Title

Rule-encoding neurons in prefrontal and auditory cortex of rats performing a task similar to the cocktail party problem

Author

Rodgers, Christopher

Publication Date

2013

Peer reviewed|Thesis/dissertation

Rule-encoding neurons in prefrontal and auditory cortex of rats performing a task similar to the cocktail party problem

by

Christopher Rodgers

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Neuroscience

and the Designated Emphasis

in

Computational Science and Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:
Professor Michael DeWeese, Chair
Professor Daniel Feldman
Professor Jonathan Wallis
Professor Michel Maharbiz

Fall 2013

Abstract

Rule-encoding neurons in prefrontal and auditory cortex of rats
performing a task similar to the cocktail party problem

by

Christopher Rodgers

Doctor of Philosophy in Neuroscience

and the Designated Emphasis in
Computational Science and Engineering

University of California, Berkeley

Professor Michael DeWeese, Chair

The human auditory system easily solves the "cocktail party problem" – that is, even when multiple people are speaking at once, we can easily select and pay attention to a single voice while ignoring the others. Though this seems easy to do, the problem is known to be quite computationally complex. It requires identifying the important sound, selecting it for special processing, and using information from it to make behavioral decisions; meanwhile, the other voices must not be allowed to distract us.

How does the brain do this? In chapter 1, I review previous approaches to this question and motivate the choices I made in designing my experiments. In chapter 2, I present the data and conclusions I obtained in collaboration with my advisor, Dr Michael DeWeese. (We are submitting this chapter for publication separately.) In chapter 3, I present a detailed protocol for repeating our behavioral results.

The final chapter, Chapter 4, is broader in scope. I discuss how our models and results relate to existing models of prefrontal control over other brain regions. Finally, I consider what my results have taught me about the scientific process of investigating neural function and ruminate on where this field may be headed next.

Acknowledgements

I would like to acknowledge Mike for always believing in me, for his kind words, and for entrusting me with an enormous amount of freedom to do the things I wanted to do. He invariably encourages free thinking in his students and his tireless faith in science and scientists is an inspiration.

My thesis committee has been very supportive over the years. Dan never failed to provide cool-headed and incisive advice on any topic. His commitment to the graduate program and its students goes the extra mile and more and is much appreciated. Jon provided useful guidance on many topics, first by gently reminding me that science is as much about proposing hypotheses as falsifying them, and later for training us in the ways of psychology. Michel brought a refreshingly frank candor to our meetings; he also always came to our meetings, something I learned not to take for granted.

I would like to thank all of the other graduate students and postdocs in our program. The lateral support provided by all of them is incredible. Without the gracious help and support I received from the students above me (especially Mike Goard, Natalia Caporale, and John Long), I could never have finished grad school and would undoubtedly now be churning out code at some dull social media startup. I tried to learn from their example.

Finally I would like to thank my parents for teaching me that science and computers were cool, and Anita, not only for sharing this journey but for often leading the way.

Table of Contents

Chapter 1. Introduction

What is attention? William James believed the answer was obvious. He defined it via an appeal to the common human experience:

> Everyone knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. (James 1890)

Over a century later, a prolific researcher in the field, Jeffrey Schall, evoked those remarks but proposed a rather more concrete definition. He directly equated attention with the activation of neurons in one particular region of the brain, the frontal eye field (FEF).

> While everyone may know what attention is, the description of attention in the neuroscience literature is rather confused…. [Attention] need be no more than the selective differential activation of neurons in the appropriate network that includes FEF. (Schall 2004)

It is a commendably explicit and falsifiable definition, well-matched to the empirical and reductionist times in which we live. As an auditory neurophysiologist, I note one surprising corollary: because the FEF is a region that exists only within the visual system of primates, this definition excludes the possibility of attention in any non-primate or in any modality other than vision.

In fact we and others (Ding 2012, Mesgarani 2013) believe that selective attention does exist in the auditory system and that it is the mechanism by which the brain solves the "cocktail party problem". The cocktail party problem (Cherry 1953, Sayers 1957) may be described in this way: imagine you are at a cocktail party, surrounded by a din of voices, yet you only care about the voice of the person with whom you are speaking. The attended speaker's voice becomes more salient, more clear, while the ignored voices fade away.

Humans solve the cocktail party problem easily but this is an illusion: it is known to be quite computationally difficult and no algorithms can rival the human brain at this task (McDermott 2009). Clearly our brains are performing some fundamental computation that we do not understand. A desire to identify that computation was our first motivation for joining the search for the cocktail party solution. Our second motivation was that this behavior provides an entry point into probing how the brain controls the flow of information.

Somewhere in the brain, attention must "throw a switch", allowing some sensory representations to rise to the level of consciousness and suppressing others. More concretely, some sensory stimuli (the attended voice) influence behavior, while others (the ignored voices) do not. The switch may be thrown extremely rapidly: in a fraction of a second one may begin ignoring one's partner at the cocktail party in favor of eavesdropping on the conversation behind. Yet there is no known mechanism for how this may occur.

Synapses connect neurons in the brain; could synaptic rewiring be the underlying mechanism of this switch? As is so often the case in neuroscience, there is a large gap between what we know the brain can do, and the specific mechanisms that have been characterized in single neurons or circuits of neurons. Synaptic plasticity, thought be the basis of learning and memory, requires many pairs of pre- and post-synaptic spikes, occurring on the timescale of minutes (Feldman 2012), much slower than the timescale of attention.

So how is the switch thrown so quickly? Despite decades of intensive research into selective attention, we still do not know. We began our investigations, perhaps over-ambitiously, by developing a completely new approach to the question. Selective attention has been almost exclusively studied thus far in primates, human and otherwise. Moreover the bulk of the work was in visual attention; data on other senses are much more scarce. We chose to work with rats because they are amenable to some techniques that are not feasible in primates, but we first had to develop a behavioral task for rats that shared at least some features with existing primate attention tasks. We also chose to work in the auditory system, even though much less is known about it than about vision, because it is the most important sense for our motivating cocktail party problem.

I first briefly review what is known about selective attention – the preferential processing of important stimuli at the expense of unimportant stimuli – before discussing our results.

The canonical task

To empirically investigate a cognitive phenomenon, one must begin by defining a behavioral task and asserting that it requires the cognitive phenomenon of interest. If we wanted to study the neural basis of mathematics, why not just ask subjects to perform mathematical operations in their heads while we measure some aspect of their physiology? The problem is that we have very little control over the subject's internal state, even in the best of circumstances.[1] Only when we are sure that the subject is performing our task, consistently and at the limits of his ability, can we even begin to identify our measurements with the cognitive phenomenon we wish to study.

Selective attention has most often been investigated using the Posner task (Posner 1980). The subject is presented with two visual stimuli, for instance, a grating on the left and a grating on the right. After a random delay period, one of them will be subtly changed – for instance, it may undergo a slight change in orientation. The subject must detect and report when he detects this *target change*, perhaps by pressing a lever. Critically, before the trial begins the subject is cued in some way that one of the two stimuli is more likely to undergo the shift. The subject must report the change no matter where it occurs, but it is more likely to occur at the cued location than at the uncued location. The fact that the cued location is more likely to contain the shift is

[1] "The mind is its own place, and in it self / Can make a Heav'n of Hell, a Hell of Heav'n." John Milton, *Paradise Lost*, London: 1674. I:244-245. Print. These lines are spoken by Satan.

the critical difference that makes this a selective attention task, rather than a simple sensory discrimination.

Assuming the subject is highly motivated to perform well, he will attempt to maximize his performance on the task. But the task can be made arbitrarily difficult by the use of increasingly subtle target changes; eventually, the subject will not be able to detect targets perfectly. Can the subject transfer cognitive resources from processing one stimulus to the other? If so, then it is to his advantage to do so -- to process preferentially the cued ("attended") stimulus at the cost of the uncued ("ignored") stimulus.

In fact it is absolutely the case that human and primate subjects process preferentially the cued stimulus (Cohen 2009, Reynolds 2004). We know this because the performance on *valid* trials (during which the target occurred at the cued location) is significantly higher than the performance on the less common *invalid* trials (during which the target occurred at the uncued location). This classic result that has now been replicated many times under many conditions, and this behavioral effect is now identified with the psychological concept of selective attention by most researchers in the field.

Note the logical flow: a cognitive phenomenon was subjectively observed; scientists invented a task that seemed likely to produce such a phenomenon in a subject and hypothesized that subjects performing such a task would show certain behavioral metrics; those metrics were observed. No step in this argument is rigorously proven, and it goes without saying that we never really know whether the personal and subjective experience of a human subject, let alone a non-human primate, mirrors our own private experience of choosing to pay attention to a stimulus. However, no one has identified a behavioral metric that falsifies the assertion that the Posner task does produce selective attention in the subject, and so for now it remains accepted theory.

In part due to the success of this model, it is now difficult, or perhaps impossible, to propose an alternative task to probe selective attention that differs in any substantial way from the Posner task.

- The target must be subtle, for instance a slight change in orientation. If the change were not subtle, then the subject would still detect it easily even at the uncued location. There would be no difference in performance between valid and invalid trials, which is the defining metric. Might the subject still be attending the cued location? We cannot know.[2]
- The cue cannot be perfectly reliable, because then we would have no invalid trials to measure. Yet it stands to reason that the allocation of cognitive resources to the target

[2] If the subject is human, the experimenter can simply ask what he is attending. This is often proposed as an advantage of human research over animal research. This is not rigorous, because there is no way to know if the subject is reliably reporting his internal state. Quantitative performance metrics must be used in order to prove (*e.g.* by measuring performance on invalid trials) that a cognitive state has been produced.

stimulus should increase with the reliability of the cue. In this case, we would "know" that the subject is using selective attention, yet we would be completely unable to prove it!

Thus, one particular issue with the Posner task is that we can never know if the observed results are generalizable to selective attention, or particular to the Posner task (for instance, to the problem of subtle targets and imperfect cues). Another issue is that it has been consistently difficult to dissociate the allocation of attention from the planning of saccades (eye motion), since saccades are typically used to indicate behavioral choice in this task. In fact, more recent results have suggested that these two effects are less separable than previously thought (Zénon 2012).

Our task

A more troubling problem is that we may have strayed too far from the original motivating question. When asked what attention is, the man on the street would probably propose a situation like the following: a schoolchild is supposed to be taking notes on his teacher's lecture, yet cannot help but listen to his friend beside him who is making plans for the weekend. It's not clear that this maps onto the Posner task. Is the student's performance contingent on detecting a subtle variation in the teacher's voice? Is there any uncertainty in which person's voice is the more important?

In fact it seems that the student in such a situation faces a far more discrete choice. He is faced with two dichotomous choices: take notes, or talk with his friend. In either case both voices enter his auditory system, but the student's behavior differs depending on which voice he selects.

Motivated by such thoughts, we set out to design a synthetic, controlled version of the student's problem. In our reduced preparation, we present the subject (a rat) with two simultaneous sounds on every trial (Sound A + Sound B). The rat has been previously trained on each sound individually and knows the correct behavioral response for each one (A means X; B means Y). These two behavioral responses may conflict. We resolve the uncertainty by cueing the subject -- telling it which of the two sounds it should select. The subject's behavior reports whether it correctly selected the sound.

The rat in this task is performing a non-trivial feat. On any given trial, the sensory stimuli presented are identical. Yet, depending on which sound we tell it to select, it correctly performs either action X or Y. That is, the same stimulus produced a different response. How does the brain produce this behavior? Somewhere, a "switch must be flipped" in order to route the same incoming sensory information into opposing motor production circuits.

Expected neural effects

Returning to the existing literature, the predominant effect observed in the neural activity of subjects performing selective attention tasks is amplification of the target representation and suppression of the distractor (*i.e.*, ignored stimulus) representation. This has been characterized in a number of different ways.

Non-parametric approaches

One approach is to build decoders: computer algorithms that are trained to predict the target stimulus in terms of the recorded neural activity. Increased performance of the decoder is taken as evidence that the target stimulus was more strongly represented, or amplified, in the neural activity. This has proven successful in the auditory cortex of humans solving the cocktail party problem (Mesgarani 2013, Zion-Golumbic 2013).

The decoders may be implemented in any number of ways, and it is rarely or never claimed that the brain itself is using the same particular implementation. Rather, the claim is that the decoder's increased performance is a metric of attentional amplification. A strength of this approach is that it provides a proof of principle that the target stimuli could be read out with greater fidelity than the distractor by at least one technique. A disadvantage is that it is rarely clear what exactly has changed in the neural activity, let alone how the neurons computed that change.

Parametric approaches: changes in receptive field

In this approach, the experimenter first parameterizes the stimulus space by fitting a model between the neuron's firing rate and the luminance at every location on the retina while presenting a wide battery of task-irrelevant probe stimuli. The coefficients in this model are taken as the *receptive field* of the neuron – the area in stimulus space to which it responds. The receptive fields measured while the subject attends stimulus A or B are compared. The typical finding is that the receptive field more closely matches the attended stimulus' properties (Reynolds 2000, Cohen 2011, David 2008).

The parametric approach is satisfying in that it produces clear pictures of receptive fields and statistically quantifiable metrics of their change. The implicit argument is that attention works by changing receptive fields to better match the stimuli. But I dispute this argument.

- If the probe stimuli were task-irrelevant, then they were not attended and the response to them cannot be a measure of attention.
- If the probe stimuli were the task-relevant targets, then it is tautological to say that the neurons respond more to the target stimuli because their receptive fields changed to match the targets. Both claims arise from the same basic finding: that target stimuli elicit higher responses.

The problem is that receptive fields are derived quantities that do not physically exist in the brain. They are a way to visualize the effects of attention, but they cannot in themselves be a

mechanism. If it were the case that the receptive field changes occurred because of axon rewiring or synaptic plasticity on the neuron's inputs, then this could be a mechanism of attention. But such an explicit claim is never made, probably because it seems highly implausible that such rewiring could occur with the timescale and specificity required by the task.

Just as with the Posner task behavior, I feel the argument has become circular. The question (how does attention work?) begat a method (estimating receptive field changes); then the method became the question – the field began to focus narrowly on cataloguing and parameterizing receptive field changes.

The topographic / gaze control model

Perhaps the most plausible model, from a mechanistic and evolutionary point of view, is the topographic model (Reynolds 2004). This model explicitly relies on the retinotopic organization of visual cortex, which means that neural representations of the target and distractor stimuli are spatially segregated. The attended location is encoded in frontal areas. Topographic projections from frontal areas to sensory cortex may be activated such that the area of sensory cortex encoding the target location is receiving additional excitatory input, thus amplifying the stimulus representations there. An important result (Moore 2003) lends enormous credence to this model: electrically stimulating the area of frontal cortex corresponding to a certain location on the retina produces behavioral effects that are very similar to cueing that area in space.

One problem with this model is that it would seem to predict an increased firing rate in sensory cortex even in the absence of explicit stimuli, yet such pre-stimulus firing has rarely been reported. Another is that it relies on topography of representation and of projection. How may it be applied to tasks that are not spatial, such as attending a certain color of object wherever it appears?

More problematic is that it is unclear how this could be applied to the auditory case. In auditory cortex, maps are much coarser (Rothschild 2010); vocalization stimuli overlap extensively in acoustic frequency and thus in stimulus representation within the cortex (Ding 2012). It defies reason that a certain top-down projection from frontal to auditory cortex exists for every possible human's voice that the subject may choose to attend. One could imagine that the projection could be calculated on the fly, but this is hardly more mechanistic than simply asserting that neurons encoding the target sound are activated.

Conclusions

What is our contribution to this? We developed a new, purely auditory, task for rats that shares some features with the cocktail party problem. Because our task differs from the Posner task, we refrain from claiming that it requires selective attention. Instead, we propose the term *stimulus selection* for our task. We discuss in detail later the formal psychological relationship between our task and other related tasks.

Instead of attempting to measure receptive fields at all, we focused exclusively on measuring the changes in neural response to our particular task stimuli. Finally, in order to investigate the role of brain regions outside of sensory cortex, we also recorded from prefrontal cortex.

I next discuss in greater detail our experimental paradigm and results.

Chapter 2. Rule-encoding neurons in prefrontal and auditory cortex in a novel rodent model of the cocktail party problem[3]

Abstract

Animals can selectively respond to a target sound in the presence of simultaneous distractors, similar to the way in which humans can respond to one person's voice at a 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 pre-stimulus activity in mPFC encoded the selection rule — the sound to which the rat would respond. Moreover, electrically disrupting activity in mPFC significantly impaired performance. Surprisingly, pre-stimulus and stimulus-evoked activity in A1 also encoded the selection rule, a cognitive variable typically considered the domain of prefrontal regions. However, stimulus tuning was not strongly affected. We suggest a model in which activation of a specific network of neurons underlies the selection of an imminent 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 even 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 et al., 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, 2001). Humans use selective attention, the cognitive process of selecting and responding to a single target stimulus amongst simultaneous distractors (Desimone and Duncan, 1995), to solve the cocktail party problem (Ahveninen et al., 2011). Experiments in visual selective attention reveal that prefrontal cortex sends top-down "bias signals" to sensory cortex (Miller and Cohen, 2001; Moore et al., 2003) in order to select the target stimulus and subsequently enhance its neural representation, while suppressing the representation of distractors. Similar mechanisms may be at work in the auditory cortex: electrocorticographic (Mesgarani and Chang, 2013; Zion Golumbic et al., 2013) and magnetoencephalographic (Ding and Simon, 2012) recordings show that brain activity is dominated by the attended voice. Without recordings from single neurons it is difficult to ascertain what changes on the single-neuron level give rise to these effects.

Towards this goal, we have developed a new behavioral task for rats with three key properties. First, on each behavioral trial the subject hears a pair of simultaneous sounds, each drawn from a different category (e.g., white noise bursts vs. warbles). Second, the experimenter can indicate which sound the subject should select in order to receive a reward. Third, the subject

[3] This chapter was co-written with my advisor, Dr Michael DeWeese, and will be submitted for publication separately.

then selects and responds to the correct sound from the pair. This means that the subject must be capable of selecting either of the two sounds upon demand, and in fact must be able to switch multiple times during a single behavioral session. This task requires cognitive flexibility because the same pair of stimuli demands a different behavioral response ("same stimulus; different response") depending on which sound we require the subject to select. We are aware of no purely auditory single-unit studies in any animal satisfying these three conditions. 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). An alternative name for this type of task is stimulus feature selection, since two simultaneously presented sounds may be perceived as a single sound with two features. 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). Our behavioral paradigm shares attributes with all of these, but for consistency we will refer to our task as stimulus selection below.

Although monkeys are the traditional model organism of choice for complex cognition (Gold and Shadlen, 2007), rodents are capable of sophisticated decision-making, in some ways very similar to humans (Raposo et al., 2012; Brunton et al., 2013; Zariwala et al., 2013). Rodents also show behavioral flexibility under the control of the prefrontal cortex (Karlsson et al., 2012; Kvitsiani et al., 2013; Young and Shapiro, 2011), even though this region is not necessary for simple sensory discriminations (Pai et al., 2011). The mPFC in particular appears to be critical for task switching (Birrell and Brown, 2000; Floresco et al., 2008; Durstewitz et al., 2010; Ragozzino et al., 1999). For example, when rats learn to switch the navigational strategy they use to solve a maze, the mPFC encodes this switch and inactivating the area severely disrupts performance (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 and Everling, 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 an important role in directing flexible auditory processing in A1 (Fritz et al., 2010). In that paradigm, ferrets were trained to detect tones at a specific target frequency, which resulted in rapid task-related plasticity (tuning changes) in A1 and increased functional connectivity with frontal areas (Fritz et al., 2010; Fritz et al., 2003). Moreover, the ferrets could rapidly switch between different auditory tasks and the character of the observed tuning changes matched the demands of each task (Fritz et al., 2005, Fritz et al., 2010). These experiments shed light on the mechanisms of attending to acoustic frequency and revealed A1 to be surprisingly dynamic for a primary sensory area, but the behaviors used

were not stimulus selection according to our criteria. In our task, stimuli are always presented simultaneously rather than sequentially, and a stimulus used as a distractor on one trial can be the target on subsequent trials.

We are unaware of any single-unit studies of purely auditory stimulus selection in any animal. A model of this ability in rodents would be especially useful because of the relative ease and speed with which they can be trained on cognitively demanding tasks (Carandini and Churchland, 2013) and as a first step toward more complex behaviors, such as selective attention, which has traditionally been studied only in primates. Finally, many models of visual selection are not obviously applicable to the auditory modality — for instance the idea that visual attention co-opted the neural mechanisms for shifting gaze over evolutionary time. Establishing an auditory selective attention paradigm could shed more light on whether the known mechanisms of visual selection are universal or specific to one modality.

Toward these goals, we recorded from individual neurons in both mPFC and the primary auditory cortex (A1) of rats performing our task. We found that the pre-stimulus, anticipatory activity of our recorded neurons in mPFC encoded which sound would be selected. Surprisingly, we also found this pre-stimulus effect in a sizable fraction of our recorded neurons in A1. Finally, stimulus-evoked activity in both brain regions was similarly modulated, although 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 for at least 250 ms — the "hold period." This triggered speakers on the left and right to play in stereo 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 a broadband noise burst from either the LEFT or RIGHT speaker, and either a HIGH- or LOW-pitched warble (frequency-modulated tone). 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 (Methods).

On each trial, one of the sounds in the stimulus pair (the "target") indicated the correct response but the other sound (the "distractor") was uninformative. To indicate which sound the rat should select, 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). Each block consisted of 80 trials (Figure 1C), 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. Behavioral controls (Figures S2B, S2C) showed that the rats responded to the target sound, not to the target/distractor combination.

We refer to this task as auditory stimulus selection, by which we mean that the rat selectively responds to one of two simultaneous sounds on any given trial. Importantly, a stimulus selection task requires the subject to be able to select either of the two sounds, depending on which one the experimenter designates as predicting reward. This designation could be accomplished with cues presented before the start of each trial, but in our task we use a block design with no explicit pre-trial cue, so the rat must use its recent reward history to determine which sound it should select.

Stimulus selection — selectively responding to a behaviorally relevant target in the presence of distractors — is one component of selective attention, a broader and more complex ability that also includes perceptual enhancement (Knudsen, 2007; Reynolds and Chelazzi, 2004; Pestilli et al., 2011). We feel that tasks requiring sustained tracking (Mitchell et al., 2007) or the enhanced detection of faint stimuli (Cohen and Maunsell, 2009) are the gold standard of selective attention research. Nonetheless, our task represents an important step forward; we are aware of no other paradigms to study stimulus selection in rodents, nor any single-unit studies of purely auditory stimulus selection in any animal.

Rats perform the task well above chance

We ensured that the rats were in fact selecting the correct target sound by verifying that their behavioral response was driven by the target sound, significantly above chance and significantly more than it was driven by the distractor sound, and also that they were not using the same stimulus/response strategy in both blocks. Some strategies allow 50% performance without using any information from the target, such as always going to the choice port for the current block even in response to a nogo target, a strategy that we commonly observed in rats before they were fully trained. For this reason we verified that performance was significantly and consistently greater than 50% in both blocks, and also that the animals were responding to the target sound and not the distractor (Methods), before and after implanting the recording electrodes. This typically required about 40 one-hour training sessions, for up to eight weeks. Our best rats’ typical performance during recording sessions was approximately 85% in both blocks (Figure 2). In general, the rats performed well above chance, rapidly and correctly changing which sound they selected after each block change.

Because our task associates a different choice/reward port with each block, a different distribution of motor responses is required in localization (50% go left, 50% nogo) than in pitch discrimination (50% go right, 50% nogo). We note two consequences of this. First, this allows us to identify an interesting type of error trial on which the rat appeared to respond to the wrong sound. On such “interference” trials, the rat heard a “go” distractor (i.e., a sound to which the rat should respond with a GO response in the other block) and incorrectly went to the choice port associated with that distractor, instead of doing what the target sound indicated. We later

analyze the neural correlates of this error. Second, it is plausible that the rat's motor plan differs between the blocks. There is a similarity in this sense between our task and some blocked visual spatial attention tasks, in which 80% of the trials require a saccade in the same direction (Cohen and Maunsell 2009). It can be difficult to tease apart response selection from stimulus selection (but see Erlich et al., 2011; Sato and Schall, 2003; Steinmetz and Moore 2012). We return to this issue later.

Anticipatory neuronal activity in mPFC encodes the selection rule

We next asked what differences in neuronal activity between blocks correlated with the selection of the target. We implanted tetrodes into A1 and/or mPFC and recorded single-unit action potentials (spikes) from multiple neurons during behavior. By analogy with the rule-encoding neurons in primate prefrontal cortex, we hypothesized that mPFC would encode the selection rule. That is, we expected that the firing rates of single mPFC neurons would differ significantly between localization and pitch discrimination trials. We first confined our analysis to the hold period, the interval before stimulus onset while 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. An example unit (Figure 3A) fired significantly more ($p < 0.001$, Mann-Whitney U-test) in the hold period during localization trials (mean: 12.1 Hz) than it did during pitch discrimination trials (mean 7.2 Hz). A different but simultaneously recorded single unit in mPFC (Figure 3B) fired significantly more during pitch discrimination (mean 5.4 Hz) than during localization (mean 2.7 Hz). In both cases the effect persisted across the entire session of over 1300 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 more during) localization and 40 preferred pitch discrimination; neither preference was significantly more common (binomial test, $p > 0.05$).

