Neuron
Article

Neural Correlates of Task Switching in Prefrontal Cortex and Primary Auditory Cortex in a Novel Stimulus Selection Task for Rodents

Chris C. Rodgers[1,3,*] and Michael R. DeWeese[1,2,3]
[1]Helen Wills Neuroscience Institute, University of California, Berkeley, 132 Barker Hall, Berkeley, CA 94720, USA
[2]Department of Physics, University of California, Berkeley, 132 Barker Hall, Berkeley, CA 94720, USA
[3]Redwood Center for Theoretical Neuroscience, University of California, Berkeley, 132 Barker Hall, Berkeley, CA 94720, USA
*Correspondence: xrodgers@gmail.com
http://dx.doi.org/10.1016/j.neuron.2014.04.031

SUMMARY

Animals can selectively respond to a target sound despite simultaneous distractors, just as humans can respond to one voice at a crowded cocktail party. To investigate the underlying neural mechanisms, we recorded single-unit activity in primary auditory cortex (A1) and medial prefrontal cortex (mPFC) of rats selectively responding to a target sound from a mixture. We found that prestimulus activity in mPFC encoded the selection rule—which sound from the mixture the rat should select. Moreover, electrically disrupting mPFC significantly impaired performance. Surprisingly, prestimulus activity in A1 also encoded selection rule, a cognitive variable typically considered the domain of prefrontal regions. Prestimulus changes correlated with stimulus-evoked changes, but stimulus tuning was not strongly affected. We suggest a model in which anticipatory activation of a specific network of neurons underlies the selection of a sound from a mixture, giving rise to robust and widespread rule encoding in both brain regions.

INTRODUCTION

Humans can select and respond to one person's voice while many others are speaking at the same time. We do this effortlessly, yet no known algorithm can solve this "cocktail party problem" in realistic settings, perhaps because we do not fully understand the relevant computations performed in the brain (Cherry, 1953; Sayers and Cherry, 1957; Ding and Simon, 2012; McDermott, 2009). Other social animals such as birds and rodents demonstrate a similar ability (Bee and Micheyl, 2008); for instance, mother mice respond to distinct pup calls when several are calling at once (Geissler and Ehret, 2002). Humans use selective attention, the cognitive process of selecting and responding to a single target stimulus among simultaneous distractors (Desimone and Duncan, 1995), to solve the cocktail party problem (Ahveninen et al., 2011). Experiments in visual selective attention have revealed that the prefrontal cortex (PFC) sends top-down "bias signals" to sensory cortex (Miller and Cohen, 2001; Moore et al., 2003) to select the target stimulus, enhancing its neural representation while suppressing the representation of distractors. Similar mechanisms may be at work in the auditory cortex: electrocorticographic (Mesgarani and Chang, 2012; Zion Golumbic et al., 2013) and magnetoencephalographic (Ding and Simon, 2012) recordings show that brain activity is dominated by the attended voice. Ultimately, recordings from single units (individual neurons) will be needed to understand the circuit. In addition, many models of visual selection are not obviously applicable to the auditory modality—for instance, the idea that visual attention co-opts the neural mechanisms for shifting gaze (Moore et al., 2003). Establishing an animal model of auditory selective attention could shed more light on whether the known mechanisms of visual selection are universal or specific to one modality.

Nonhuman primates are the traditional model organism for studying complex cognition (Gold and Shadlen, 2007) but rodents are also capable of sophisticated decision-making (Raposo et al., 2012; Brunton et al., 2013; Zariwala et al., 2013). The behavioral flexibility of rodents is thought to be mediated by the PFC (Karlsson et al., 2012; Kvitsiani et al., 2013), even though this region is not necessary for simple sensory discriminations (Pai et al., 2011). The medial PFC (mPFC) in particular is critical for task switching (Birrell and Brown, 2000; Floresco et al., 2008; Durstewitz et al., 2010; Ragozzino et al., 1999), such as switching the navigational strategy used to solve a maze (Rich and Shapiro, 2009). Rodent mPFC thus appears to maintain a representation of the current task rule, analogous to the rule-encoding neurons observed in primate PFC (Wallis et al., 2001; Asaad et al., 2000; Johnston et al., 2007), although large parts of the monkey PFC appear to be functionally and anatomically unique to primates (Wise, 2008).

Frontal areas have been shown to play a role in directing flexible auditory processing in the primary auditory cortex (A1). For example, training ferrets to detect tones at a specific target frequency (Fritz et al., 2003) produces rapid tuning changes in A1 (i.e., the neurons responded more to the target frequency) and modulates functional connectivity between A1 and frontal areas (Fritz et al., 2010). Moreover, when these ferrets switch between different auditory tasks, the observed tuning changes match the demands of each task (Fritz et al., 2005). These experiments

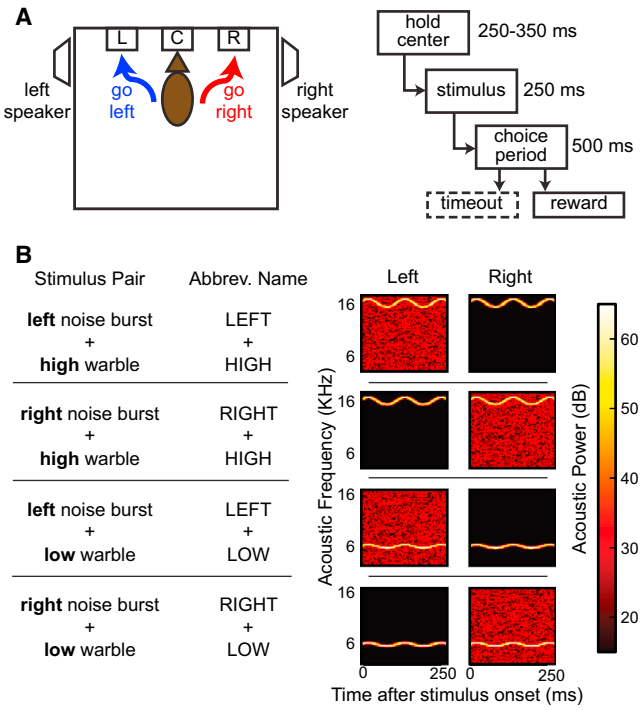

Figure 1. Behavioral Paradigm

(A) Left: a schematic of the behavioral arena with left (L), center (C), and right (R) ports (or nose-pokes), and left and right speakers. Right: timeline of each trial. The rat initiates a trial by nose-poking the center port as shown. After a hold period, an auditory stimulus is played. Following this, the rat may choose to go to the left port (blue arrow), go to the right port (red arrow), or do neither of those (a "nogo" response).

(B) Task stimuli (left, description; right, spectrogram of the auditory waveform). On each trial, the rat hears one of four possible auditory stimulus pairs: LEFT+HIGH, RIGHT+HIGH, LEFT+LOW, or RIGHT+LOW. Each is a simultaneous combination of a broadband noise burst played from either the left or right speaker, and a low-pitched or high-pitched warble played with equal intensity from both speakers.

(C) Task rules. The session consists of alternating localization and pitch discrimination blocks. In localization blocks, the rat must go left for sounds containing LEFT and it must nogo for sounds containing RIGHT; the low- or high-pitched warble is an irrelevant distractor. In pitch discrimination blocks, the rat must go right if the stimulus pair contains LOW and it must nogo if the stimulus pair contains HIGH; the noise burst is an irrelevant distractor.

revealed A1 to be surprisingly dynamic; however, subjects were not required to select a target sound in the presence of simultaneous distractors, a critical aspect of the cocktail party problem (McDermott, 2009; Ding and Simon, 2012).

In this study, we have taken advantage of the relative ease and speed with which rodents can be trained on demanding tasks (Carandini and Churchland, 2013) to develop a novel behavioral task in which rats hear two simultaneous sounds, but select and respond only to one. This requires cognitive flexibility because, depending on which sound the experimenter instructs the subject to select, the same pair of sounds can elicit a different behavioral response ("same stimulus; different response"). The subjects must alternate which sound they select multiple times within each session. We are aware of no purely auditory single-unit studies in any animal with these properties. The analogous ability in vision—to respond to a behaviorally relevant stimulus in the presence of competing distractors—has been referred to as stimulus "selection" (Knudsen, 2007; Reynolds and Chelazzi, 2004; Pestilli et al., 2011); following this, we refer to our task as auditory stimulus selection.

Similar visual and cross-modal tasks have been termed set shifting (Stoet and Snyder, 2004), task switching (Sasaki and Uka, 2009), and selective attention (Moran and Desimone, 1985; Hocherman et al., 1976; Otazu et al., 2009). Other studies have investigated "response selection": how decisions are translated into appropriate motor actions, following stimulus selection or even in the absence of an explicit stimulus (Young and Shapiro, 2011; Turken and Swick, 1999). We also note a similarity between our task and the Wisconsin Card Sorting Task for diagnosing disorders of executive function (Monchi et al., 2001).

We recorded from individual mPFC and A1 neurons in rats performing our task. We found that the prestimulus, anticipatory activity of our recorded neurons in mPFC encoded the selection rule—which sound the subject should select. Surprisingly, we also found this prestimulus effect in a sizable fraction of the neurons we recorded in A1. Disruption of mPFC through electrical microstimulation significantly impaired task performance. Finally, although changes in prestimulus baseline correlated with changes in stimulus-evoked activity in both brain regions, this did not appear to alter tuning properties in a way that would be obviously beneficial for responding to the selected sound.

RESULTS

A Novel Behavioral Task for Rodents: Auditory Stimulus Selection

We developed an auditory stimulus selection task for rats, in which the subject was trained to respond to either of two simultaneously presented sounds. The rat initiated each trial (Figure 1A) by holding its nose in the center port of a three-port behavior box—the "hold period." This triggered speakers on the left and right to play one of the following four equally likely stimulus pairs: LEFT+HIGH, RIGHT+HIGH, LEFT+LOW, or RIGHT+LOW (Figure 1B). Each stimulus pair was a simultaneous combination of (1) a broadband noise burst, played from either the LEFT or RIGHT speaker; and (2) either a HIGH- or LOW-pitched warble (frequency-modulated tone), played from both speakers simultaneously. After the onset of stimulus presentation, the rat could then choose to "go left" (poke its nose in the left port), "go right" (poke its nose in the right port), or "nogo" (not poke either side). Correct pokes into the side ports were rewarded with water; incorrect pokes were penalized with a 2–6 s timeout (see Supplemental Experimental Procedures available online).

