Prefrontal deep projection neurons enable cognitive flexibility via persistent feedback monitoring

Spellman T,[1][*] Svei M[1], Liston C[1]

[1]Feil Family Brain and Mind Research Institute, Weill Cornell Medicine, New York, NY 10021

Corresponding Author
Conor Liston (col2004@med.cornell.edu); Timothy Spellman (tis2013@med.cornell.edu)

Summary

Cognitive flexibility, the ability to alter one's strategy in light of changing stimulus-response-reward relationships, is critical for acquiring and updating learned behavior by trial and error. Successful performance of attentional set-shifting, a test of cognitive flexibility, has been shown to critically involve the prefrontal cortex (PFC). It is unclear, however, whether PFC neurons support set-shifting behavior by guiding action selection or by monitoring feedback from completed actions. Using optogenetics and 2-photon calcium imaging, we demonstrate that PFC activity critically encodes trial feedback information, rather than guiding action selection, to enable set-shifting. Moreover, we find that, in two distinct PFC projection populations, representations of trial feedback form a topological gradient, with cells more strongly selective for feedback information located further from the pial surface, regardless of their projection targets. Together, these findings indicate that deep PFC projection neurons enable set-shifting through monitoring of behavioral feedback.

Keywords: prefrontal cortex, cognitive flexibility, calcium imaging, optogenetics, projection neurons, executive functioning, feedback monitoring

Introduction

Cognitive flexibility, the ability to rapidly respond to changing external contingencies, is critical for any organism that must navigate through a dynamic environment and learn to adapt its behavior to changing environmental contingencies. A form of cognitive flexibility frequently employed in clinical assessments is *attentional set-shifting*, a kind of task-switching behavior that requires a subject to ignore a previously relevant stimulus feature and instead attend to a previously irrelevant feature to execute a correct response (Heisler et al., 2015; Tait et al., 2014).

Impairment in set-shifting is a prevalent cognitive symptom in a range of psychiatric disorders (Ceaser et al., 2008; Disner et al., 2011; Halleland et al., 2012; Jazbec et al., 2007; Murphy et al., 2012) and often persists after treatment in both depression and schizophrenia, despite remission of other symptoms (Bortolato et al., 2016; Gonda et al., 2015; Harvey et al., 2004). Prefrontal cortex (PFC) plays a critical and well-established role in supporting set-shifting behavior in human (Kim et al., 2012; Milner, 1963) and rodent models (Birrell and Brown, 2000; Bissonette et al., 2013; Brigman et al., 2005), but the circuit elements and circuit-level mechanisms that enable set-shifting remain poorly resolved. A functional dissection of the prefrontal circuitry underlying set-shifting is of particular translational interest because pharmacotherapeutic interventions in this circuitry can have complex or paradoxical effects by acting on competing circuit elements (Miller, 2016), which may account for the deleterious cognitive side effects that may accompany SSRIs and antipsychotics and SSRIs (Hill et al., 2010; Sayyah et al., 2016).

Neural activity in primate prefrontal cortex has been shown to encode abstract, context or rule-related information (Hymana et al., 2012; Mante et al., 2013; Rigotti et al., 2013), and in rodents performing set-shifting tasks, such rule-related representations can shift flexibly with repeatedly changing stimulus-reward contingencies (Durstewitz et al., 2010; Malagon-Vina et al., 2018). In light of the critical role of PFC in cued

modulation of attention (Wimmer et al., 2015), the idea that PFC activity might guide action selection during set-shifting tasks remains a well-predicated hypothesis.

Alternatively, because set-shifting performance is based on continuous trial-and-error learning, and rules are not explicitly cued, the critical role of PFC in set-shifting may lie in monitoring feedback from actions once performed. In addition to encoding context, prefrontal activity has been shown to represent feedback signals associated with trial outcome (Bissonette and Roesch, 2015; Starkweather et al., 2018), and recent evidence suggests that such feedback-related activity is critical for performing task-switching behavior (Biró et al., 2019; Ellwood et al., 2017).

It is also unclear how distinct circuit elements interact within PFC to support set-shifting. A major set of questions in behavioral neuroscience in recent years has centered on whether the functional identity of a given neuron is defined by its connectivity profile or by its laminar location (Adesnik and Naka, 2018; Harris and Shepherd, 2015). One difficulty in resolving this question lies in the fact that laminar location and connectivity are often correlated. While findings from numerous data sets implicate both laminar location and connectivity in determining the circuit contribution of cortical neurons (Otis et al., 2017; Senzai et al., 2019), to disambiguate these two factors requires an experimental preparation that controls for both. We therefore employed a virally-mediated expression approach that labeled two PFC projection populations across largely overlapping regions of PFC.

The critical contribution of PFC activity to set-shifting may be mediated by a number of output pathways, but two PFC projection targets of particular interest are the projection to ventromedial striatum (PFC-VMS) and to the medio-dorsal thalamus (PFC-MDT). In rodents, both target structures have established roles in set-shifting behavior (Aoki et al., 2015; Block et al., 2007; Floresco et al., 2006; Kato et al., 2018), and PFC projections to both structures underlie behavior that relies on cognitive flexibility (Marton et al., 2018; Nakayama et al., 2018), however it is unknown whether PFC projections to these downstream targets convey similar or distinct task-critical information during set-shifting.

To examine the roles of PFC output neurons in set-shifting performance, we developed and validated a novel set-shifting task for head-fixed mice to enable 2-photon imaging during behavior. Muscimol infusion in PFC impaired performance in both initial and overtrained set-shifting sessions. Attentional set shifts were acquired gradually and exhibited characteristic switching costs, and performance depended upon the recency of rule-dependent trials. Signals encoding the selected response and trial outcome were decoded from neural population activity only *after* trial completion, supporting a role for PFC cells in feedback monitoring rather than action selection. This feedback signal persisted for four subsequent trials (up to 55 seconds). Separate analysis of PFC-VMS and PFC-MDT neurons revealed robust representations of all task-related features in both cell types, with strikingly similar trial responsiveness and temporal dynamics. Unexpectedly, optogenetic inhibition of either cell type had no effect on performance when delivered during stimulus delivery and response epochs, but inhibition during the post-trial epoch, impaired performance. Moreover, this impairment was limited to performance on rule-dependent trials, and only in cases when photo-inhibition was administered following rule-dependent trials, providing further support for a critical role for PFC cells in feedback monitoring rather than action selection in this task. Analysis of the topological distribution of response and outcome-selective cells revealed a laminar gradient, with more informative cells located further from the pial surface, indicating a prominent role for projection neurons in deeper cortical layers in conveying feedback-related information to the downstream regions.

Results

A cross-modal set-shifting task for head-fixed mice

To image PFC activity during set-shifting behavior in mice, we developed a head-fixed paradigm leveraging the sensitivity of olfactory and whisker somatosensory systems. Water-restricted mice were presented on each trial with one of two possible whisker stimuli (e.g. 35Hz vs 155Hz vibration presented bilaterally) and two possible odor stimuli (e.g. almond oil vs olive oil), in randomized combinations, with left and right lickspouts serving as behavioral readouts (**Fig. 1A**. As in most set-shifting tasks, animals underwent a standardized series of transitions in task structure involving different types of stimulus discriminations, response reversals, and

Figure 1: Validation of a prefrontal-dependent serial set-shifting task for head-fixed mice
A. Task schematic. Top: diagram of stimulus delivery and lick response configuration. Middle: illustration of relevant and irrelevant stimuli and rule shifting. Bottom: definition of congruent (CG) and incongruent (ICG) trial.
B. Training curves for the seven learning stages. Y-axis denotes the proportion of animals at a given learning stage after x completed trials (concatenated sessions; mode = 350 trials per session). SD = Simple Discrimination, CD = Compound Discrimination, IDS1 = Intradimensional Shift (new stimuli, initial modality rule), Rev = Reversal (initial modality rule), EDS = Extradimensional Shift, IDS2 = Intradimensional Shift (second modality rule), SEDS = Serial Extradimensional Shifting. N=53 animals.
C. Left: Number of trials to criterion, Intradimensional vs Extradimensional Shift, n=53 animals. Signed rank z = -2.9, p = 0.0034, for IDS1/EDS; z = -4.0, p = 6x10⁻⁵ for IDS2/EDS, z = 1.4, p = 0.15 for IDS1/IDS2. Right: Mean trials to criterion during SEDS, whisker rule vs odor rule. N=115 animals, signed rank z = 0.14, p = 0.9.
D. Left: trials to criterion in EDS sessions during transcranial infusion (SAL = saline; MUS = muscimol). N= 12, 13 mice (SAL, MUS). Rank sum z = 2.4, p = 0.02. Right: Number of trial blocks reaching criterion performance in SEDS sessions following 10 rule shifts. N=12; median blocks (BL/MUS/SAL): 4, 1.5, 4. Signed rank p = 0.0005 (BL/MUS), 0.001 (SAL/MUS), 0.68 (BL/SAL). Median total trials completed: SAL = 644; MUS = 651; Signed rank p = 0.52.
E. Incongruent trial performance by recency of previous incongruent trial during SEDS sessions (mean ± SEM), shown for trial groups relative to rule shifts. N = 693 sessions in 131 animals. ANOVA for ICG trial recency vs ICG trial performance: t = -4.68, p = 3x10⁻⁶.
F. SEDS trial performance relative to rule shift. Left: proportion of trials correct. Right: anticipatory licking (correct – incorrect direction). Same sessions as in **E**.

attentional set shifts (**Fig. 1A-B,** see Methods for details), culminating in a Serial Extradimensional Shifting task (SEDS), in which the relevant modality (olfactory vs. whisker somatosensory) switched upon reaching a criterion performance of 80% correct within a 30-trial moving window. As expected for a set-shifting task, animals took longer to perform an extradimensional shift (EDS) between modalities than to acquire a new set of exemplars

Figure 2 Prefrontal populations represent task-critical features as activity trajectories

A. Top left: schematic of coronal prism implantation field. Bottom left: hSyn-GCaMP6f fluorescence in fixed tissue (DAPI in blue). Right: standard deviation image from 115-min. hSyn-GCaMP6f recording (~3Hz, N = 953 units). Note: not same animal as bottom left.