Anticipatory neuronal activity in A1 also encodes the selection rule

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 in the pre-stimulus 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 pitch discrimination; binomial test, $p > 0.05$). Since A1 is known to encode many types of sounds in a sparse fashion (DeWeese et al., 2003; Hromádka et al., 2008; Carlson et al., 2012), we were not surprised to observe that only some of our recorded neurons in A1 significantly responded to our task stimuli (Supp. Info.). However, rule encoding was approximately equally widespread in both stimulus-responsive (14/49) and non-responsive (22/50) neurons. This finding is reminiscent of human imaging results suggesting that neurons in auditory cortex may carry top-down attention signals even in the absence of stimulus information (Ahveninen et al., 2011).

These effects were strong: among the significantly rule-encoding neurons, the median increase in firing rate during the preferred block was 74.7% in mPFC and 99.7% in A1. We controlled for the possibility that these results in either brain region could be explained by firing rate drift over the course of the session or by spike sorting errors arising from small differences in spike waveform shape between blocks (Supp. Info.). We did not observe clustering or any other topographic organization of neurons preferring the same block, which implies that these effects could have been obscured in multi-unit recordings. In sum, these results demonstrate widespread and robust encoding of selection rule in the pre-stimulus activity of both mPFC and A1 neurons.

Motor preparation

The mPFC regulates cognitive state, but it also plays a role in motor planning (Erlich et al., 2011). A classic result (Euston and McNaughton, 2006) showed that PFC neural activity, apparently related to working memory, could actually be well-explained solely in terms of behavior variability. We analyzed video of our subjects and found evidence of preparatory changes in head angle that differed between blocks (Supp. Info.). However, unlike the results of McNaughton and colleagues, we found that the block rule explained more of the variability than head angle did in the vast majority of rule-encoding neurons (Figure S3G-J, S4G-J).

We propose an alternative hypothesis: instead of behavioral context driving postural changes that in turn drive mPFC activity, it could be that behavioral context drives both mPFC activity and adaptive postural changes. Our analysis lends credence to this hypothesis. However, the potentially confounding role of motor planning remains an important consideration for interpreting recordings from prefrontal and even, given our results in A1, sensory cortex.

Error trial analysis

In the previous sections we considered only correct trials. We next considered interference trials, during which the rat erroneously chose the port associated with the other block, suggesting that it was selecting the wrong sound from the mixture. If encoding of selection rule in the anticipatory activity is important for successful stimulus selection, then the encoding should be weaker or even reversed when the rat selected the wrong sound.

In mPFC, the encoding of selection rule was significantly weakened on interference trials, versus correct trials (Figure 3D). In A1, we observed a more extreme effect (Figure 4D): the rule encoding was actually reversed on interference trials, meaning that firing rates were greater during the non-preferred block on such trials. These observations are consistent with the idea that anticipatory activity predicts which sound the rat will respond to, even for trials on which the rat appears to respond to the distractor by going to the wrong choice port. Although the activity thus predicts a motor response to the block-irrelevant port, it does not differ between trials where the rat ultimately goes to the block-relevant port (correctly or incorrectly) or chooses the nogo response (Figures S3C, S4C).

Within-trial timescale of the encoding of the selection rule

We next asked how long before the stimulus the encoding emerged, and for how long afterwards it persisted. For each rule-encoding neuron, we compared across blocks the smoothed firing rates in every 50 ms bin before and after the stimulus onset, up to plus or minus 3 seconds from the stimulus onset. We thereby determined the largest interval of time around the hold period during which the neural activity significantly encoded the selection rule. Across the dataset, the median inter-trial interval was 4.0 s (inter-quartile range: 2.7 s to 5.3 s) and so this time range (plus or minus 3 s) will overlap 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, B). For some neurons, rule encoding was strictly confined to the hold period: their firing rate was modulated only in the immediate pre-stimulus period. Other neurons showed significant encoding at all time bins tested: their firing rate was persistently elevated throughout 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 pre-stimulus (IQR: 0.15 to 1.2 s); in PFC the median was 0.625 s pre-stimulus (IQR: 0.34 to 1.0 s). That is, the majority of rule-encoding neurons developed this property well before the rat initiated a trial by center-poking.

To examine the typical dynamics within each population and to determine which brain area first encodes the selection rule, we averaged the normalized activity (mean: 0, variance: 1) of all rule-encoding neurons in both brain regions 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). The activity in mPFC was first significantly elevated 2.4 seconds before stimulus onset, while population activity in A1 became elevated 0.78 seconds before stimulus onset. That the effect occurs first in mPFC is consistent with its hypothesized role as the origin of top-down bias signals to sensory cortex (Miller and Cohen, 2001); however, we emphasize that the wide variability in timescale within both regions, and the fact that only a minority of our dataset was collected in simultaneous A1/PFC recordings, complicates a direct comparison between brain regions.

Encoding of behavioral choice

We found a prominent difference between the firing rates on GO and NOGO trials, beginning around the time the rat left the center port and continuing for several seconds (Figure 5D). During each rule-encoding unit's preferred block, its firing rate remained elevated on NOGO trials for several seconds, during which time the rat was generally beginning the next trial. In contrast, on GO trials, the unit's firing rate rapidly fell, and in fact remained below its long-term mean firing rate 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. In this model, once the animal perceives the GO stimulus, the rule-encoding

disappears and the GO response is produced. Another interpretation is that the rule-encoding is persistent on NOGO trials because the animal is already preparing to begin the next trial less than a second later, whereas on GO trials the animal no longer needs to encode the rule because it is simply moving to the reward port to consume water.

The data do argue against one particular model. The observed anticipatory effect are unlikely to be a pre-emptive encoding of motion to the block-relevant port (*e.g.*, go left during the localization block) because the firing rates differ greatly pre- and post-stimulus on GO trials. It could be the case that the encoding of motor plan is inversely correlated with encoding of the subsequent motor act, but this is inconsistent with the traditional notion of motor plan.

Changes in baseline activity correlate with similar changes in evoked activity

Given that the pre-stimulus 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 of time after stimulus onset during which the firing rate was significantly elevated above the pre-stimulus rate (Supp. Info.). 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 (Figure 6A, B). Such neurons were rarer in PFC than in A1, though not significantly so (PFC: 31/90, A1: 49/99; $p > 0.05$, Fisher's exact test). Evoked responses were significantly weaker in PFC (Figure S6).

Based on the previous results, in which we found that the increased firing rate during the preferred block often persisted for a period of time after stimulus onset, we expected that the evoked firing rate would also be higher during the preferred block. In both regions, this is indeed the case: an increase in pre-stimulus firing rate during one block correlates with an approximately equal increase in evoked firing rate during the same block (Figure 6C, D; exemplar: Figure 4B). 21% (9/43) of A1 neurons and 24% (4/17) of PFC neurons showed a significant elevation of evoked response during their preferred block. We note that this analysis had less power than the pre-stimulus analysis because of the additional variability introduced by the stimulus dependence (Supp. Info.).

We next used an ideal decoder analysis (Methods, Figure 6E) to ask whether the recorded neurons encoded the identity of the noise burst or warble with greater fidelity in either block, either due to changes in stimulus tuning, the baseline elevation described above, or some other effect. We can decode the identity of both the noise burst and the warble from the evoked responses in A1 (n=57 neurons in 22 simultaneous ensembles) and PFC (n=25 neurons, 13 ensembles). The A1 neurons provide a significantly better source of data from which to decode the sound identity, probably due to their stronger responses and tighter stimulus selectivity. However, for both brain regions and both sounds, we cannot decode the sound any more accurately from the localization trials than from the pitch discrimination trials. Other versions of

this analysis yielded generally similar results, even when we separately considered neurons preferring each block (Supp. Info., Figure S6F).

We conclude that, although neurons showing an increased pre-stimulus firing rate in one block generally showed an equivalent increase in the evoked rate during the same block, these changes in evoked rate do not obviously improve the detectability of the target sound. However, we note that our ensembles of neurons provided useful information about the identity of both sounds, and the brain has access to a pool of neurons orders of magnitude larger than our recorded population. It may be that the problem faced by cortex in this task is not to maximize the information available about the stimuli in individual neurons, but rather a wiring problem of how to flexibly re-route the relevant stimulus information to the relevant 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 the mPFC of three trained animals (rats Z1, Z2, and Z3) with extracellular stimulating electrodes. On 20% of trials ("zap" trials), we injected a 10Hz train of current pulses during both the hold period and the duration of the auditory stimulus. Such electrical stimulation drives an extremely rapid activation of nearby neurons, followed by a slower suppression of firing rates (Logothetis et al., 2010) for a few hundred milliseconds. Thus, this approach neither "silences" nor "activates" the brain region, but rather disrupts the normal firing rates and patterns.

Typically, pharmacological agents such as muscimol are used to test whether neural activity is necessary for a certain behavior. Electrical disruption allows much finer control over the strength and timing with which we can perturb the circuit. We were able to stimulate during a desired subset of trials, for a certain time range (i.e., throughout the center-poke hold and auditory stimulus presentation), while sparing activity the rest of the time. This provided us with a statistically powerful within-session control.

Across all three animals, electrical disruption tended to impair performance (Figure 7) in both localization (mean impairment 5.4% in Z1, 12.4% in Z2, and 27.5% in Z3) and pitch discrimination (19.1% in Z1, 18.7% in Z2, 13.0% in Z3). This impairment was significant across sessions ($p < 0.05$, binomial test) for pitch discrimination in 3/3 rats (Z1, Z2, and Z3) and for localization in 2/3 rats (Z2 and Z3). Electrical disruption largely, though not exclusively, affected performance on NOGO trials. All rats were impaired on pitch discrimination NOGO trials in almost all sessions (Z1: 6/6, Z2: 8/8, Z3: 7/8 sessions). Some rats also exhibited additional impairments: Z3 was impaired on localization NOGO (8/8 sessions) and Z2 was impaired on localization GO trials (Z2, 8/8 sessions). These effects were generally quite strong within individual sessions (Figure 7B) even though they varied between rats. Taken together, these data suggest that, in the absence of normal mPFC 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 stimulus selection in particular was impaired, as opposed to impulse control, or some other aspect of the task (Supp. Info.).

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, which we have elaborated into a quantitative simulation as a proof of principle. The model 1) requires only random stimulus tuning in A1; 2) does not require tuning changes or synaptic reweighting after the initial training phase; 3) uses only excitatory connections, consistent with the observation that most long-range projections in the brain are excitatory.

The model (Figure 8) consists of a population of N neurons in A1, randomly tuned for each of our four stimulus pairs. The activation of each A1 neuron was subject to additive Gaussian noise. The “sensory SNR”, defined as the ratio of the strength of this noise to the strength of the stimulus tuning, is a free parameter. Half of the neurons are arbitrarily assigned to each of the two tasks. Each subpopulation projects to two command neurons encoding the two possible behavioral responses during that block (e.g., go left and nogo). Each projection is trained to activate the correct command neuron using a least-squares fit constrained to use only positive (excitatory) weights. 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 a new set of test stimuli are presented. To produce the block-appropriate response, a “task signal” is added to the activations of the neurons in the appropriate A1 subpopulation for the current block, as indicated by either the red or blue neurons in Figure 8. Because all feed-forward weights are positive, adding this task signal translates into an excitatory boost to the premotor neurons appropriate for that block. Thus, even without any synaptic reweighting, the model will tend 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.

This demonstrates that anticipatory modulation can be part of a scheme that is capable of solving the task switching problem, even with weak sensory responses and “random but fixed” tuning in A1, in which each subpopulation consists of randomly tuned neurons that do not change their tuning between blocks.

Discussion

Auditory stimulus selection: task switching between conflicting auditory discriminations

When human listeners hear two simultaneous voices they can selectively respond to either one. This is a complex ability, and our task models one 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. The rats learned the task with less than eight weeks of training and performed many trials per session (median: 698; inter-quartile range: 507 to 912). To our knowledge this is the first published example of rodents performing such a stimulus selection task in any sensory modality.

Previous studies have identified critical roles for mPFC in behavioral flexibility in several contexts. For example, elegant work (Rich and Shapiro, 2009) established not only that the mPFC encodes the switches in navigational strategies ("go east" vs "turn right") that rats use to solve a maze, but also that inactivating this region impairs severely and selectively their ability to perform the switch. Other studies of task switching in rodents required them to switch between a sensory discrimination and a (potentially habitual) fixed response ("follow the light" vs "always go left"; Floresco et al., 2008; Durstewitz et al., 2010). Many researchers are interested in extending these results to task switching between sensory discriminations, but it is often challenging to induce the switch when it requires ignoring a previously trained stimulus. Even in cross-modal switching, where 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, while our study requires multiple switches per session. We believe our task advances the study of task switching in rodents to be much closer to the standard set by human and non-human primate studies.

Despite its clinical and computational relevance (Ding and Simon, 2012), the auditory cocktail party problem remains less studied than comparable visual tasks. Even in primates we are not aware of any single-unit studies of purely auditory stimulus selection. A multi-unit study (Lakatos et al., 2013) required monkeys to sustain attention to streams of pure tones; however, the researchers found that the monkeys were unable to ignore the distractor stream if it was within 1.5 octaves of the target stream. Human voices, even those with very different pitch, are much closer than this and actually overlap extensively in acoustic frequency (McDermott, 2009). For this reason, we believe animal models of this ability should use stimuli that, like ours, overlap at least partially in frequency and require solutions not based purely on frequency separation. 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., the use of viral vectors expressing light-gated ion channels in specific brain regions or genetically-identified cell types).

Anticipatory activity in both mPFC and A1 encodes the selection rule

We found that rodent mPFC robustly encode the subject's selection rule, analogous to the rule-encoding role of primate prefrontal cortex (Asaad et al., 2000; Wallis et al., 2001; Johnston and Everling, 2007). Rule encoding develops in our recorded mPFC population over 2.5 seconds before the stimulus onset, as the rat is planning to initiate a trial or even finishing the previous trial. The widespread nature of the encoding and the broad timescales over which it persists 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. One possibility is that this persistent activity represents a memory trace of the selection rule (Funahashi et al., 1989), meaning that it densely and persistently encodes cognitive variables like selection rule. In fact, the cortex may shift to a completely different network state (Karlsson et al., 2012) depending on which stimulus the rat plans to select.

We also observed anticipatory encoding of the selection rule in primary auditory cortex (A1), a surprising result since encoding of selection rule in the absence of sensory stimulation has traditionally been considered the domain of prefrontal areas. However, attention is known to modulate the pre-stimulus activity of single neurons in monkey V2 and V4, although not in V1 (Luck et al., 1997; Reynolds et al., 2000). At a larger spatial scale, visual attention can produce a similar increase in pre-stimulus baseline in V1, as assessed both with fMRI in humans (Pestilli et al., 2011) and with voltage sensitive dye in monkeys (Chen and Seidemann, 2011). Higher visual cortex also shows pre-stimulus modulation by attention in humans (Pestilli et al., 2011; Kastner et al., 1999; Thut et al. 2006). More generally, single neuron activity in primary sensory cortex can anticipate reward (Shuler and Bear, 2006) or a motor response (Niwa et al., 2012), and anticipation of a visual stimulus can trigger a hemodynamic response in V1, though without a corresponding change in neural activity (Sirotin and Das, 2009). In this light, perhaps it is not surprising that primary sensory cortex could also encode the selection rule for an imminent stimulus. In this way both the information about the stimulus and the information about how that stimulus should be interpreted are encoded in the same neurons, providing a possible locus for the behavioral decision to be made.

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 (Luck et al., 1997). In contrast, our results show that rat A1 already robustly encodes a non-sensory variable, very similar to mPFC. This could be a difference between rats and monkeys, or between auditory and visual cortex, or both. Disambiguating these possibilities will be an important direction for future work.

Comparison with studies of selective attention and task-relevant plasticity

This pre-stimulus change in baseline contributed in an additive way to the strength of the sensory-evoked responses in both A1 and PFC; however, we found limited evidence for any additional modulation of sensory-evoked responses in A1. For example, the neurons did not appear to encode the target stimulus with any greater fidelity than the distractor stimulus. This

is consistent with some, but not all, previous studies of auditory task switching. Although neuronal activity in A1 is robustly modulated in the aroused/engaged behavioral state versus the passive/idle state (Otazu et al., 2009; Lee and Middlebrooks, 2011), the neuronal effects of shifting between different engaged behaviors tend to be weaker or even non-existent. For instance, switching between an auditory task and an olfactory or visual task does not change evoked spiking auditory responses in A1 (Otazu et al., 2009; Lakatos et al., 2009), and switching between temporal and spatial auditory discriminations does not significantly change spatial tuning in A1 (Lee and Middlebrooks, 2011). Nonetheless, the fact that these studies (and ours) found no evidence of tuning changes in A1 does not mean that they do not exist under some circumstances.

In fact, a series of pioneering experiments demonstrated task-relevant plasticity in A1 of ferrets trained to detect a target frequency (Fritz et al., 2003; Fritz et al., 2010). One important methodological difference is that their study, unlike ours, made use of a large battery of probe stimuli and was therefore optimized to detect receptive field changes, including those affecting only task-irrelevant stimuli. Intriguingly, this plasticity was nuanced: it could induce facilitation or, alternatively, significant suppression at the task-relevant frequency. Facilitation was more common than suppression, but the use of a different reinforcement paradigm reversed this (David et al., 2012). More studies of complex auditory behaviors will be necessary to better understand the factors that determine whether a given behavioral paradigm produces task-related modulation of evoked spiking responses in auditory cortex.

The lack of evidence for tuning modulation in our data is a surprising result, given that visual selective attention enhances target representations and suppresses 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 (Knudsen, 2007; Reynolds and Chelazzi, 2004; Pestilli et al., 2011). Target-enhancing modulation of evoked responses is believed to mediate perceptual enhancement (although see Zénon and Krauzlis, 2012), as assessed behaviorally by a lower threshold or steeper psychophysical curves (Cohen and Maunsell, 2009; Moore et al., 2003). This predicts that only tasks that require perceptual enhancement will produce such effects.

In contrast, stimulus selection is often investigated with easily detectable stimuli far above threshold (Hocherman et al., 1976; Stoet and Snyder, 2004) and such studies, like ours, often find no modulation of evoked responses in sensory cortex (Sasaki and Uka, 2009; Mante et al., 2013). It may be that stimulus selection is the dominant computational challenge in such tasks and perceptual enhancement is therefore less important. Similarly, the cocktail party problem is often difficult because all voices are of competing intensity, not because the target voice is barely audible. Additive, pre-stimulus baseline increases have been observed in V1 during attention-demanding tasks and may lead to efficient stimulus selection (Chen and Seidemann 2012; Pestilli et al., 2011); our data support a similar hypothesis in A1. In summary, while the mechanisms by which selective attention mediates perceptual enhancement remain an

important area of inquiry, enhancement of sensory evoked responses in A1 may not be necessary for our task or other similar stimulus selection paradigms.

Stimulus selection via activation of latent circuits for each target

Based on our results, we propose a model for stimulus selection based on task-specific activation of latent circuits, rather than task-specific adaptation of a single circuit. We found subpopulations of neurons in both A1 and mPFC — one activated during the localization block, the other during the pitch discrimination block. The signature of this activation is increased baseline activity. However, they do not change their tuning for specific stimuli. We hypothesize that the difference between the circuits is their downstream connectivity: each circuit may project to separate circuits in a downstream effector region, perhaps the striatum since the corticostriatal projection plays an important role in auditory decisions (Znamenskiy and Zador, 2013). In this model, only one circuit is activated at a time, via feedforward excitation perhaps originating in mPFC, and only this circuit has sufficient baseline activity to drive behavior.

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. Second, neurons in A1 and in striatum showing the same block preference should be more strongly connected than those showing the opposite block preference. Finally, specific activation of one of the subpopulations in mPFC, A1, or striatum should specifically bias behavior toward the block preferred by that subpopulation. However, such a manipulation would require a means of stimulating only those neurons that can be functionally identified by their anticipatory firing rate, perhaps by expressing light-gated ion channels in the appropriate populations. Possibly activity-dependent promoters such as cFos or other immediate early genes could be used to this effect.

In some ways, this model is more parsimonious than the traditional tuning change model of auditory attention, 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 the fastest that they can be estimated from those data), it is unclear how known synaptic mechanisms could mediate task-specific tuning changes on a sub-second timescale. By contrast, our model requires only circuits with essentially fixed stimulus tuning, and the selection mechanism occurs by activating one of these circuits, rather than by changing the tuning of any of the neurons. This reflects the challenge of the task, which does not require amplifying the neural representation of a faint stimulus but rather a discrete change in sensorimotor mapping.

Might the pre-stimulus effects we describe be due to a difference in motor plan? Because each block is associated with a different choice port, it is plausible that the rat plans a different response in each block (go left versus go right). Rodent frontal (Erlich et al., 2011) and primate primary auditory (Niwa et al.; 2012) cortex have both been reported to encode imminent motor actions. In those studies the neural encoding of motor plan was observed in a delay period

between stimulus and response, and it correlated with the subject's response. Our neural effects differ in a few ways. First, motor plans are usually characterized in the context of a task with a delay period between the stimulus and the "go" cue (Erlich et al., 2011). During this delay period there is no uncertainty about the correct response. In our task, we report on effects that precede the stimulus, and the rat ultimately chooses to go or not go with roughly equal probability. Moreover, the neural effects are the same regardless of whether the rat produces a GO or NOGO response (Figure S3C, S4C). This is inconsistent with the most straightforward meaning of motor plan. It is more similar to what has been called a "countermanding" or "stop signal" task (Schall et al., 2000), in which a motor plan is formed but is ultimately discarded on certain catch trials. Lesion studies in rodents demonstrate the involvement of striatum and orbitofrontal cortex in such tasks, though the mPFC does not appear to be necessary (Eagle and Robbins, 2003; Eagle et al., 2008).

Second, the timecourse of the neural effects we observed was quite protracted, even persistent throughout the block, in many of our recorded neurons. The baseline increase precedes the initiation of the trial, and, in some neurons, even overlaps the previous trial. Throughout this period, the rat is engaged in various motor actions, such as reward consumption. Nevertheless, the effects we have observed may still reflect with the rat's motor plan. Importantly, our task requires remapping sensory stimuli to motor responses, and it is reasonable to expect rule encoding to incorporate both the sensory and motor aspects of this remapping.

Taken together, our results are consistent with a distributed processing model in which contextual information from PFC modulates activity in A1 in order to increase the fidelity with which the appropriate motor action can be read out. This idea was proposed in the context of tuning changes (Fritz et al., 2010; David et al., 2012; Blake et al., 2002), but we demonstrate that it could also operate by activating a separate circuit without retuning neurons to task-relevant stimuli. Alternative models based on visual selection (Gilbert and Shallice, 2001; Mante et al., 2013) propose that stimulus selection occurs in frontal areas, not sensory cortex. Our data are similar to theirs in the sense that we do not observe tuning changes in sensory cortex (Mante et al., 2013) but different in the sense that we do not observe strong representations of the stimuli in PFC. We do observe encodings of the motor choice in both areas (Figure 5D). Whether these differences reflect a distinction between auditory and visual processing, or can perhaps be unified, remains a question for future work.

Our 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 have presented what we believe to be the first single-unit results in any animal performing an auditory stimulus selection task and we have found widespread and robust rule encoding in mPFC and A1, though 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.

Methods

Behavior training

We used male Long-Evans rats and began training them when their body mass reached 150g-200g, 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 water rewards. After each session they were given ad lib access to water for one hour. We monitored body weight and other markers to ensure they remained healthy.

We used a typical "shaping" procedure to train the rats. First they learned the localization task and pitch discrimination tasks separately and without a distractor. Next they learned to alternate between the tasks. Finally they learned to respond to the mixed stimulus containing target and distractor based on the block. Human intervention was required to determine when the rats were ready to progress to the next stage of training (generally, at least 80% hit rate). Human intervention was also required to discourage certain unwanted response strategies using the following tools: 1) increasing error timeout; 2) temporarily enforcing "all GO" or "all NOGO" trials (and dropping such trials from analysis); 3) giving water rewards out of the left or right port even in the absence of good performance in order to maintain motivation or encourage a task switch. Once the rats were sufficiently well-trained that little or no human intervention was required, they were implanted with the drive. Some rats required "retraining" after implantation using the techniques listed above; any trials thus affected were discarded from analysis. The entire training process takes about 10 weeks.