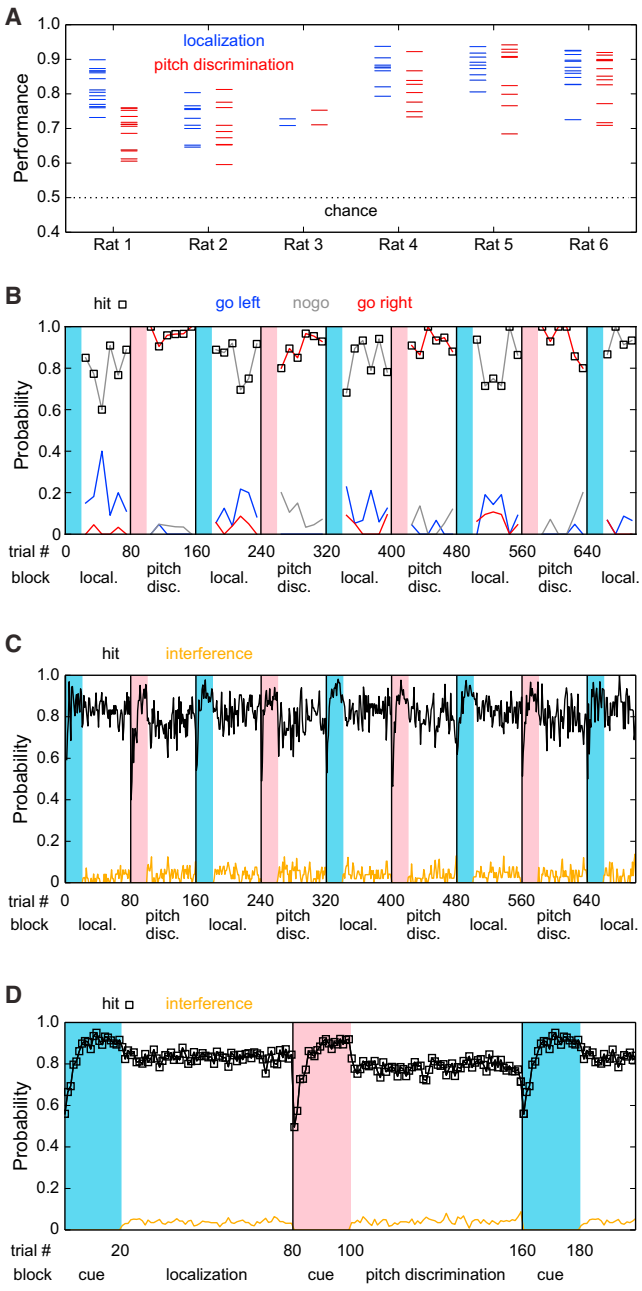

Figure 2. Trained Rats Select and Respond to the Target Sound, Not the Distractor

(A) Behavior performance during recording sessions. Each hash mark is the performance during localization (blue) or pitch discrimination (red) in a single recording session. Performance is well above chance (black dotted line; see Supplemental Experimental Procedures).

(B) Distribution of behavioral responses to an example stimulus pair (RIGHT+LOW) over the course of an average session. We averaged all sessions from a single rat (rat 5) and binned the trials into groups of ten. The x-axis shows both trial number and block type. The correct response to this stimulus pair is to go right during pitch discrimination and to nogo during localization. Each trace shows the probability that the rat will go right (red), nogo (gray), or go left (blue); black open squares mark the correct response for that block. The rat responds correctly most of the time, even though the required action changes abruptly at the block boundaries. This stimulus pair does not occur during cue trials, which begin each block and are shaded in cyan and pink throughout this figure.

(C) Similar to (B), but averaged over all sessions, rats, and stimuli. Most trials are correct (black trace). Interference trials (orange; see text) are rare.

(D) Performance briefly dips during cue trials at the beginning of a block but recovers within a few trials. All localization blocks from (C) are averaged together, as are all pitch discrimination blocks. To emphasize block transitions, the x axis repeats itself after trial 160; the cyan shaded areas are identical because the block structure is cyclical.

On each trial, one of the sounds in the stimulus pair (the "target") indicated the correct response; the other sound (the "distractor") was uninformative. The behavioral session alternated between "localization" blocks of trials, during which the noise burst was the target, and "pitch discrimination" blocks, during which the warble was the target (Figure 1C). Each block consisted of 80 trials, the first 20 trials of which were reserved to indicate the block change. During these 20 "cue trials," the rat heard only target sounds without any distractor.

The entire training process required approximately 10 weeks. Trained rats performed many trials per session (median, 698). We verified that the rats were performing significantly above 50% in both blocks, which meant that their behavioral response was driven by the target sound, rather than by the distractor or by a combination of target and distractor (Figures S1B, S1C, and Supplemental Experimental Procedures). Our best rats' typical performance during recording sessions was approximately 85% in both blocks (Figure 2). After each block change, rats rapidly and correctly switched to selecting the new target sound. Performance was typically better on go trials than on nogo trials (Figure S1A).

Anticipatory Neuronal Activity in mPFC and A1 Encodes the Selection Rule

We next asked whether neuronal activity changed according to which target the rats selected. We implanted tetrodes into the brain, targeting A1 and/or the prelimbic region of mPFC, and recorded single-unit action potentials (spikes) during behavior. By analogy with the rule-encoding neurons in the primate PFC, we hypothesized that the firing rates of single mPFC neurons would differ significantly between localization and pitch discrimination blocks. We first confined our analysis to the hold period, the interval before stimulus onset when the rat is holding its nose in the center port and presumably preparing to select the target sound from the imminent stimulus pair.

We found that the hold period activity of a majority of mPFC neurons robustly encoded the selection rule on "correct trials," those trials on which the rat gave the correct response. An example unit (Figure 3A) fired significantly more in the hold period during localization trials than it did during pitch discrimination trials. A different but simultaneously recorded single unit in mPFC (Figure 3B) fired significantly more during pitch discrimination than during localization. In both cases, the effect persisted across the entire session of over 1,300 trials, alternating with each block just as the behavior did. Across our recorded population of mPFC neurons, 63% (76/121) of the neurons individually and significantly encoded the selection rule during the hold period (Figure 3C). Of these, 36 neurons preferred (i.e., fired

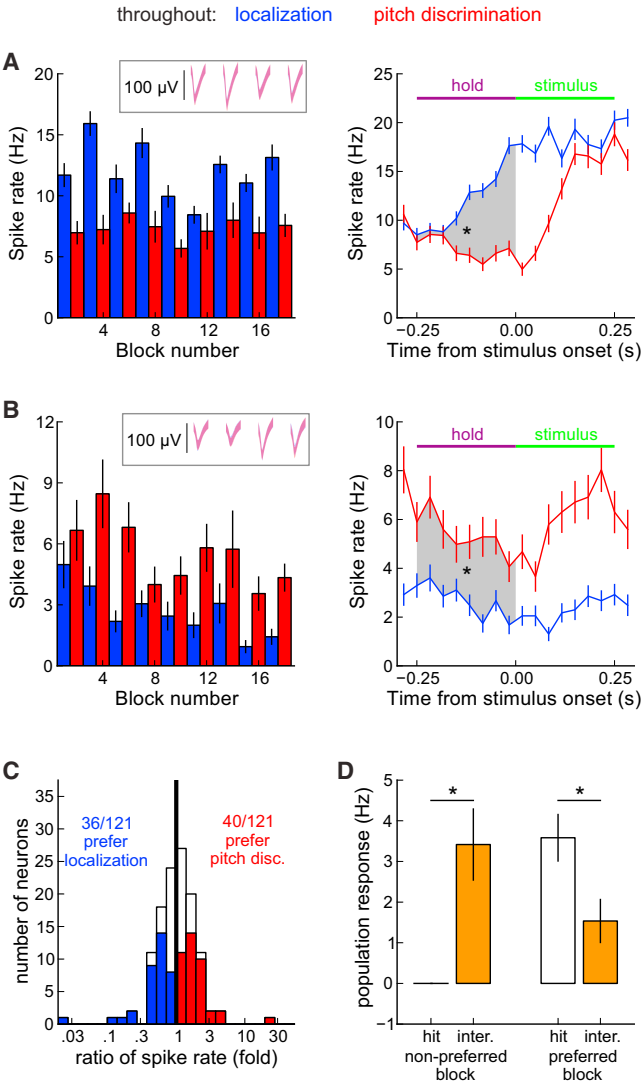

Figure 3. Prestimulus Activity in mPFC Encodes the Selection Rule

(A) Left: An example mPFC single unit that fires more during the hold period for localization (blue bars throughout this figure) than for pitch discrimination (red bars). For all figures, error bars represent SEM unless otherwise noted. Inset: Extracellular waveforms (mean ± SD) with duration of 0.8 ms on each channel of the tetrode. The waveforms are colored red and blue based on the block in which they were recorded, but are almost entirely overlapping (purple). Right: peristimulus time histogram (PSTH) of the same unit, averaged over all correct trials from each block. During the hold period (gray), the firing rate is significantly ($p < 0.001$, unpaired Mann-Whitney U-test) higher on localization (mean 12.1 Hz, n = 483) than on pitch discrimination (mean 7.2 Hz, n = 295) trials.

(B) Another example mPFC single unit, this one preferring pitch discrimination. The hold period firing rate is significantly ($p < 0.001$) higher on pitch discrimination (mean 5.4 Hz) than on localization (mean 2.7 Hz) trials. Trial counts are the same as the simultaneously recorded unit in (A). This neuron's firing rate is persistently elevated at all points plotted.

(C) Stacked histogram of the ratio of hold period firing rate (pitch discrimination over localization) for all mPFC neurons. Red and blue bars represent significantly modulated neurons. We used an unpaired Mann-Whitney U-test for all neurons and controlled for multiple comparisons using the Benjamini-Hochberg false discovery rate.

(D) Rule encoding is diminished on interference trials. We averaged together the firing rates of each rule-encoding neuron during either correct (white) or interference (orange) trials in their preferred and nonpreferred blocks. Firing rates are normalized by subtracting the firing rate on correct trials in the nonpreferred block. The population response on interference trials is significantly increased during the nonpreferred block and decreased during the preferred block (n = 57 neurons, paired Mann-Whitney test).

more during) localization and 40 preferred pitch discrimination; neither preference was significantly more common (binomial test, $p > 0.05$).