B. Top: Thresholded/rasterized events from deconvolved calcium traces. N=100 units recorded over the course of 10 trial blocks.

C. Left: mean ± SEM of z-scored activity from an example neuron during trials with four stimulus combinations. Stimulus presented from -2.5s to 0s. Response window is from 0 sec to 1.5 sec. Earliest subsequent trial onset at 8 sec. Right: trial-separated z-score values for left/right whisker stimuli in a 3-unit space, with illustrated classification boundary.

D. SVM decoder accuracy for five task-related features, temporally aligned to the response window. Cross-validation was performed on trials (50%) held out from training. N=4740 units from 21 animals. Shaded areas are 90% confidence intervals, and equalizing trial numbers for each feature value ensured that performance on shuffled data was ~50% for each feature (Methods).

E. Temporal auto-correlation of β-weights from decoders for the features plotted in **d**. Warmer colors correspond with higher correlation coefficient for each pair of response-aligned timepoints. Pearson coefficients were obtained from partial correlations, which, for each of the five features shown, controlled for the β-weights of the other four features. Low correlation coefficients for non-adjacent time bins signify that independent neural populations determine the decoder at non-adjacent times.

F. Temporal trajectories of first three principal components for odor-rule and whisker-rule trials. PCA performed separately on rule-separated trial sets. Units are derived from z-scored activity. N = 1770 neurons in 8 animals.

within either modality (termed an intradimensional shift or IDS; **Fig. 1C**), and there was no difference between modalities in trials to criterion (Fig. 1C). Infusion of muscimol within PFC increased trials to criterion in the initial EDS shift, as well as after completing 10 shifts of SEDS (**Fig. 1D**), indicating that set-shifting in this paradigm is dependent on PFC activity, both initially and after repeated shifts.

Trials belonged to two classes: those in which whisker and odor stimuli would indicate the same response direction (*congruent* trials: CG), and those in which the two stimuli would indicate opposing directions (*incongruent* trials: ICG). Therefore, ICG trials, but not CG trials, required application of the modality rule, and feedback from ICG trials alone carried information about modality rule. Animals used prior ICG trial information to guide behavior, evidenced by a negative correlation between ICG performance and the number of intervening CG trials that had been presented since the most recent ICG trial, and this held true in trials occurring both early and late within trial blocks (**Fig. 1E**). Performance on ICG trials, but not CG trials, dropped below chance following rule switches, revealing perseverance of the previous rule (**Fig. 1F**). Sub-chance performance persisted over multiple ICG trials, regressing gradually to chance performance before new the new rule was acquired

Prefrontal populations represent task-critical features as activity trajectories

To assess task-related activity in PFC neurons, GCaMP6f-mediated 2-photon calcium imaging was performed through a coronally-implanted microprism (Andermann et al., 2013; Low et al., 2014), allowing for a field of view that preserved cortical laminar structure and included both prelimbic and infralimbic areas (**Fig. 2A-B**). This field of view allowed for monitoring of a 1500μM X 760μM area, with a median of 450 active neurons per animal in a single session (registering cells across consecutive sessions resulted in a median of 179 – see Methods). GCaMP was expressed pan-neuronally (hSyn-GCaMP6f) and in projection neurons (hSyn-DIO-GCaMP6f and rAAV2-Cre in PFC-VMS and PFC-MDT). Although these cell types would be analyzed separately, we began by characterizing the task-responsiveness of all labeled neurons. Examination of activity traces averaged over feature-matched trials revealed neurons whose response profiles displayed compound selectivity for multiple task features

, such as whisker stimulus and odor stimulus (**Fig. 2C**, left). While neurons with individual feature selectivity like the one shown in Fig. 2c (cell 244) were relatively rare, feature selectivity could be readily decoded from the combined firing rates of groups of neurons, (Fig. 2C, right). We therefore used high-dimensional population analysis to examine the coding properties of the recorded neurons.

Analysis of SEDS sessions with maximum-margin linear decoders (Christianini and Shawe-Taylor, 2000; Meyers et al., 2008) revealed robust representations of whisker and odor stimuli, response, trial outcome, and task rule (**Fig. 2D**). SVM decoders shown in Fig. 2 and in subsequent figures were iteratively trained on half of trials and tested on the held-out half (Methods). To assess whether these features were represented by persistent activity or by temporally dynamic trajectories, we performed a temporal autocorrelation analysis of the β-weights from these decoders, coefficients whose magnitude denotes the contribution of each cell to the discrimination of the feature (**Fig. S2**). While high correlation (~0.5) was seen between adjacent time-points (346ms time bins), these correlations dropped to near zero for non-adjacent time-points, revealing a coding scheme by which activity in different subsets of neurons contributed to feature discrimination at different moments of the trial and post-trial interval (**Fig. 2E**). Unlike the other four features plotted, the task rule was a relatively stable variable which rarely changed from trial to trial (median = 93 trials), yet this feature also exhibited trajectory-like representation, rather than representation by persistent activity. Principal component analysis performed on averaged traces (Churchland et al., 2012; Heetderks, 1978; Wang et al., 2018) from whisker-rule and odor-rule trials showed distinct trajectories for the first three components, indicating that a large portion of the population variance was accounted for by the task rule (**Fig. 2F**).

Task-critical action and outcome representations in post-trial activity

We next addressed how trial feature information is represented in the trial and post-trial epochs. Specifically, we wanted to know whether the response signal that appeared prior to lickport selection and trial termination in

Figure 3: Task-critical action and outcome representations in post-trial activity
A. Left: schematic of cross-validation approach, training and testing across cue directions. Upper right: SVM decoder accuracy for response trained/tested across cue direction (gray: CD; red: SEDS). Lower right: decoder accuracy for outcome trained/tested across cue direction (gray: CD; red: SEDS). N=4740 units from 21 animals.
B. Response-aligned values of correlation coefficients between β-weight values for whisker stimulus, odor stimulus, and response. Same population as **a**.
C. Means ± 95% confidence intervals for decoders tested on past (t-...) and future (t+...) trial features. Top: cue direction. Middle: response direction. Bottom: trial outcome. Red dashes: chance performance.
D. Histograms: fraction of cells modulated by task features over trial-aligned timepoints and through learning stages (see Fig. 1 and Methods). Top row: stimuli (solid: initially relevant stimulus; dashed: initially irrelevant stimulus). Middle row: response direction. Bottom row: trial outcome.

Fig. 2d constituted an action selection signal or a feedback signal. Resolution of this question required decoupling stimulus signal from response signal. We therefore trained SVM decoders on subsets of trials in which the cue indicated one direction (e.g. lick left) and tested on subsets of trials with the opposing cue direction (**Fig. 3A**). We first trained and tested these decoders on trials from earlier compound discrimination sessions, which required the same overt behavior as the set-shifting sessions (discrimination-guided licking) but that lacked a cognitive flexibility requirement. Here we observed significantly *below-chance* performance, indicating strong correlation in the neural activity associated with stimulus direction and response direction, and an absence of a detectable, independent response signal (Fig. 3A, upper panel). When the same procedure was performed on trials from SEDS sessions, a similar below-chance decoding accuracy preceded the response selection, while after the response, a significantly above-chance decoding accuracy emerged. This result shows a robust, independent response signal that is retrospective of the action it represents, and which is present only during

set-shifting. The same procedure was repeated for trial outcome and revealed a similar retrospective signal that was limited to set-shifting sessions (Fig. 3A, lower panel), again consistent with the interpretation that in this context, PFC cells encoded an action feedback signal, not an action selection signal.

Feedback monitoring in the SEDS task would, we reasoned, require not only a representation of response but also of the stimuli that preceded it. We therefore asked whether the neurons responsible for the response feedback signal seen in the post-trial epoch also showed compound selectivity for whisker and/or odor stimulus. To answer this question, we tested for correlation of the decoder β-weights for response, whisker stim and odor stim during and after the trial. We found that response β-weights were highly (Pearson ρ>0.4) correlated with whisker and odor β-weights across the neural population (**Fig. 3B**), meaning that a neuron highly selective for response was more likely to be selective for whisker or odor stim, particularly during the post-trial epoch. Such an elevated correlation was not seen between whisker and odor stimulus β-weights during the post-trial epoch, indicating that stimulus-response compound selectivity did not result from a subpopulation encoding all three trial features, but rather from subpopulations that encoded [whisker+response] and [odor+response] independently.