Trial timings

In three rats (Rats 1-3) the hold period was drawn from a uniform distribution on 0-100 ms; after pilot results indicated pre-stimulus effects, the hold period duration was increased to 250-350 ms in the other three rats. All trials with a hold period <50 ms were discarded for the analyses in Figure 3 and Figure 4. Hold period response was counted in the minimum window that applied to all trials: 50 ms for the first 3 rats and 250 ms for the rest.

The duration of the choice period differed between sessions, but was fixed within a session (or if it was changed slightly within a session, then the trials before the change were discarded from analysis). Correct entries into the choice port on go trials were rewarded with water from the same port. Incorrect entries into the choice port on nogo trials results in a 2-6 s timeout. Correct nogo responses were not explicitly rewarded with water, although the rat avoided a timeout with this response. Poking neither port on a go trial was not explicitly punished with a timeout, other than a lost opportunity for reward.

Chance performance on the task

In order for the rat to perform significantly above chance within a session, its behavior had to satisfy three criteria: 1) the rat performed significantly above 50% in each block, meaning that it must be using some information from the target sound (which is the only possible source of information on the correct response) to decide whether to go or nogo; 2) the rat is significantly more likely to perform the action indicated by the target than the action indicated by the

distractor; 3) the rat is not using a "fixed strategy", that is, the same mapping from stimulus pair to behavioral response in each block. (Because the target and distractor swap roles in each block, satisfying the second criterion is sufficient to satisfy the third.)

The first criterion rules out strategies like "always go left during localization", which was a common strategy while first learning the task. We used a binomial test to compare the proportion of hits to 0.5 in both blocks and discarded any sessions that were not significantly ($p > 0.05$) above 50% in either block. The second criterion rules out certain hypothetical strategies such as always getting the congruent-nogo stimulus (RIGHT+HIGH) correct, and otherwise guessing randomly between the correct choices for that stimulus pair in each block. This fixed strategy yields 62.5% in both blocks but it uses information equally from both target and distractor; thus, it fails the second criterion. To test this, we used a paired Mann-Whitney U-test to compare whether the action on each trial was correct for the target versus correct for the distractor. In practice, none of the rats actually adopted such a hypothetical strategy: although some sessions failed the first criterion and were discarded ($p > 0.05$ for 4/55 sessions), no sessions failed the second criterion ($p < .005$ for all sessions). Therefore the first criterion (performance above 50% in both blocks) is actually the most relevant, and we mark the chance level on the plot as 50%.

Construction and implantation of tetrode microdrives

We constructed tetrodes by cutting lengths of 12.5 micron nichrome wire coated with partially annealed polyimide insulation (Kanthal Palm Coast), twisting them, and heating with a heat gun until the 4 individual strands melted together. The tetrodes were then routed through polyimide guide tubes and glued to the moveable plastic tab within a potentiometer. We pinned the individual wires using gold pins (Neuralynx) into a custom printed circuit board (custompcb.com, beta-layout.com) that we designed. To reduce the Johnson-Nyquist noise at the electrochemical interface, which scales with the square root of impedance, electrode wires were gold-plated to 0.3 megaohm before implantation. For animals with dual implants, we built separate drives to target A1 and PFC. These were connected with a custom-designed adapter to a 32-channel preamp/headstage (Triangle BioSystems International).

Standard surgical techniques were used (Supp. Info.). Craniotomies were performed directly dorsal to the target areas (A1: 5.25 mm posterior and 6.5 mm left from bregma; prelimbic (PL) region of mPFC: 3.0 mm anterior and 1.0 mm left from bregma) and the tetrodes subsequently lowered downward into the target areas before recording.

For stimulation experiments: we implanted low impedance (about 0.1 megaohm) platinum-iridium stimulating electrodes into bilateral mPFC of three animals. In one rat (Z1), a single pair of electrodes (FHC) was located in the dorsal portion of the prelimbic region in each hemisphere; in the other two rats, an array of three stimulating electrodes (MicroProbes) was placed in each hemisphere to span the anterior-posterior and dorsal-ventral extent of the prelimbic region.

Recording and signal processing

The electrodes were lowered by approximately 100-200 microns before most recording sessions by turning the potentiometer's screw. Before recording, we waited 30 minutes to allow the tetrodes to fully adjust. Broadband data were acquired at 30KHz and digitized and stored using a neural signal processor from Blackrock Microsystems. After the behavior, white noise bursts were presented passively to the animal to detect field and/or multi-unit auditory responses. Strong, low-latency auditory responses indicated that the electrodes were in A1 (in combination with the stereotactic coordinates used during implantation and, when possible, post-mortem histological reconstruction of electrode tracks). We only considered sessions in which we believed the electrodes to be in the correct brain regions.

We filtered the data offline to separate LFP (<200Hz) and spikes (>3 KHz). Butterworth, non-causal (temporally symmetric) filters were used to ensure that no phase distortion occurred. We used a detection threshold of 4.5 sigma (calculated using the more robust median absolute deviation) and a short window of 0.8 ms in order to minimize collisions between detected spikes. We extracted spike waveforms using our own contributions to the open-source OpenElectrophy software suite, reduced the dimensionality with principle component analysis, clustered with KlustaKwik, and manually reclustered as necessary with Klusters (Hazan et al., 2004) while blind to the experimental variables. Single units were identified based on the existence of a refractory period and minimal cluster overlap with other putative single units or noise.

We analyzed the data with Python and the modules numpy, scipy, scikits-learn, rpy2, statsmodels, and pandas, as well as custom-written data analysis code. Except where otherwise noted in the text, we observed consistent results across all subjects and therefore pooled the data.

Decoder analysis

An ideal decoder was trained on the evoked rates, including baseline. We implemented this decoder using the LogisticRegression object in scikits-learn and assessed its performance by the number of trials on which the identity of the noise burst and/or warble was correctly predicted. Figure 6E shows the results on ensembles of simultaneously recorded neurons, but the results (no effect of block) were similar for individual neurons.

Figures

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, in the position shown on the left. After a hold period, an auditory stimulus plays in stereo. 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. The warble is always played with equal intensity from both speakers.

C) Task rules. The session consists of alternating localization and pitch discrimination blocks of 80 trials each. Left: 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. Right: 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 localized noise burst is an irrelevant distractor. Good performance depends on selecting and responding to the target sound, not the distractor sound.

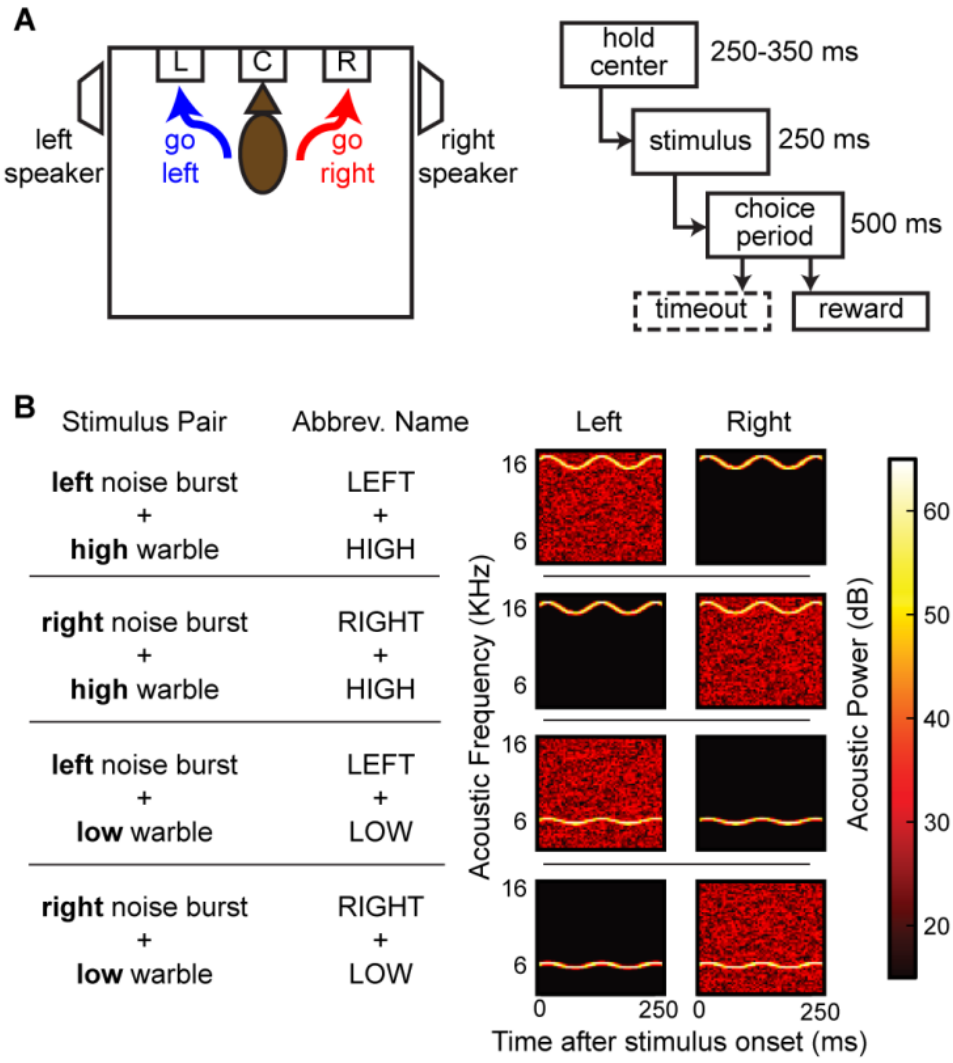
A
L
C
R
left speaker
go left
go right
right speaker
hold center
250-350 ms
stimulus
250 ms
choice period
500 ms
timeout
reward
B
Stimulus Pair
Abbrev. Name
left noise burst + high warble
LEFT + HIGH
right noise burst + high warble
RIGHT + HIGH
left noise burst + low warble
LEFT + LOW
right noise burst + low warble
RIGHT + LOW
Left
Right
Acoustic Frequency (KHz)
16
6
0
250
Time after stimulus onset (ms)
Acoustic Power (dB)
60
50
40
30
20
C
Localization
LEFT + HIGH
RIGHT + HIGH
LEFT + LOW
RIGHT + LOW
Correct Response
go left
nogo
go right
Pitch Discrimination
LEFT + HIGH
RIGHT + HIGH
LEFT + LOW
RIGHT + LOW

FIGURE 1

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) and pitch discrimination (red) in a single recording session. Performance is well above chance (black dotted line, see Methods).

B) Distribution of behavioral responses to an example stimulus pair (RIGHT+LOW) over the course of an average session. We averaged across all sessions from a single rat (CR21A) and binned the trials into groups of 10 to smooth the traces. 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 (see Figure 1C). 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. Cue trials, during which this stimulus pair does not occur, begin each block and are shaded in cyan and pink throughout this figure.

C) Combined performance, similar to (B) but averaged over all sessions, rats, and stimuli. Correct responses (black trace) are the most common outcome. Performance is consistently high throughout, except immediately after a block change. The orange trace shows the probability of an “interference” trial (see text).

D) Analysis of performance immediately after block changes. All localization blocks from (C) are averaged together as are all pitch discrimination blocks. (In order to emphasize block transitions, the x-axis repeats itself after trial 160; the block structure is cyclical and so the cyan shaded areas are identical.) Immediately after the beginning of a new block (cyan and pink areas), performance decreases briefly but recovers within a few trials.

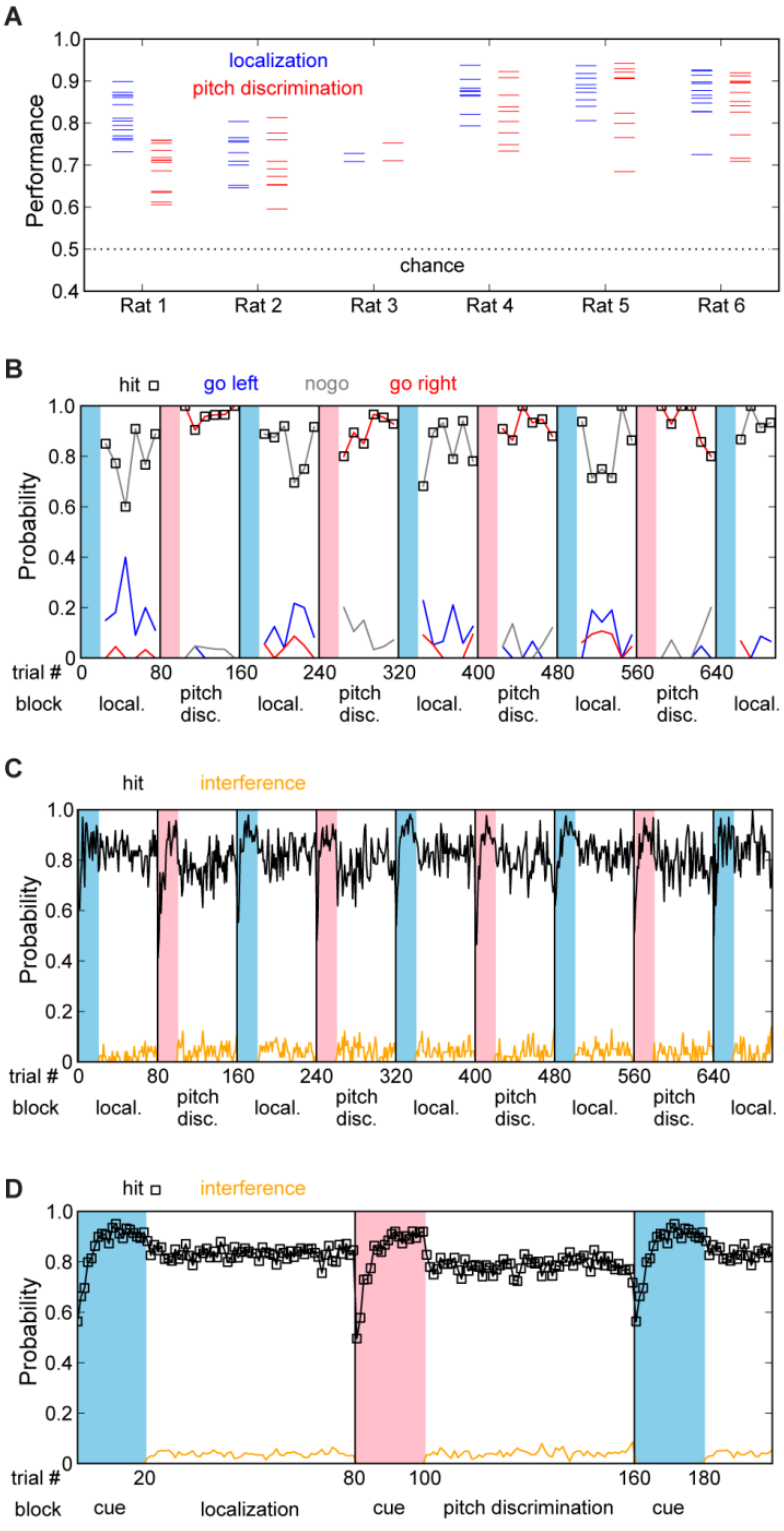
A
localization
pitch discrimination
Performance
chance
Rat 1
Rat 2
Rat 3
Rat 4
Rat 5
Rat 6
B
hit
go left
nogo
go right
Probability
trial #
block
local.
pitch disc.
C
hit
interference
Probability
trial #
block
local.
pitch disc.
D
hit
interference
Probability
trial #
block
cue
localization
pitch discrimination

FIGURE 2

Figure 3. Pre-stimulus 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); error bars SEM. Inset: Extracellular waveforms (mean plus or minus standard deviation) on each channel of the tetrode, duration 0.8 ms. The waveforms are colored red and blue based on the block in which they were recorded, but are almost entirely overlapping (purple). Right: peri-stimulus time histogram (PSTH) of the same unit, averaged over all correct trials from each block. The firing rate is significantly higher ($p < 0.001$) during the hold period (gray shading) for localization (mean 12.1 Hz, n=483 trials) vs. pitch discrimination (mean 7.2 Hz, n=295 trials). We assessed significance for all neurons with the Mann-Whitney U-test and controlled for multiple comparisons with the Benjamini-Hochberg false discovery rate.

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

C) Stacked histogram of the ratio of hold period firing rate (pitch discrimination over localization) for all mPFC neurons. Red and blue bars are significantly modulated neurons.

D) Rule encoding during the hold period is diminished on interference trials. We averaged together the firing rates in the preferred and non-preferred blocks of each rule encoding neuron, after normalizing by subtracting the firing rate on correct trials in the non-preferred block. Error bars: SEM; orange bars: interference trials; white bars: correct trials. The population response on interference trials is significantly decreased during the preferred block and increased during the non-preferred block. Thus, the encoding of selection rule is diminished on trials on which the rat may be selecting the wrong sound. Significance was assessed with a paired Mann-Whitney test (n=57 neurons), which is invariant to the subtractive normalization performed.

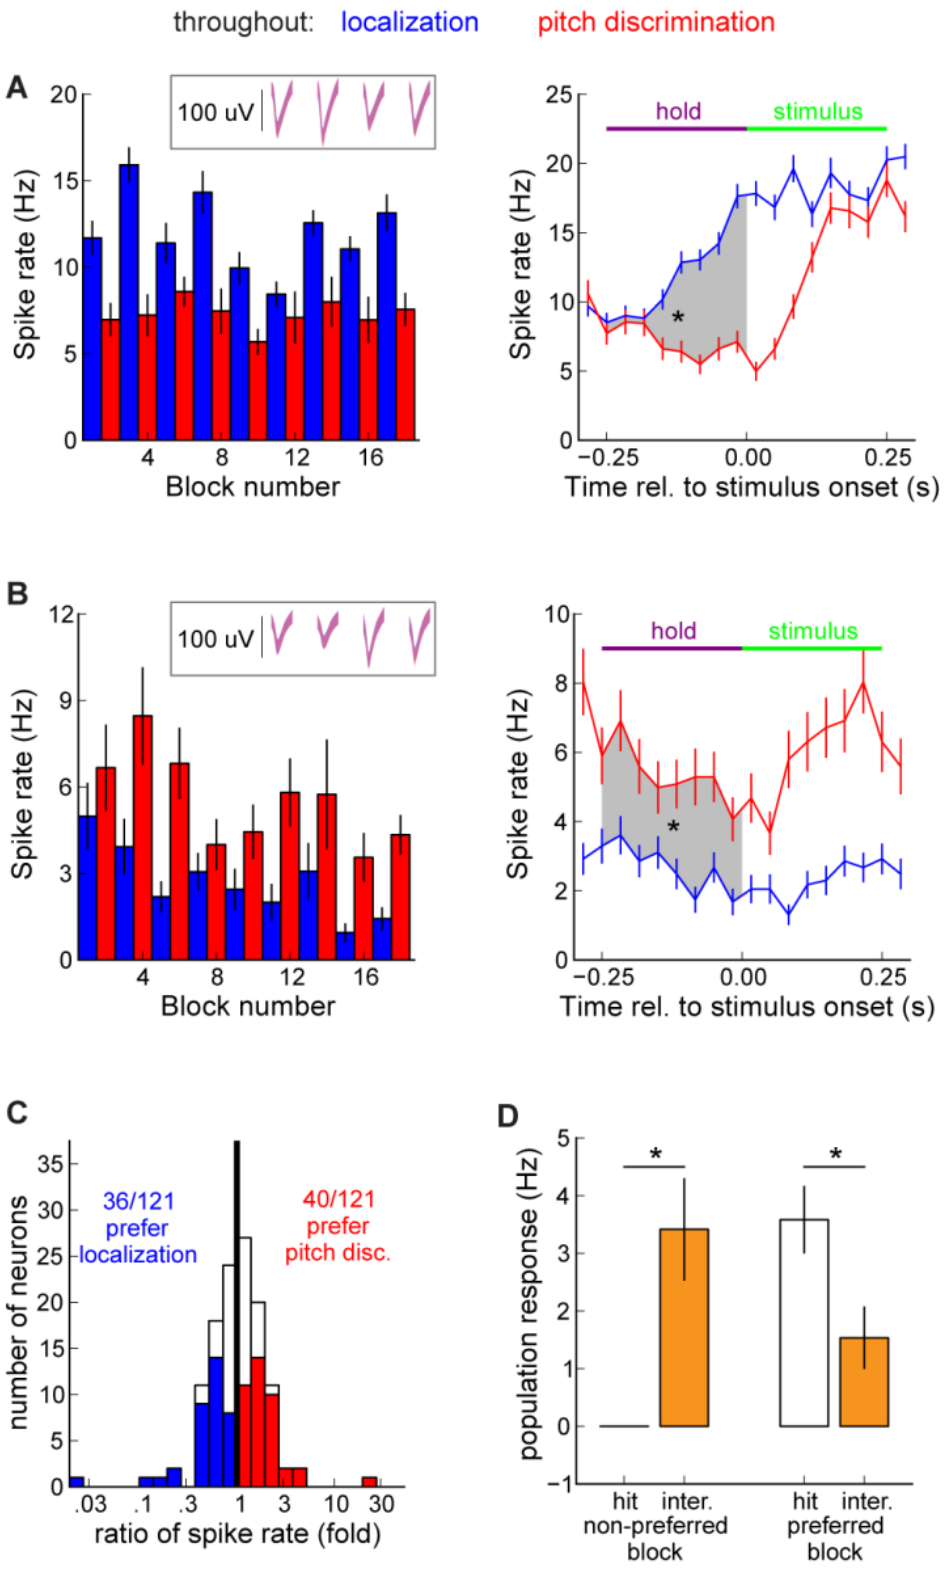
throughout: localization pitch discrimination
A
100 uV
Spike rate (Hz)
Block number
hold
stimulus
Time rel. to stimulus onset (s)
B
C
number of neurons
36/121 prefer localization
40/121 prefer pitch disc.
ratio of spike rate (fold)
D
population response (Hz)
hit inter. non-preferred block
hit inter. preferred block

FIGURE 3

Figure 4. Pre-stimulus activity in A1 also encodes the selection rule.

A) An example neuron recorded in primary auditory cortex (A1). This neuron responds significantly more ($p < 0.001$) during localization (8.0 Hz, n=312 trials; blue throughout this figure) than during pitch discrimination (4.8 Hz, n=253; red). Note the peak following stimulus onset, which was used to analyze the evoked response (Figure 6). Throughout this figure, we use the same conventions and statistical procedures as in Figure 3.

B) Another simultaneously recorded example A1 neuron that encoded the selection rule. This neuron 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 for A1 neurons. Same conventions as Figure 3D, but the effect is stronger here. The population response on interference trials (orange bars) is significantly greater during the non-preferred block than during the preferred block ($p < 0.05$, n=16 neurons, paired Mann-Whitney U-test), opposite to the encoding on correct trials (white bars).

