Title

Competing Neural Ensembles in Motor Cortex Gate Goal-Directed Motor Output

Journal

Neuron, 88(3)

ISSN

0896-6273

Authors

Zagha, Edward
Ge, Xinxin
McCormick, David A

Publication Date

2015-11-01

DOI

10.1016/j.neuron.2015.09.044

Peer reviewed

HHS Public Access

Author manuscript

Neuron. Author manuscript; available in PMC 2016 November 04.

Published in final edited form as:

Neuron. 2015 November 4; 88(3): 565–577. doi:10.1016/j.neuron.2015.09.044.

Competing neural ensembles in motor cortex gate goal-directed motor output

Edward Zagha[1], **Xinxin Ge**[1], and **David A. McCormick**[1,2]

[1]Department of Neurobiology, Yale School of Medicine, New Haven, Connecticut, USA

[2]Kavli Institute for Neuroscience, Yale School of Medicine, New Haven, Connecticut, USA

Summary

Unit recordings in behaving animals have revealed the transformation of sensory to motor representations in cortical neurons. However, we still lack basic insights into the mechanisms by which neurons interact to generate such transformations. Here, we study cortical circuits related to behavioral control in mice engaged in a sensory detection task. We recorded neural activity using extracellular and intracellular techniques and analyzed the task-related neural dynamics to reveal underlying circuit processes. Within motor cortex, we find two populations of neurons that have opposing spiking patterns in anticipation of movement. From correlation analyses and circuit modeling, we suggest that these dynamics reflect neural ensembles engaged in a competition. Furthermore, we demonstrate how this competitive circuit may convert a transient, sensory stimulus into a motor command. Together, these data reveal cellular and circuit processes underlying behavioral control, and establish an essential framework for future studies linking cellular activity to behavior.

Introduction

The cerebral cortex modulates behavior according to momentary goals. In the context of sensory detection, this requires selectively linking sensory stimuli with appropriate motor responses. The neural correlates of sensory detection in the sensorimotor system have been studied extensively by Romo and colleagues in non-human primates (de Lafuente and Romo, 2005, 2006). These investigators show that neural activity evolves from representing the stimulus to representing the decision and action, and that this evolution unfolds both in time and cortical space, with frontal cortical regions linking the stimulus to the decision. However, many questions remain about how such neural representations are generated and

Correspondence should be addressed to D.A.M. (david.mccormick@yale.edu).

Author Contributions

The experiments were designed by EZ and DAM. EZ and XG conducted the recordings, analyses and computational modeling. EZ and DAM wrote the manuscript. All authors commented on the manuscript.

The authors declare no competing financial interests.

transformed within and between cortical areas. In this study, we explore the neural mechanisms in mouse motor cortex that gate behavior during a sensory detection task.

Foundational stimulation and observational studies identified motor cortex as the region most directly related to the initiation and representation of movements. Since then, many details have emerged about the relationships between cortical activity and movement parameters (Bruce and Goldberg, 1985; Georgopoulos et al., 1982; Kalaska et al., 1989). In the rodent whisker system, neuronal spiking in primary motor cortex (M1) correlates with the initiation and slow modulation of whisking and licking (Carvell et al., 1996; Huber et al., 2012). However, even during non-movement periods motor and pre-motor cortices generate robust and specific activity, which have been considered ‘preparatory’ or ‘orienting’ (Churchland et al., 2006; Erlich et al., 2011; Tanji and Evarts, 1976). Motor cortex neurons with different activity patterns may independently represent different task or movement parameters. However, we explore an alternative possibility, that cortical neurons with different activity patterns are interacting components of a larger functional circuit that controls the initiation of behaviorally relevant movements.