Response and outcome representations were seen to persist through the post-trial epoch, with decoder accuracies remaining elevated above chance until the onset window of the following trial (Fig. 2D, 3A). Given that animals required ~20 incongruent trials to abandon a rule after experimenter-imposed rule-changes (Fig. 1F), we next looked for evidence of outcome-related evidence accumulation that spanned multiple trials. We therefore ran SVM classifiers on whisker and odor stimuli, response, and outcome for past trials and found that, while whisker and odor representations were no longer detectable after one trial, response and outcome representations persisted for up to four trials, or up to 55 seconds (**Fig. 3C**). Importantly, the inability to decode these features at above-chance accuracies from future trials served as a negative control against autocorrelative behavioral effects.

To test whether response and outcome representations resulted from the performance of the overt behavior required for the task (discrimination-guided licking) or whether it resulted from the demands of the task related to cognitive flexibility, we tracked response and outcome signals across learning stages. While signal corresponding with the relevant stimulus was present from the earliest recorded session (compound discrimination) and remained throughout all subsequent task stages, signal corresponding with the irrelevant stimulus emerged only after the EDS session (**Fig. 3D**. While few neurons were modulated by response or outcome in early (e.g. compound discrimination) sessions, response and outcome-modulated activity emerged over the course of multiple task transitions (Fig. 3D.

Recent trial outcomes are colinearly multiplexed in a shared activity subspace

That trial feedback information could be decoded for up to four trials in the past led us to question how these past trial outcomes were being encoded. The temporal trajectory structure seen within trials (Fig. 2E-F) might play out across trials, with different groups of neurons inheriting representations of different past trials. Alternatively, these trial histories might be multiplexed by the same neurons, whose activity could be incrementally (linearly) modulated with each correct and incorrect trial result. To test these possibilities, we correlated β-weights from outcome decoders for successive past trials and found that they were indeed supported by correlated populations, whereby, for example, a neuron selective for the outcome of trial t-1 was relatively more likely to be selective for trial t-2 (**Fig. 4A**). As a further illustration of this phenomenon, the path of each neuron's incorrect/correct ratio over successive past trials tended to display a linear regression toward the unity line (**Fig. 4B**). An analysis of the distances and angles of pairs of past outcomes in neural activity space (**Fig. 4C**) revealed that more similar outcome combinations (e.g. incorrect-correct vs correct-incorrect) were closer to one another than more different outcome combinations (e.g. incorrect-incorrect vs correct-correct), and that all pairs of outcome vectors were more colinear (i.e. lower vector angles) than with vectors from shuffled trials (**Fig. 4D, S7**). Together, these results depict a coding scheme in which trial outcome influences the activity of a subset of the population in an incremental, colinear way with decreasing amplitude as the corresponding

Figure 4: Recent trial outcomes are colinearly multiplexed in a shared activity subspace
A. Trial-aligned values of correlation coefficients between β-weight values for trial t-1 outcome and t-2-6 outcomes. N=4740 cells from 21 animals.
B. Regression of past outcome preference from trials t-1 to t-4. Same cells as **a**. Insets: Mean z-scored activity differences (incorrect – correct) for example cells preferring correct (top) and incorrect (bottom) outcomes in trials t-1 to t-5.
C. Schematic diagram of Euclidean distance and vector angle computed from trial-averaged outcome (correct vs incorrect) conditions in neural activity space.
D. Left: Euclidean distances for all combinations of trial t-1 and t-2 outcomes. Z-scored activity differences are scaled by the mean β-weights for trial t-1 and t-2 outcome decoders. Bar heights and error bars correspond with mean ± 95% CI intervals over 500 data sub-samplings with replacement (same data as **b**). Right: angles between pairs of trial outcome vectors shown at left. All vectors use [incorrect, incorrect] as vertex. Bar heights and error bars are mean ± 95% CI as at left.

trial recedes into the past, rather than different groups of neurons representing trial outcome at different latencies from the corresponding trial.

Similar task responsiveness in two major PFC projection neuron populations

Is the task-related activity of a PFC neuron a function of its long-range connectivity profile or of its laminar location? While not mutually exclusive, these two proposed axes of functional specialization within cortex have motivated numerous studies in recent years, and noteworthy findings have lent support to both hypotheses.

Pioneering work by Otis et al, for example, found dissociated response properties in cortico-striatal and cortico-thalamic neurons in a Pavlovian conditioning task, with the former activated and the latter inhibited by reward-predicting conditioned stimuli (Otis et al., 2017). However, because the projection populations examined in that study were located in different cortical layers (layer 2/3-5 and layer 6, respectively), the contribution of laminar depth to coding differences was not directly addressed. Recent work by Liu et al demonstrated that ventral hippocampal afferents within mouse infralimbic cortex show both projection and layer specificity in their axonal targets, preferentially driving cortico-cortical projection neurons over other projection subtypes, and preferentially driving layer 5 neurons over layer 2/3 neurons (Liu and Carter, 2018). Thus, projection specificity and laminar depth may both contribute to functional specialization of subpopulations within PFC.

To examine the contributions of projection specificity and laminar depth to the performance of set-shifting behavior, we next focused our analysis on PFC-VMS and PFC-MDT neurons, two projection-defined PFC output populations whose target structures have both been implicated in supporting cognitive flexibility in prior work (Block et al., 2007; Floresco et al., 2006; Marton et al., 2018). Labeling of both populations within the same animals (**Fig. 5A**) using Cav2 and rAAV2 virus-mediated expression (Soudais et al., 2001; Tervo et al., 2016) revealed two largely non-overlapping cell types that were also spatially intermingled, with PFC-VMS neurons largely present in layers II/III and V, and PFC-MDT neurons largely present in layer V (**Fig. 5B**). Despite their being spatially interspersed, the two populations showed little (<5%) overlap (Fig. 5B).

Contrary to our expectations, the two populations showed a striking degree of overall similarity in their task-responsiveness. Both exhibited robust representation over similar time-courses for all five examined task features (**Fig. 5C**). Both populations showed modulation by trial outcome that emerged through successive learning stages (**Fig. 5D**). Principle component analysis yielded similar trial-aligned temporal profiles for the main components (**Fig. 5E**), revealing similar underlying structures accounting for the largest proportions of variance across the populations. Additionally, similar fractions of neurons in each population were excited vs inhibited by trial onset (**Fig. 5F**) and the two cells types each showed roughly equal proportions of correct-preferring and incorrect-preferring neurons (**Fig. 5G**).

The finding of durable representations of trial feedback signals over successive trials raised the possibility that the task-critical role of the PFC might be feedback monitoring, a prediction consistent with results from visually-guided rule-switching experiments (Biró et al., 2019). The robust representation of stimuli and task rule, particularly during stimulus presentation, raised the (not exclusive) possibility that PFC activity might be critical to guiding rule-dependent action selection in the task. The fact that PFC-VMS and PFC-MDT populations exhibited similar physiological task responsiveness suggested that both cell types might critically support set-shifting performance. We next tested these possibilities using projection-targeted optogenetics.

Interference with feedback monitoring in PFC projection populations impairs set-shifting

We used the soma-targeted anion-conducting channelrhodopsin stGtACR2 (Mahn et al., 2018) to inhibit PFC activity during set-shifting in one of three temporally-controlled regimes (**Fig. 6A**). Photoactivating light stimulation was delivered during trials (0.5s trial-ready cue period, 2.5s stimulus presentation, and ≤1.5s response window), or during inter-trial intervals (8-10s epoch triggered on response lick) following CG or ICG trials (**Fig. 6B**).

Trial blocks (beginning with rule switch and ending with the animal reaching criterion for the new rule) alternated light off/on. Photoactivation of pan-neuronally-expressed stGtACR2 during the inter-trial interval following incongruent trials impaired performance on ICG trials (N=16, p=0.0005 Fig. **6C**) but not on CG trials (N=12, p=0.5, Fig. **6D**). No effect of light was seen for control tdTomato-expressing animals (N=13, p=0.4 and p=0.9 for ICG and CG trials, respectively, Fig. 6c-d).

This effect was also seen in animals expressing stGtACR2 in PFC-VMS and PFC-MDT projection neurons (VMS: N=17, p=0.007, 0.1 for ICG, CG respectively; MDT: N=18, p=0.01, 0.1 for ICG, CG respectively, Fig. 6c-d). When light was delivered following CG trials, no impairment was seen on ICG trial performance for either projection population (N=18, 18, p=0.1, 0.1 for VMS and MDT respectively, Fig. **6E**), indicating that the impairment seen with post-ICG inhibition reflected an interference with prior trial feedback, rather than preparation for the subsequent trials. This finding of selective impairment on ICG trials following inhibition during the inter-trial interval following ICG trials, but not CG trials, also provides further support for the conclusion that

Figure 5: PFC-VMS and PFC-MDT projection populations exhibit similar task response properties

A.   Diagram and *in vivo* images of dual projection neuron labeling approach. Top: PFC-VMS neurons. Bottom: PFC-MDT neurons. Left: diagram of viral targeting. Right: isosbestic GCaMP6f images for projection-specific neural populations.

B.   Dual PFC->VMS and PFC->MDT projection labeling in a single brain. Top: green is rAAV2-CAG-mNeuroGFP in MDT; red is AAV1-hSyn-DIO-nls-tdTomato in PFC and CAV2-Cre in VMS (DAPI in blue). Bottom left: relative green, red, and overlapping cell counts. Bottom right: median relative overlap (N = 8 sections, 2 animals).