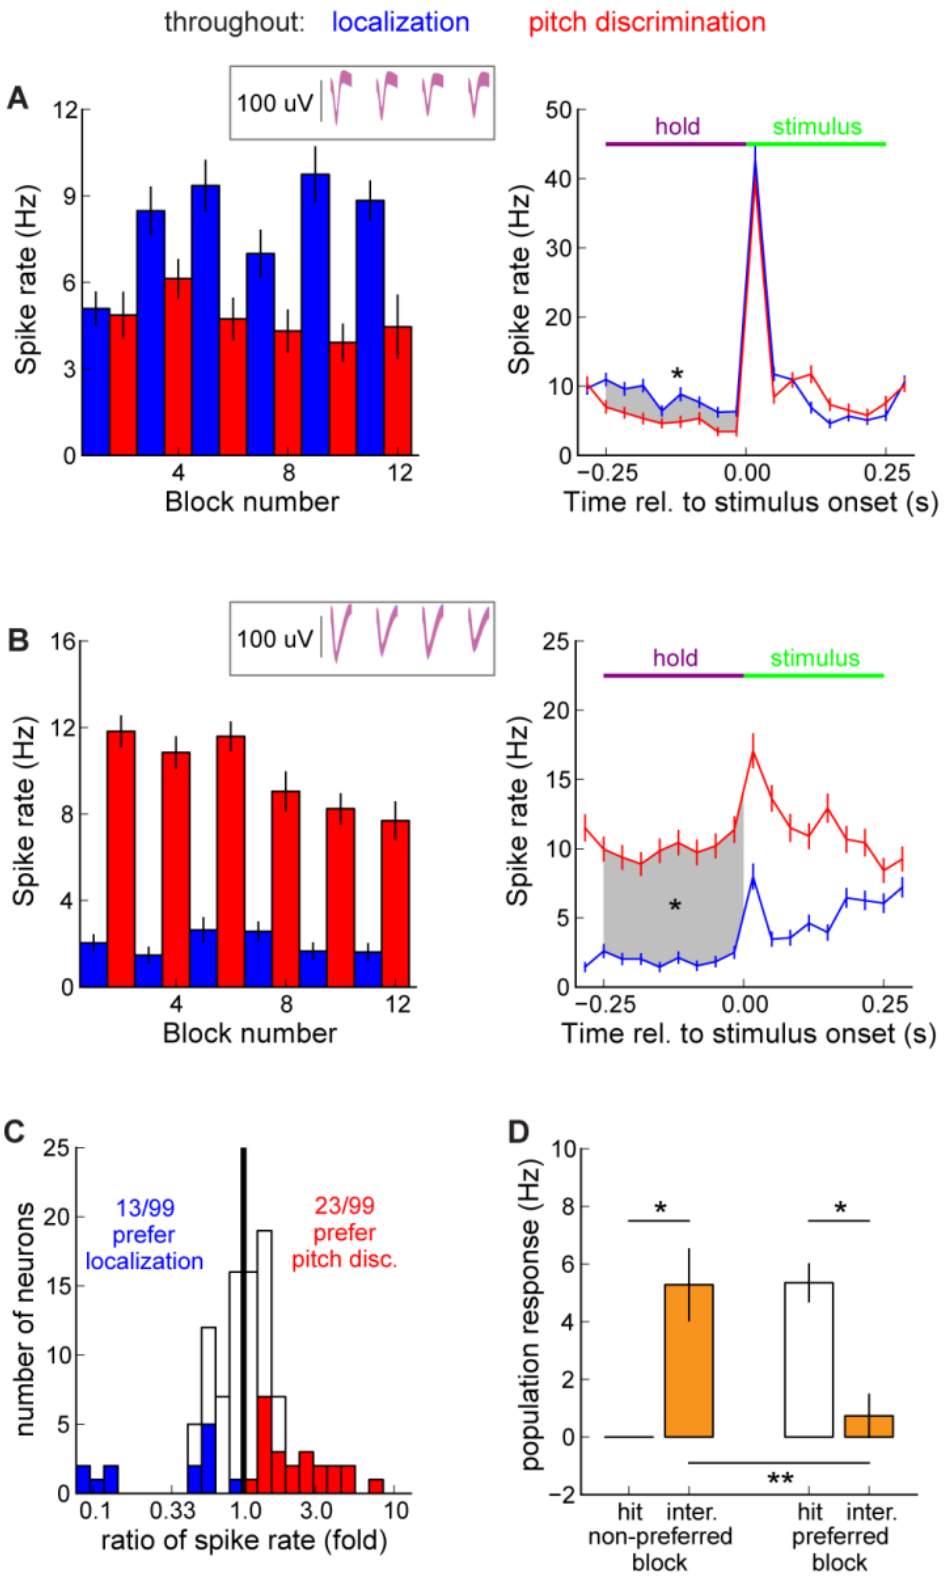
throughout: localization pitch discrimination
A
100 uV
Spike rate (Hz)
Block number
hold
stimulus
Time rel. to stimulus onset (s)
B
100 uV
Spike rate (Hz)
Block number
hold
stimulus
Time rel. to stimulus onset (s)
C
number of neurons
13/99 prefer localization
23/99 prefer pitch disc.
ratio of spike rate (fold)
D
population response (Hz)
hit inter. non-preferred block
hit inter. preferred block

FIGURE 4

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 broader than in previous figures. Firing rates are smoothed with a 50ms Gaussian kernel, normalized to equal variance, and locked to stimulus onset at time 0 ms. The time interval containing the hold period during which the traces significantly diverge is shaded gray. Although these neurons were identified based on a difference in firing rate during the hold period, the traces often diverge for much longer than that. We observed a wide variety of timescales and dynamics in the block-specific anticipatory modulation. The first neuron effectively fires persistently more in one block. The third and fourth neurons demonstrate that the firing rate can either rise during the preferred block, or, less commonly, drop during the non-preferred block. The fifth neuron shows that the anticipatory effect can be limited to just the hold period alone.

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 three seconds preceding the stimulus (first neuron), to very brief activation well under 1 second (last neuron). The third neuron shows increased baseline firing and increased stimulus-evoked firing (peak immediately after time zero) in the same block. This was typical of our dataset (see Figure 6).

C) Population timecourse: the curves represent the average response during the preferred block across all rule-encoding neurons in mPFC (purple) and A1 (orange). The firing rates of all rule-encoding neurons were normalized (mean: 0, variance: 1) and then averaged together. Only the response during the preferred block is shown. Traces are mean response (plus or minus SEM) across neurons. Thick mean trace: timepoints during which the population response significantly exceeds zero, the mean firing rate ($p < 0.05$, one-sample t-test across neurons). In both populations, the firing rate in the preferred block shows a gradual increase, 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 2.4 s before stimulus onset, n=76) than the A1 population (first significantly activated 0.78 s before stimulus onset, n=36), 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. Left: peri-event time histograms (PETHs), locked to the post-stimulus exit from the center-port, from rule-encoding mPFC neurons during their preferred block. Right: same as left panel, but for A1 neurons. Each PETH is aligned to the post-stimulus exit from center-port. Trials are grouped according to correct GO responses (green) and correct NOGO responses (gray). As in (C), PETHs were normalized to unit variance and zero mean before averaging across neurons; only the response during the preferred block is shown. Trace thickness indicates SEM across neurons. On NOGO trials, the firing rate remains 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 falls below baseline and remains there as the rat moves to the choice port and drinks a reward (which

always required at least several seconds). To illustrate this last point, along the lower edge of the figure we also plot the distribution of latencies to the relevant trial events: stimulus onset (black), reward delivery (green, GO trials only), and center-poke beginning the next trial (gray, NOGO trials only; for GO trials the beginning of the next trial would not be visible in this time range due to the time necessary to consume reward).

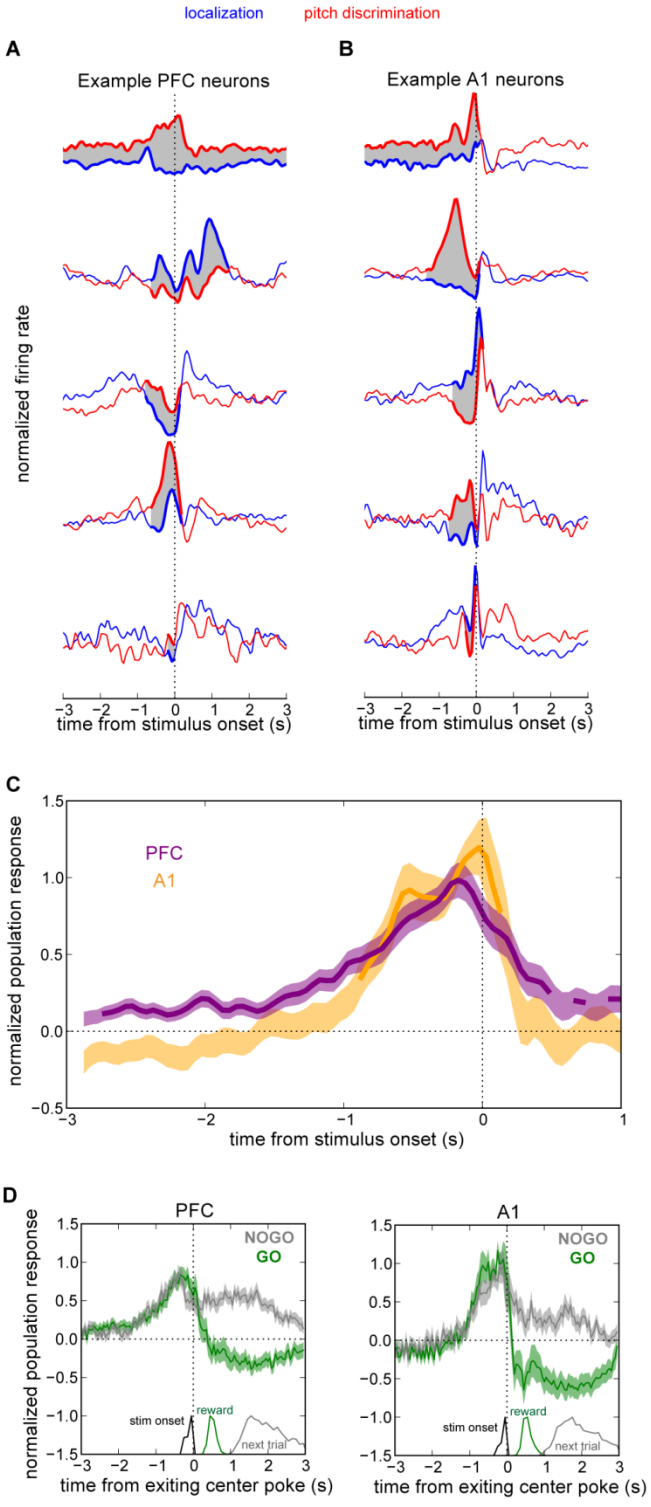

Figure 6. Limited evidence for modulation of stimulus-evoked activity.

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 triangle: stimulus onset; shaded area: response window for this neuron.

B) An example mPFC neuron that responds to the task stimuli with a low-latency response. Auditory responses were weaker in mPFC neurons compared with A1 neurons (Figure S6).

C) For A1 neurons, increase in hold period activity during one block correlates with increased evoked response during that block. For each neuron, the change in evoked response (driven spikes in pitch discrimination vs. localization) is plotted against the change in hold period firing rate (anticipatory spikes in pitch discrimination vs. localization). The trend line (n=43 neurons, r=0.52) has a slope of 0.98, suggesting that much of the modulation of evoked strength is due to anticipatory modulation (example: Figure 4B).

D) Following the conventions of (C), but for auditory-responsive PFC neurons. Again, a change in baseline activity correlates closely with a change in evoked activity (n=17 neurons, slope=1.46, r=0.85).

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) can be decoded from the trial-by-trial responses of simultaneously recorded ensembles of auditory-responsive cells in either A1 or PFC. It can be decoded significantly better ($p < 0.001$) from A1 cells (n=22 ensembles of 57 neurons total) than from mPFC cells (n=13 ensembles of 25 neurons total), but it cannot be decoded significantly better during either block. The chance decoding level, attainable by a neuron with no information about the stimulus, is 0.5. The mean and SEM over the ensembles is shown. Significance was assessed with a 3-way ANOVA on brain region, target sound, and block.

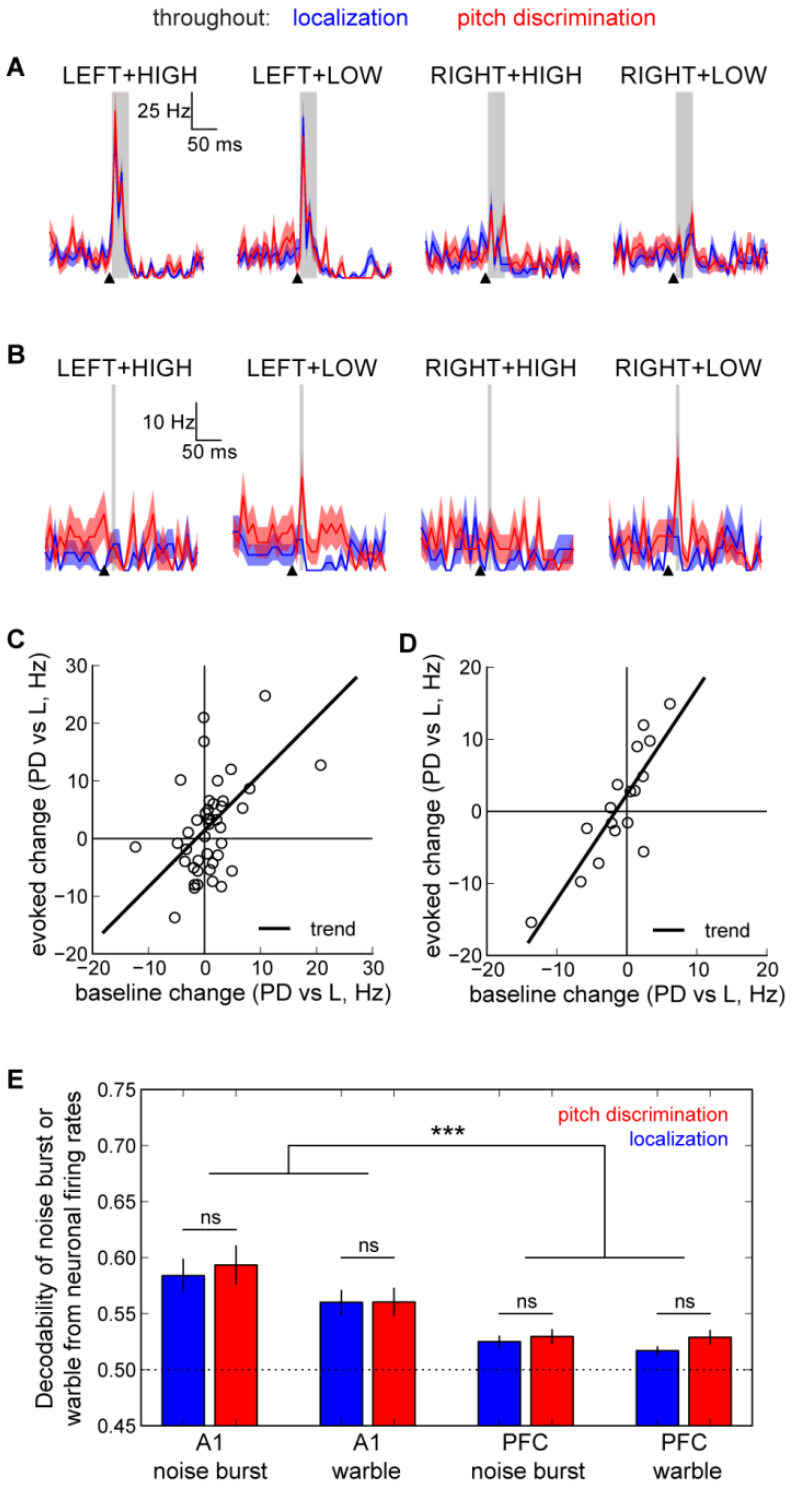
throughout: localization pitch discrimination
A
LEFT+HIGH
LEFT+LOW
RIGHT+HIGH
RIGHT+LOW
25 Hz
50 ms
B
LEFT+HIGH
LEFT+LOW
RIGHT+HIGH
RIGHT+LOW
10 Hz
50 ms
C
evoked change (PD vs L, Hz)
baseline change (PD vs L, Hz)
trend
D
evoked change (PD vs L, Hz)
baseline change (PD vs L, Hz)
trend
E
Decodability of noise burst or warble from neuronal firing rates
pitch discrimination
localization
***
ns
ns
ns
ns
A1 noise burst
A1 warble
PFC noise burst
PFC warble

FIGURE 6

Figure 7. Disruption of mPFC robustly impairs performance on the task.

A) Electrical disruption of mPFC significantly impacted task performance during localization trials (left panel) and/or pitch discrimination trials (right panel) in most sessions. Each point represents the performance within a single session on control (x-coordinate) vs. zap trials (y-coordinate). Plus symbols represent sessions during which the performance was significantly impaired ($p < 0.05$, Fisher's exact test). Throughout this figure, different colors represent different rats (red: Z1, yellow: Z2, green: Z3).

B) Example session from each rat. Performance is shown for each trial type (GO and NOGO in each block). Solid bars represent control trials; open bars represent zap trials. Error bars: 95% confidence intervals using Pearson-Klopper binomial fit. Asterisks indicate trial types for which electrical disruption significantly impairs performance (Fisher's exact test). The effect is robust within each example session, but variable across rats. See Supp. Info. for the data for all sessions.

C) Impairment on each trial type for each rat, across sessions. Error bars: SEM across sessions. All rats showed a significant impairment on NOGO trials in one block or the other ($p < 0.05$, binomial test on the number of sessions demonstrating impairment). One rat (Z2) also showed a significant impairment on localization GO trials.

FIGURE 7

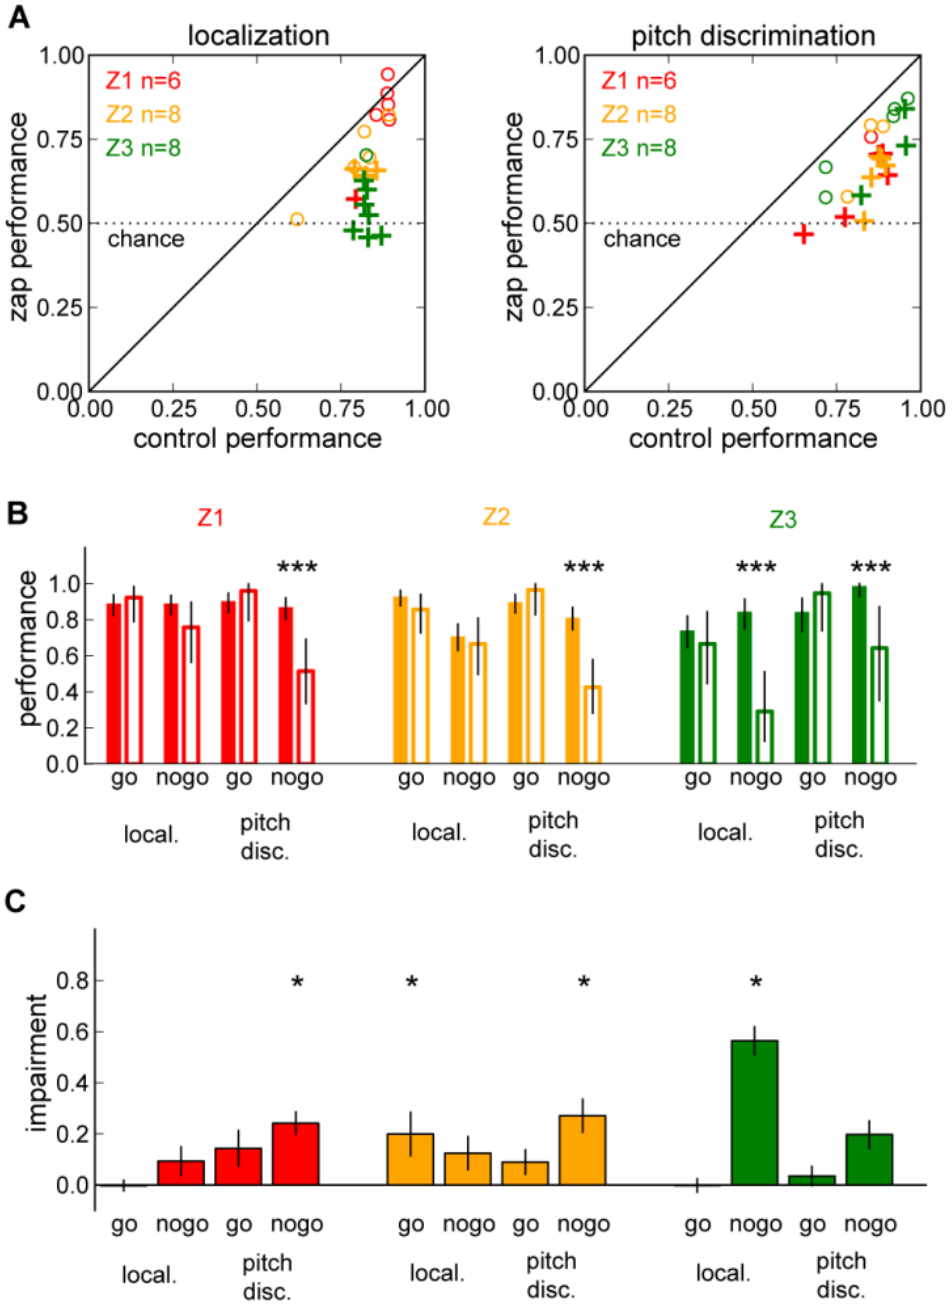

FIGURE 7

Figure 8. A simulated network model demonstrates how anticipatory modulation could solve the stimulus selection problem.

A) Network connectivity. A1 consists of a population of N neurons, each with random tuning for the four task stimuli and subject to additive Gaussian noise. Red and blue neurons in A1 are differentially activated in one block or the other, based on an excitatory "task signal" projection, hypothetically originating in PFC. Each subpopulation in A1 connects to a set of premotor command neurons encoding the possible responses in that block. The model's choice is determined by which command neuron is the most active via a winner-take-all mechanism. The weights W1 and W2 are constrained to be excitatory (non-negative) and are separately optimized during an initial supervised training phase, then fixed.

B) Performance of the model for N=320 neurons on task 1 (left panel) and task 2 (right panel). 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. We plot the probability of a correct choice versus the strength of the task signal. For 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 subpopulation corresponding to the incorrect block.) As the task signal increases in strength, the sensory input is eventually drowned out and the model's performance falls to chance (50%). This problem is especially pronounced at the lowest SNR, near 0.015. However, larger networks can still perform well at such SNRs (Supp. Info.).

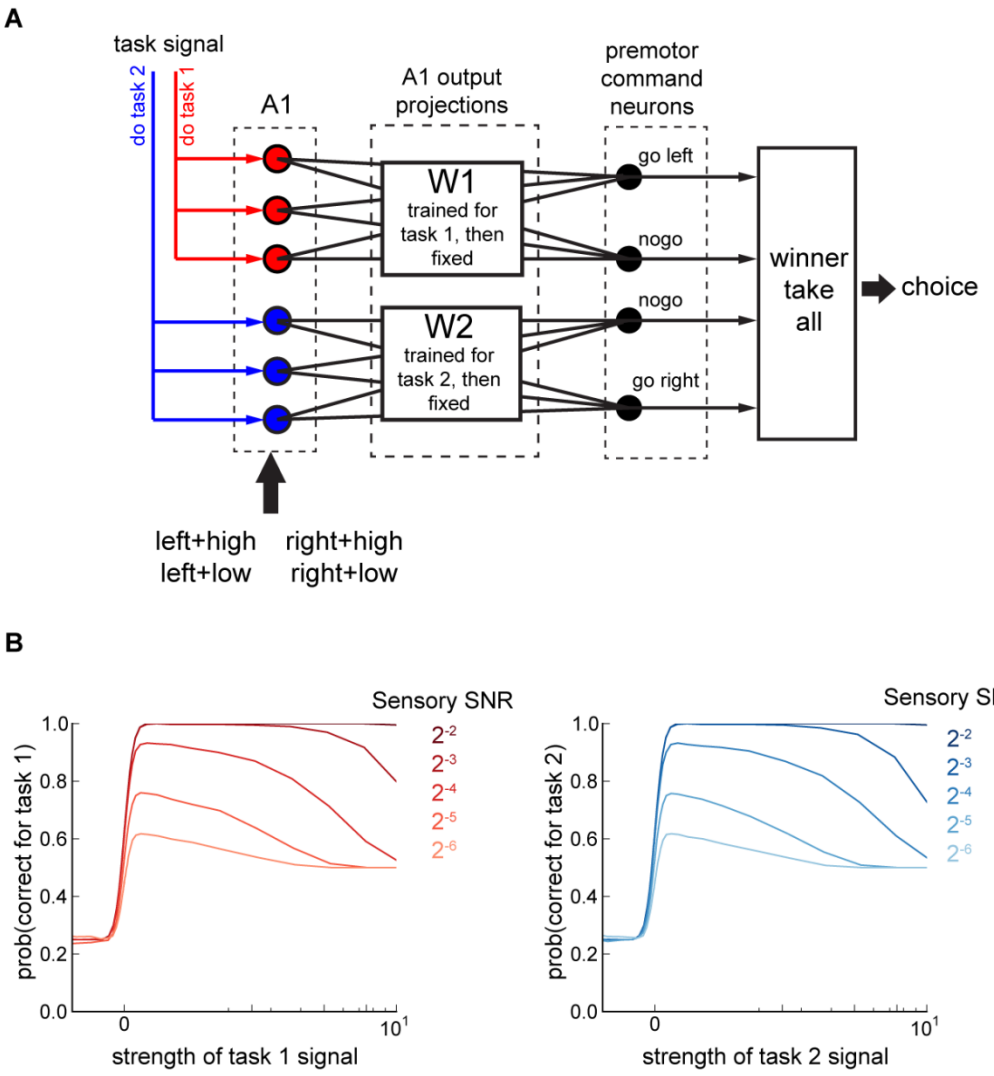
A
task signal
do task 2
do task 1
A1
A1 output projections
premotor command neurons
W1
trained for task 1, then fixed
W2
trained for task 2, then fixed
go left
nogo
nogo
go right
winner take all
choice
left+high
left+low
right+high
right+low
B
Sensory SNR
prob(correct for task 1)
strength of task 1 signal
Sensory SNR
prob(correct for task 2)
strength of task 2 signal

FIGURE 8

Supplementary Figures

Each figure is linked to the main figure with the same number. Not all main figures have associated supplementary figures.

Figure S2. Behavioral performance, related to Figure 2

A) Performance of each rat in greater detail, with go and nogo trials separately considered in each block. Rats generally did better on GO than on NOGO trials (first and third columns above second and fourth). Some rats did better on localization than on pitch discrimination (first and second columns above third and fourth). Error bars show SEM across sessions.