Surprisingly, we also found a similar effect in A1 (Figure 4). Although encoding of selection rule was our hypothesized result in mPFC, this was unexpected in A1, especially given the absence of auditory stimulation during the hold period. Across our recorded population, 36% (36/99) of A1 neurons encoded selection rule. As with mPFC, neither population was significantly larger (13 preferring localization, 23 preferring pitch discrimination; binomial test, $p > 0.05$). Because A1 encodes sounds sparsely (DeWeese et al., 2003; Hromádka et al., 2008; Carlson et al., 2012), we were not surprised to observe that only some (49/99) of our recorded neurons in A1 significantly responded to our task stimuli (Figures S5A–S5E). Rule encoding was not significantly more widespread in either stimulus-responsive (14/49) or nonresponsive (22/50) neurons ($p > 0.05$, Fisher's exact test). This finding is reminiscent of evidence from human imaging that attention affects strongly stimulus-driven regions of auditory cortex less than it affects other, more poorly tuned regions (Petkov et al., 2004).

These effects were strong: among rule-encoding neurons, the median increase in firing rate during the preferred block was 74.7% in mPFC and 99.7% in A1. These results are unlikely to be due to firing rate drift over the course of the session or spike sorting errors arising from small differences in spike waveform shape between blocks (see Supplemental Experimental Procedures). In sum, these results demonstrate widespread and robust encoding of selection rule in the prestimulus activity of both mPFC and A1 neurons.

Error Trial Analysis

In the previous section, we analyzed only correct trials. We next considered "interference" trials, during which the rat appeared to select the wrong sound. On such trials, the rat heard a "go" distractor (i.e., a sound that should have elicited a go response, had it been presented during the other block) and incorrectly went to the choice port associated with that distractor instead of doing what the target sound indicated. If anticipatory encoding of selection rule is important for successful stimulus selection, then this encoding should have been weaker or even reversed when the rat selected the wrong sound.

Indeed, in mPFC, the encoding of selection rule was significantly weakened on interference trials, as compared with correct trials (Figure 3D). In A1, we observed a more extreme effect (Figure 4D): the rule encoding was actually reversed on interference trials (i.e., firing rates were greater during the nonpreferred block on such trials). These observations are consistent with the idea that anticipatory activity predicted which sound the rat would select, even for trials on which the rat appeared to respond to the distractor by going to the wrong choice port.

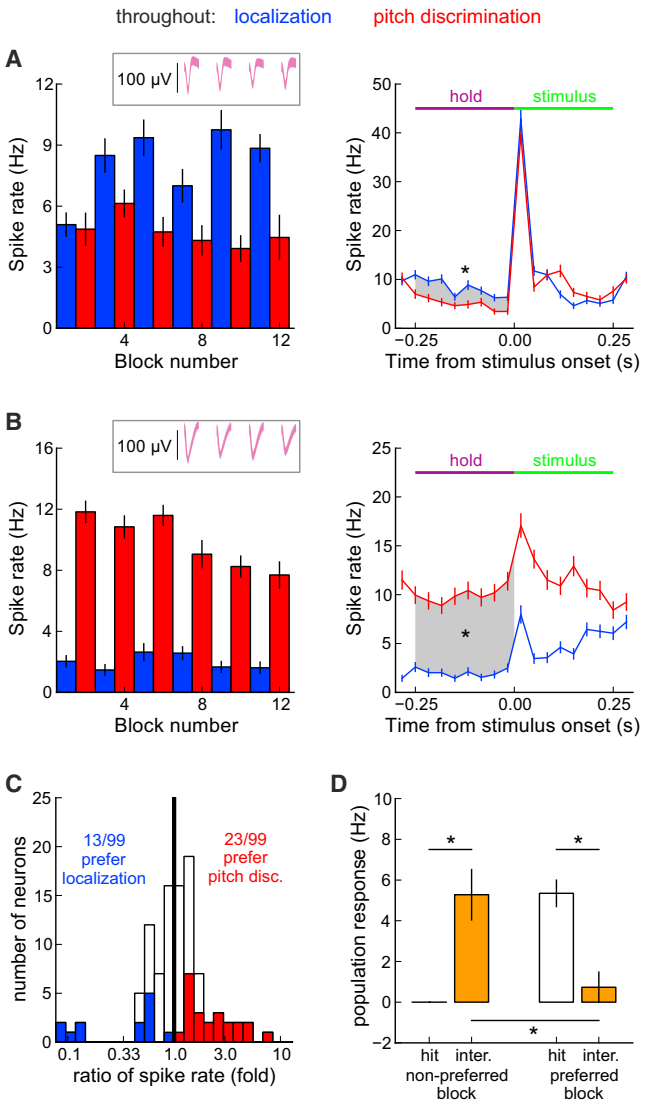

Figure 4. Prestimulus Activity in A1 also Encodes the Selection Rule

(A) An example A1 neuron that responds significantly more ($p < 0.001$) during localization (8.0 Hz, n = 312 trials) than during pitch discrimination (4.8 Hz, n = 253). Throughout, conventions and statistical procedures are as in Figure 3. Note the peak following stimulus onset, which was used to analyze the evoked response (Figure 6).

(B) Simultaneously recorded A1 neuron that significantly ($p < 0.001$) prefers pitch discrimination (10.1 Hz, n = 312 trials) over localization (2.0 Hz, n = 253).

(C) Stacked histogram of the ratio of hold period firing rate (pitch discrimination over localization) for all A1 neurons.

(D) Rule encoding during the hold period is inverted on interference trials. The population response on interference trials (orange bars) is significantly greater during the nonpreferred block than during the preferred block ($p < 0.01$, n = 16 neurons, paired Mann-Whitney U-test). In contrast, on correct trials (white bars) the firing rate is higher during the preferred block than the nonpreferred block, by definition.

Potential Role of Posture

The mPFC regulates cognitive state, but it also encodes body position and plays a role in motor planning (Erlich et al., 2011; Euston and McNaughton, 2006). We analyzed video of the rats and found evidence of preparatory head positioning that differed between blocks (see Supplemental Experimental Procedures), presumably a behavioral strategy that the rat used to prepare for the differing motor actions required in each block (i.e., go left in localization and go right in pitch discrimination).

This raised the question of whether neurons were encoding this postural difference, rather than selection rule. We found that, in the vast majority of rule-encoding neurons in both brain regions, the selection rule explained more of the variance in firing rates than did head angle (Figures S2I–S2T and S3I–S3Q). In addition, the rule encoding was largely maintained on a subset of "posture-equalized" trials, selected so that the mean head angle was the same in each block (Figures S2M–S2O and S3M–S3O). Finally, as we discuss below, the long duration of the neural effects we observed further argues against the possibility that changes in posture were the underlying cause.

Within-Trial Timescale of the Encoding of the Selection Rule

We next asked how soon before the stimulus the rule encoding emerged, and for how long afterward it persisted. For each rule-encoding neuron, we compared across blocks the smoothed firing rates in every 50 ms bin up to 3 s before and after stimulus onset. We thereby determined the largest interval around the hold period during which the neural activity significantly encoded the selection rule. Across the data set, the median intertrial interval was 4.0 s (interquartile range, 2.7–5.3 s) and so this range (± 3 s) overlapped with the previous and/or next trial in many cases.

The temporal dynamics of the encoding varied widely across neurons in both regions (Figures 5A and 5B). For some neurons, rule encoding was strictly confined to the hold period. Other neurons showed significant encoding during all time bins tested: their firing rates were persistently elevated during the preferred block. We found neurons spanning this range of timescales in both brain areas. In A1, the median rule-encoding unit first developed a significant block preference 0.55 s prestimulus; in PFC, the median was 0.625 s prestimulus. Thus, the majority of rule-encoding neurons developed this property well before the rat initiated a trial by center-poking.

To examine the population dynamics of rule encoding, we averaged the normalized activity of all rule-encoding neurons during their preferred block. On average, the population activity ramped up gradually before stimulus onset, over a timescale of several seconds, and then fell relatively quickly afterward (Figure 5C; Figures S4B–S4E). The population activity in mPFC was significantly elevated earlier than in A1, consistent with its hypothesized role as the origin of top-down bias signals to sensory cortex (Miller and Cohen, 2001). However, we note that the wide range of timescales within both regions, and the fact that only a small fraction of our data set consists of simultaneous recordings from A1 and PFC, complicates a direct comparison between brain regions.

Finally, we asked whether the rule encoding reflected an increased firing rate in one block (as compared with the

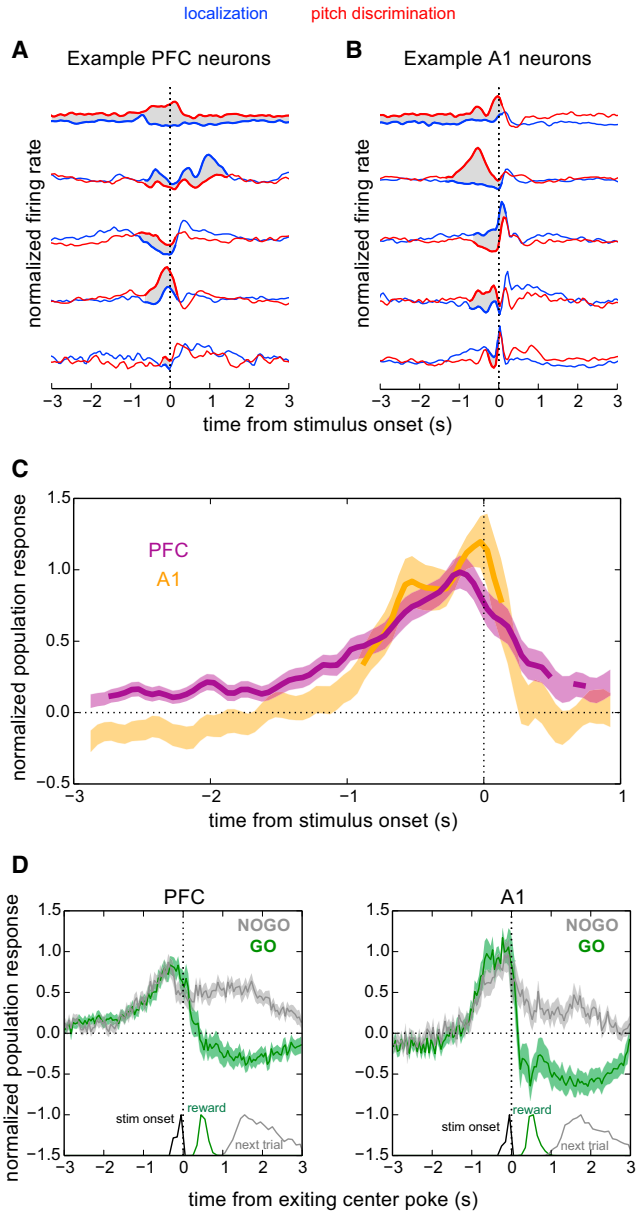

Figure 5. Within-Trial Timescale of the Encoding of Selection Rule

(A) PSTHs from example rule-encoding mPFC neurons in each block (blue, localization; red, pitch discrimination). Note that the timescale is much longer than that in previous figures. Firing rates are smoothed with a 50 ms Gaussian kernel and normalized to equal variance across neurons. Gray shading represents the maximum time interval, containing the hold period, during which the traces significantly diverge. Although these neurons were identified based on a difference in firing rate during the hold period, the traces often diverge for much longer. We observed a wide variety of timescales and dynamics. The first neuron effectively fires persistently more in one block. The third and fourth neurons demonstrate that the firing rate can either decrease during the nonpreferred block or, more commonly, increase during the preferred block.