C.   Trial-aligned SVM decoder accuracy for whisker stim, odor stim, response, outcome and rule. N=1770 units from 8 animals (PFC-VMS, top) and 1115 units from 9 animals (PFC-MDT, bottom).

D.   Histograms: proportion of cells modulated (rank sum p<0.01) by trial outcome over trial-aligned timepoints and through learning stages for PFC-VMS (red) and PFC-MDT (green) units (same cells as in **c**).

E.   First four principle components for trial-averaged activity histograms in PFC-VMS (left) and PFC-MDT (right) cells.

F.   Trial-averaged activity histograms for PFC-VMS (left) and PFC-MDT (right) cells. Each cell's mean trace is normalized to its peak value, and cells are sorted by time of peak excitation (top) or inhibition (bottom).

G.   Mean outcome difference (incorrect – correct) traces for PFC-VMS and PFC-MDT neurons. Each cell's mean difference trace is normalized to its peak value, and cells are sorted from most strongly preferring correct outcomes (top) to incorrect outcomes (bottom).

PFC activity supports set-shifting performance by encoding trial feedback information, rather than guiding action selection.

Strikingly, neither PFC-VMS nor PFC-MDT activity was necessary for execution of the rule-guided response, as no impairment was seen with photoactivation during ICG trials (N=10, 10, p=0.4, 0.4 for VMS and MDT respectively, Fig. **6F**).

The impairment on ICG trials that resulted from photoactivation following ICG trials was seen both early and late in trial blocks (first 20 incongruent trials: p=0.02, 0.03; incongruent trials after 20: p=0.01, 0.007 for VMS, MDT respectively). Inhibition during the beginning (first 4.5 seconds) or end (beginning 4.5 seconds into the 8-10 second post-trial interval) alone was insufficient to impair behavior (N=13,15; beginning: p=0.4, 0.2; end: p=0.3, 0.06 for VMS and MDT, respectively), which indicated that the post-trial activity encoding feedback information was robust to transient perturbation (**Table 1**).

Together, these behavioral results illustrate a critical role for post-trial activity in both PFC-VMS and PFC-MDT neurons in enabling dynamic rule-guided responding.

Feedback-related activity follows an anatomical gradient independent of projection subtype.

The surprising similarity of both task-related activity and the task-critical function of the PFC-VMS and PFC-MDT pathways left open the question of whether any laminar topological organization is shared between the two populations. We therefore examined whether the relative locations of neurons within the PFC imaging field affected task-related activity. As has been demonstrated in sensory cortex(Smith and Kohn, 2008), temporal correlations across pairs of simultaneously-recorded neurons decayed with distance (**Fig. 7A**), demonstrating an association between anatomical location and temporal activation within the context of the task.

Because neurons exhibited both increased and decreased activity profiles in response to trial onset (Fig. 5f), we quantified the population variance of trial-averaged waveforms in order to capture the magnitudes of neuron responses, including those both excited and inhibited at each time point. More deeply located subsets of neurons (further from pial surface) showed greater population variance (**Fig. 7B**, exhibiting more diverse trial responses than more superficially located subsets. While more superficial neurons exhibited elevated activity during the post-trial interval (**Fig. 7C**), increased response magnitude (absolute value of z-scored activity) was seen in deeper neurons from both projection populations, and with similar spatial correlation (**Fig. 7D**). Thus, while superficial neurons maintained relatively uniform, elevated mean activity during the post-trial interval, deeper neurons exhibited a wider diversity of responses.

In addition to trial responsiveness, task-related information was heterogeneously distributed across the cortical laminar axis, with more deeply situated neurons contributing greater β-weights to linear decoders for response and outcome during the inter-trial interval (**Fig. 7E**). General linear models for the β-weights of response and outcome yielded significant coefficients for neuron depth but not for projection type (**Fig. 7F**). Because the distribution of β-weights is non-normal, with a few neurons contributing most of the high values, we compared sample distributions across groups of neurons from four quartiles of pial depth and found that sample distributions differed significantly across more distance quartiles for both response and outcome weights in both projection types (**Fig. 7G**). Confirming the differential contribution of neurons from deep and superficial areas to decoder accuracy, binned subpopulations from spatial quartiles showed differential accuracies, with decoders comprised of deeper neurons showing greater accuracy for response and outcome during the ITI than those comprised of more superficial neurons (**Fig. 7H**).

Discussion

We combined projection-targeted *in vivo* calcium imaging with optogenetics to identify the relative contributions of two major PFC projection subtypes to a validated behavioral model of cognitive flexibility in mice. We found that PFC activity persistently tracked recent trial histories, sustaining representations of response and outcome, but not stimulus identity or cue direction, over multiple trials. The behavioral importance of these feedback signals was confirmed by temporally precise inhibition, which revealed that PFC activity following the rule-informative ICG trials, but not following rule-uninformative CG trials, impaired behavior, and that neither population was critical to the execution of rule-dependent responses in real-time. PFC-VMS and PFC-MDT

Figure 6: Dynamic rule-guided performance requires PFC-VMS and PFC-MDT activity selectively following rule-informative trials

A. Cre-mediated expression of gtACR-FusionRed in PFC-VMS neurons, and fiber tracks from chronically implanted bilateral optical fibers (fixed dissue, DAPI in blue).

B. Schematic diagram of light delivery conditions.

C. Effect of light stimulation following incongruent trials on incongruent trial performance (fraction of trials correct). Left: viral control animals (td-Tomato expressed in PFC); second from left: animals with pan-neuronal expression of stgtACR in PFC; third from left: animals with stgtACR expressed in PFC-VMS projection neurons; right: animals with stgtACR expressed in PFC-MDT projection neurons. N = 13, 16, 17, 18 resp. Sign rank p = 0.4, 0.0005, 0.007, 0.01, resp.

D. Effect of light stimulation during ITIs following incongruent trials on congruent trial performance (proportion of trials correct). Left to right: as in (a) above. N = 13, 16, 17, 18, resp. Sign rank p = 0.9, 0.5, 0.9, 0.07, resp.

E. Effect of light stimulation during ITIs following congruent trials on incongruent trial performance. N = 18, 18, resp. Sign rank p = 0.1 and 0.1, resp.

F. Effect of light stimulation during trials (stimulus presentation + response window) on incongruent trial performance. N = 10 and 10, resp. Sign rank p = 0.4 and 0.4, resp.

projections were found to play parallel roles in the task, exhibiting similar, task-critical activity, and both exhibited a spatial functional gradient, with more deeply located neurons in both groups displaying more task-critical information during the feedback epoch.

While the primary motivation behind the development of this task was to functionally emulate clinical assessments of cognitive flexibility, there is also an extensive literature within human psychology that concerns task-switching and the study of switching costs (Kiesel et al., 2010). Indeed, we note that, as with previous rodent behavioral models of task-switching (Rodgers and DeWeese, 2014), mice here exhibit switching costs that persist well beyond the first trial (Fig. 1F), that they appear to use model-free learning, and that this differentiates them from their human counterparts. We also note our use of the term *serial extradimensional set-shifting*, where the term *task switching* is more often used to refer to paradigms with multiple, bidirectional changes in stimulus-response-reward contingencies. We do so to differentiate this task from the reversal, intradimentional, and first-time extradimensional transitions, which here form distinct learning stages and which are also forms of task-switching. Nevertheless, increased trials to criterion for extradimensional vs. intradimensional shifts (Fig. 1C) and elevated reaction times on incongruent vs. congruent trials (Meiron and Kessler, 2008) are hallmarks of shifting across feature dimensions.

While the finding of a critical role for PFC output neurons in supporting feedback monitoring in set-shifting was in keeping with our initial hypothesis, the lack of behavioral impairment seen when inhibiting these pathways during trials was striking and unexpected, particularly given the robust activation of stimulus-selective representations seen across PFC activity during the stimulus presentation and response selection epoch. The observation of strong task-related activity, which is nevertheless evidently not critical for task-critical performance, leaves unanswered the question of what adaptive information-processing purpose this activity may serve and reinforces the need to assess the role of neural signals through modulation, rather than assuming a critical role for observed activity.

The PFC-MDT population targeted in these experiments is confined largely to layer 5. This distinguishes it from layer 6 cortico-thalamic neurons, also found in PFC and elsewhere in cortex. It should be noted that these projections were labeled using a rAAV2 vector, which has selective tropism for layer 5 PFC-MDT neurons, rather than canine adenovirus (Cav2), which preferentially infects axons of layer 6 neurons (Collins et al., 2018) (Fig. S4). This differential labeling, along with pronounced task differences, may account for differential effects of projection neuron inhibition relative to other PFC projection targeting studies (Otis et al., 2017).

Feedback signals within PFC are modulated by expectation and by cognitive states, including attention (Smout et al., 2019). Extensive research in feedback-related negativity (Hauser et al., 2013) and reward prediction error (Garrison et al., 2013) support the idea that feedback-related signals in prefrontal areas are dynamically mediated by ongoing task demands and are disrupted in disorders in which set-shifting is also disrupted (Toyomaki et al., 2017). The overt behavior performed in the compound discrimination task is identical to that performed during set-shifting: the animal is presented with cues comprised of two stimulus modalities and responds to one with a left or right lick. The fact that dedicated response and outcome signals are absent during early compound discrimination sessions and emerge as changing task rules require inferential learning (Fig. 3A,D), lends strong support to the idea that these signals are not innate but are brought about by the demands of feedback monitoring. Future work will seek to elucidate potential neuromodulatory and plasticity-related mechanisms behind these signals.