B) Performance (fraction of correct responses) of one rat that performed a slightly modified "catch trial" task on the last day of recordings. This task was designed to probe whether the rats learn to respond to a unified stimulus pair, or whether they learn to respond just to the target sound regardless of the identity of the distractor. On a small proportion (15%) of trials, we replaced the distractor with a neutral sound, to which the rat had never been trained. For example, on catch trials during localization the rat heard the same target as always (LEFT or RIGHT) with a novel mid-range warble of no behavioral relevance. If the rats had memorized each of the four possible stimulus pairs and were unable to generalize, they should perform at chance on these novel stimulus pairs. The performance on catch trials (red) and standard trials (white) for each trial type is shown, with 95% bootstrapped confidence intervals. The rats perform just as well on catch trials as on standard trials (unpaired Mann-Whitney U-test on the outcome of each trial, $p > 0.05$ in all cases). This suggests that the rats are selecting the target stimulus, not memorizing a fixed set of four stimulus pairs.

C) Same as panel B, but for a different rat.

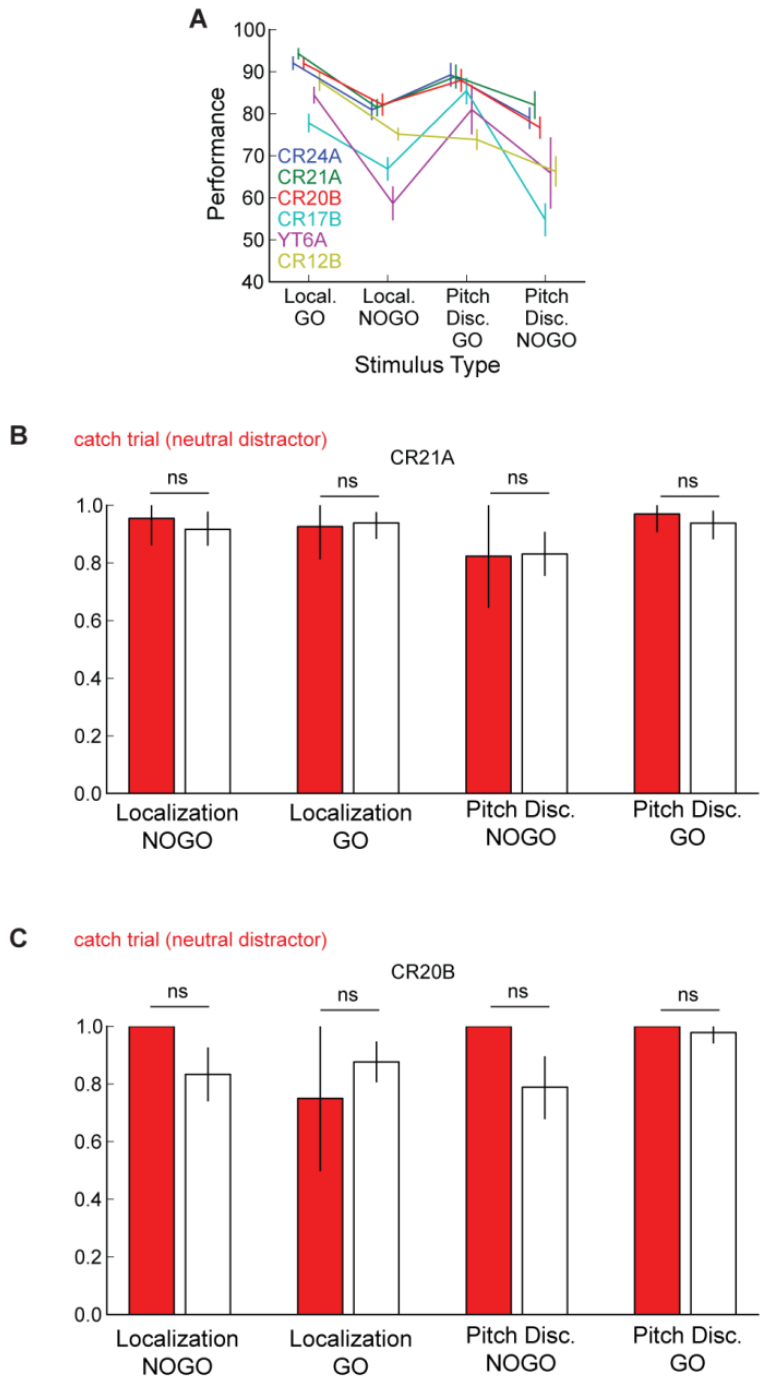
A
Performance
100
90
80
70
60
50
40
CR24A
CR21A
CR20B
CR17B
YT6A
CR12B
Local. GO
Local. NOGO
Pitch Disc. GO
Pitch Disc. NOGO
Stimulus Type
B
catch trial (neutral distractor)
CR21A
ns
1.0
0.8
0.6
0.4
0.2
0.0
Localization NOGO
Localization GO
Pitch Disc. NOGO
Pitch Disc. GO
C
catch trial (neutral distractor)
CR20B
ns
1.0
0.8
0.6
0.4
0.2
0.0
Localization NOGO
Localization GO
Pitch Disc. NOGO
Pitch Disc. GO

FIGURE S2

Figure S3. PFC hold period, related to Figure 3

A) Stacked histogram of the difference, rather than ratio, of hold period firing rate between blocks for all mPFC neurons.

B) Alternative presentation of the hold period effect across mPFC neurons. The hold period firing rate in each block is shown as the x- and y-coordinate of each point (red and blue: significant block preference; gray: not significant). Note the logarithmic scaling, necessary to avoid crowding the points with low firing rate. Error bars are 95% confidence intervals obtained by bootstrapping and were truncated at the edge of the plot. Significance was assessed with a Mann-Whitney U-test as described in the text.

C) Analysis of the hold period effect in mPFC on various types of correct and error trials. The trials are grouped by the meaning of the target sound, the rat's response, and the meaning of the distractor sound. Neurons are grouped by their preferred block. White bars represent correct trials, on which the rat's response matches the target sound. The gray bar represents "go-on-nogo" error trials when the target meant nogo but the rat went to the choice port anyway. (The opposite error, nogo-on-go, was too rare to include in this analysis. We only analyzed neurons from sessions with at least 3 trials of each type.) The orange bar represents "interference" trials on which the rat heard a distractor sound meaning go and went to the choice port associated with that distractor (WP, or wrong port). To aid in visualization, firing rates were normalized by subtracting the mean response on correct trials during the non-preferred block, and then averaged across neurons. There is no significant difference in the hold period activity between correct go, correct nogo, and incorrect go-on-nogo trials. But if the hold period activity encoded a simple plan to go to the choice port regardless of the upcoming stimulus, then it should be lower on correct nogo trials where the rat did not perform this action. Thus, if the hold period activity represents a motor plan, it must be subject to change ("countermanding") after the stimulus. However there is a significant difference between interference trials and all other trial types in that block. That is, when the rat gives the response that would be appropriate in the other block, the anticipatory activity is higher in the non-preferred block and lower in the preferred block — *i.e.*, the block modulation is attenuated, trending toward reversed. Significance between each pair of bars was assessed with a paired Mann-Whitney U-test across neurons, which is invariant under the subtractive normalization performed, and the p-values were Bonferroni corrected.

D) Histogram of the number of neurons preferring pitch discrimination (red), localization (blue), and neither (black). In some rats (marked N/A), no mPFC neurons were recorded.

E) Proportion of rule-encoding neurons in mPFC is consistent across rats. The data from (D) are now expressed a percentage of total neurons. No percentages are plotted for rats with fewer than 8 neurons total recorded in mPFC.

F) On average, firing rates are bidirectionally modulated in both blocks, versus the spontaneous rate. We defined the spontaneous rate as the average rate during epochs more than 2 s from

the nearest stimulus onset. Localization-preferring neurons (left) fire significantly more than spontaneous in localization and significantly less than spontaneous in pitch discrimination. A similar statement holds for pitch discrimination-preferring neurons (center). Neurons that do not prefer either block (right), that is, neurons for which the pre-stimulus firing rate does not significantly differ between blocks, tend to be suppressed versus spontaneous in both blocks, though this effect was not quite significant during pitch discrimination ($p = 0.06$). The bars show the average and SEM across neurons of the hold period firing rates minus the spontaneous firing rate for that neuron. We assessed significance with a paired Mann-Whitney test.

G) The azimuthal (left/right) angle of the rat's head during center poke differs by approximately 30 degrees between blocks. Red: pitch discrimination. Blue: localization. Values are normalized such that the average head angle across all trials is zero. These data are from an example session but all analyzed sessions yielded similar results. This preparatory motor activity is presumably an adaptive behavioral strategy in response to the fact that the choice port differs between blocks.

H) Example PFC neuron. Each point shows the square root of the number of spikes fired and head angle on a single trial. (Square root is a normalizing transformation for Poisson-like counts.) Across all trials, these variable are significantly correlated (black trend line). However, this correlation between firing rate and head angle is almost entirely explained by block type. Within each block, there is no such correlation – the red and blue trend lines are not significantly different from horizontal.

I) Fraction of explainable variance (FEV) in the spike counts (again square root-normalized) that the least-squares linear fit attributes to block (yellow), head angle (black), and interaction between block and head angle (purple), individually for each rule-encoding PFC neuron during the analyzed video sessions. Red horizontal line shows 50%. Most bars are mostly yellow, indicating that block is the major explanatory factor in the spike count. Only one bar is more than 50% black, corresponding to a neuron whose firing rate was mostly explained by head angle.

J) Summary plot of the data in panel (I). The distribution of FEV across neurons is plotted for each factor (block, head angle, and interaction). Individual points represent individual rule-encoding PFC neurons: pluses when that factor was a significant ($p < 0.05$, ANOVA) predictor, open circles where that factor was not significant. The red line shows the median; the blue box outlines the inter-quartile range. Across the population, most of the variance (median: 67.1%) is explained by block; in contrast, only a small amount is explained by head angle (median: 5.6%) and the rest (12.2%) is explained by the interaction between the two. (These values do not sum to 100% because of the use of the median.)

K) No evidence for a correlation between the pre-stimulus firing rate of PFC neurons and the animal's reaction time, defined as the time between stimulus onset and withdrawal from center port. We plot here the distribution of correlation coefficients obtained in both blocks, with red representing significantly correlated neurons ($p < 0.05$ after correction with the false

discovery rate). Only a small minority of neurons showed a significant correlation, and the overall distribution is not significantly different from zero ($p > 0.05$, one-sample t-test). Similar results were obtained when considering only rule-encoding neurons, *i.e.*, those that with a significantly increased pre-stimulus firing rate in one block or the other.

L) Similar to (K), but now correlating the anticipatory firing rate with the "motion time" – the time necessary for the rat to move to the choice port on successful GO trials. Again, only a small minority of neurons show an individually significant correlation and the population distribution is not significantly different from zero. Similar results were obtained when considering only rule-encoding neurons.

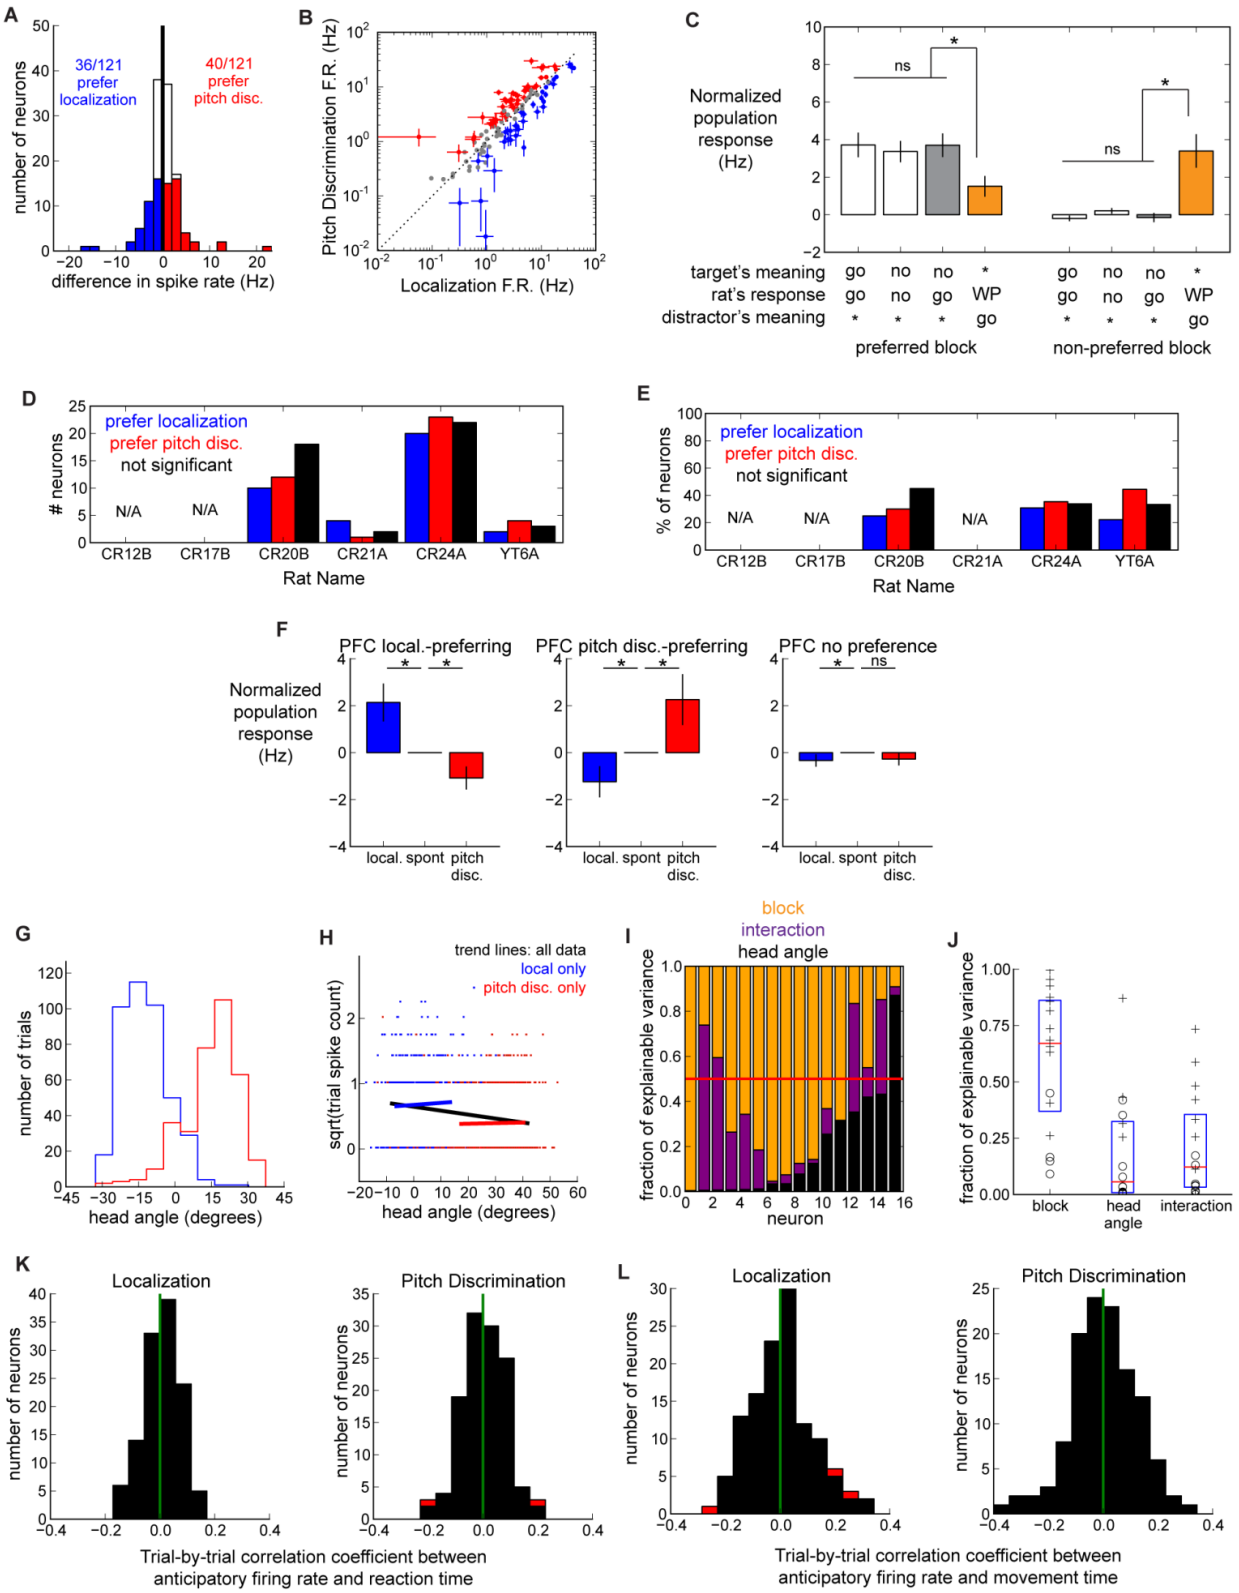
A
number of neurons
36/121 prefer localization
40/121 prefer pitch disc.
difference in spike rate (Hz)
B
Pitch Discrimination F.R. (Hz)
Localization F.R. (Hz)
C
Normalized population response (Hz)
ns
*
target's meaning
rat's response
distractor's meaning
preferred block
non-preferred block
D
neurons
prefer localization
prefer pitch disc.
not significant
N/A
CR12B CR17B CR20B CR21A CR24A YT6A
Rat Name
E
% of neurons
F
PFC local.-preferring
PFC pitch disc.-preferring
PFC no preference
local. spont pitch disc.
G
number of trials
head angle (degrees)
H
trend lines: all data
local only
pitch disc. only
sqrt(trial spike count)
I
block
interaction
head angle
fraction of explainable variance
neuron
J
block
head angle
interaction
K
Localization
Pitch Discrimination
Trial-by-trial correlation coefficient between anticipatory firing rate and reaction time
L
Trial-by-trial correlation coefficient between anticipatory firing rate and movement time

FIGURE S3

Figure S4. A1 hold period, related to Figure 4

A, B) Same as Figure S3A,B, but for A1 neurons instead of mPFC neurons.

C) Same as Figure S3C (correct trials: white; go-on-nogo errors: gray; interference trials: orange), but for A1 neurons instead of mPFC neurons. The effects are similar to those for mPFC, but the effect on interference trials is stronger. Now the direction of hold period modulation is significantly reversed -- the firing rate is higher during such trials in the non-preferred block than in the preferred block. However, there is no difference between the other trial types, regardless of whether the rat performed a go or nogo response.

D) Histogram of the number of neurons preferring pitch discrimination (red), localization (blue), and neither (black). In one rat, no A1 neurons were recorded (marked N/A).

E) Proportion of rule-encoding neurons in A1 is consistent across rats. The data from (D) are now expressed a percentage of total neurons. No percentages are plotted for rats with fewer than 8 neurons total recorded in A1.

F) On average, firing rates are bidirectionally modulated in both blocks, versus the spontaneous rate. We defined the spontaneous rate as the average rate during epochs more than 2 s from the nearest stimulus onset. Localization-preferring neurons (left) fire significantly more than spontaneous in localization and significantly less than spontaneous in pitch discrimination. A similar statement holds for pitch discrimination-preferring neurons (center). Neurons that do not prefer either block (right), that is, neurons for which the pre-stimulus firing rate does not significantly differ between blocks, showed no change in their firing rates versus spontaneous. The bars show the average and SEM across neurons of the hold period firing rates minus the spontaneous firing rate for that neuron. We assessed significance with a paired Mann-Whitney test.

G) The azimuthal (left/right) angle of the rat's head during center poke differs by approximately 30 degrees between blocks. Red: pitch discrimination. Blue: localization. Values are normalized such that the average head angle across all trials is zero. These data are from an example session but all analyzed sessions yielded similar results. This preparatory motor activity is presumably an adaptive behavioral strategy in response to the fact that the choice port differs between blocks.

H) Example A1 neuron. Each point shows the square root of the number of spikes fired and head angle on a single trial. (Square root is a normalizing transformation for Poisson-like counts.) Across all trials, these variable are significantly correlated (black trend line). However, this correlation between firing rate and head angle is almost entirely explained by block type. Within each block, there is no such correlation – the red and blue trend lines are not significantly different from horizontal.

I) Fraction of explainable variance (FEV) in the spike counts (again square root-normalized) that the least-squares linear fit attributes to block (yellow), head angle (black), and interaction between block and head angle (purple), individually for each rule-encoding A1 neuron during the analyzed video sessions. Red horizontal line shows 50%. Most bars are mostly yellow, indicating that block is the major explanatory factor in the spike count. Only one bar is more than 50% black, corresponding to a neuron whose firing rate was mostly (50.2%) explained by head angle.

J) Summary plot of the data in panel (I). The distribution of FEV across neurons is plotted for each factor (block, head angle, and interaction). Individual points represent individual rule-encoding A1 neurons: pluses when that factor was a significant ($p < 0.05$, ANOVA) predictor, open circles where that factor was not significant. The red line shows the median; the blue box outlines the inter-quartile range. Across the population, most of the variance (median: 71.5%) is explained by block; in contrast, only a small amount is explained by head angle (median: 14.1%) and the rest (17.8%) is explained by the interaction between the two. (These values do not sum to 100% because of the use of the median.)

K) No evidence for correlation between pre-stimulus firing rate of A1 neurons and the animal's reaction time, defined here as the time between stimulus onset and withdrawal from center port. We plot here the distribution of correlation coefficients obtained in both blocks, with red representing significantly correlated neurons ($p < 0.05$ after correction with the false discovery rate). Only a small minority of neurons showed a significant correlation, and the overall distribution is not significantly different from zero ($p > 0.05$, one-sample t-test). Similar results were obtained when considering only rule-encoding neurons, *i.e.* those with a significantly increased pre-stimulus firing rate in one block or the other.

L) Similar to (K), but now correlating the anticipatory firing rate with the "motion time" – the time necessary for the rat to move to the choice port on successful GO trials. Again, only a small minority of neurons show an individually significant correlation and the population distribution is not significantly different from zero. Similar results were obtained when considering only rule-encoding neurons.