Behavioral control requires the correct timing of both movement initiation and movement suppression. The neural activity underlying movement initiation is most commonly represented as accumulation-to-bound in higher-order and motor cortical areas (Cook and Maunsell, 2002; Gold and Shadlen, 2007; Hanes and Schall, 1996; Schall and Thompson, 1999). In this framework, motor command neurons trigger movement initiation once their activity surpasses a spike rate threshold. Alternatively, inhibitory control is the ability to suppress movement generation at times when inappropriate or non-productive. Dysfunctions of inhibitory control contribute to impulsivity which is prominent in many psychiatric and neurologic disorders (Robbins et al., 2012; Schachar et al., 1995). In laboratory settings, many behaviors require inhibitory control by including periods in which animals must withhold behavior (e.g. No-Go or delay between stimulus acquisition and response window). While behavioral and imaging studies reveal an essential role for frontal cortex in inhibitory control (Sebastian et al., 2014), the neural mechanisms underlying this executive function are largely unknown.

In this study, we trained mice to respond (lick) selectively to somatosensory (target) stimuli and withhold licking at all other times. In expert mice, we characterized the role of motor cortex in this task. We find that suppressing motor cortex alters task behavior by increasing inappropriately timed responses (e.g. a decrease in response suppression). From extracellular and intracellular neural recordings, we find two populations of neurons that enhance or suppress their activity in anticipation of movement. Rather than acting independently, our analyses suggest that the anti-correlated spiking between these populations emerges from competition via bidirectional inhibition. Furthermore, this competition enables the circuit to suppress inappropriately timed responses, yet also act as a switch that converts a transient, sensory stimulus to a motor command.

Results

Behavioral task

We progressively trained wild type mice in a somatosensory detection task (Figure 1, see Methods). Water restricted mice were head-fixed in a behavioral setup (Figure 1A). Piezo-controlled paddles were placed in the whisker fields bilaterally, each contacting multiple whiskers. Target whisker deflections consisted of rapid (10 ms) caudal deflections of either paddle. To receive a sugar water reward, mice had to lick a lickport following a whisker stimulus within 1.5 seconds (hit). To discourage impulsive behavior, all licking outside of the post-target interval was punished by aborting the current trial. In addition, in 80% of trials we first presented an auditory stimulus (tone, Figure 1C). If mice withheld licking following the tone (correct rejection) they were presented with a target stimulus and opportunity for reward 1–3 seconds later. Licking to the tone (false alarm) was punished by aborting the trial. We trained 22 mice to expert performance over the course of this study (Table S1). Initially, mice were highly impulsive and would frequently lick to both tone and target stimuli. With sufficient training over 2–3 weeks, licking behavior and task performance improved significantly (d′, 2.5+/−0.1; hit rates, 86+/−2%; false alarm rates, 10+/−1%; n=22) (Figures 1E and 1F). Mice were considered expert once they achieved a d′ greater than 2, corresponding to an average hit and false alarm rate of, e.g., 85% and 15%, respectively.

Expert mice converged to a similar behavioral strategy to solve the task. Mice withheld whisking prior to the target stimulus, and received the stimulus as a passive whisker deflection. In a subset of mice we quantified whisker pad EMG during the task. We observed an 8 fold increase in EMG amplitude post-target compared to pre-target ($p<10^{-13}$ for each mouse, n=3) (Figures 1D and S1A). Following whisker deflection we observed a gradual ramp up of EMG activity beginning 117+/−21 ms after target onset, which preceded repetitive whisking and licking of the lickport. Average reaction time, defined as the period from target onset to first lick contact, was 519+/−28 ms (n=22) and was highly stereotyped for each expert mouse within individual behavioral sessions (Figure S1B).

To determine the contribution of motor cortex for this task, we suppressed motor cortex bilaterally during task performance. Motor cortex suppression significantly altered task performance by increasing false alarms (149+/−27% increase in false alarm rate, $p<0.01$, n=7 mice, Figure 1G). Motor cortex suppression was accomplished both by muscimol application (155+/−33% increase in false alarm rate, $p<0.05$, n=5 behavioral sessions in n=5 mice, Figure S1C, D) and optogenetic activation of PV-containing GABAergic interneurons (132% increase, n=7 behavioral sessions in n=2 mice, $p<0.05$ for each mouse, Figure S1J). Thus, motor cortex appears to be necessary for inhibitory control by withholding action when it is not rewarded. In contrast, suppression of primary somatosensory cortex (S1, unilateral muscimol application) resulted in reductions in hit rate for contralateral whisker stimuli (60+/−8% average reduction in hit rate, $p<0.05$, n=3 mice, Figure S1E–G). Hit rates for ipsilateral stimuli and false alarm rates were unaffected, suggesting specific deficits from S1 suppression in detection rather than arousal or task engagement.