(B) Example neurons from A1, following the conventions of (A). Again, the neurons exhibit a wide variety of dynamics, from essentially persistent block-specific activation for over 3 s preceding the stimulus (first neuron), to very brief activation well under 1 s (last neuron).

(C) Population time course: the traces represent the mean response, ± SEM, during the preferred block in mPFC (purple, n = 76) and A1 (orange, n = 36). The firing rates of all rule-encoding neurons were normalized (mean, 0; variance, 1) and then averaged together. The bold mean trace represents time points during which the population response significantly ($p < 0.05$, one-sample t test) exceeds zero, the mean firing rate. In both brain regions, the firing rate in the preferred block gradually increases, peaking around the time of stimulus onset, and then decreases more quickly back to baseline. The PFC population increases its response earlier (first significantly activated 2.7 s before stimulus onset) than the A1 population (first significantly activated 0.88 s before stimulus onset), consistent with the hypothesized role of PFC as the source of top-down modulation.

(D) Population time course plotted separately for go and nogo trials, from rule-encoding mPFC (left) and A1 (right) neurons during their preferred block. The peri-event time histograms (PETHs) are locked to the poststimulus exit from the center-port. As in (C), PETHs are mean ± SEM, and were normalized to unit variance and 0 mean before averaging across neurons. On nogo trials (gray), the firing rate remained elevated above baseline for at least several seconds, during which time the rat typically had already initiated the next trial. On go trials, the firing rate was suppressed below baseline as the rat moved to the choice port and consumed a reward, which always required at least several seconds. Latency distributions of trial events are plotted along the lower edge: stimulus onset (black), reward delivery (green, go trials only), and center-poke beginning the next trial (gray, nogo trials only). The next trial after a go trial would not be visible on this timescale due to the time required to consume the reward.

spontaneous rate while the rats were not performing the task), a decreased firing rate in the other block (versus spontaneous), or something else (for instance, a low spontaneous rate, an elevated rate for one block, and an even higher rate for the other block, which might reflect an encoding of task difficulty). Individual neurons exhibited a diversity of effects and we observed single units showing each of these possibilities (Figures 5A and 5B). However, across the population of rule-encoding neurons, the firing rate was significantly higher than spontaneous during the preferred block and significantly lower than spontaneous during the nonpreferred block (Figures S2F and S3F). These data argue against a model in which neurons encode task difficulty and instead suggest that each block actively engages two different populations of neurons, increasing the firing rate in one population and suppressing it in the other.

Encoding of Behavioral Choice

We found a prominent difference between the firing rates during go and nogo trials (Figure 5D). During a typical rule-encoding unit's preferred block, its firing rate remained elevated on nogo trials for several seconds after the center-poke, during which time the rat was often beginning the next trial. In contrast, on go trials, the typical unit's firing rate rapidly fell as the rat left the center port and remained suppressed for several seconds, during which time the rat was typically moving to the reward port and consuming reward.

One interpretation of this result is that rule encoding is particularly important for producing the nogo response, consistent with previous reports of enhanced encoding of nogo stimuli (David et al., 2012; Fritz et al., 2003). Another interpretation is that rule encoding persisted on nogo trials simply because the animal was already preparing to begin the next trial, whereas on go trials the rat was moving to the reward port to consume water and thus no longer needed to represent the stimulus selection rule.

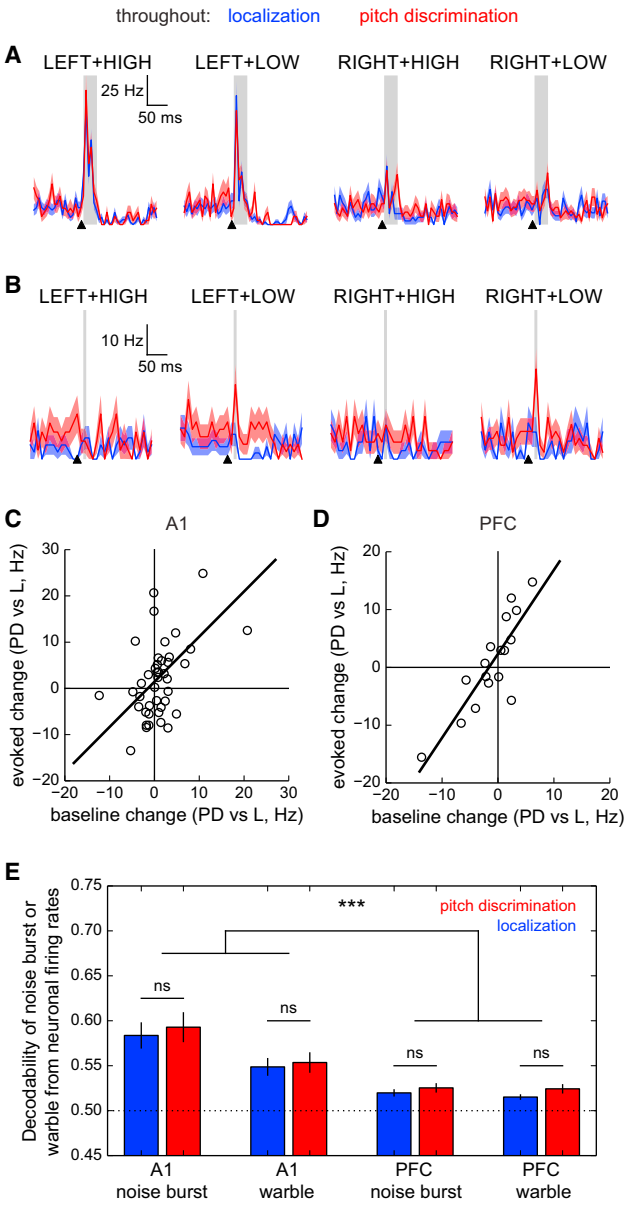

Figure 6. Limited Evidence for Modulation of Stimulus Tuning

(A) An example A1 neuron exhibiting a preference for some acoustic stimuli (LEFT+HIGH, LEFT+LOW) over others (RIGHT+HIGH, RIGHT+LOW), but no change in this tuning with block (localization, blue; pitch discrimination, red). Black triangles represent stimulus onset; shaded areas represent the response window for this neuron.

(B) An example auditory-responsive mPFC neuron. Evoked responses were weaker for mPFC neurons than for A1 neurons (Figures S5A–S5E).

(C) For A1 neurons (n = 43), increased anticipatory firing during one block significantly correlated ($p < 0.001$) with increased evoked responses during the same block. Each circle shows the change in evoked response (y-axis) versus the change in hold period firing rate (x-axis) for each neuron, quantified as mean pitch discrimination firing rate minus localization firing rate. The slope of the trend line is close to 1, suggesting that most of the evoked modulation across neurons is due to anticipatory modulation.

(D) Following the conventions of (C), but for auditory-responsive mPFC neurons (n = 17). Again, the change in anticipatory activity across neurons correlated closely with the change in evoked response ($p < 0.001$).

(E) No evidence for tuning changes that increase the decodability of the target sound. The identity of the noise burst (LEFT or RIGHT) or the warble (LOW or HIGH) could be decoded from the trial-by-trial responses of simultaneously recorded ensembles of auditory-responsive cells in either A1 or PFC. It could be decoded significantly more accurately ($p < 0.001$ for the main effect of brain region) using A1 responses (n = 21 ensembles of 49 neurons total) than using mPFC responses (n = 17 ensembles of 31 neurons total). However, for both sounds and both brain regions, the decoding was not significantly more accurate during either pitch discrimination (red) or localization (blue) trials ($p > 0.05$ for each pair of bars using a paired t test). The chance decoding level, attainable by a neuron with no information about the stimulus, was 0.5. Error bars represent SEM over ensembles. We used a three-way ANOVA on brain region, target sound, and block.

Changes in Baseline Activity Correlate with Changes in Evoked Activity

Given that the prestimulus activity encoded the selection rule, we next assessed whether the stimulus-driven activity in A1 differed between blocks. We first defined the evoked response window of each neuron as the period after stimulus onset during which the firing rate was significantly elevated above the prestimulus rate (Figures S5A–S5E). The evoked response on each trial was then defined as the number of spikes emitted during this window. We analyzed the mPFC neurons in the same way and found a population of neurons showing auditory responses to our task stimuli that were low-latency and tightly locked to stimulus onset, similar to A1 (Figures 6A and 6B). Auditory-responsive neurons in mPFC were significantly rarer (mPFC, 31/121; A1, 49/99; $p < 0.001$, Fisher's exact test) than in A1, and their responses were weaker (Figures S5C and S5D). Their median response latency was also significantly longer, though only slightly (19.75 ms versus 16.75 ms in A1; Figure S5E).

Our results show that prestimulus activity is modulated by selection rule, often persistently, consistent with a model in which rule-encoding neurons receive a higher level of tonic excitatory input during one block, for example. We expected that this task-specific modulation of baseline activity might correspond to an increased response during the stimulus-evoked response as well. In both A1 and mPFC, this was indeed the case: across the population, an increase in prestimulus firing rate during one block correlated with a comparable increase in evoked firing rate during the same block (Figures 6C and 6D; example cell: Figure 4B). However, after accounting for changes in prestimulus activity, we found very little evidence of any block-specific change in evoked firing rate (see Supplemental Experimental Procedures), which suggests that evoked activity is not strongly altered beyond an additive effect of baseline.