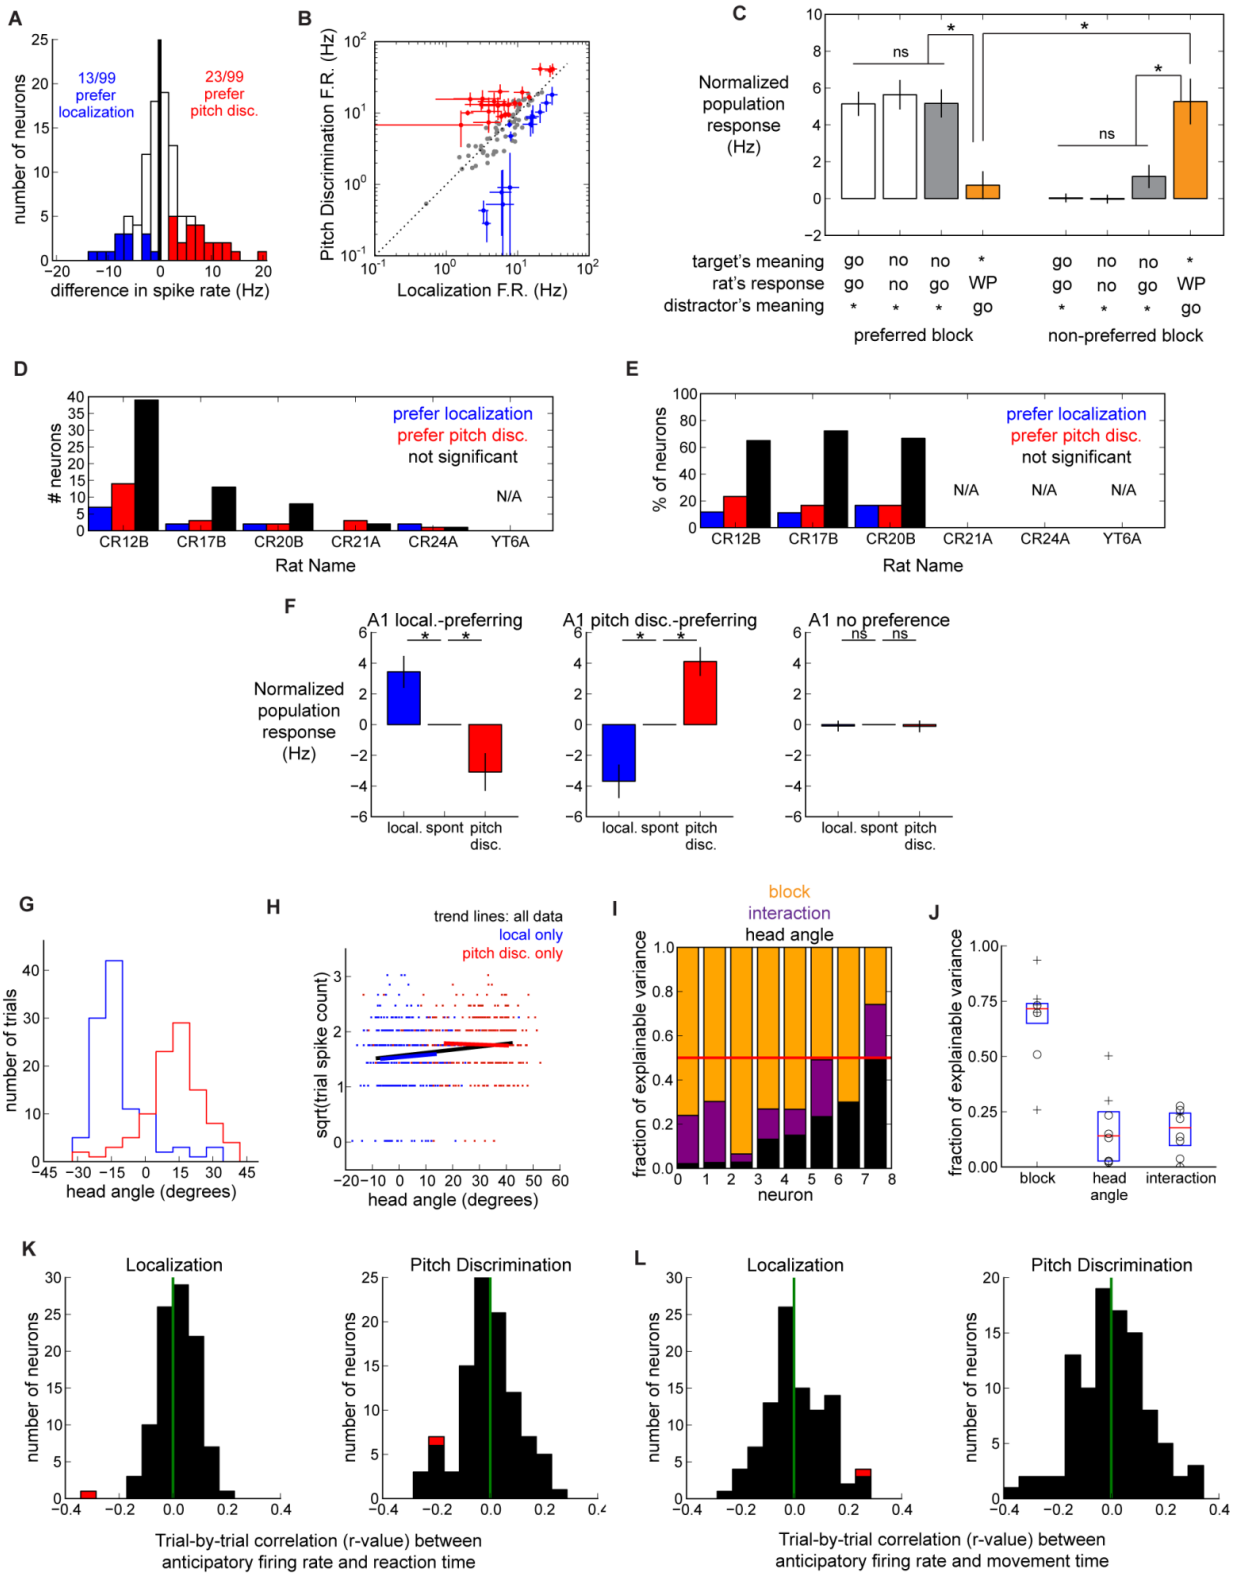
A
number of neurons
13/99 prefer localization
23/99 prefer pitch disc.
difference in spike rate (Hz)
B
Pitch Discrimination F.R. (Hz)
Localization F.R. (Hz)
C
Normalized population response (Hz)
ns
*
target's meaning
rat's response
distractor's meaning
preferred block
non-preferred block
D
neurons
prefer localization
prefer pitch disc.
not significant
N/A
CR12B
CR17B
CR20B
CR21A
CR24A
YT6A
Rat Name
E
% of neurons
F
A1 local.-preferring
A1 pitch disc.-preferring
A1 no preference
Normalized population response (Hz)
local. spont
pitch disc.
G
number of trials
head angle (degrees)
H
trend lines: all data
local only
pitch disc. only
sqrt(trial spike count)
I
block
interaction
head angle
fraction of explainable variance
neuron
J
block
head angle
interaction
K
Localization
Pitch Discrimination
number of neurons
Trial-by-trial correlation (r-value) between anticipatory firing rate and reaction time
L
Trial-by-trial correlation (r-value) between anticipatory firing rate and movement time

FIGURE S4

Figure S5. More information on timecourse, related to Figure 5

A) Similar to figure 5C but now locked to entry into the center port that initiated the trial, rather than the subsequent stimulus onset. This shows the population time course during the preferred block, averaged over all rule-encoding neurons in each region (purple: mPFC; orange: A1). Throughout this figure, firing rates were first normalized to zero mean and unit variance and then averaged over neurons; the thickness of the trace represents SEM over neurons. This demonstrates more clearly that the increase in anticipatory activity definitely precedes the center-poke entry.

B) Population time course during the preferred block for rule-encoding neurons in mPFC (left) and A1 (right). In this panel, we include only trials *following* a successful NOGO. Along the lower edge of the figure we also plot the distribution of latencies to the end of the previous trial, defined as the withdrawal from the center port following the NOGO stimulus onset. The increased activity persists between the trials, even though the rat is typically remaining motionless near the center port during this time. This demonstrates that the anticipatory effect is not due to the rat's motion toward the center port.

C) Similar to Figure 5D in the main text, but now showing activity in the non-preferred block (green) in addition to activity during the preferred block (purple); also, only correct NOGO trials are included to demonstrate the persistence of the effect. The population time course is locked to exit from center-port after the stimulus onset. Before exiting the center port, average activity in the non-preferred block is suppressed below baseline; average activity in the preferred block is increased above baseline. This remains true for at least several seconds after the rat exits the center port, during which time the rat typically had already initiated the next trial. To illustrate this last point, along the lower edge of the figure we also plot the distribution of latencies to the relevant trial events: stimulus onset (black) and the center-poke beginning the next trial (gray).

D) Similar to panel (C), except we now include only correct GO trials. After exiting the center port, the firing rate during the preferred block (again averaged over neurons) falls below baseline and remains there as the rat moves to the choice port and drinks a reward (which always required at least several seconds). In contrast, during the non-preferred block A1 neurons are actually activated above baseline during this period, thus inverting the usual rule encoding. Along the lower edge of the figure we also plot the distribution of latencies to the relevant trial events: stimulus onset (black) and reward delivery (green). We finally note that these results depict the population average; individual neurons displayed a wide variety of dynamics (Figure 4).

E) Similar to (C) and (D), but now separately plotting the population activity during localization (blue) and pitch discrimination (red) blocks from neurons preferring localization (left column) and pitch discrimination (right column) on correct NOGO trials (top two rows) and correct GO trials (bottom two rows). The dynamics of the pitch discrimination-preferring and localization-preferring neurons are quite similar (compare left and right columns), other than the fact that

the red and blue traces are roughly reversed (by definition of preferred block). Along the lower edge of the figure we also plot the distributions of latencies to stimulus onset (black), reward delivery (green, GO trials only), and center-poke beginning next trial (gray, NOGO trials only).

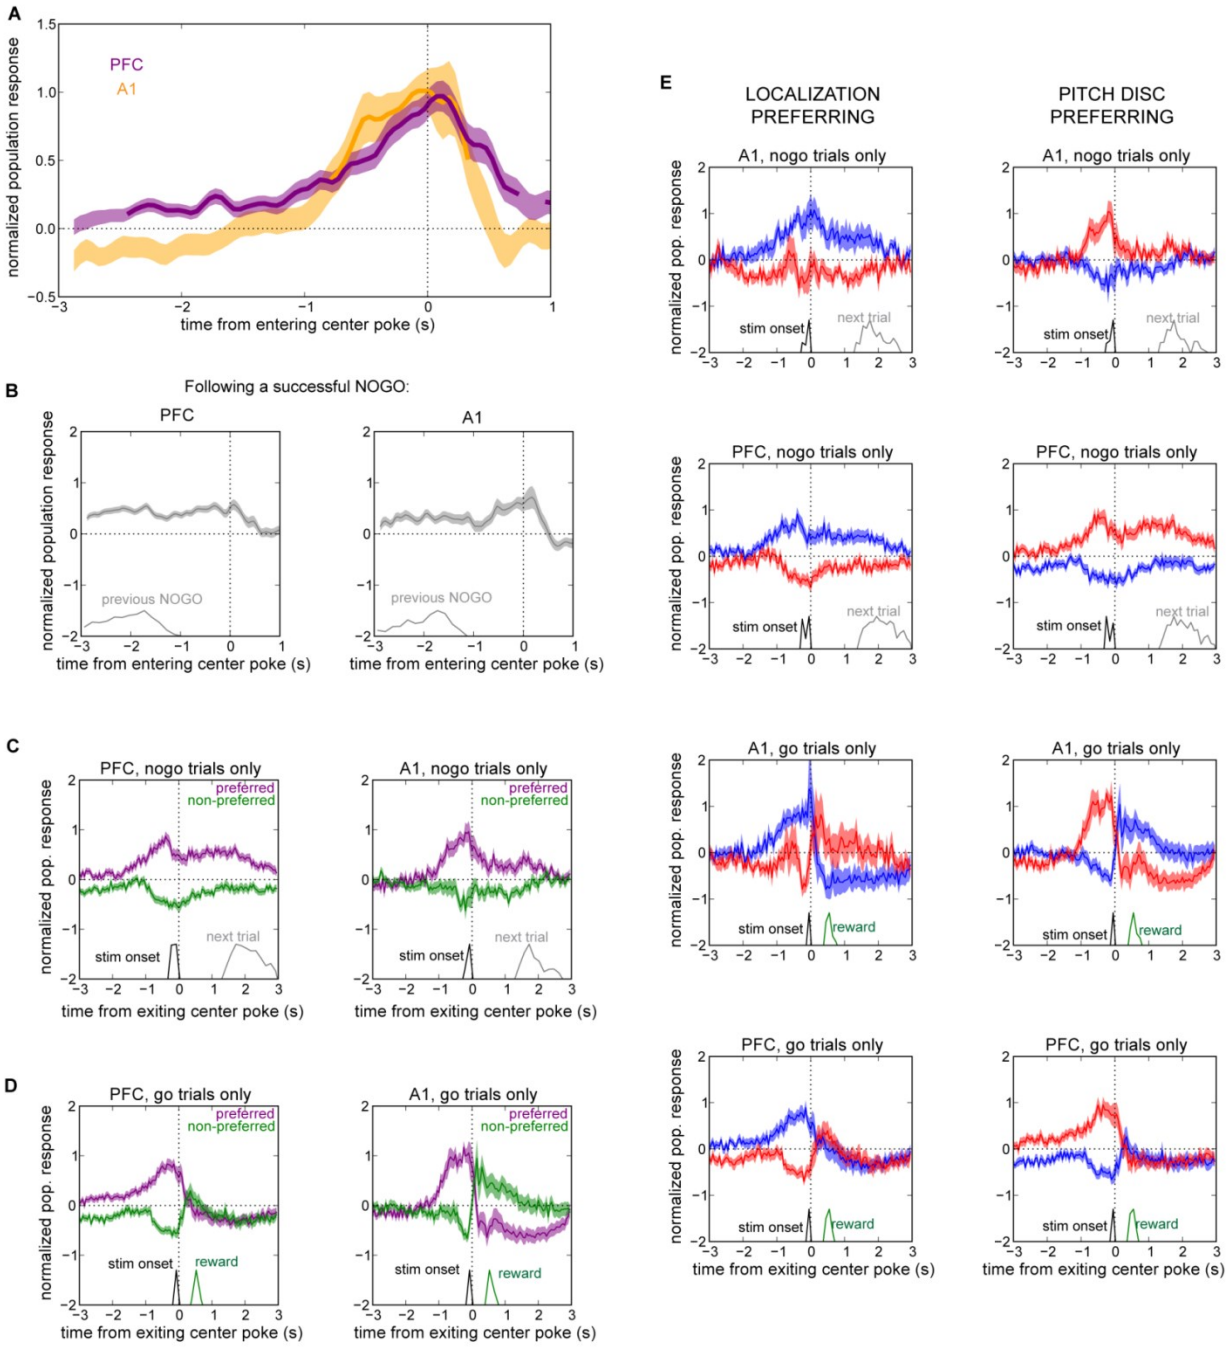

FIGURE S5

Figure S6. More information on evoked activity, related to Figure 6

A) PSTH of a typical A1 neuron, to illustrate the notion of "evoked response strength." This neuron's response strength is near the median of the A1 population. All stimuli and trials are included in this PSTH. Note that the response to the onset of the sound is rapid, short-latency, tightly stimulus-locked, and brief. The onset window is shaded. This was defined as the continuous set of time bins post-onset over which the firing rate was significantly greater than the rate preceding the stimulus. Some mPFC units also showed similar (though weaker) stimulus-driven responses.

B) An example A1 neuron showing one of the strongest and most sustained recorded responses.

C) Distribution of onset response strengths across n=49 auditory-responsive A1 neurons (blue) and n=31 auditory-responsive PFC neurons (orange). The response strength is expressed as the average number of additional spikes (over baseline) during the onset response window for that neuron. All trials and stimuli are included, and the baseline rate is calculated over the 50 ms preceding the stimulus onset. The responses are significantly stronger in A1 neurons (median: 0.11 spikes) than in PFC neurons (median: 0.02 spikes), $p < 0.05$, unpaired Mann-Whitney U-test. Note the long tail of the distribution: a small subpopulation fires much more strongly than the median.

D) Alternative presentation of onset response strength. The data are the same as in (C), but now expressed as percentage of baseline firing rate. Because the response window is so brief, a small number of additional spikes over baseline typically represents a large (many-fold) increase in rate. Again the responses are stronger in A1 neurons (median: 209% of baseline) than in PFC neurons (median: 171% of baseline), $p < 0.001$, unpaired Mann-Whitney U-test.

E) Distribution of latencies to center of onset response window across the same populations as (C) and (D). The PFC latencies are significantly longer (median: 19.75ms) than the A1 latencies (median: 16.75ms). $p < 0.05$, unpaired Mann-Whitney U-test.

F) No correlation between stimulus tuning and the change in pre-stimulus rate for auditory-responsive A1 neurons. We calculated the strength of the tuning for the noise burst (LEFT versus RIGHT) or the warble (LOW versus HIGH) using the difference in firing rate to those stimuli when they were presented on cue trials. Similar to Figure 6E in the main text, we quantified tuning using the performance of an ideal decoder. This metric has a value of 0.5 for identical responses, and 1.0 for perfectly discriminable responses. We found no correlation between how well the neurons were tuned for either stimulus and the change in pre-stimulus rate across blocks. Although these results were obtained using the cue trials to measure tuning, similar results were obtained when we calculated the tuning using the responses to the stimuli containing distractors (data not shown).

G) Same as (F), but for auditory-responsive PFC neurons. Note the difference in the scale of the x-axis vs (F), reflecting the fact that PFC neurons tend to be more poorly tuned for the stimuli than A1 neurons. Again there is no significant correlation, though there is a weak, non-significant trend in the following direction: neurons that are well-tuned for the warble tend to show increased anticipatory firing in pitch discrimination; similarly, neurons that are well-tuned for the noise burst tend to show increased anticipatory firing during localization.

H) Across the population of auditory-responsive neurons in both brain regions, all four stimulus pairs elicit equally small and non-significant changes across blocks in their evoked firing rate. Each box and whiskers plot shows the distribution across neurons of the difference in evoked firing rate between pitch discrimination and localization, for each stimulus pair. (We used square root as a normalizing transform but this did not affect the results.) The box's top and bottom show the inter-quartile range (IQR); the tapered area of each box shows the 95% bootstrapped confidence intervals on the median. Pluses indicate outliers, defined as more than 1.25 times the IQR from the median; these points were still included in the analysis. We also repeated the analysis after subtracting the pre-stimulus rate from the evoked rate in both blocks (labeled "baseline-subtracted"). A one-way Kruskal-Wallis test revealed no significant difference between the stimulus pairs. We also assessed whether the change to each individual stimulus pair was different from zero using the Wilcoxon signed-rank test, correcting for multiple comparisons within each group (*e.g.,* each group of four boxplots) using the false discovery rate. We found was no significant difference from zero for any individual stimulus pair, indicating that there was no consistent trend toward increased evoked firing rates in either block for any stimulus pair.

I) Similar to (H), but here we took the absolute value of all of the data. We reasoned that some stimulus pairs might elicit increased firing rates during pitch discrimination in some neurons and in localization in other neurons, resulting in no net effect. Under this hypothesis, we would expect that the absolute value of the difference between blocks, across neurons, would be significantly greater for that stimulus pair versus the other stimulus pairs. However, a one-way Kruskal-Wallis test across stimulus pairs again revealed no significant difference: all stimulus pairs yield equivalent absolute differences in evoked firing rate across neurons.

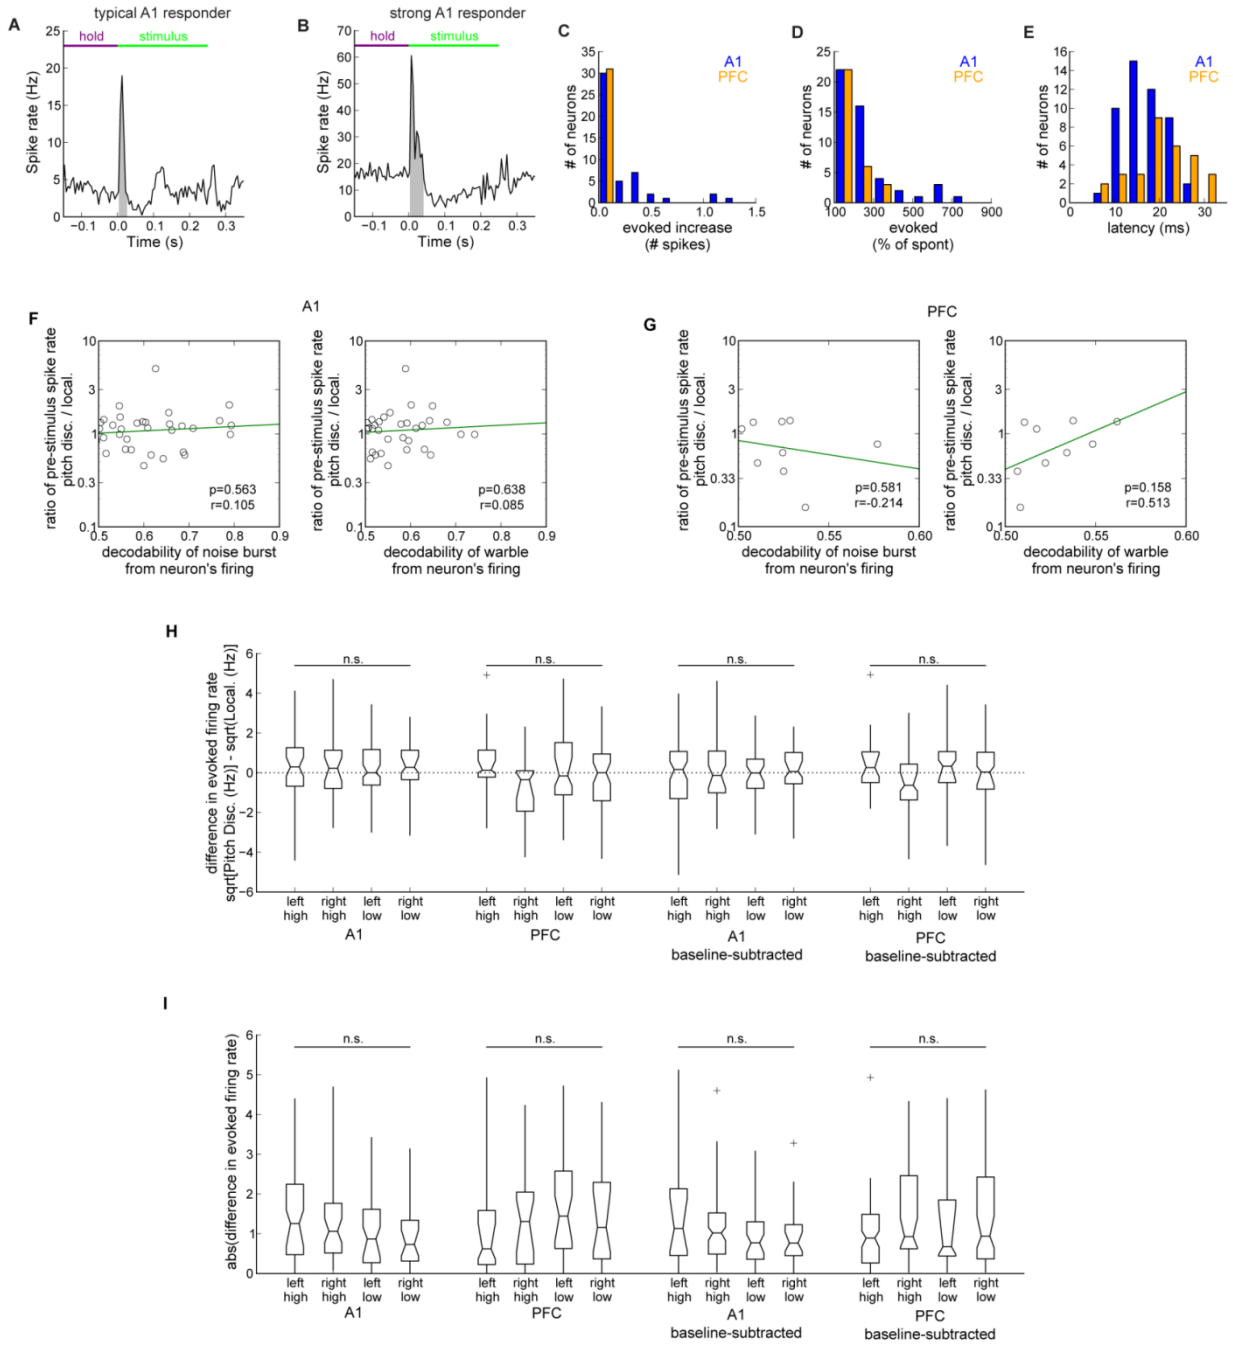
A
typical A1 responder
hold
stimulus
Spike rate (Hz)
Time (s)
B
strong A1 responder
C
# of neurons
A1
PFC
evoked increase
(# spikes)
D
evoked
(% of spont)
E
latency (ms)
F
A1
ratio of pre-stimulus spike rate
pitch disc. / local.
p=0.563
r=0.105
decodability of noise burst
from neuron's firing
p=0.638
r=0.085
decodability of warble
from neuron's firing
G
PFC
p=0.581
r=-0.214
p=0.158
r=0.513
H
difference in evoked firing rate
sqrt[Pitch Disc. (Hz)] - sqrt[Local. (Hz)]
n.s.
left high
right high
left low
right low
A1
PFC
A1
baseline-subtracted
PFC
baseline-subtracted
I
abs(difference in evoked firing rate)

FIGURE S6

Figure S7. More information on electrical disruption of mPFC, related to Figure 7

A) Effect of electrical disruption of mPFC on performance for all individual sessions from three rats (red: Z1, yellow: Z2, green: Z3). Each line connects the performance on control trials (left) and zap trials (right) within the same session. Plus marks indicate a significant difference, which was a decrease in every case ($p < 0.05$, Fisher's exact test). Trials are grouped according to GO and NOGO in both blocks. Asterisks indicate trial types for which the effect of disruption was significant across sessions ($p < 0.05$, binomial test on the number of sessions showing impairment).

B) Impairment caused by electrical disruption by stimulus pair, averaged across sessions for each rat. Error bars: SEM across sessions. Impairment is defined as the difference between performance on control trials and zap trials. Colors represent individual rats, following (A). Two rats (Z1 and Z2) showed a significant decrease for RIGHT+HIGH during pitch discrimination across sessions ($p < 0.05$, binomial test on the number of sessions showing impairment). One rat (Z3) showed a deficit for both NOGO sounds during localization.