Single unit spike patterns

To probe the neural signals related to behavioral control, we recorded neural activity in expert animals in motor cortex during the task (centered on M1, see Methods). We recorded spiking activity from well-isolated single units (101 units total, 92 with sufficient spiking activity for further analyses [firing rate > 0.1 Hz], 44 recorded by loose patch and 48 by multielectrode array; all recordings were from layer 5). One subset of M1 neurons rapidly (within 20 ms) enhanced spike rate following the target stimulus (Figures 2A and 2C). Increases in spike rate of these neurons preceded the initiation of whisking and licking, similar to previous recordings of M1 neurons (Carvell et al., 1996). We refer to these neurons as the 'post-target enhanced' population. In contrast, a second subset of M1 neurons showed a complimentary spiking pattern. This population of neurons exhibited a rapid (within 40 ms) spike rate suppression following target onset (Figures 2B and 2D). Spike rate suppression typically lasted hundreds of milliseconds and until the onset of licking. We refer to these neurons as the 'post-target suppressed' population. From simultaneous recordings (Figure 2E) we find that these spike patterns co-occur during behavior. Furthermore, these two response patterns were stable, such that enhanced (suppressed) neurons continued to show post-target enhancement (suppression) for the duration of each recording.

We analyzed changes in spiking in multiple ways. First, we determined the proportion of neurons with significant changes in firing rate across the target on hit trials, as compared to other periods during the task. We found that the firing rate of 61% (n=56/92) of all active neurons were significantly ($p<0.05$) modulated across the target, 41% (n=38) enhanced and 20% (n=18) suppressed. This was much greater than the proportion of these same neurons with significant firing rate modulation during the inter-trial interval ('spontaneous' 9% modulated, $p<0.0001$ compared to target) or across the tone for correct rejection trials (18% modulated, $p<0.0001$ compared to target). To further characterize these activities, we calculated the relative change in firing rate (FR) of all active neurons across the target stimulus [peri-target ΔFR index = $(FR_{post}-FR_{pre})/(FR_{post}+FR_{pre})$] (Figure 2F–I). We found that the distribution of peri-target ΔFR indices can be fit by the sum of two Gaussians (Ashman's D=2.81) (Figure 2I). This is in contrast to the singly peaked distributions during the inter-trial interval (D=0.24) or across the tone (D=0.09) (Figures 2F and 2G). These analyses indicate that on hit trials the target stimulus evokes significant spike rate enhancement or suppression across most layer 5 M1 neurons.

We further characterized the spike shape, spike rates and laminar profile of the modulated units. We identified putative fast spiking units (FSU) based on spike waveform (n=13, 14% of total; Figures S2A and S2B). FSUs had higher pre-target spike rates than putative regular spiking units (RSU) (FSU: 21.2+/−5.1 Hz, RSU: 9.9+/−1.0 Hz, $p<0.001$), which is a general characteristic of FS interneurons *in vivo* (Beloozerova et al., 2003; Gentet et al., 2010). FSUs had peri-target ΔFR indices ranging from −1 to 0.9 (Figure S2C), suggesting that these interneurons are components of both populations (n=2 significantly suppressed, n=8 significantly enhanced). Amongst the RSUs, neurons that were significantly enhanced or suppressed across the target had similar pre-target spike rates (suppressed: 8.9+/−1.4 Hz, n=16; enhanced: 12.9+/−2.3 Hz, n=30, p=0.23). Following the target, there was a large divergence of average firing rates between these populations (suppressed: 2.8+/−0.8 Hz;

enhanced: 29.9+/−4.3 Hz; p<0.0001) (Figure S2G). Enhanced and suppressed populations contained units with broadly distributed spike widths (Figures S2E and S2F), indicating that both of these populations contain fast spiking and regular spiking neurons. Additionally, enhanced and suppressed neurons were intermixed within layer 5 (Figure S2D), and often recorded simultaneously from the same electrode.