We also asked whether selection rule modulated stimulus tuning, for instance, to enhance the representation of the target sound. Such an effect might have been obscured in our analyses thus far, which averaged over stimuli in order to detect a change in overall response strength. To ascertain directly whether the target sound was better represented in the neural activity, we used an ideal linear decoder analysis (see Supplemental Experimental Procedures; Figure 6E) to quantify how well the stimulus-evoked activity in both brain regions encoded the identity of the noise burst or warble. As expected, the identity of each sound may be decoded more accurately from the activity of A1 neurons than from mPFC neurons, likely due to their stronger responses

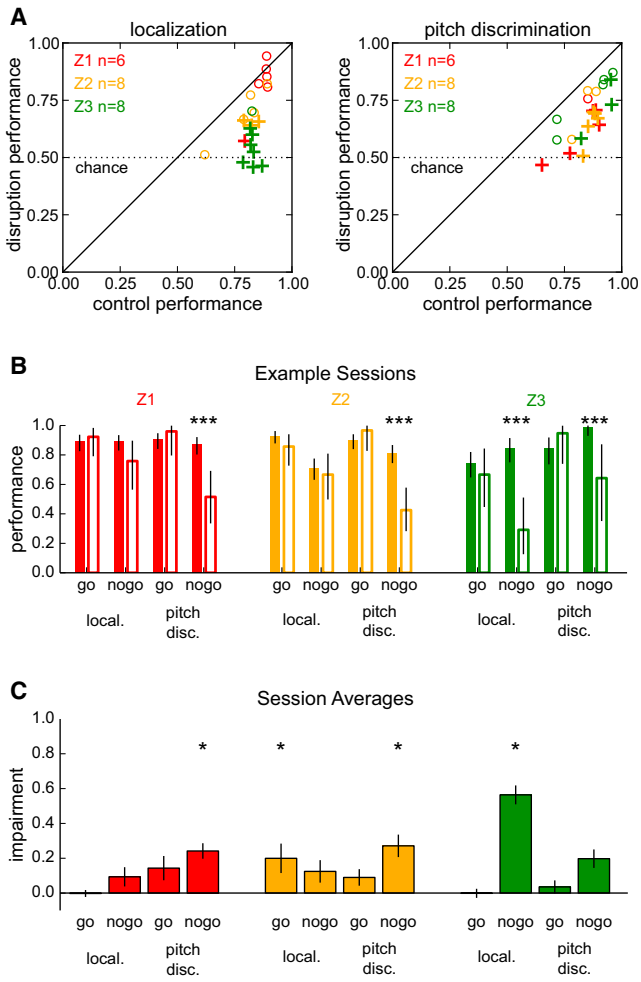

Figure 7. Disruption of mPFC Robustly Impairs Performance

(A) Electrical disruption of mPFC significantly impaired task performance (i.e., fraction of correct trials) during localization trials (left) and pitch discrimination trials (right) in most sessions. Each point represents the performance within a single session on disruption (y-axis) versus control (x-axis) trials. + represents sessions during which the performance was significantly impaired ($p < 0.05$, Fisher's exact test). Throughout this figure, colors represent different rats (red, Z1; yellow, Z2; and green, Z3).

(B) Example session from each rat. Performance is shown for each trial type (go and nogo in each block) on control (solid bars) and disruption (open bars) trials. Error bars represent 95% confidence intervals using Pearson-Klopper binomial fit. Asterisks indicate trial types for which electrical disruption significantly impairs performance (Fisher's exact test, $p < 0.001$ for all significant comparisons shown). The effect is robust within each example session, but varies between rats.

(C) Pattern of impairment (i.e., the difference in performance between control and disruption trials) across sessions for each rat. Error bars represent SEM across sessions. Asterisks show trial types that were significantly impaired ($p < 0.05$, binomial test on number of impaired sessions) for each rat. All rats were significantly impaired on nogo trials during one block or the other. One rat (Z2) also showed a significant impairment on localization go trials.

See also Figure S6.

and tighter stimulus selectivity. However, for both brain regions and for both noise bursts and warbles, we cannot decode the sound any more accurately from the responses on localization trials than on pitch discrimination trials. Moreover, we did not observe any correlation between each neuron's change in anticipatory activity and its tuning for the stimuli (Figures S5F and S5G) or any indication that some stimulus pairs (e.g., those requiring different responses in each block) elicited a greater response or a stronger modulation of that response (Figures S5H and S5I).

Thus, whereas neurons showing an increased prestimulus firing rate in one block generally showed a corresponding increase in the stimulus-evoked rate during the same block, these changes in evoked rate did not enhance the representation of the target sound, at least in any way that our decoding analysis could detect. However, as we discuss further below, it may be that in this task the brain does not need to maximize the information available about the target sound via tuning changes. (After all, even the small ensembles of neurons we recorded provided useful information about the identity of both sounds; given access to ensembles on the scale of auditory cortex, the brain should be able to decode the stimuli with virtual certainty.) Rather, the challenge of stimulus selection in a task such as ours may be to flexibly re-route the relevant stimulus information to the appropriate motor neurons at every block change.

Disruption of mPFC Significantly Impairs Task Performance

mPFC has been shown to be required for many task switching paradigms, which prompted us to ask whether it is required for our task. To answer this question, we developed an electrical disruption technique, inspired by transcranial magnetic stimulation (TMS) in humans (Dayan et al., 2013). We first implanted mPFC of three trained animals with extracellular stimulating electrodes. On 20% of trials ("disruption" trials), we injected a 10 Hz train of current pulses (see Experimental Procedures) throughout both the hold period and the auditory stimulus. Such electrical stimulation drives an extremely rapid activation of nearby neurons (Histed et al., 2009), followed by a slower suppression of firing rates (Butovas and Schwarz, 2003; Logothetis et al., 2010) for a few hundred milliseconds. Thus, the primary effect of this approach is neither to silence nor activate the brain region, but rather to disrupt the normal firing rates and patterns. Moreover, because we did not observe any spatial clustering of neurons preferring one task or the other, it is unlikely that such microstimulation would preferentially activate neurons of either preference, even if the stimulation protocol were purely excitatory.

Across all three animals, electrical disruption tended to impair performance (Figure 7) in both localization and pitch discrimination. This impairment was significant across sessions ($p < 0.05$, binomial test on the number of impaired sessions) during pitch discrimination for all (three of three) rats and during localization for most (two of three) rats. Electrical disruption largely, although not exclusively, affected performance on nogo trials. All rats were impaired on pitch discrimination nogo trials in almost all sessions. Some rats were additionally impaired on go or localization nogo trials. These effects were generally quite strong within individual sessions even though they varied between rats (Figures 7B and 7C). Taken together, these data suggest that in the absence of normal mPFC

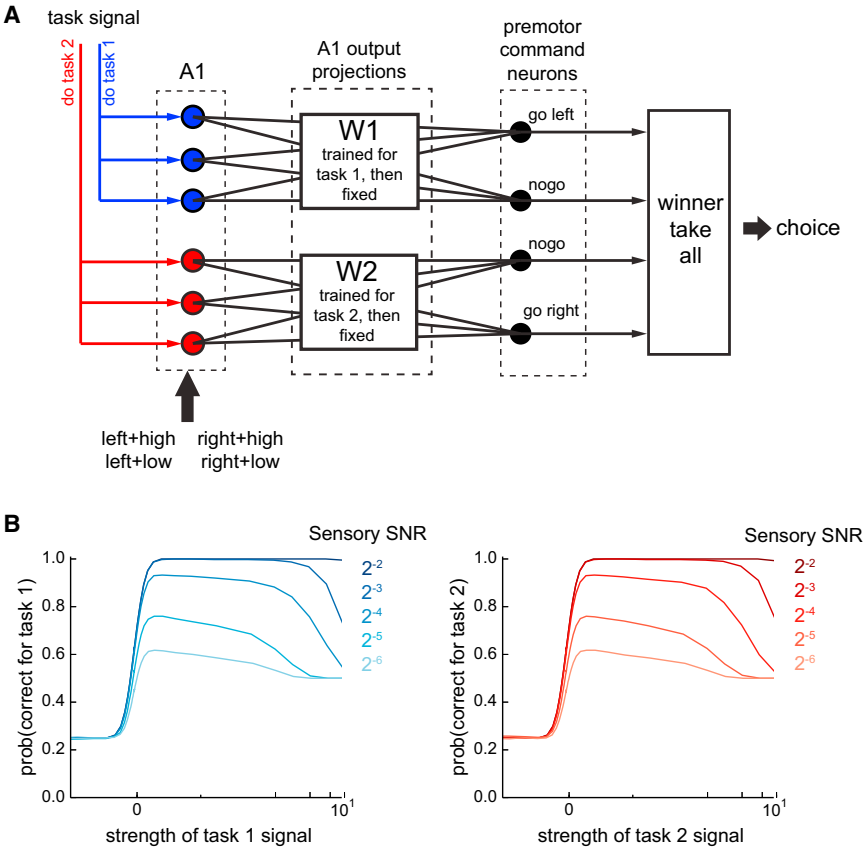

Figure 8. A Simulated Network Model Using Anticipatory Modulation Can Perform Stimulus Selection

(A) Network connectivity. Simulated A1 consists of N neurons, each with random tuning for the four task stimuli and subject to additive Gaussian noise. Red and blue populations are activated in one or the other block by an excitatory "task signal" projection. Each population connects to a set of premotor command neurons encoding the possible responses in that block. The projection weights W1 and W2 are optimized independently during an initial supervised training phase and constrained to be excitatory. The most active command neuron determines the network's choice.

(B) Performance of the model for N = 320 neurons on task 1 (left) and task 2 (right). We tested a range of values for the sensory signal-to-noise ratio (SNR), defined as the ratio of the tuning for sensory stimuli to the strength of the additive Gaussian noise in each A1 neuron. At the highest SNR of 0.25 (darkest trace), the model produces 100% correct responses for virtually any positive task switch signal. (Negative task signals correspond to activating the incorrect population for the current task.) At low SNRs, as the task signal increases in strength, the sensory input is eventually drowned out and the model's performance falls to chance (50%).

See also Figure S7.

activity, each rat resorts to its default strategy (typically "always go") in one or both blocks. Normal activity in mPFC is therefore important for good performance in our paradigm, but the strong impairment on nogo trials in particular made it difficult to ascertain whether the primary effect was on stimulus selection, as opposed to impulse control or some other aspect of the task (Figure S6).