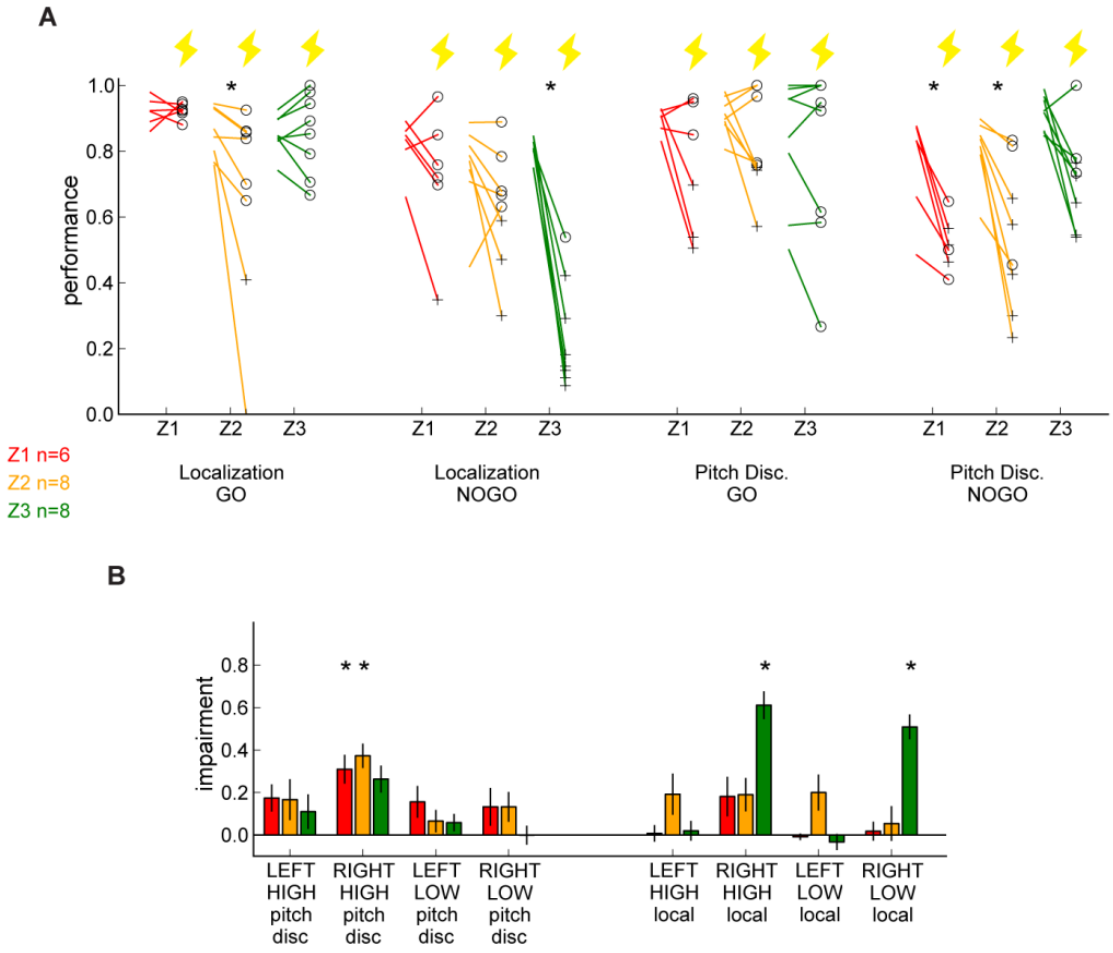

FIGURE S7

Figure S8. Behavior of the model over parameter space, related to Figure 8

A) The model described in Figure 8 operates well over a wide range of parameters. We quantify its performance as the strength of the task signal (x-axis in Figure 8B) necessary for the model to reach a criterion performance of 80%. As in Figure 8B, we normalize this task signal strength to the sensory noise (the variance of the Gaussian noise that was added to each A1 neuron's activation during the simulation). This plot shows a color-coded map of the strength of the task signal necessary to reach criterion, for various values of N (x-axis) and of the sensory signal-to-noise ratio (y-axis), defined as the strength of the tuning of each neuron to the variance of the same Gaussian noise mentioned above. Darker colors are good: they indicate that a weak task signal was sufficient to reach criterion; lighter colors indicate that a stronger signal was necessary. It is desirable for the model to operate with as weak a task signal as possible, because when the task signal becomes quite strong it overwhelms the sensory input and the performance drops again. We found that larger networks invariably performed better, as expected. For a given network size (any column of the figure), the performance was stable across a certain range of SNRs, but quickly degraded below a certain SNR cutoff. In the white region (small network sizes and/or very poor sensory signal-to-noise ratio), the model never reaches criterion performance.

B) Similar to panel A, but now we show the maximum magnitude of the task signal before performance falls below criterion again. Here, lighter colors (as in the upper right corner) are good, because they indicate that the network performs well over a wide range of task signal strengths; this is associated with large network sizes and/or high sensory signal-to-noise ratios.

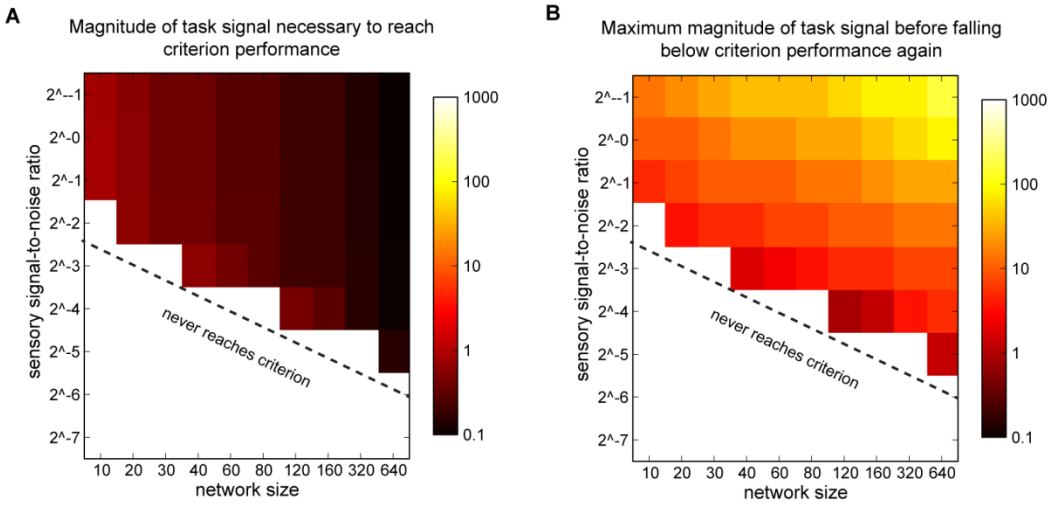

FIGURE S8

Supplementary Experimental Procedures

Behavioral parameters

Each stimulus was 250 ms in duration. The warbles were frequency-modulated tones, centered at 6 KHz and 16 KHz with a 10 Hz modulation frequency of amplitude 0.07 octaves, and presented with equal intensities (65 dB SPL) from both speakers. The white noise bursts were of approximately equal power at all frequencies between 5KHz and 50KHz, decaying rapidly outside of this range; the total acoustic power over this range was 55 dB SPL, delivered from only one speaker at a time (LEFT or RIGHT). We used a Lynx L22 sound card to convert digital signals to analog voltages, and Fostex FT17H tweeters to produce the sound.

In a subset of sessions, we presented task-irrelevant natural sounds during epochs when the rat was not initiating trials for the purpose of probing receptive fields. These probe sounds were terminated as soon as the rat began performing the task again. We did not observe any correlation between the neural results presented in this study and the receptive fields thus obtained.

Surgical implantation

Rats were anesthetized using ketamine/xylazine and isoflurane as necessary to maintain a deep anesthesia, assessed using toe pinch. Skin and fascia were resected from the midline and the skull cleaned. Titanium screws (Small Parts) were inserted into each cranial plate. Two stainless steel screws with a wire soldered onto the head were inserted into the occipital plate above the cerebellum. These were later soldered onto the reference and ground inputs on the microdrive.

Craniotomies were performed directly dorsal to the target areas (A1: 5.25 mm posterior and 6.5 mm left from bregma; PL: 3.0 mm anterior and 1.0 mm left from bregma). The dura was removed and the tetrodes gradually inserted into each region. The craniotomy was filled with agar, which surrounded and protected the tetrodes. Methyl methacrylate (Teet's, Henry Schein) was used to affix the entire drive to the skull and screws.

Finally the tissue was flushed thoroughly with sterile saline and sutures were used if necessary to seal the skin around the implant. Aseptic technique was maintained throughout and the tetrodes themselves were disinfected before implantation. Post-operatively, the rat was given buprenorphine and/or meloxicam to provide analgesia and its health and weight monitored twice daily. Once the rat was fully recovered, the behavioral task was resumed, now concurrent with electrophysiological recording.

Analysis of possible confounds: waveform variation and firing rate drift

In addition to the standard spike-sorting procedure for identifying stable units, we also used an extra, stricter check on the quality of our data to ensure that the hold period effect could not be due to sorting errors arising from small variations in waveform between blocks. For each neuron, we identified the sub-cluster of sorted spikes occurring just during the localization hold period, and calculated the Mahalanobis distance (in the first four PCA feature dimensions) between this subcluster and the entire cluster of spikes from this unit. We assessed the

significance of this distance by randomly permuting the labels on the subcluster and the full cluster 2000 times, and calculating the probability of observing a distance less than or equal to the true distance. We repeated this analysis for the other block (pitch discrimination). We discarded the neuron from analysis entirely if the subcluster in either block was significantly more separated from the full cluster than the permuted subclusters were ($p < 0.05$, permutation test). We also repeated the analysis with simulated Gaussian sub-clusters of the same size as the actual sub-clusters and derived the distribution of mean-squared distance from the cluster center, again rejecting any neuron whose sub-cluster that exceeded the 95th percentile of this distribution in either block.

Additionally, we also considered the possibility that a spurious hold period effect could arise from a slow increase or decrease in firing rate over the entire session, perhaps due to drift or motivation, even though the multiple switches between blocks within each session made such a possibility unlikely. We reasoned that, if this were true, then when taking block number into account the difference between block types should no longer be significant. We fit a linear model to the square root of the spike count in the hold period on each trial, using both block type (localization or pitch discrimination) and block number (1, 2, 3, …) as predictors. (Square root is a variance-stabilizing transform for Poisson counts.) We assessed the significance of each predictor with ANOVA. Any neuron that showed a hold period effect according to the analysis described in the text, but that failed to show a significant effect of block type or failed the overall F-test ($p > 0.05$) was discarded from the analysis. 8/231 neurons (combined across brain regions) were discarded for this reason.

For this ANOVA we used the type-III sum of squares. For all ANOVA analyses in this study, we avoided using type-I sum of squares because we found it to be much more sensitive to unequal trial counts (e.g., more hits in one block than in the other).

Power analysis

We analyzed the statistical power of our methods (unpaired Mann-Whitney U-test on the spike counts across blocks) on simulated Poisson counts. We determined the total spike count had to be at least 20 spikes to detect a change between blocks; therefore neurons with fewer total spikes than this in the hold period of all trials combined were discarded from analysis. For the typical trial counts in our dataset, we would not be able to detect any less than a doubling of firing rate for a neuron with this minimum firing rate (though the method becomes much more sensitive at higher firing rates). For this reason, selection rule encoding could be even more common than we have shown. Also, due to the fact that some of our A1 data was collected in earlier animals for which the hold periods were shorter and the trial counts lower, we have less statistical power in that portion of the dataset.

Calculation of evoked responses

For each neuron, the spike times on each trial were smoothed with a Gaussian kernel with 1 ms standard deviation. For every 0.5 ms time bin after stimulus onset, the distribution of smoothed spike counts was compared to the combined distribution of all 0.5ms time bins in the 50 ms

preceding stimulus onset with a Mann-Whitney U test. The first window of contiguous time bins that were all significantly greater than the spontaneous rates was defined as the onset response window. Windows of less than 1 ms were discarded because these neurons emitted far too few spikes to analyze statistically. A few neurons showing atypical auditory onset responses (*e.g.*, 4/108 showed a slow build rather than a short-latency peak) were discarded because we were concerned that their activity might be driven by the decision rather than the stimulus.

For the average evoked response plotted on the y-axis of Figures 6C and 6D, we used a bootstrap procedure to draw equally from each stimulus, separately for each block. This procedure accounts for differences in the proportion of each stimulus type across blocks, arising from random chance or from better performance on some stimuli than others (*e.g.*, better performance on go than on nogo, Figure S2A) since only correct trials were included for this analysis.

Changes in evoked response not explained by changes in baseline

We asked whether there were any additional changes in evoked response, above and beyond what could be explained by pre-stimulus effects, by subtracting the block-specific baseline firing rate from the evoked response on each trial and then repeating the bootstrap procedure described immediately above. We found that a small population (6/43, or 14%) of neurons increased their evoked response significantly ($p < 0.05$ from the overlap of the bootstrapped distributions in each block), above and beyond any baseline changes. (Another analysis in which we directly compared across blocks the number of spikes emitted in response to each stimulus individually yielded similar results, as did a stimulus*block ANOVA on each neuron.) However, unlike the other results in the paper, these neurons were largely (4/6) observed in a single animal (Rat 1), the rat that had the most difficulty with pitch discrimination trials, and in these neurons the firing rate was higher during pitch discrimination. One possibility is that the greater difficulty this rat had with one block led to this block-specific increased in evoked rate; our other rats were more evenly matched in performance between blocks.

Disruption of mPFC by electrical stimulation

Electrical microstimulation of cortex produces first a strong and synchronous activation of nearby neurons and later a slower, long-lasting suppression (Logothetis et al., 2010). This rebound is thought to be a homeostatic response, either network-level or cell-autonomous.

We injected a train of 1 ms current pulses at 10 Hz into mPFC on a subset of trials (“zap trials”) and only during the center-poke hold and stimulus presentation. In one animal we used a single pair of electrodes centered near the dorsal portion of the prelimbic region in each hemisphere; in the other animals we used an array of three electrodes spanning the anterior-posterior and dorsal-ventral extent of the prelimbic region in each hemisphere. Stimulation was always bipolar, first positive and then negative, with respect to a cranial ground screw. Two rats (Z1 and Z3) were also implanted with recording drives in auditory cortex.

We began with a very low current, around 10uA per electrode, which was typically too low to produce any behavioral effect. We used pilot sessions to increase the amount of current until performance on the task became moderately impaired. This typically occurred at around 25 uA per electrode (although in rat Z2, current levels of twice this were required to have any effect). We wanted to use a minimal perturbation to ensure that the effects were as localized as possible in both time and space. We emphasize that these current levels are far lower than are typically used, for instance to evoke a sensory percept or motor response, demonstrating that the mPFC is particularly sensitive to disruption at least during our behavior.

We sometimes used the same stimulation protocol during epochs in between trials when the rats were not behaviorally engaged in the task. We never noticed any overt behavioral response to stimulation under these conditions. Besides the impairment described in the main text, the only additional effect we observed during behavior was that rats appeared to have more trouble completing the center poke. The sound does not play until the rat holds for a random duration between 250ms and 350ms. Shorter ("failed") center-pokes do not initiate a trial. Consistent with the proposed role of the mPFC in estimating temporal duration, we noticed that rats exhibited more failed center-pokes during disruption (data not shown), especially at higher current levels. Typically they did not go to the choice port after a failed center-poke (that is, the stimulation did not directly elicit a choice motion); they simply repeated the center-poke until successfully initiating a trial.

For some individual sessions, the disruption caused a significant increase in the number of "wrong-port" responses, that is, trials in which the rat went to the choice port associated with the other block (data not shown). This suggested a possible specific deficit in stimulus selection, or in the memory of the current block, but the effects were insufficiently consistent to draw firm conclusions. We only included sessions for which the performance on control trials was significantly above chance, using the same definitions of chance as we did previously for the non-stimulated animals (see: "Chance performance on the task" in the Methods section of the main text).

We observed that rats appeared to be particularly impaired on the "congruent" NOGO stimulus RIGHT+HIGH in one or both blocks (Figure S7B). This is interesting because this stimulus should, in theory, be the least ambiguous stimulus of all: it always means NOGO. For this reason, there was no significant increase in the proportion of trials on which the rat gave the response that would have been appropriate in the other block, as one might have expected were stimulus selection the only cognitive ability that was affected. Future experiments will be needed to disentangle the role of the mPFC in holding the center port, interpreting the stimulus, and producing the correct motor act.

Simulated network model

The weights from the A1 neurons to the command neurons were trained using a non-negative least squares algorithm (scipy.optimize.nnls). Each subpopulation was trained separately – the first subpopulation, which projects to two command neurons, was trained to produce the

responses appropriate for the first block; the process was repeated for the second subpopulation and the second block.

For a given SNR, as the task switch signal increases in strength, the model's performance eventually drops to 50%. This is because the size of the task switch signal begins to dominate the sensory input (the "overdriven" regime). In this regime, the model still produces responses that are appropriate to the block (*i.e.*, it does not "go to the wrong port") but the responses are no longer related to the sensory input.

When the task signal corresponding to task 1 is activated, but the model is assessed on its performance on task 2, it performs poorly (as expected). These data are shown in Figure 8B of the main text: negative values of the task signal correspond to activation of the "wrong" network. This corresponds to doing localization instead of pitch discrimination in the real task; in such a case, only 25% of trials will be correct: those which present the RIGHT+HIGH stimulus which always means NOGO.

For a network size N = 640, and at very low SNRs <1%, our model reached the overdriven regime before it ever produced good performance. We measured the minimum and maximum values of the task switch signal that produced good performance (>80%) over a range of network sizes and SNRs. Increasing the size of the network increases the working range of the model by decreasing the minimum task signal that is necessary to reach criterion performance and increasing the maximum acceptable task signal before reaching the overdriven regime (Figure S8).

Video analysis of preparatory head positioning

We recorded video of all behavioral sessions using an infrared camera. We hypothesized that the rat might use a different posture in each block since the reward ports were different. We took the video frame closest to the center of the hold period on each trial and manually scored the position of each ear. We did this by asking a human observer (CR), who was blind to the block and outcome of each trial, to click on each ear in the frame from every trial and recorded the position of the clicks. This manual scoring was time-consuming so we analyzed only the three sessions during which we recorded the most rule-encoding neurons (one session from Rats 2, 4, and 6). Using the position of each ear, and knowing that the nose was located in the center port during this interval, we were able to construct the center position and azimuthal angle of the head. We found a prominent correlation between the head angle and the block in all analyzed sessions (*e.g.,* Figure S3G, S4G).

Because head angle is correlated with block, and because we analyzed rule-encoding neurons for which, by definition, the firing rate correlated with block, it stands to reason that head angle correlates with firing rate. In an example neuron (Figure S3H, S4H), the firing rate is highly significantly correlated with block (black trend line). However, for this neuron the firing rate is not correlated with head angle within each block separately (red and blue trend lines), suggesting that this is a mere side effect of the fact that the neuron is encoding block.

We obtained these results by regressing the square root of the spike counts on each trial onto a 2-way linear model including both head angle and block as factors. We used ANOVA to calculate the explainable variance and statistical significance of head angle and block for each neuron. Importantly, this class of linear models explicitly accounts for the correlation between explanatory factors by inverting the covariance matrix. Thus, the explainable variance obtained is a measure of the information uniquely available from each factor in the model. Across our population of rule-encoding neurons, far more of the variability was explained by block than by head position or by the interaction (cross-term) between these two variables (Figure S3I, J; S4I, J). Moreover head angle was rarely a statistically significant predictor of firing rate (4/16 PFC neurons, 2/8 A1 neurons). The fact that some neurons do encode head angle, not block, according to this analysis is consistent with the known role of some neurons in this brain region and it is also a proof of principle that, in at least some cases, the head angle scoring procedure is sufficiently sensitive to uncover these effects.

We chose to focus on head angle because this was the most prominent and easily quantifiable postural difference between blocks and is therefore likely to be correlated with preparatory motor activity in general. No analysis can rule out the possibility that PFC is actually encoding some unknown, subtle difference in behavior between blocks. However, the effects we observe favor the hypothesis that activity of PFC neurons primarily encodes the current task. It is possible that a side effect of this difference in cognitive state is a block-specific difference in motor planning and execution; these motor effects would thus correlate with PFC activity even though they may not be directly encoded by PFC activity. In fact, even our neurons that seem to be encoding primarily head angle may actually be encoding cognitive state: trials on which the rat most strongly plans to perform localization might also produce the strongest preparatory actions.

We also asked whether the effects could be explained by the motion history of the animal preceding center-poke entry. We reasoned that the rat's movement history before each trial would strongly correlate with the choice on the previous trial. In particular, if the rat had just completed a successful NOGO trial, it was likely to have remained relatively motionless in the center port; on trials following a successful GO trial, it was likely to have just moved from the choice port to the center port. When we analyzed only trials following a successful NOGO and found that the rule-encoding remained the same: the firing rate was still elevated during the preferred block and suppressed during the non-preferred block (Figure S5B). This demonstrates that the rule encoding cannot be purely an effect of the motion to the center port.

Acknowledgements

The authors would like to thank Mohammad Dastjerdi and Christian Fernandez for developing the initial rodent behaviors in the lab and managing behavioral training; Alfonso Apicella for helping set up the lab; Tony Zador, Uri Livneh, Erin Rich, Nicholas Steinmetz, Anita Devineni, 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.

Chapter 3. Detailed procedure for training rats on stimulus selection

We describe here the procedure we developed for training rats to perform stimulus selection. It is presented in stages. Each stage may take multiple, in some cases many, sessions until the rat is ready to move on to the next stage. Begin each day where the previous day left off, or perhaps a small step backward if necessary.

Materials

We used a three-port behavioral chamber with three nose-pokes (left, center, and right) and two speakers mounted such that they were on the left and right side of the rat's head when it poked in the center. The left and right nose-pokes could also deliver water rewards. We used 10 uL of water per reward. We found that one hour was a reasonable amount of time to train each rat. Our rats received *ad lib* water after every behavioral session to satiation. They were then deprived of water until the next behavioral session, approximately 24 hours later.

Our entire system, both hardware and software, was modeled on that used by Dr Carlos Brody's lab at Princeton University. They have graciously shared an extensive amount of code and guidelines on their website (brodylab.princeton.edu).

Stimulus parameters

We use Fostex FT17H tweeters as speakers. These can be calibrated to produce sound between 5KHz and 50KHz.

The LEFT and RIGHT white noise bursts were 250 ms in duration. Each was played entirely from a single speaker (left or right); the other speaker remained silent. We began with random samples (*e.g.*, the "rand" function in Matlab) and applied a filter that we calculated to counteract the frequency response of the tweeter. The end result was roughly flat over the range 5KHz to 50KHz and contained very little power outside of this range. The total amount of power over this range was 60 dB SPL.

The LOW and HIGH warbles were also 250 ms in duration. Each was an FM modulated tone with amplitude 0.07 octaves, modulation frequency 10 Hz, and carrier frequency 6KHz (LOW) and 16KHz (HIGH). Each was presented at 65 dB SPL.

The stimulus pairs used in the task (LEFT+HIGH, LEFT+LOW, RIGHT+LOW, RIGHT+HIGH) were simultaneous combinations of the above sounds: the stereo waveforms were simply added together.

Pilot experiments demonstrated that rats' predilection for performing pitch discrimination or localization could be titrated by adjusting the relative volumes of these sounds. For instance, by

increasing the volume of the tones by 5 dB SPL, the rats were noticeably better at pitch discrimination and correspondingly worse at localization. The volumes stated above were chosen to make the tasks approximately equal in difficulty. Nonetheless, it was quite common for rats to be better at localization. We speculate that this is an innately easier task for them; alternatively, this may have been because we always trained the rats on localization first.

Stage 0: Initial handling

Goal: The rats should become comfortable with the experimenter.

Spend five to ten minutes getting to know the rats and letting them become accustomed to being picked up. The rats should not be water deprived until they are comfortable with the experimenter.

Stage 1: Nose-poke training

Goal: The rats should learn how to nose-poke; that the nose-pokes produce water; and that persistent nose-poking is sometimes required.

Cover the center and right nose pokes, leaving only the left exposed. Every time the rat pokes once into the left nose-poke, a water reward is dispensed. Slowly increase the number of required nose-pokes in order to receive reward, to about five or six. Once the rat has consumed approximately 1 mL of water (100 rewards), cover the left port and uncover the right port. Repeat the process.

If the rat does not realize how to nose-poke, dispense water rewards manually from the nose-pokes until it discovers the water.

Stage 2: Center-out training

Goal: The rats should learn the "center-out" structure of the task; that all trials begin by poking the center port; and that subsequently poking the side port will produce reward

Cover the right port. Leave the left and center ports uncovered. Turn on a white LED in the center port. For the duration of the stages, the white LED means that the rat can poke the center port to initiate a trial. The rats cannot yet know this, but they will hopefully be curious about the LED and will poke the center port.

As soon as they poke the center port, the LED should turn off. The rats needs to learn that this means it now needs to make a behavioral choice, *i.e.* go to the choice port. If it pokes to the left port within a certain time after the end of the stimulus (referred to hereafter as the maximum choice time, or MCT), then it will receive water at the left port. Set the MCT to 30 s at the beginning of this stage.

This can be the most frustrating stage for all concerned because the rats cannot know to perform the center-out structure until they have randomly done it by accident enough times to realize a pattern. We speculate that this notion of "action at a distance" -- that is, that center-poking enables water delivery at the choice port -- is not an innate concept for the rats.

We often found it useful to give a reward to the choice port as soon as the rat center-pokes. We also would reward successful center-out motions with multiple rewards. Typically rats perform only a few correct trials per session for the first day or two.

For particularly recalcitrant learners, water can be delivered to the center port before the rat performs a trial. This will encourage the rat to investigate the center port and increase its chances of discovering the center-out pattern.