Acknowledgements

T.S. is supported by NIH (1K99MH117271), the Brain and Behavior Research Institute, and the Leon Levy Foundation. C.L. is supported by NIH (R01 NS095441, MH109685) and by the Rita Allen Foundation, Hope for Depression Research Foundation, Pritzker Neuropsychiatric Disorders Research Foundation, and the Foundation for OCD Research. The authors thank S. Fusi, F. Stefanini, and L. Grosenick for consultation on the planning and implementation of data analyses.

Author Contributions

CL and TS designed the behavioral task and planned the study. TS and MS performed the injection and implantation surgeries and gathered the behavioral data. TS and MS performed the neural recordings. TS analyzed the behavioral and neural data. TS and CL wrote and edited the manuscript.

Figure 7: Post-trial feedback-related activity in PFC-VMS and PFC-MDT follows a common topological gradient

A.   Spatial distance of all simultaneously recorded cell pairs vs. temporal correlation (N = 1562607 pairs from 4740 cells in 21 animals). White plot is mean ± SEM for 50µM bins.

B.   Cross-cell variance of trial-averaged activity, over trial-aligned timepoints and relative distance from pial surface. Warmer colors correspond with higher cross-unit variance per spatiotemporal bin. Units are means of z-scored activity. Same cells as in a.

C.   Correlation of population activity by pial distance over trial-aligned timepoints. Asterisks denote significant for Bonferoni-corrected rank sum scores (p<0.01).

D.   Correlation of absolute values of z-scored activity with relative pial distance during ITI for PFC-VMS and PFC-MDT neurons. PFC-VMS Spearman ρ = 0.16, p = 8e-79, 1770 neurons, 8 animals; PFC-MDT Spearman ρ = 0.12, p = 3e-52, n=1155 neurons, 9 animals.

E.   Correlations of decoder β-weights with pial distance at trial-aligned timepoints. Larger circles, Spearman p<0.01. PFC-VMS in red, PFC-MDT in green (same units as in d).

F.   GLM relating SVM β-weights for response and outcome during ITI to pial distance and cell type (same units as in d, e). Left, response: SVM β-weights by pial distance (i, t = -5.07, p = 4e-7), and cell type (ii, t = -0.98, p = 0.33). Interaction: t = -1.08, p = 0.28. Right, outcome: SVM β-weights by pial distance (iii, t = -4.68, p = 3e-6), and cell type (iv, t = 1.0, p = 0.31). Interaction: t = 0.53, p = 0.6.

G.   Distributions of β-weights for response and outcome during ITI for PFC-VMS and PFC-MDT populations, binned by quartile of relative pial distance. Red cross-bars are means, black cross-bars are medians. Checkered plots are coefficients from 2-sample Kolmogorov-Smirnov distribution, comparing distribution similarity across all pairs of quartiles. Red asterisks are comparisons with Bonferroni-corrected significance.

H.   Mean ± 95% confidence intervals for response and outcome decoders by relative pial distance quartiles. Same cells as d-g.

Declaration of Interests: Competing interests: The authors declare no competing interests.

Table 1: Summary of behavioral statistics for optogenetic experiment (signed rank results)

photoactivation	rials	alue	tdt	hSyn	FC-VMS	FC-MDT
post-ICG	ICG	N	13	16	17	18
		rank	58	135	133	142
		p	0.4	**0.0005**	**0.007**	**0.01**
post-ICG	CG	N	13	16	17	18
		rank	47	55	74	44
		p	0.9	0.5	0.9	0.07
post-CG	ICG	N			18	18
		rank			123	123
		p			0.1	0.1
during ICG	ICG	N			10	10
		rank			37	36
		p			0.4	0.4
early ITI post-ICG	ICG	N			13	15
		rank			59	85
		p			0.38	0.17
late ITI post-ICG	ICG	N			13	15

		rank			53	93
		p			0.3	0.06
first 20 block trials, post-ICG	ICG	N			17	18
		rank			126	136
		p			**0.02**	**0.03**
trials after first 20, post-ICG	ICG	N			17	18
		rank			130	147
		p			**0.01**	**0.007**

METHODS

LEAD CONTACT AND MATERIALS AVAILABILITY: This study did not generate new unique reagents.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Male C57BL/6 mice (Jackson Labs) were used for all experiments, aged 6-10 weeks at first use. Mice were housed in a Weill Cornell Medical College facility and were maintained on a 12-hour light-dark cycle. Except when water-restricted for the purpose of behavioral training and testing, all mice were given *ad libitum* access to food and water. Littermates   underwent prism or fiber implant surgeries within the same week, and mice were group-housed with littermates with the same surgery status. All procedures were approved by the Weill Cornell Medicine IACUC. Sample sizes for each experiment were determined using power analysis estimates computed in Matlab, based on anticipated effect sizes that were estimated from previously published reports whenever they were available, and were powered to detect moderate, biologically meaningful effect sizes.

METHOD DETAILS

Surgery

Animals were placed inside a flow box and anesthetized with isoflurane gas (2%) until sedated, at which point they were placed in a stereotax and maintained on 0.5% isoflurane for the duration of the surgery. Scalp hair was trimmed away, and a midline incision was made using fine surgical scissors (Fine Science Tools), exposing the skull. The periosteum was bluntly dissected away and bupivacaine (0.05 mL, 5 mg/kg) was topically applied. For prism implantation, a large rectangular craniotomy was made above left PFC, extending from 1.5mm anterior to 3.7mm anterior, and from 2.0mm lateral (left) to 0.2mm lateral (right, across midline).

A 0.5-mm burr (Fine Science Tools) and a high-speed hand dental drill (Osada) were used, taking great care not to compress brain tissue or damage the sagittal venous sinus. In the event of venous bleeding, Gelfoam (Pfizer) was applied to the dura surface to accelerate clotting. Gentle irrigation with phosphate-buffered saline (137 mM NaCl, 27 mM KCl, 10 mM phosphate buffer, VWR) was used to clear debris at regular intervals. The dura beneath the craniotomy was removed using the tip of a 26g insulin syringe (VWR) and fine forceps (Fine Science Tools).

Chronically implanted microprisms (1.5mm X 1.5mm X 3mm; M/L,A/P,D/V), from OptoSigma (BK7 borosilicate glass with aluminum hypotenuse and silicon dioxide coating), were implanted at a depth of 2.3mm ventral to brain surface using a stereotaxic micromanipulator (Kopf). During implantation, the prism was held in place using vacuum suction via an 18G blunt needle. As in previous studies (Andermann et al., 2013; Low et al., 2014) minimal reactive gliosis was seen in the coronal imaging field, and maximum calcium-mediated fluorescence was seen 50-150μM past the prism face; therefore imaging planes were confined to this depth.

For PFC-VMS and PFC-MDT projection targeting, additional craniotomies were made at 1.25mm/1.25mm A/L and 1.2mm/0.35mm P/L, respectively. All head-fixed animals received custom-machined stainless steel head plates affixed to skull surface with Metabond dental cement. Head plates featured a circular central aperture centered around the imaging field (9mm I.D.), with right and left securing arms (25mm total width) that accommodated 0-80 socket screws (0.38g in total). Sterile eye lubricant (Puralube, FischerSci) was administered to prevent corneal drying, and a microwavable heating pad (Snugglesafe) was used to maintain body temperature. Metacam (1 mg/kg, i.p.) was administered after surgery as a prophylactic analgesic.

Viral Transduction

AAV1 of titer exceeding 10^{12}vg/ml (Vector Biolabs, UNC Vector Core and Addgene) was used to package the plasmids. For imaging experiments, AAV1-hSyn-GCaMP6f or AAV1-hSyn-DIO-GCaMP6f was targeted to PFC. Injection coordinates for PFC were 1.75mm anterior. Two parallel injection tracks were made at 0.2mm and 0.5mm lateral, and in each of these tracks, two D/V sites received infusions, at 2.0mm and 1.5mm ventral to brain surface. For rAAV2-Cre projection targeting, VMS injections were delivered at 1.25mm/1.25mm/4.7mm A/L/V, and MDT injections at 1.2mm/0.35mm/3.2mm P/L/V. Hamilton syringes and beveled 36G or 33G NanoFil needles (WPI) were used, and at each site, the needle was allowed to sit 5min to allow for tissue settling before infusion. Virus was infused at a rate of 50nL/min, for a total of 250nL per site.

For the dual labeling tracing experiment (Fig. 5B), rAAV2-mNeuroGFP was injected in MDT (500nL), rAAV2-Cre was injected in VMS (500nL), and AAV1-NLS-tdTomato was injected in PFC (500nL).

Muscimol silencing

Animals received chronically implanted 26G bilateral stainless steel guide cannulae (Plastics One), implanted at 1.75mm/0.35mm/1mm A/L/V. After undergoing the training and task transition sequence up to and including Reversal, animals underwent muscimol infusion. In a familiar cage, bilateral internal infusion cannulae were inserted into the guide cannulae, protruding 0.5mm from the end of the guide cannulae, and were left for 5min to allow tissue to settle. Muscimol ($1\mu g/\mu L$) or physiological saline (0.9%) was infused at a rate of 50nL/min, for a total of $0.25\mu L$. Five minutes after completion of the infusion, the internal cannulae were removed and them mouse underwent behavior al testing.

