessions. **D.** Averaged performance on low-value, novel scenes, across trial number. Yellow lines are 95% confidence intervals.

Table 1

Performance on all scene types.

Reference scenes	High Value	Low Value
Saccade Left	.97	.84
Saccade Right	.99	.87
Novel scenes	High Value	Low Value
Saccade Left	.66	.64
Saccade Right	.67	.53
Addition scenes	High Value	Low Value
Saccade Left	.65	.60
Saccade Right	.62	.54