Once the rats begin reliably producing the center-out pattern, gradually taper the MCT to 1 s. Ensure the rats are still successfully receiving reward at least 90% of the time before continuing.

Stage 3: Localization

Goal: The rats should learn that auditory stimuli will play upon poking the center port; that this is not something to be feared; that LEFT means "go left"; and that RIGHT means "nogo"

Begin the session as before, with 10-20 trials of center-out without any sound at MCT of 1 s. Now, raise the MCT to 3 s and enable the sound. Ensure that only go trials, that is, LEFT stimuli, are presented. No LOW or HIGH sounds should be presented at any point during this stage.

The first time the rat hears the sound it will be startled. The higher MCT gives it time to receive its reward even though it was startled. It is important that it knows center-out structure very well. Once it does 2-3 trials with sound, it will no longer be scared of the sound. It may be necessary to deliver extra rewards during this time. Some rats will attempt to partially center-poke in order to avoid triggering a trial. Do not reward this behavior.

Once the rat is consistently doing trials, lower the MCT again to 1 s and ensure the rat can still perform well.

At this point we introduce NOGO stimuli. Raise the MCT again, to 1.5 s or so, before doing this. Then begin to present LEFT and RIGHT stimuli randomly on each trial. We call this the "mixed" mode, as opposed to the "forced-go" or "forced-nogo" mode in which only one stimulus is presented. We do not recommend using any other mode, *e.g.* 80%/20% go/nogo or structures in which the next trial is contingent on the current trial. This will encourage the rats to learn spurious strategies. The rats should learn that all trials are independent and equally likely.

The first time the rat hears the RIGHT it may be startled. It will soon grow accustomed to it and will attempt to give the same "go left" response regardless of the sound. Here for the first time we introduce the error timeout. Set this initially to a very low amount, like 0.5 s. After each go-on-nogo error, the rat must wait this long before the center LED turns on again and the next trial may be initiated. There is no need to punish nogo-on-go errors with a timeout: the lack of reward is punishment enough.

Rats may sometimes grow discouraged that the "go left" response no longer invariably yields a reward. Maintain motivation by alternating between forced-go and mixed modes. Continue until the rats can consistently work and do not give up after nogo trials.

Next, let the rats perform this task until they begin to learn the nogo response by trial and error. The first sign is that the rats will move visibly more slowly to the choice port on nogo trials, because they know they are not going to be rewarded. Eventually they will learn that there is no need to move away from the center port at all; they simply should wait until the MCT has expired and the next trial can begin.

We generally found it necessary for the MCT to be somewhere between 1 s and 2 s. Note that decreasing the MCT will increase the difficulty of go trials (because the rat must move more quickly) and decrease the difficulty of nogo trials (because the rat doesn't have to wait as long and/or responses that are slightly too slow will be scored as correct). Do not vary the MCT more than necessary. Only lower it if the rats are performing over 85% on go trials, performing very badly on nogo trials, and are not improving. Do not lower it so far that the rats perform under 80% on go trials or they will lose motivation. Remember that the rats primarily care about the number of rewards received, not the numerical performance.

We found that most rats would learn the localization task to a reasonable level (over 80% overall, over 70% on nogo trials) in a few days to a week. Sometimes it was necessary to increase the error timeout to 2 s or 4 s, if the rats were doing large numbers of trials (over 350 in an hour) and all of the above manipulations did not work. We view increasing the error timeout as the last resort because it decreases the yield of trials.

Stage 4: Pitch discrimination

Goal: the rats should learn to go the right port; that LOW means "go right"; and that HIGH means "nogo".

This stage is quite similar to the previous stage. Begin by covering the left port only, with no sound playing, a high MCT, and with only "go right" trials.

Once the rats have learned to go right, follow the above protocol for turning the sound on, mixing in nogo trials, and gradually increasing performance.

Stage 5: Task switching

Goal: the rats should learn to switch between localization and pitch discrimination; and to achieve very high and equally high performance on both

Uncover all ports. Set the MCT a bit higher than was necessary for the individual tasks, perhaps 2 s. Alternate the trials in blocks: 40 localization, 40 pitch discrimination, *etc*. At the block changes, some rats will instantly switch because they remember the sounds. Others will persist at the previous task. If necessary, switch to "all go" mode to encourage a block switch, then return to mixed when possible.

Continue this stage until the rats are quite good at both tasks. They will need to know these basic stimuli, as well as the go/nogo structure, quite well in order to succeed at the stimulus selection task.

Follow the protocol outlined in Stage 3 for adjusting the MCT and error timeout parameters. Keep performance on go trials high, but if it is very high (above 95%) and nogo performance is very low, the MCT should probably be decreased. Do not let the go performance fall below 75%. Increase error timeout if the rats are doing many trials with poor nogo performance. We found that increasing error timeout beyond 6 s was counterproductive.

Use the "forced nogo" mode when all other strategies fail. Present only nogo stimuli until the rat stops responding. As soon as the rat gives 1-3 correct nogo responses in a row (out of frustration), then immediately give a few go trials to keep motivation up. But do not give too many go trials in a row as this is the opposite of what the rat should learn.

The general philosophy is a 50/50 compromise between adjusting the parameters based on the rat's behavior, and leaving the parameters fixed so that the rat adjusts to them. A certain amount of intuition and user intervention is necessary, though this should be minimized. Certainly the parameters should not be "overcorrected" and wildly varied within a session. Eventually they should be fixed within each session.

Stage 6: Introducing the distractor

Goal: the rats should learn how to perform the task with a distractor; to switch between the blocks even with the distractor; and to achieve an equally high performance in both blocks

We found this stage to be surprisingly easy for the rats to learn. Use a trial structure of 20 cue trials (no distractor), followed by 60 trials with distractor. Then switch to the other task, again using 20 cue trials and 60 trials with distractor.

If the rats persistently do the wrong task or go to the wrong port, switch to "all go" mode or (in extreme circumstances) give water manually to the port they are avoiding. At first, be generous

with the “all go” mode and help them to switch between the tasks. Later, avoid the “all go” mode and let the rats learn to switch by trial and error.

We found that rats very often adopted a strategy in which they would perform well at one task, then switch to an “always go” strategy in the other block. This is a simple, but undesired, form of task switching. Use the techniques above (increasing error timeout, forcing nogo) to combat this. Very recalcitrant rats can be switched to a block structure in which they only perform only their non-preferred task until they begin to perform well at it. However, we found it useful to have the rats do at least one block of each task, every day, so that they still have to switch every day.

As rats become very highly trained on this task, we typically had to lower MCT gradually. In the most extreme cases an MCT of 0.4 s was necessary because rats simply would not wait longer than this on nogo trials under any conditions.

Our best rats were able to learn to perform both tasks above 90% correct and to switch almost instantly, reaching plateau performance within 2 trials of a block change.

Why the complicated go-left/nogo/go-right structure?

As social animals who live in burrows, rats probably have the capability to identify and select important sounds (*e.g.*, vocalizations from other rats) in order to preferentially respond to them. We have found that the bottleneck in developing new behaviors is not the rat’s ability, but the researcher’s ability to incentivize the rat to perform the desired task. We have found rats to be excellent at making very difficult sensory discriminations; however, it is much more difficult to train them to perform tasks requiring cognitive flexibility, especially if there is an alternative, suboptimal strategy that works pretty well (such as always-go, as described above).

We first conceived of the task as a “two-alternative forced choice” paradigm. In that paradigm, the rat has two behavioral choices on every trial: to poke its nose in the left port, or to poke its nose in the right port. Each sound in the pair could mean either “go left” or “go right”. The problem we encountered was that rats would always select the same sound (*e.g.*, always do localization), regardless of the block. That meant that they performed at 100% in one block and at 50% in the other block, by random chance. The overall performance was 75%, more than acceptable to the rat.[4] We tried a wide variety of workarounds to avoid this but found no success. Even when we carefully selected the stimulus pairs such that the rat would perform at 0% if it selected the wrong sound, it still took several daily sessions before the rat began to select the other sound.

Our solution was to modify the task to the form described above. We found that rats learned this task much more quickly, even though the only difference was the motor act required to report a decision, an implementation detail of seemingly incidental importance. One possible

[4] Also sufficient to pass a course at Berkeley.

explanation is that, with the current version of the task, the rat never receives reward for performing the action associated with the distractor. Another possibility is that the association of a separate reward port with each task makes it cognitively easier for the rat to switch between them.

Chapter 4. Interpretation and perspective

Models of prefrontal control

I began my research with a "feedforward" hypothesis in mind. Under this hypothesis, information flows from sensory organs in the periphery, through sensory cortex, and eventually reaches the "apex" region: prefrontal cortex. At this point, a final decision is made based on all available information. The correct motor action is determined and transmitted downwards through premotor and motor areas, eventually entering the periphery again and activating musculoskeletal effectors. This hypothesis remains popular among some authors (Mante 2013, Gilbert 2002)

Lesion studies (Rich 2009, Pai 2011) showed that animals are still able to make sensory discriminations even without prefrontal cortex (though not necessarily on all tasks, especially those requiring cognitive flexibility). I modified my feedforward model with an exception for simple sensory discriminations: for instance, in tone detection, the information might follow a tighter loop directly from auditory cortex to striatum and from there to the motor effectors. This could occur without any input from PFC. I suggest the striatum here because activating this projection is sufficient to bias pitch discrimination decisions (Znamenskiy 2013).

After having taken the data, I now take a different view. I now propose that the decision is always made in sensory cortex, regardless of the task. When necessary, contextual information (*e.g.*, the current task and appropriate rules) flows downward from PFC and mixes with bottom-up sensory information in sensory cortex. In my case I suggest that these streams converge in A1 and the decision is produced there. Downstream regions like the striatum read out this decision from the activity in A1 and signal the appropriate motor act to the peripheral motor system.

I cannot definitively identify the regions involved in this model because I did not record in other structures, such as striatum, motor cortex, or higher sensory cortex. The specific proposal that I am making is that information is flowing primarily from PFC to sensory cortex, and only weakly in the other direction. The most obvious piece of evidence in support of these ideas is that I observed only weak and poorly tuned sensory information in PFC. This seems incompatible with the notion that PFC not only receives sensory information but performs the final computation on it.

Secondly, I observed a surprising amount of cognitive and motor signals in A1. In fact, I found A1 and PFC to be extremely similar in the variables they encoded. There were some exceptions: A1 encoded sounds more strongly; PFC encoded the task more strongly; but to a first pass the regions were more similar than they were different.

I proposed a quantitative model to make these ideas more explicit. In this model, contextual information from PFC is simply an additive excitatory signal that activates the appropriate subnetwork in A1. One important limitation of this model is that it can only activate

subnetworks that already exist in A1. New subnetworks can only be formed via behavioral training. A question for future work is how this model could be improved to more closely resemble the cocktail party problem. Clearly a subject can direct attention to a new speaker, whom he has never heard before, rapidly and without extensive retraining. To understand such a behavior, I would need to design a task in which the subject must generalize to a new stimulus perhaps even on every trial.

How do we know what we know in neuroscience?

Consider a simple thought experiment[5]. The subject is a trained mathematician. The task is to view a long and complex mathematical proof and to evaluate its correctness. The subject reports his decision by pressing a button if and only if he believes the proof to be true. Clearly this cognitive ability is quite elaborate and we would like to understand how the brain can perform this.

Now imagine that I have found a neuron in the subject's brain with the following characteristics:

1) It is highly correlated with the task -- i.e., the neuron fires when the proof is true; the neuron does not fire when the proof is false

2) We have the ability to perturb this neuron. When we cause the neuron to fire, the subject reports that the proof is true. When we silence it, the subject reports that the proof is false.

By the standards of the field, we have found a neuron of critical importance to mathematical proving. A variety of experiments might be suggested to understand what computations it is performing.

Now I reveal the identity of this neuron -- it is the motor neuron controlling the mathematician's index finger which pushes the button. Suddenly the results seem trivial and uninteresting. But why?

One simple answer is that we did not do certain control experiments. For instance, if a follow-up task had required the subject to report using his big toe, then our results would have disappeared. But it is easy to imagine that there is a neuron one single step upstream from the putative motor neuron, that would exhibit the above properties regardless of what effector was used. In fact, such a neuron almost certainly exists, simply because mathematicians can be trained to report their answers with any appendage without any additional effort.

[5] The original version of this was posed to me by Jack Gallant. In those days perturbational techniques were less common and I have updated it to include them.

Another control experiment is to consider error trials, but this would not change anything either. When the mathematician makes a mistake, the neuron would correlate with his response; perturbing the neuron would still affect his response. It seems extremely unlikely that we would discover an "oracular" neuron that encodes the veracity of the proof regardless of the mathematician's response.

These follow-up experiments and many more are *ad hoc*: we know that they are important, but they don't fit into the standard approach for understanding neural functioning, which is to search for physiological correlates, then to silence or to activate the relevant neurons and observe, respectively, an abolishment and a rescue of the behavior. These experiments are the accepted standard for proving sufficiency and necessity, *i.e.*, logical truth.

So in any experiment, once we have satisfied the physiological and perturbational conditions, what do we do next in order to determine what role the neuron plays in behavior? I would argue that the next step is to characterize quantitatively the computations at every point in the circuit which we believe to be necessary. We must be able to re-implement the same computations in a simulation and produce generally the same results (at least on average across trials). Only then can we answer the underlying question ("how is the computation done") in addition to the questions answered by physiology ("what is the brain doing") and perturbation ("where is the computation done").

Throughout my experiments, presented in Chapter 2, I struggled with this problem. Do our effects represent stimulus selection? Motor planning? Memory of the current block? This problem is not specific to our study. I am evoking here the words of J Erlich, M Bialek, and C Brody (Erlich 2011), who raised a very similar question about their own work, an extremely careful investigation into decision making in the rat frontal orienting field (FOF).

> There are several possible interpretations as to what component(s) of response preparation FOF neurons might encode: do they represent a motor plan? A memory of the identity of the motor plan? Attention? Intention? Our data do not discriminate between these possibilities. (Erlich 2011)

Psychology is a useful framework for understanding human (or rat) behavior by conceptualizing the aspects of cognition (attention, motor planning, stimulus processing, *etc*.), but there is no reason to assume that these concepts map on physiology in a one-to-one fashion. In attempting to do so, we risk falling into the trap of being unable to dissociate effects that may be different aspects of the same neural processing.

Instead, I propose an alternative approach. (Unfortunately it is at the limits of what is technically feasible presently.) Not only must we identify what neurons are involved (by perturbational techniques) and what activity patterns they produce (by physiological recordings); we must also identify the computations they perform. This will require measuring their inputs, their outputs, and potentially their entire internal state, and finally entering all of the data simultaneously into a model of the entire system. The model must produce the same

result as the actual brain, at least on average. This is based on the engineering doctrine: if you really understand something, you can take it apart and put it together again (and it will still work).

Assuming we could actually do this, what form would our understanding of the neural circuit take? I do not think it would change the way we conceptualize behavior, which will remain rooted in a psychological/ethological model. Nor would it change our beliefs about what individual neurons are doing, which we can already investigate with physiology. Instead it would answer an entirely different question at an intermediate level of abstraction: how do thoughts form?

It remains to be seen whether we could even understand the computations performed at such an intermediate level. It may be that there is no simpler description of the system, other than a complete one. Even with purpose-built neural networks, designed by humans *ab initio* without attempting to recapitulate an actual brain, it can be extremely difficult to understand how they work, just that they do. Nonetheless, I believe a full description of the computation is a prerequisite for understanding it.

Ruminations on perturbational techniques

Throughout the history of neuroscience, experimenters have both observed the brain and interfered with it, sometimes in the same experiment. In decades past, commonly used techniques for interfering with the brain were lesioning it (perhaps reversibly), stimulating it, applying drugs to it, and so on. The invention of *optogenetics* has made it possible to interfere with much greater resolution and control, by expressing light-gated ion channels in genetically-identified neuronal cell types, and has correspondingly increased the popularity of studies which manipulate neuronal activity and measure the effect on behavior.

I am not aware of an accepted term for the general approach of interfering with the brain. It is sometimes just called optogenetics, but this is a technique, not an approach, just as extracellular recording is a single technique within physiology. It is sometimes called “causal manipulation”, in the sense that it causally affects the brain, but I consider this a misnomer because it seems to imply that other manipulations somehow violate causality, a logical impossibility. I prefer the term *perturbation* (after Erlich 2011, though it has surely been used before).

Correlational physiology explores the relationship between neural activity and behavior under the endogenous state, when every attempt is made to interfere as little as possible. Perturbation probes this relationship by manipulating the state and observing the results. Both are critical for understanding the functioning of the circuit.

What is the logical basis of the perturbational approach? It is typically understood in terms of *activation* and *silencing* which are understood to be opposing manipulations that should yield opposing behavioral results. If activating a neuron (or set of neurons, or brain area) N produces

a behavior B, this is taken to be evidence that "the neuron is sufficient for the behavior" (N implies B). If silencing a neuron abolishes a behavior then this is taken as evidence that "the neuron is necessary for the behavior" (not N implies not B)[6]. I call this the *classical interpretation* of perturbation experiments. The combination of these two logical statements is B⇔N, that is, a logical identity between the firing of a neuron and the behavior.[7]

The classical interpretation is initially satisfying. Yet there are some extremely simple examples in which it completely breaks down. Consider these three examples:

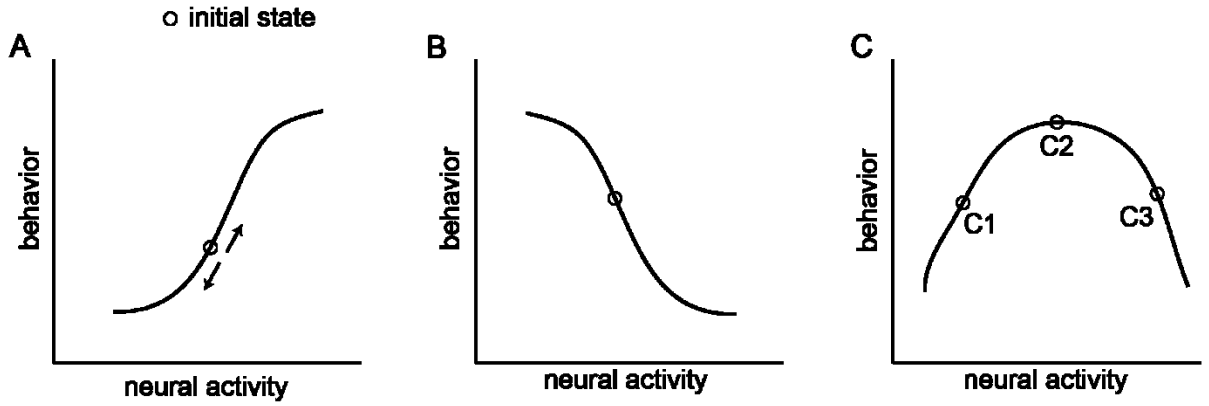

Figure 4.1 - Three example relationships (cases A, B, and C) between neural activity and behavior. Circles represent the initial (endogenous) state of the system. Perturbing the neural activity moves the state rightward or leftward along the curve.

Case A: The simple case. Behavior is positively correlated with neural firing. In this case, activating the neuron increases behavior (sufficiency); silencing the neuron suppresses it (necessity). Conclusion: high firing in this neuron is both necessary and sufficient for the behavior.

Case B: Only slightly more complex. Behavior is negatively correlated with neural firing. In this case, activating the neuron suppresses the behavior and silencing the neuron enhances it. What do we conclude? It stands to reason that we should conclude the opposite of (A): low firing is both necessary and sufficient for behavior.

The experiment that abolishes behavior must always be the necessity test, so in case (B) the activation experiment is the necessity test, in apparent violation of the classical interpretation. This is not at all an unusual case however; for instance, this neuron may be inhibiting the neuron in case A.

---

[6] This is logically equivalent to B implies N. The phrasing may suggest that the behavior preceded the neural activity but this is incorrect: logical implications do not have a temporal order.

[7] Typically ignored in these arguments is another variable C: the entire *context* in which the experiment occurs. This includes the subject's environment, but also the functioning of every other neuron and the rest of the body. Strictly speaking, no neuron can be sufficient for any behavior, because firing this neuron in a paralyzed subject (or isolated in a dish) would never produce the behavior. It is also likely that vast swathes of the brain and body are necessary for almost every behavior, for instance the brainstem respiratory systems.

Case C: The *inverted-U* relationship. In this case, the results of the experiment depend on the initial conditions. At point C1, the results match case (A); at point C3, the results match case (B); at point C2, both perturbations suppress the behavior. What do we conclude?

Clearly medium firing is both necessary and sufficient for behavior, but we could only know this if we had access to preparations beginning at all three points. In the worst case, if we had access only to preparations near point C2, then all manipulations would perturb behavior and we could only conclude that the endogenous firing patterns were necessary for behavior; never being able to enhance the behavior, we could not conclude anything about sufficiency.

Any more complex relationship between neuronal firing and behavior can only make these matters worse. Moreover, it is reasonable to expect (by both evolutionary and efficiency arguments) that the brain typically operates close to a local optimum similar to the troublesome point C2. In such cases, it is unlikely that activating or silencing neurons will move the system toward a more optimal state. Both manipulations are likely to "disrupt" activity – to move it further from its optimal state. (Unfortunately, such results are likely to be perceived as "null" or "conflicting" and remain unpublished.)

The common phrase "the neurons are necessary and sufficient for the behavior" is shorthand for "increased firing of the neurons is necessary and sufficient for behavior". And in cases like (A) above where the classical interpretation holds true, that conclusion is indeed valid. These are exactly the circumstances under which perturbation techniques have been the most illuminating (Lin 2011). This argument does not invalidate those results by any means; it simply shows that there are quite simple examples where the endogenous firing pattern of the neuron is both necessary and sufficient for behavior but the classical interpretation will yield incorrect or confusing results.

In summary:

1) Silencing is not the only way to show that a neuron's firing is necessary for a behavior. Any manipulation that abolishes behavior (silencing, activation, even random disruption) are equally valid.[8]
2) A neuron's observed firing pattern can be sufficient for the behavior, but in such a way that neither stimulation nor any other manipulation could reveal it.[9]

Throughout, I have only considered the case where the neuron is actually controlling behavior. Under the null hypothesis, perturbation techniques would produce a null result. This is a very useful conclusion, subject of course to the usual caveats of null results.

[8] Of course, the type of manipulation constrains our interpretation of what exactly about the neuron's firing is necessary. But such interpretations must be drawn carefully anyway. For instance, we may mistakenly conclude we are in case A when we are actually at point C1.

[9] If we already know the full dynamics of the network and could calculate the optimum exactly, we might be able to do a very impressive sufficiency test. This would be equivalent to knowing that we were close to point C2, measuring the exact location and curve, and delivering exactly the right perturbation to reach C2 exactly.

I make one final technical note. Thus far I have assumed that activating and silencing neurons do exactly that. However, perturbing a system typically results in a rebound in the opposite direction, possibly because the network is attempting to regain its desired state. For instance, activating a system with electrical or optogenetic perturbations can paradoxically produce widespread and long-lasting suppression a bit later. These are typically viewed as undesirable side effects to be controlled via clever experimental design, but if one is willing to consider all of these manipulations as necessity tests via disruption, it becomes a non-issue or even a strength.

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. U. S. A. 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. U. S. A. 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.

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. U. S. A. 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. U. S. A. 2012, 5–10.

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.

Feldman, D.E. (2012). The Spike-Timing Dependence of Plasticity. Neuron 75, 556–571.

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

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.

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

Gilbert, S.J., and Shallice, T. (2002). Task switching: a PDP model. Cogn. 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. Meth. 155, 207–216.

Hocherman, S., Benson, D.A., Goldstein, M.H., 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, 14.

James, W. (1890). The principles of psychology (Vols. 1 & 2). New York Holt 118, 688.

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.

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.

Lakatos, P., O'Connell, M.N., Barczak, A., Mills, A., Javitt, D.C., and Schroeder, C.E. (2009). The leading sense: supramodal control of neurophysiological context by attention. Neuron 64, 419–430.

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.

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

Lin, D., Boyle, M.P., Dollar, P., Lee, H., Lein, E.S., Perona, P., and Anderson, D.J. (2011). Functional identification of an aggression locus in the mouse hypothalamus. Nature *470*, 221–226.

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. Current Biology: CB 19, R1024–7.

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.
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.
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.
Posner, M.I. (1980). Orienting of attention. Q. J. Exp. Psychol. 32, 3–25.
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.
Rothschild, G., Nelken, I., and Mizrahi, A. (2010). Functional organization and population dynamics in the mouse primary auditory cortex. Nat. Neurosci. 13, 353–360.
Sasaki, R., and Uka, T. (2009). Dynamic readout of behaviorally relevant signals from area MT during task switching. Neuron 62, 147–157.
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., 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. Neuro. 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.

Zénon, A., and Krauzlis, R.J. (2012). Attention deficits without cortical neuronal deficits. Nature 489, 434–437.

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.