Optogenetic implantation and stimulation

Bilateral fibers were implanted over PFC (Thorlabs dual fiber cannulae, $700\mu M$ center-to-center spacing, $200\mu M$ core) at 0.35mm lateral, 1.75 anterior, 1.2mm ventral to brain surface. Light was delivered via Thorlabs M470F3 fiber-coupled LED, 1.5-3mW light power from each fiber. In the trial-concurrent stimulation condition, photoactivation was initiated at the same time as the 500ms trial-ready cue and persisted through the 2.5s stimulus presentation epoch until either a lick terminated the trial or the 1.5s response period expired. In the full inter-trial interval stimulation conditions, photoactivation was initiated concurrently with lickspout selection and persisted through the 8-10s post-trial epoch, terminating with the onset of the following trial-ready cue. In the experiment testing the effects of inhibition during the beginning of the inter-trial interval, the photoactivation was initiated concurrently with lickspout selection and persisted for 4.5 seconds before terminating; in the experiment testing the effects of inhibition during the end of the inter-trial interval, photoactivation was initiated beginning at 4.5 seconds following lickspout selection, persisting for the remainder of the 8-10s inter-trial interval. Light was delivered in alternating trial blocks, with no light delivered in the initial block. Sessions in which animals failed to reach criterion on the initial block were re-run until the animal reached criterion.

2-photon imaging

Imaging was performed via an Olympus 10x 0.6NA objective, with 8mm working distance. All images were acquired using a commercial two-photon laser-scanning microscope (Olympus RS) equipped with a scanning galvanometer and a Spectra-Physics Mai Tai DeepSee laser tuned to 920nm, operating at 300-500mW. Fluorescence was recorded through gallium arsenide phosphide (GaAsP) detectors using the Fluoview acquisition software (Olympus) using a green light emission bandpass filter (Semrock). All imaging experiments began by obtaining an isosbestic anatomical scan (810nm 2P excitation light) to aid in relocating the same sites over multiple sessions. Calcium signals were acquired at 256 X 130 pixel resolution, covering a $1500\mu M$ X $760\mu M$

field of view, with a μM/pixel ratio of 5.85. The scan time was 346ms, with a frame rate of 2.89Hz. All calcium imaging experiments occurred in awake mice. For analysis of SEDS sessions, sessions from each animal were concatenated using CellReg (Sheintuch et al., 2017) and non-rigid transformation. Neurons were modeled with a maximal centroid distance of 15μM and a threshold correlation of 0.65. Z-scoring of deconvolved activity traces was performed prior to concatenating across sessions. Data sets for each of 21 animals represent a median of 3 sessions (2.5,4 upper, lower quartiles), 1462 trials (1034, 1667 upper, lower quartiles), and 11 set shifts (8.5, 17.5 upper, lower quartiles).

<u>Image processing</u>

Videos were motion-corrected using NoRMCorre (Pnevmatikakis and Giovannucci, 2017) in Matlab (Mathworks), and a constrained non-negative matrix factorization-based source extraction method was used to denoise, deconvolve and demix the videos to produce neural traces (using the extensively validated CNMFE package with OASIS signal deconvolution (Pnevmatikakis et al., 2016)). The resulting traces were deconvolved calcium traces corresponding with estimated event rates, which were then z-scored over the full session to normalize. Calcium signals from sequential sessions were concatenated using nonrigid co-registration of spatial cell footprints using CellReg (Sheintuch et al., 2017), (Fig. S2).

<u>Behavioral training and testing protocols</u>

Animals recovered 14 days from surgery before being placed on water restriction for four days, after which they typically drank 1-1.2mL/day. Behavioral training and testing procedures were carried out 5 days per week, with each animal undergoing one session per day. Animals underwent two days of hand-feeding, in which they were handled for up to ten minutes by the experimenter while receiving water from a 1mL syringe with a rounded stainless steel gavage needle. The following day they underwent a *Habituation-1* session in the behavioral apparatus, which consisted of an aluminum restraint tube with dual lickspouts positioned at one end. This tube was 27mm in diameter, a width calibrated to allow the animal to groom and adjust its posture during sessions but which prevented significant lateral or vertical body movement. During *Habituation-1*, lickspouts were alternately armed so that a single lick would trigger delivery of a 3uL water bolus, and the identity of the armed lickspout changed every 1-4 trials, with a 1.5s timeout after each bolus delivery. This alternating schedule forced the animal to explore both lickspouts equally to maximize rewards. Animals would periodically venture out of the restraint tube to explore the behavior chamber, at which point the experimenter would guide them back into the tube. The session terminated when the animal reached satiety, and the animal was considered to have passed this stage when it had consumed 500μL in a session. The *Habituation-2* stage consisted of the same lick/dispense schedule, but during this session the animal was head-restrained for the first time. Here again, the animal needed to consume 500μL to pass. After this stage animals underwent *Habituation-3*, in which lickspouts alternated every 20 lick/delivery trials, to establish left/right trial blocks.

After this, animals underwent a *Shaping-1* session, which introduced the trial structure to be used for the rest of the sessions: a 500ms white noise trial-start cue was presented, followed by a 2.5s stimulus (whisker or odor, depending on the training sequence, which was counter-balanced across animals), followed by a 1.5s response window, during which the correct lickspout was armed. When the response rate fell to between 0.3 and 0.6 within a 10-trial moving window, reward would be dispensed regardless of whether the animal licked, in order to promote licking. Trials were blocked (20 right, 20 left), and a lick to the incorrect side would not terminate the trial (the animal as allowed to continue until it licked the correct spout). At the start of each trial block, a single 3uL reward was dispensed at the now-rewarded spout. Criterion for passing *Shaping-1* was 500μL consumed. *Shaping-2* differed from *Shaping-1* in that trials were randomized between sides rather than blocked. Criterion here, and for all remaining sessions up to SEDS, entailed reaching 80% correct within a 30-trial moving window, and simultaneously performing above 50% on both left and right trials, at any point within the session. Following this session, animals underwent *Simple Discrimination*, in which a lick to the wrong side during the response window would terminate the trial. After reaching criterion, animals underwent *Compound Discrimination* (CD), *Intradimensional Shift-1* (IDS1), *Reversal* (Rev), *Extradimensional Shift* (EDS), *Intradimensional Shift-2* (IDS2), and finally *Serial Extradimensional Shifting* (SEDS). Sessions CD through IDSB consisted of up to 350 trials, while SEDS sessions continued until the animal reached 1100uL in rewards. In all sessions, failure to respond

in ten consecutive trials would trigger the end of the session. Behavioral data was sent from an Arduino board to a desktop computer through serial input and logged as rich text format files.

To counteract animals' tendency to periodically exhibit stereotyped direction bias, left/right stimulus probability was dynamically computed as a function of the animal's concurrent direction bias coefficient. A 10-trial moving window was populated with a -1 for each incorrect left response, a 1 for each incorrect right response, and a 0 for each correct response. For each trial, the coefficient was the mean over the 10 previous trials. A coefficient absolute value between 0 and 0.25 resulted in 50% left-right stimulus probability; an absolute value between 0.25 and 0.75 resulted in a 33% chance of a stimulus toward the direction of bias; an absolute value between 0.75 and 1 resulted in a 17% chance of a stimulus toward the direction of bias. Trial runs exhibiting lick bias coefficients greater than 0.75 were therefore excluded from analysis.

Out of 21 mice used in calcium imaging, 14 underwent whisker discrimination first, while 7 underwent odor discrimination first; no differences in task-related activity were observed between the two groups. All mice used in the muscimol and optogenetic experiments underwent whisker discrimination first. Whisker stimulus pairs consisted of either 35Hz vs. 155Hz pure tones, or a 210Hz tone vs. a Poisson click train, all delivered bilaterally. Odor cue pairs consisted of sesame oil vs. olive oil, or almond oil vs. orange extract. Clean air flowed through the odor presentation port (0.2L/min through 1/32" Teflon tubing), and clean air was completely replaced by air channeled through odorant bottles within 50ms from stimulus onset.

QUANTIFICATION AND STATISTICAL ANALYSIS

Linear decoding with SVM maximum margin classifiers

Population decoding was performed with maximum-margin linear decoders (Matlab fitcsvm). At each iteration data were separated into equally sized training and testing sets, and in each set, trials were separated into equal numbers from each binary value for the feature of interest. To combine inputs from neurons across multiple sessions, and to remove inter-trial variability each of these trial groups was than randomly separated into 16 trial groups, each of which was averaged into one *supertrial* to serve as a decoder input (16 supertrials for one value of the feature of interest and 16 from the opposing value). This number was calibrated to ensure robust trial subsampling within training and testing sets while removing sufficient trial variability to reveal accurate decoding. Confidence intervals were obtained by iterating over 500 subsamplings and identifying centiles. Decoder accuracies performed on data from inter-trial intervals were run on binned activity values over the 4-8 second post response epoch unless otherwise specified.