A Simulated Network Model Demonstrates How Modulation of Anticipatory Activity Could Solve the Stimulus Selection Problem

Our data suggest a simple model of how the brain might perform stimulus selection: the PFC sends tonic excitation, perhaps indirectly, to a specific population of neurons for each task (e.g., populations for localization and pitch discrimination), which increases both their anticipatory and stimulus-evoked activity but does not affect their tuning. These populations connect to downstream motor regions that produce the appropriate response for that block; however, only the population with the increased firing rates can control behavior. We produced a quantitative simulation of this model to show that it can indeed solve the problem of stimulus selection. The simulation: (1) requires only random stimulus tuning in A1, (2) does not require tuning changes or synaptic reweighting after the initial training phase, and (3) uses only excitatory connections, consistent with the observation that most long-range projections in the brain are excitatory (Logothetis et al., 2010).

The model (Figure 8; see Supplemental Experimental Procedures for details) consists of a population of N neurons in A1, randomly tuned for each of our four stimulus pairs and subject to Gaussian noise. Half of the neurons are arbitrarily assigned to each task. Each population projects to two command neurons encoding the two possible behavioral responses during that block (e.g., go left and nogo during localization); this projection is trained to activate the correct command neuron for each stimulus pair. The actual behavioral choice is determined by which command neuron is the most active ("winner-take-all").

After the training phase, the synaptic weights are fixed and the model is tested on its ability to produce the correct response in each block. To simulate the rule encoding we observed in our data set, a "task signal" is added to the activations of the neurons in the appropriate population for the current block. Because all feed-forward weights are positive, adding this task signal translates into an excitatory boost to the premotor neurons receiving input from that population. Thus, even without any synaptic reweighting, the model tends to choose the response appropriate for the current block and stimulus. With 320 neurons, the network performs above 80% correct even with a signal-to-noise ratio (SNR) as low as 0.0625 (i.e., very weak sensory responses in each neuron relative to its internal noise). Increasing the network size can lower this SNR limit even further (Figure S7). Thus, our model demonstrates that a network can perform stimulus selection by task-specific activation of neurons, even without task-specific adaptation of their tuning.

DISCUSSION

Auditory Stimulus Selection: Task Switching between Conflicting Auditory Discriminations

When human listeners hear two simultaneous voices, they can selectively attend and respond to either one. This is a complex ability, and the rodent task we have developed models part of it—selecting and responding to one of two simultaneous sounds. Our subjects can voluntarily switch which sound they select, and do so at each block change within a single recording session.

Although previous studies of task switching in rodents did not require stimulus selection, they did require subjects to switch the navigational strategy they use to solve a maze (Rich and Shapiro, 2009) or between a sensory discrimination and a fixed response (e.g., "follow the light" versus "always go left"; Floresco et al., 2008; Durstewitz et al., 2010). It has been challenging to extend these results to task switching between distinct sensory discriminations, perhaps because this requires ignoring a previously trained stimulus. Even in cross-modal tasks, in which the targets and distractors come from entirely different modalities, strong cueing mechanisms (violating our "same stimulus; different response" condition) have been used to induce the switch: introducing novel stimuli (Birrell and Brown, 2000), deleting distractors (Otazu et al., 2009), or changing the behavioral arena completely (Haddon and Killcross, 2007). Finally, most previous studies required rats to shift no more than once per session, sometimes just once per lifetime, whereas our study requires multiple switches per session.

Despite its clinical and computational relevance (Ding and Simon, 2012), the auditory cocktail party problem remains less studied than comparable visual tasks. One multi-unit study (Lakatos et al., 2013) required primates to select a target stream of tones; however, the subjects were unable to ignore any distractor stream within one octave of the target. Human voices typically overlap extensively in acoustic frequency (McDermott, 2009), which contributes to the difficulty of the cocktail party problem, and we thus designed our stimuli to overlap in frequency. In sum, we believe our task represents an important first step toward understanding the cocktail party problem in rats, paving the way toward future studies with the modern tools available in rodent models (e.g., viral vectors for manipulating genetically identified cell types).

Anticipatory Activity in both mPFC and A1 Encodes the Selection Rule

We found that rodent mPFC robustly encodes the selection rule, analogous to the rule-encoding role of the primate prefrontal cortex (Asaad et al., 2000; Wallis et al., 2001; Johnston et al., 2007). Rule encoding develops in the mPFC population over 2.5 s before the stimulus onset, as the rat is planning to initiate a trial or even finishing the previous trial. Thus, we find that the prefrontal cortex densely and persistently codes for cognitive variables (cf. Rigotti et al., 2013), in contrast to the sparse coding of stimuli typical of sensory cortex (Hromádka et al., 2008; Olshausen and Field, 1996). This dense and widespread coding of selection rule in our data are perhaps surprising because only one bit of information needs to be encoded—pitch discrimination or localization—and this information is only necessary while making a decision on each trial. This persistent activity may represent a memory trace of the selection rule (Funahashi et al., 1989) or it may represent a shift to a completely different network state (Karlsson et al., 2012) depending on which stimulus the rat plans to select.

We also observed rule encoding in A1, a surprising result because this has traditionally been considered the domain of prefrontal areas. However, attention is known to induce anticipatory activity in sensory areas (Luck et al., 1997; Reynolds et al., 2000; Chen and Seidemann, 2012; Kastner et al., 1999; Thut et al., 2006). More generally, single neuron activity in primary sensory cortex can predict a motor response (Niwa et al., 2012) or an expected reward (Shuler and Bear, 2006), and anticipation of a visual stimulus can trigger a hemodynamic response in V1, although without a corresponding change in spiking (Sirotin and Das, 2009). Therefore, perhaps it is not surprising that primary sensory cortex could also encode selection rule. In this way, both the stimulus and the information about how that stimulus should be interpreted are encoded in the same neurons, providing a potential locus for the behavioral decision to be made.

Finally, we observed a surprising amount of similarity between A1 and mPFC, both of which showed robust encoding of the selection rule and of behavioral choice (Figure 5D). In monkeys, attention effects become more prominent higher in the visual hierarchy (Reynolds and Heeger, 2009). In contrast, our results show that rat A1 already encodes a nonsensory variable. This could be a difference between rats and monkeys, or between auditory and visual cortex, or both.

Comparison with Studies of Selective Attention and Task-Relevant Plasticity

In this study, we found limited evidence for any modulation of sensory-evoked responses in A1 beyond an additive effect of increased baseline. In particular, the neurons did not appear to change their tuning to encode the target stimulus with greater fidelity. This is consistent with some, but not all, previous studies of auditory task switching. For instance, switching between temporal and spatial auditory discriminations does not significantly change spatial tuning in A1 at the population level (Lee and Middlebrooks, 2011), although switching between passive and engaged states robustly modulates neuronal sensitivity (Otazu et al., 2009; Lee and Middlebrooks, 2011).

However, a series of pioneering experiments did demonstrate task-relevant tuning changes in A1 of ferrets trained to detect a target frequency (Fritz et al., 2003, 2010). One important methodological difference is that their study, unlike ours, made use of a large battery of probe stimuli and was therefore better optimized to detect receptive field changes. This plasticity was nuanced: it could induce both facilitation and, intriguingly, suppression at the task-relevant frequency; whether facilitation or suppression was more prevalent depended on whether positive or negative reinforcement was used (David et al., 2012). Further studies of complex auditory behaviors will be necessary to better understand the factors that determine whether a given behavioral paradigm produces tuning changes in auditory cortex.

Visual selective attention has been shown to enhance target representations and suppress distractors in V4 and other visual areas (Cohen and Maunsell, 2011; David et al., 2008; Mitchell et al., 2007; Reynolds and Heeger, 2009). However, selective attention consists of two component processes with separate behavioral measures: stimulus selection and perceptual enhancement (cf. Knudsen, 2007; Reynolds and Chelazzi, 2004; Pestilli et al., 2011). Perceptual enhancement is typically measured using faint stimuli to probe psychophysical thresholds (Cohen and Maunsell, 2009). In contrast, studies such as ours that use easily detectable stimuli far above threshold (Hocherman et al., 1976; Stoet and Snyder, 2004) often report limited evidence for enhanced representations of target stimuli in sensory cortex (Sasaki and Uka, 2009; Mante et al., 2013; Pestilli et al., 2011). In such tasks, the dominant challenge is not detecting the stimuli, but rather selecting the relevant target, which may rely on changes of baseline (Pestilli et al., 2011) perhaps due to anticipatory modulation (Chen and Seidemann, 2012). Similarly, the cocktail party problem is often difficult because all voices are of competing intensity, not because the target voice is barely audible. Thus, the nature of the task may determine whether stimulus representations are modulated or remain relatively fixed.

Some models of visual selection (Gilbert and Shallice, 2002; Mante et al., 2013) propose that stimulus selection occurs in frontal areas, not sensory cortex. Our data are similar in the sense that we do not observe tuning changes in sensory cortex (Mante et al., 2013) but different in that we do not observe strong representations of the stimuli in PFC, similar to a recent observation in primate PFC (Lara and Wallis, 2014). Our results are more consistent with a distributed processing model in which contextual information from PFC modulates activity in A1 to increase the fidelity with which the appropriate motor action can be read out (Fritz et al., 2010; David et al., 2012; Blake et al., 2002).

The Potential Roles of Motor Planning and Posture

We considered the potential roles of both posture—the angle of the rat's head relative to the behavior box in particular—and motor planning in driving the observed task-specific modulation of neuronal activity. Because each block is associated with a different choice port, it is plausible that the rats adopted a different default motor plan for the two blocks: go left for one task and go right for the other. Moreover, we observed a difference in head angle between blocks, presumably a behavioral strategy that rats used to prepare for the differing motor actions required. We note a similarity with some blocked visual spatial attention tasks, in which 80% of the trials require a saccade in the same direction (Cohen and Maunsell, 2009). In all such tasks, it can be difficult to disambiguate response selection and stimulus selection (Erlich et al., 2011; Sato and Schall, 2003; Steinmetz and Moore, 2012).

We found that some rule-encoding neurons, especially in mPFC, also encoded head angle to some extent (Figures S2K and S3K). This is consistent with previous mPFC data (Euston and McNaughton, 2006) and the idea that single prefrontal neurons simultaneously encode disparate sensorimotor and cognitive signals (Rigotti et al., 2013). However, we found that the firing rate of most neurons was better explained by block than by head angle and that rule encoding persisted even when we controlled for head angle by trial selection (Figures S2I–S2T, S3I–S3Q). These results suggest that cognitive context (i.e., task rule) drives both the observed neuronal activity and the adaptive posture, rather than posture directly driving the neuronal activity.