For comparison of temporal autocorrelations in feature coding properties, Pearson's partial correlation was performed using Matlab's native *partialcorr* function, in order to account for an elevated floor of β-weight correlations, owing to the outsized contribution of more active neurons to feature coding. This analysis was repeated using simple Spearman's and Pearson's correlations, with similar results (for each feature, high correlation with adjacent timepoints and low correlation with non-adjacent timepoints).

Principle components analysis (PCA)

PCA was used to identify the trial-aligned activity patterns responsible for the greatest variance within the activity space of the neural traces. For the analysis in Figure 2f, the data were divided into odor rule and whisker rule trials, then trials were averaged for each neuron, resulting in a neuron x time matrix for each rule. PCA was performed using Matlab's native *pca* function, and the first three resulting components were plotted through trial-aligned time. For the analysis performed in Figure 5e, PCA was performed separately on trial-averaged traces from PFC-VMS and PFC-MDT sets (neuron x time).

Analysis of covariance (ANCOVA)

To quantify the fraction of cells at each learning stage modulated by response direction, we adopted an analysis of covariance approach to control for two covariates: stimulus direction and the previous trial's response direction. ANCOVA was performed using Matlab's *anovan* function, to which these two covariates were passed

as continuous variables. Because the introduction of two additional degrees of freedom resulted in the dilution of explanatory variance across negative cofounders, we relaxed the threshold p-value to 0.02 for this analysis.

General Linear Model

Matlab's *fitglm* function was used to fit a GLM model to projection cell type and topological location (depth from pial surface). β-weight was used as the dependent variable, and cell type (PFC-VMS vs PFC-MDT) and relative pial depth (microns from the most medial neuron in the field of view) were passed as independent variables, along with their interaction.

$$\hat{Y} = \beta_0 + \beta_1 p + \beta_2 d + \beta_3 p \cdot d$$

Euclidean distance and vector angles

Euclidean distance between pairs of trial averaged conditions was computed as the square root of the sum of the differences between the two conditions across all neurons in the set.

$$d_{xy} = \sqrt{\sum_{i=1}^{n}(x_i - y_i)^2}$$

Euclidean distance was also computed on activity vectors scaled by the vector weights of their corresponding SVM decoders in order to scale the amplitude of feature-selective cells relative to noise.

Scaling by β-weights proved useful for comparing feature-averaged activity vectors in activity space, as noisy or randomly selected vectors in a sufficiently high-dimensional space tends to approach orthogonality. Vector angles were computed as the dot product of the difference vectors, divided by the product of their norms. This value was converted into an angle of degrees by taking its inverse cosine.

$$\theta_{xy} = \cos^{-1}(x \cdot y) \div (\|x\| \cdot \|y\|)$$

DATA AVAILABITY: All data and Matlab analysis scripts are available from the corresponding author upon reasonable request.

Figure S5: Retrograde labeling of PFC-MDT neurons by CAV2-FLP-mediated tdTomato (red) and AAV2Retro-GFP (green). Both viruses were infused in MDT as a mixture by a single injection. FRT-tdTomato as infused in PFC in a separate injection.

Figure S6: Cross-animal distribution of β-weights from feature decoders. Box plots are medians and quartiles taken from the epochs of peak decoding accuracy for each feature (stimulus delivery epoch for whisker and odor stims, post-trial epoch for response and outcome, both epochs for task rule).

Figure S7: Euclidean distances between incorrect and correct outcomes for trials t-1 through t-7.

References

Adesnik, H., and Naka, A. (2018). Cracking the Function of Layers in the Sensory Cortex. Neuron *VOLUME 100, ISSUE 5, P1028-1043*.

Andermann, M., Gilfoy, N., Goldey, G., Sachdev, R., Wölfel, M., McCormick, D., Reid, R., and Levene, M. (2013). Chronic Cellular Imaging of Entire Cortical Columns in Awake Mice Using Microprisms. Neuron *80*, 900–913.

Aoki, S., Liu, A.W., Zucca, A., Zucca, S., and Wickens, J.R. (2015). Role of Striatal Cholinergic Interneurons in Set-Shifting in the Rat. J. Neurosci. *35*, 9424–9431.

Biró, S., Lasztóczi, B., and Klausberger, T. (2019). A Visual Two-Choice Rule-Switch Task for Head-Fixed Mice. Front Behav. Neurosci. *06 June*.

Birrell, J., and Brown, V. (2000). Medial frontal cortex mediates perceptual attentional set shifting in the rat. J. Neurosci. *20*, 4320–4324.

Bissonette, G., and Roesch, M. (2015). Neural correlates of rules and conflict in medial prefrontal cortex during decision and feedback epochs. Front Behav. Neurosci. *9: 266.*

Bissonette, G., Powell, E., and Roesch, M. (2013). Neural structures underlying set-shifting: roles of medial prefrontal cortex and anterior cingulate cortex. Behav. Brain Res. *250*, 91–101.

Block, A., Dhanji, H., Thompson-Tardif, S., and Floresco, S. (2007). Thalamic-prefrontal cortical-ventral striatal circuitry mediates dissociable components of strategy set shifting. Cereb. Cortex *17*.

Bortolato, B., Miskowiak, K., Köhler, C., Maes, M., Fernandez, B., Berk, M., and Carvalho, A. (2016). Cognitive remission: a novel objective for the treatment of major depression? BMC Med. *14, Article number: 9*.

Brigman, J., Bussey, T., Saksida, L., and Rothblat, L. (2005). Discrimination of multidimensional visual stimuli by mice: intra- and extradimensional set shifts. Behav. Neurosci. *119*, 839–842.

Ceaser, A., Goldberg, T., Egan, M., McMahon, R., Weinberger, D., and Gold, J. (2008). Set-shifting ability and schizophrenia: a marker of clinical illness or an intermediate phenotype? Biol Psychiatry *64*, 782–788.

Christianini, N., and Shawe-Taylor, J. (2000). An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods (Cambridge, UK: Cambridge University Press).

Churchland, M., Cunningham, J., Kaufman, M., Foster, J., Nuyujukian, P., Ryu, S., and Shenoy, K. (2012). Neural population dynamics during reaching. Nature *July 5; 487(7405): 51–56.*

Collins, D., Anastasiades, P., Marlin, J., and Carter, A. (2018). Reciprocal Circuits Linking the Prefrontal Cortex with Dorsal and Ventral Thalamic Nuclei. Neuron *98, 366–379*.

Disner, S.G., Beevers, C.G., Haigh, E.A.P., and Beck, A.T. (2011). Neural mechanisms of the cognitive model of depression. Nat. Rev. Neurosci. *12*, 467–477.

Durstewitz, D., Vittoz, N., Floresco, S., and Seamans, J. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. Neuron *66*, 438–448.

Ellwood, I., Patel, T., Wadia, V., Lee, A., Liptak, A., Bender, K., and Sohal, V. (2017). Tonic or Phasic Stimulation of Dopaminergic Projections to Prefrontal Cortex Causes Mice to Maintain or Deviate from Previously Learned Behavioral Strategies. J Neurosci *37(35): 8315–8329.*

Floresco, S., Ghods-Sharifi, S., Vexelman, C., and O, M. (2006). Dissociable Roles for the Nucleus Accumbens Core and Shell in Regulating Set Shifting. J. Neurosci. *26*, 2449–2457.

Garrison, J., Erdeniz, B., and Done, J. (2013). Prediction error in reinforcement learning: a meta-analysis of neuroimaging studies. Neurosci Biobehav Rev *Aug;37(7):1297-310*.

Gonda, X., Pompili, M., Serafini, G., Carvalho, A.F., Rihmer, Z., and Dome, P. (2015). The role of cognitive dysfunction in the symptoms and remission from depression. Ann. Gen. Psychiatry *14*, 27.

Halleland, H., Haavik, J., and Lundervold, A. (2012). Set-shifting in adults with ADHD. J Int Neuropsychol Soc *18*, 728–737.

Harris, K., and Shepherd, G. (2015). The neocortical circuit: themes and variations. Nat. Neurosci. *18, pages170–181*.

Harvey, P., Green, M., Keefe, R., and Velligan, D. (2004). Cognitive Functioning in Schizophrenia: A Consensus Statement on Its Role in the Definition and Evaluation of Effective Treatments for the Illness. J. Clin. Psychiatry *65*, 361–372.

Hauser, T., Iannaccone, R., Staempfli, P., and Drechsler, R. (2013). The Feedback-Related Negativity (FRN) revisited: New insights into the localization, meaning and network organization. NeuroImage *84*.

Heetderks, W. (1978). Principal component analysis of neural population response of knee joint proprioceptors in cat. Brain Res *Volume 156, Issue 1*, Pages 51-65.

Heisler, J.M., Morales, J., Donegan, J.J., Jett, J.D., Redus, L., and O'Connor, J.C. (2015). The Attentional Set Shifting Task: A Measure of Cognitive Flexibility in Mice. J. Vis. Exp. JoVE.

Hill, S.K., Bishop, J.R., Palumbo, D., and Sweeney, J.A. (2010). Effect of second-generation antipsychotics on cognition: current issues and future challenges. Expert Rev. Neurother. *10*, 43–57.

Hymana, J., Maa, L., Balaguer-Ballesterb, E., Durstewitzb, D., and Seamansa, J. (2012). Contextual encoding by ensembles of medial prefrontal cortex neurons. PNAS *109*, 5086–5091.