Even if the rule-encoding activity we observed does not simply reflect postural differences, it is possible that it represents an internal motor plan (which could be covertly present even in the absence of a measurable behavioral parameter like head angle). It is difficult to disambiguate motor planning from rule encoding because the task itself requires different sensorimotor mappings in each block. However, the time course of the rule encoding was quite protracted in many neurons, in some cases even persistent throughout the block (Figure 5), during which time the rat was engaged in various motor actions such as moving to or from the center port and harvesting reward (see example neurons in Figures S2P and S3P). It seems unlikely that neurons would continue to represent the specific action of moving from the center port to the choice port on such a long timescale. In addition, in our paradigm any default motor plan is subject to cancellation on nogo trials (sometimes called "countermanding," cf. Schall et al., 2000; Eagle and Robbins, 2003; Eagle et al., 2008), because the animal does not know during the anticipatory period whether it will be signaled to perform a go response or not. Finally, we did not observe any correlation between the anticipatory firing rate and the reaction or movement time (Figures S2G, S2H, S3G, and S3H). To summarize, our task requires remapping sensory stimuli to motor responses, and it is reasonable to expect rule encoding to incorporate both the sensory and motor planning aspects of this remapping. Anticipatory modulation may encode both the selection rule and, therefore, the motor plan required to implement that rule.

Stimulus Selection via Activation of Latent Circuits for Each Target

In light of our results, we propose a model for stimulus selection based on activation of separate, task-specific circuits. In this model, there are two neuronal populations in both A1 and mPFC—one activated during the localization block, the other during the pitch discrimination block. These populations show increased prestimulus and stimulus-evoked activity during their preferred block but do not change their tuning for specific stimuli. We hypothesize that the difference between these populations is their downstream connectivity: each may project to separate targets in a downstream effector region such as the striatum (Znamenskiy and Zador, 2013), forming distinct circuits for each task. In this model, only one circuit is activated at a time, via feed-forward excitation perhaps originating in mPFC, and only this circuit has sufficient baseline activity to drive behavior. In some ways, this model is more parsimonious than the traditional tuning change model, which requires that prefrontal (or other) brain regions be able to modulate the tuning of many A1 neurons as quickly as the subject shifts the focus of attention. Although attention does produce tuning changes (David et al., 2008; Fritz et al., 2003)

over minutes (which is typically the fastest that they can be estimated), it is unclear how known synaptic plasticity mechanisms could mediate task-specific tuning changes on a subsecond timescale.

Our model makes several testable predictions. First, there should exist "premotor" neurons (possibly in the striatum) receiving input from A1 that also show a block-dependent anticipatory modulation. In addition, specific activation of one of the subpopulations in mPFC, A1, or striatum should bias behavior toward the block preferred by that subpopulation. Such a manipulation would require targeting specific neurons based on their anticipatory firing rate, a challenging experiment that might nonetheless be feasible using activity-dependent promoters to drive light-gated ion channels, for example.

In conclusion, these results establish the rat as a model organism for auditory stimulus selection, paving the way for future investigations of the cocktail party problem with emerging optical and genetic tools amenable to rodents. We found widespread and robust rule encoding in mPFC and A1, although we observed little change in the stimulus tuning of evoked responses. We propose a simple model to explain these results: task-specific activation of latent circuits, rather than task-specific adaptation of a single circuit.

EXPERIMENTAL PROCEDURES

All procedures were approved by the Animal Care and Use Committee at the University of California, Berkeley. We used male Long-Evans rats (Harlan), housed in pairs. Training began when their body mass reached 150–225 g, at approximately 45–60 days old. Rats were given restricted access to water in the day before the training session so that they would be motivated to obtain a water reward. After each session, they were given ad lib access to water for 1 hour. We monitored body weight and other markers to ensure they remained healthy. We used standard behavioral shaping and surgical implantation techniques (see Supplemental Experimental Procedures).

Electrical Disruption Protocol

We began with a very low current, ~10 uA per electrode, which was typically too low to produce any behavioral effect. We wanted to use a minimal perturbation to ensure that the effects were as localized as possible in both time and space, and so we used pilot sessions to increase the amount of current until performance on the task became moderately impaired. During the testing sessions that we report in the main text, the mean currents used were 37 uA, 41 uA, and 23 uA per electrode for Z1, Z2, and Z3 respectively. See Supplemental Experimental Procedures for further details and comparisons with other studies.

Data Analysis

We preprocessed the data using the open-source OpenElectrophy software suite (Garcia and Fourcaud-Trocmé, 2009) built on the Neo object model (Garcia et al., 2014). We used KlustaKwik (Kadir et al., 2013) and Klusters (Hazan et al., 2006), while blind to the experimental variables, to identify single units.

We analyzed the data with Python within the IPython environment (Pérez and Granger, 2007) and the modules numpy, scipy, matplotlib, scikits-learn, statsmodels, and pandas. We also conducted some statistical analyses in R using the rpy2 module. Except where otherwise noted in the text, we observed consistent results across all subjects and therefore pooled the data (Figures S2D, S2E, S3D, and S3E). All of the data and code necessary to recapitulate the analyses presented here are available online at https://github.com/cxrodgers/Rodgers2014 and at the data-sharing website CRCNS.org through link http://dx.doi.org/10.6080/K0W66HPJ.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures and seven figures and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2014.04.031.

ACKNOWLEDGMENTS

The authors thank Mohammad Dastjerdi and Christian Fernandez for developing the initial rodent behaviors in the lab; Alfonso Apicella for helping set up the lab; Anita Devineni, Tony Zador, Uri Livneh, Erin Rich, Nicholas Steinmetz, and Aaron Koralek for comments on the manuscript; Sarah Kochik and Vuong Vu for managing behavioral training; Trevor Dolinajec, Lily Lin, Danielle Jobe, and Yassi Sabahi for electrode fabrication; Yuliy Tsank, Sharon Lin, Ambika Rustagi, Amish Shah, Carol Pham, Francine Foo, Helene Moorman, Jason Murphy, Jason Zhang, Josh Chong, Kelly Clancy, Tanling Hsu, Pooja Upadhyaya, and Tommy Li for help with rat behavior; and Vivek Ayer for technical assistance.

Accepted: April 3, 2014
Published: June 4, 2014

REFERENCES

Ahveninen, J., Hämäläinen, M., Jääskeläinen, I.P., Ahlfors, S.P., Huang, S., Lin, F.-H., Raij, T., Sams, M., Vasios, C.E., and Belliveau, J.W. (2011). Attention-driven auditory cortex short-term plasticity helps segregate relevant sounds from noise. Proc. Natl. Acad. Sci. USA *108*, 4182–4187.

Asaad, W.F., Rainer, G., and Miller, E.K. (2000). Task-specific neural activity in the primate prefrontal cortex. J. Neurophysiol. *84*, 451–459.

Bee, M.A., and Micheyl, C. (2008). The cocktail party problem: what is it? How can it be solved? And why should animal behaviorists study it? J. Comp. Psychol. *122*, 235–251.

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

Blake, D.T., Strata, F., Churchland, A.K., and Merzenich, M.M. (2002). Neural correlates of instrumental learning in primary auditory cortex. Proc. Natl. Acad. Sci. USA *99*, 10114–10119.

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95–98.

Butovas, S., and Schwarz, C. (2003). Spatiotemporal effects of microstimulation in rat neocortex: a parametric study using multielectrode recordings. J. Neurophysiol. *90*, 3024–3039.

Carandini, M., and Churchland, A.K. (2013). Probing perceptual decisions in rodents. Nat. Neurosci. *16*, 824–831.

Carlson, N.L., Ming, V.L., and Deweese, M.R. (2012). Sparse codes for speech predict spectrotemporal receptive fields in the inferior colliculus. PLoS Comput. Biol. *8*, e1002594.

Chen, Y., and Seidemann, E. (2012). Attentional modulations related to spatial gating but not to allocation of limited resources in primate V1. Neuron *74*, 557–566.

Cherry, E.C. (1953). Some experiments on the recognition of speech, with one and with two ears. J. Acoust. Soc. Am. *25*, 975–979.

Cohen, M.R., and Maunsell, J.H.R. (2009). Attention improves performance primarily by reducing interneuronal correlations. Nat. Neurosci. *12*, 1594–1600.

Cohen, M.R., and Maunsell, J.H.R. (2011). Using neuronal populations to study the mechanisms underlying spatial and feature attention. Neuron *70*, 1192–1204.

David, S.V., Hayden, B.Y., Mazer, J.A., and Gallant, J.L. (2008). Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. Neuron *59*, 509–521.

David, S.V., Fritz, J.B., and Shamma, S.A. (2012). Task reward structure shapes rapid receptive field plasticity in auditory cortex. Proc. Natl. Acad. Sci. USA *109*, 2144–2149.

Dayan, E., Censor, N., Buch, E.R., Sandrini, M., and Cohen, L.G. (2013). Noninvasive brain stimulation: from physiology to network dynamics and back. Nat. Neurosci. *16*, 838–844.

Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. Annu. Rev. Neurosci. *18*, 193–222.

DeWeese, M.R., Wehr, M., and Zador, A.M. (2003). Binary spiking in auditory cortex. J. Neurosci. *23*, 7940–7949.

Ding, N., and Simon, J.Z. (2012). Emergence of neural encoding of auditory objects while listening to competing speakers. Proc. Natl. Acad. Sci. USA *109*, 11854–11859.

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

Eagle, D.M., and Robbins, T.W. (2003). Inhibitory control in rats performing a stop-signal reaction-time task: effects of lesions of the medial striatum and d-amphetamine. Behav. Neurosci. *117*, 1302–1317.

Eagle, D.M., Baunez, C., Hutcheson, D.M., Lehmann, O., Shah, A.P., and Robbins, T.W. (2008). Stop-signal reaction-time task performance: role of prefrontal cortex and subthalamic nucleus. Cereb. Cortex *18*, 178–188.

Erlich, J.C., Bialek, M., and Brody, C.D. (2011). A cortical substrate for memory-guided orienting in the rat. Neuron *72*, 330–343.

Euston, D.R., and McNaughton, B.L. (2006). Apparent encoding of sequential context in rat medial prefrontal cortex is accounted for by behavioral variability. J. Neurosci. *26*, 13143–13155.