Jazbec, S., Pantelis, C., Robbins, T., Weickert, T., Weinberger, D., and Goldberg, T. (2007). Intra-dimensional/extra-dimensional set-shifting performance in schizophrenia: Impact of distractors. Schizophr. Res. *89*, 339–349.

Kato, S., Fukabori, R., Nishizawa, K., Okada, K., Yoshioka, N., Sugawara, M., Maejima, Y., Shimomura, K., Okamoto, M., Eifuku, S., et al. (2018). Action Selection and Flexible Switching Controlled by the Intralaminar Thalamic Neurons. Cell Rep. *27;22(9):2370-2382*.

Kiesel, A., Wendt, M., Jost, K., Steinhauser, M., Falkenstein, M., Philipp, A., and Koch, I. (2010). Control and Interference in Task Switching—A Review. Psychol. Bull. *Vol. 136, No. 5, 849–874*.

Kim, C., Cilles, S.E., Johnson, N.F., and Gold, B.T. (2012). Domain General and Domain Preferential Brain Regions Associated with Different Types of Task Switching: A Meta-Analysis. Hum. Brain Mapp. *33*, 130–142.

Liu, X., and Carter, A. (2018). Ventral Hippocampal Inputs Preferentially Drive Corticocortical Neurons in the Infralimbic Prefrontal Cortex. J Neurosci *Aug 15;38(33):7351-7363*.

Low, R.J., Gu, Y., and Tank, D.W. (2014). Cellular resolution optical access to brain regions in fissures: Imaging medial prefrontal cortex and grid cells in entorhinal cortex. Proc. Natl. Acad. Sci. *111*, 18739–18744.

Mahn, M., Gibor, L., Patil, P., Malina, K., Oring, S., Printz, Y., Levy, R., Lampl, I., and Yizhar, O. (2018). High-efficiency optogenetic silencing with somatargeted anion-conducting channelrhodopsins. Nat. Commun. *9:4125*.

Malagon-Vina, H., Ciocchi, S., Passecker, J., Dorffner, G., and Klausberger, T. (2018). Fluid network dynamics in the prefrontal cortex during multiple strategy switching. Nat. Commun. *9:309*.

Mante, V., Sussillo, D., Shenoy, K., and Newsome, W. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature *7;503(7474):78-84*.

Marton, T., Seifikar, H., Luongo, F., Lee, A., and Sohal, V. (2018). Roles of Prefrontal Cortex and Mediodorsal Thalamus in Task Engagement and Behavioral Flexibility. J Neurosci *38(10):2569-2578*.

Meiron, N., and Kessler, Y. (2008). The Task Rule Congruency Effect in Task Switching Reflects Activated Long-Term Memory. J. Exp. Psychol. *Vol. 34, No. 1, 137–157*.

Meyers, E., Freedman, D., Kreiman, G., Miller, E., and Poggio, T. (2008). Dynamic population coding of category information in inferior temporal and prefrontal cortex. J Neurophysiol *100*, 1407–1419.

Miller, P. (2016). Dynamical systems, attractors, and neural circuits. F1000Research *5*.

Milner, B. (1963). Effects of different brain lesions on card sorting: The role of the frontal lobes. Arch. Neurol. *9*, 100–110.

Murphy, F.C., Michael, A., and Sahakian, B.J. (2012). Emotion modulates cognitive flexibility in patients with major depression. Psychol. Med. *42*, 1373–1382.

Nakayama, H., Ibañez-Tallon, I., and Heintz, N. (2018). Cell-type specific contributions of medial prefrontal neurons to flexible behaviors. J Neurosci. *10.1523/JNEUROSCI.3537-17.2018*.

Otis, J.M., Namboodiri, V.M.K., Matan, A.M., Voets, E.S., Mohorn, E.P., Kosyk, O., McHenry, J.A., Robinson, J.E., Resendez, S.L., Rossi, M.A., et al. (2017). Prefrontal cortex output circuits guide reward seeking through divergent cue encoding. Nature *543*, 103–107.

Pnevmatikakis, E., and Giovannucci, A. (2017). NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. J. Neurosci. Methods *Volume 291, Pages 83-94*.

Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., et al. (2016). Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. Neuron *89*, 285–299.

Rigotti, M., Barak, O., Warden, M., Wang, X., Daw, N., Miller, E., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature *497*, 585–590.

Rodgers, C., and DeWeese, M. (2014). Neural Correlates of Task Switching in Prefrontal Cortex and Primary Auditory Cortex in a Novel Stimulus Selection Task for Rodents. Neuron *82*, 1157–1170.

Sayyah, M., Eslami, K., AlaiShehni, S., and Kouti, L. (2016). Cognitive Function before and during Treatment with Selective Serotonin Reuptake Inhibitors in Patients with Depression or Obsessive-Compulsive Disorder. Psychiatry J. *2016*.

Senzai, Y., Fernandez-Ruiz, A., and Buzsaki, G. (2019). Layer-Specific Physiological Features and Interlaminar Interactions in the Primary Visual Cortex of the Mouse. Neuron *VOLUME 101, ISSUE 3, P500-513.E5*.

Sheintuch, L., Rubin, A., Brande-Eilat, N., Geva, N., Sadeh, N., Pinchasof, O., and Ziv, Y. (2017). Tracking the Same Neurons across Multiple Days in Ca2+ Imaging Data. Cell Rep. *21, 1102–1115*.

Smith, M., and Kohn, A. (2008). Spatial and Temporal Scales of Neuronal Correlation in Primary Visual Cortex. J Neurosci *28(48): 12591–12603*.

Smout, C., Tang, M., Garrido, M., and Mattingley, J. (2019). Attention promotes the neural encoding of prediction errors. PLoS Biol *17(2): e2006812*.

Soudais, C., Laplace-Builhe, C., Kissa, K., and Kremer, E. (2001). Preferential transduction of neurons by canine adenovirus vectors and their efficient retrograde transport in vivo. FASEB J *15:2283-2285*.

Starkweather, C., Gershman, S., and Uchida, N. (2018). The Medial Prefrontal Cortex Shapes Dopamine Reward Prediction Errors under State Uncertainty. Neuron *May 2;98(3):616-629.e6*.

Tait, D.S., Chase, E.A., and Brown, V.J. (2014). Attentional set-shifting in rodents: a review of behavioural methods and pharmacological results. Curr. Pharm. Des. *20*, 5046–5059.

Tervo, D., Hwang, B., Viswanathan, S., Gaj, T., Lavzin, M., Ritola, K., Lindo, S., Michael, S., Kuleshova, E., Ojala, D., et al. (2016). A Designer AAV Variant Permits Efficient Retrograde Access to Projection Neurons. Neuron *Oct 6. pii: S0896-6273(16)30580-3*.

Toyomaki, A., Hashimoto, N., Kako, Y., Murohashi, H., and Kusumi, I. (2017). Neural responses to feedback information produced by self-generated or other-generated decision-making and their impairment in schizophrenia. PLoS ONE *2(8): e0183792*.

Wang, J., Narain, D., Hosseini, E., and Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. Nat Neurosci. *Jan;21(1):102-110*.

Wimmer, R., Schmidt, L., Davidson, T., Nakajima, M., Deisseroth, K., and Halassa, M. (2015). Thalamic control of sensory selection in divided attention. Nature *Oct 29;526(7575):705-9*.

Figure S1: β-weight temporal cross-correlations. Pearson partial correlation values for beta-weights across trial-aligned timepoints for the five features decoded from the SEDS task. Response window initiated at time 0. All correlations controlled for β-weights from all left-out features at the given time point.

Figure S2: Validation of β-weights: A. Accuracies of SVM decoders trained and tested on trials from subsets of neurons grouped by the percentile of their β-weight values. First, a decoder was trained on all neurons (4740 from 21 animals). Neurons were then binned by their absolute β-weight percentile values (20 quantile bins, mean 237 neurons per 5 percentile bin). Decoders were then individually trained on subgroups corresponding with the percentile bins. Real data in teal, trial-shuffled data in grey. **B.** Schematic of β-weight cross-validation method. To test whether decoders weights for individual neurons depended on which other neurons were included in the decoder, subgroups were iteratively (100x) and randomly generated by dividing the population into three subgroups, training decoders on two of the three subgroups at a time, and computing correlation coefficients for neurons across the separately generated decoders. **C.** Correlation of neuron weights across decoders trainsedon separate, shuffled subgroups. Spearman p = 0 in iterations, indicating that neuron weight was independent of the identities of the other neurons in the decoder.

Figure S3: **Example cross-session cell registrations**

Figure S4: Example sagittal view of prism imaging field. hSyn-GCaMP6f in green, DAPI in blue.

Figure S5: Retrograde labeling of PFC-MDT neurons by CAV2-FLP-mediated tdTomato (red) and AAV2Retro-GFP (green). Both viruses were infused in MDT as a mixture by a single injection. FRT-tdTomato as infused in PFC in a separate injection.

Figure S6: Cross-animal distribution of β-weights from feature decoders. Box plots are medians and quartiles taken from the epochs of peak decoding accuracy for each feature (stimulus delivery epoch for whisker and odor stims, post-trial epoch for response and outcome, both epochs for task rule).

Figure S7: Euclidean distances between incorrect and correct outcomes for trials t-1 through t-7.