Floresco, S.B., Block, A.E., and Tse, M.T.L. (2008). Inactivation of the medial prefrontal cortex of the rat impairs strategy set-shifting, but not reversal learning, using a novel, automated procedure. Behav. Brain Res. *190*, 85–96.

Fritz, J., Shamma, S., Elhilali, M., and Klein, D. (2003). Rapid task-related plasticity of spectrotemporal receptive fields in primary auditory cortex. Nat. Neurosci. *6*, 1216–1223.

Fritz, J.B., Elhilali, M., and Shamma, S.A. (2005). Differential dynamic plasticity of A1 receptive fields during multiple spectral tasks. J. Neurosci. *25*, 7623–7635.

Fritz, J.B., David, S.V., Radtke-Schuller, S., Yin, P., and Shamma, S.A. (2010). Adaptive, behaviorally gated, persistent encoding of task-relevant auditory information in ferret frontal cortex. Nat. Neurosci. *13*, 1011–1019.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J. Neurophysiol. *61*, 331–349.

Garcia, S., and Fourcaud-Trocmé, N. (2009). OpenElectrophy: an electrophysiological data- and analysis-sharing framework. Front. Neuroinform. *3*, 14.

Garcia, S., Guarino, D., Jaillet, F., Jennings, T., Pröpper, R., Rautenberg, P.L., Rodgers, C.C., Sobolev, A., Wachtler, T., Yger, P., and Davison, A.P. (2014). Neo: an object model for handling electrophysiology data in multiple formats. Front. Neuroinform. *8*, 10.

Geissler, D.B., and Ehret, G. (2002). Time-critical integration of formants for perception of communication calls in mice. Proc. Natl. Acad. Sci. USA *99*, 9021–9025.

Gilbert, S.J., and Shallice, T. (2002). Task switching: a PDP model. Cognit. Psychol. *44*, 297–337.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574.

Haddon, J.E., and Killcross, S. (2007). Contextual control of choice performance: behavioral, neurobiological, and neurochemical influences. Ann. N Y Acad. Sci. *1104*, 250–269.

Hazan, L., Zugaro, M., and Buzsáki, G. (2006). Klusters, NeuroScope, NDManager: a free software suite for neurophysiological data processing and visualization. J. Neurosci. Methods *155*, 207–216.

Histed, M.H., Bonin, V., and Reid, R.C. (2009). Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation. Neuron *63*, 508–522.

Hocherman, S., Benson, D.A., Goldstein, M.H., Jr., Heffner, H.E., and Hienz, R.D. (1976). Evoked unit activity in auditory cortex of monkeys performing a selective attention task. Brain Res. *117*, 51–68.

Hromádka, T., Deweese, M.R., and Zador, A.M. (2008). Sparse representation of sounds in the unanesthetized auditory cortex. PLoS Biol. *6*, e16.

Johnston, K., Levin, H.M., Koval, M.J., and Everling, S. (2007). Top-down control-signal dynamics in anterior cingulate and prefrontal cortex neurons following task switching. Neuron *53*, 453–462.

Kadir, S., Goodman, D., and Harris, K. (2013). High-dimensional cluster analysis with the masked EM algorithm. http://arxiv.org/abs/1309.2848.

Karlsson, M.P., Tervo, D.G.R., and Karpova, A.Y. (2012). Network resets in medial prefrontal cortex mark the onset of behavioral uncertainty. Science *338*, 135–139.

Kastner, S., Pinsk, M.A., De Weerd, P., Desimone, R., and Ungerleider, L.G. (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. Neuron *22*, 751–761.

Knudsen, E.I. (2007). Fundamental components of attention. Annu. Rev. Neurosci. *30*, 57–78.

Kvitsiani, D., Ranade, S., Hangya, B., Taniguchi, H., Huang, J.Z., and Kepecs, A. (2013). Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. Nature *498*, 363–366.

Lakatos, P., Musacchia, G., O'Connel, M.N., Falchier, A.Y., Javitt, D.C., and Schroeder, C.E. (2013). The spectrotemporal filter mechanism of auditory selective attention. Neuron *77*, 750–761.

Lara, A.H., and Wallis, J.D. (2014). Executive control processes underlying multi-item working memory. Nat. Neurosci. http://dx.doi.org/10.1038/nn.3702.

Lee, C.-C., and Middlebrooks, J.C. (2011). Auditory cortex spatial sensitivity sharpens during task performance. Nat. Neurosci. *14*, 108–114.

Logothetis, N.K., Augath, M., Murayama, Y., Rauch, A., Sultan, F., Goense, J., Oeltermann, A., and Merkle, H. (2010). The effects of electrical microstimulation on cortical signal propagation. Nat. Neurosci. *13*, 1283–1291.

Luck, S.J., Chelazzi, L., Hillyard, S.A., and Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. J. Neurophysiol. *77*, 24–42.

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

McDermott, J.H. (2009). The cocktail party problem. Curr. Biol. *19*, R1024–R1027.

Mesgarani, N., and Chang, E.F. (2012). Selective cortical representation of attended speaker in multi-talker speech perception. Nature *485*, 233–236.

Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. *24*, 167–202.

Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. Neuron *55*, 131–141.

Monchi, O., Petrides, M., Petre, V., Worsley, K., and Dagher, A. (2001). Wisconsin Card Sorting revisited: distinct neural circuits participating in different stages of the task identified by event-related functional magnetic resonance imaging. J. Neurosci. *21*, 7733–7741.

Moore, T., Armstrong, K.M., and Fallah, M. (2003). Visuomotor origins of covert spatial attention. Neuron *40*, 671–683.

Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. Science *229*, 782–784.

Niwa, M., Johnson, J.S., O'Connor, K.N., and Sutter, M.L. (2012). Activity related to perceptual judgment and action in primary auditory cortex. J. Neurosci. *32*, 3193–3210.

Olshausen, B.A., and Field, D.J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature *381*, 607–609.

Otazu, G.H., Tai, L.-H., Yang, Y., and Zador, A.M. (2009). Engaging in an auditory task suppresses responses in auditory cortex. Nat. Neurosci. *12*, 646–654.

Pai, S., Erlich, J.C., Kopec, C., and Brody, C.D. (2011). Minimal impairment in a rat model of duration discrimination following excitotoxic lesions of primary auditory and prefrontal cortices. Front Syst Neurosci *5*, 74.

Pérez, F., and Granger, B.E. (2007). IPython: a system for interactive scientific computing. Comput. Sci. Eng. *9*, 21–29.

Pestilli, F., Carrasco, M., Heeger, D.J., and Gardner, J.L. (2011). Attentional enhancement via selection and pooling of early sensory responses in human visual cortex. Neuron *72*, 832–846.

Petkov, C.I., Kang, X., Alho, K., Bertrand, O., Yund, E.W., and Woods, D.L. (2004). Attentional modulation of human auditory cortex. Nat. Neurosci. *7*, 658–663.

Ragozzino, M.E., Detrick, S., and Kesner, R.P. (1999). Involvement of the prelimbic-infralimbic areas of the rodent prefrontal cortex in behavioral flexibility for place and response learning. J. Neurosci. *19*, 4585–4594.

Raposo, D., Sheppard, J.P., Schrater, P.R., and Churchland, A.K. (2012). Multisensory decision-making in rats and humans. J. Neurosci. *32*, 3726–3735.

Reynolds, J.H., and Chelazzi, L. (2004). Attentional modulation of visual processing. Annu. Rev. Neurosci. *27*, 611–647.

Reynolds, J.H., and Heeger, D.J. (2009). The normalization model of attention. Neuron *61*, 168–185.

Reynolds, J.H., Pasternak, T., and Desimone, R. (2000). Attention increases sensitivity of V4 neurons. Neuron *26*, 703–714.

Rich, E.L., and Shapiro, M. (2009). Rat prefrontal cortical neurons selectively code strategy switches. J. Neurosci. *29*, 7208–7219.

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

Sasaki, R., and Uka, T. (2009). Dynamic readout of behaviorally relevant signals from area MT during task switching. Neuron *62*, 147–157.

Sato, T.R., and Schall, J.D. (2003). Effects of stimulus-response compatibility on neural selection in frontal eye field. Neuron *38*, 637–648.

Sayers, B.M., and Cherry, E.C. (1957). Mechanism of binaural fusion in the hearing of speech. J. Acoust. Soc. Am. *29*, 973–987.

Schall, J.D., Hanes, D.P., and Taylor, T.L. (2000). Neural control of behavior: countermanding eye movements. Psychol. Res. *63*, 299–307.

Shuler, M.G., and Bear, M.F. (2006). Reward timing in the primary visual cortex. Science *311*, 1606–1609.

Sirotin, Y.B., and Das, A. (2009). Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity. Nature *457*, 475–479.

Steinmetz, N.A., and Moore, T. (2012). Lumping and splitting the neural circuitry of visual attention. Neuron *73*, 410–412.

Stoet, G., and Snyder, L.H. (2004). Single neurons in posterior parietal cortex of monkeys encode cognitive set. Neuron *42*, 1003–1012.

Thut, G., Nietzel, A., Brandt, S.A., and Pascual-Leone, A. (2006). Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. J. Neurosci. *26*, 9494–9502.

Turken, A.U., and Swick, D. (1999). Response selection in the human anterior cingulate cortex. Nat. Neurosci. *2*, 920–924.

Wallis, J.D., Anderson, K.C., and Miller, E.K. (2001). Single neurons in prefrontal cortex encode abstract rules. Nature *411*, 953–956.

Wise, S.P. (2008). Forward frontal fields: phylogeny and fundamental function. Trends Neurosci. *31*, 599–608.

Young, J.J., and Shapiro, M.L. (2011). The orbitofrontal cortex and response selection. Ann. N Y Acad. Sci. *1239*, 25–32.

Zariwala, H.A., Kepecs, A., Uchida, N., Hirokawa, J., and Mainen, Z.F. (2013). The limits of deliberation in a perceptual decision task. Neuron *78*, 339–351.

Zion Golumbic, E.M., Ding, N., Bickel, S., Lakatos, P., Schevon, C.A., McKhann, G.M., Goodman, R.R., Emerson, R., Mehta, A.D., Simon, J.Z., et al. (2013). Mechanisms underlying selective neuronal tracking of attended speech at a "cocktail party". Neuron *77*, 980–991.

Znamenskiy, P., and Zador, A.M. (2013). Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. Nature *497*, 482–485.