Considering previous studies of rodent M1 activity during whisking and licking [(Carvell et al., 1996; Hill et al., 2011; Huber et al., 2012; Zagha et al., 2013), see Discussion], we had not expected a robust population of M1 neurons undergoing prolonged suppression prior to movement. To determine whether prolonged post-target suppression was a feedforward sensory response, we delivered identical whisker stimuli to naïve, anesthetized mice while recording M1 single unit responses (n=43 neurons) (Figures S3A and S3B). M1 responses under anesthesia had a rapid onset (within 20 ms, similar to recordings from behaving mice), consistent with the anatomy of robust monosynaptic inputs from S1 to M1 (Miyashita et al., 1994). However, the sensory responses were transient, lasting approximately 70 ms, in contrast to the sustained responses observed during the task (Figures S3C and S3D). Moreover, significant responses under anesthesia were exclusively enhancing (38/43 neurons significantly enhanced, 0/43 neurons significantly suppressed). Thus, the prolonged post-target enhancement and suppression we observed in mice performing the whisker touch detection task do not simply reflect passive encoding of the target stimulus.

Next, we analyzed the average neural activity of target-modulated neurons across various task periods and conditions (Figures 3 and S3). The population neural activity was relatively stable across tone presentation for correct rejection trials (i.e. no licking response to the tone presentation; Figures 3A and 3F). This is in contrast to the rapid and robust divergence of spiking following the target stimulus on hit trials (Figures 3B and 3G). Enhancement and suppression of spike rates anticipated the motor response, as observed by aligning neural activity to the reaction time of each trial (Figures 3D and 3I). Further analyses suggest that many of these spiking signals are related to an upcoming motor response. First, post-target suppression was absent on miss trials (59% suppression on hit trials, p<0.001; 13% suppression on miss trials, p=0.17) (Figure 3H). While there was a robust post-target response in enhanced neurons on miss trials (54% enhancement, p<0.0001), this was significantly smaller than on hit trials (129% enhancement on hit trials; p<0.01 comparing enhancement on hit vs miss trials, 500 ms pre-target vs 500 ms post-target, as plotted in Figure 3C). Second, enhancement and suppression anticipated licking, regardless of whether this was post-target or off-target (false alarms and spontaneous licking) (p<0.01 for enhancement and suppression, post-target and off-target) (Figures 3E, 3J, S3C–D).

Correlations in neural activity predict competition between neural ensembles

Most of the neurons we recorded from in layer 5 of M1 undergo enhancement or suppression in anticipation of movement. These spike patterns may be generated independently, or they may reflect two components of a larger circuit. To distinguish between these possibilities, we began with the observation that the spiking of these populations appear to be anti-correlated prior to movement (Figures 3, S3C–D and S4A). In neural circuits, anti-correlated spiking may reflect anti-correlated inputs to each population

and/or local inhibitory interactions between populations (Figure 4). Unlike input correlation, lateral inhibition produces anti-correlation with a temporal lag due to synaptic delays. Therefore, we employed Granger causality, an analysis of forward prediction, to assess the possible presence of lateral inhibition.

Before analyzing the neural data, we simulated two-population mean-field circuit models (Wong and Wang, 2006) with varying input correlation and lateral inhibition (Figure 4A–D, see Methods). We verified that anti-correlated spiking can result from anti-correlated inputs or lateral inhibition. However, anti-correlated spiking and Granger causality was only observed with lateral inhibition (Table S2).

We then conducted the same correlation analyses on our simultaneously recorded units (from n=3 silicon probe recordings in three mice) to determine which model is most consistent with our neural data. We quantified the spike rate cross-correlation between pairs of simultaneously recorded enhanced (Enh) and suppressed (Supp) units across each full recording. We found that 59% of pairs showed significant spike-spike anti-correlation (mouse 1: 10 of 12 pairs; mouse 2: 11 of 21 pairs; mouse 3: 3 of 8 pairs; $p<0.05$ for each mouse), and across all pairs there was a small but significant negative correlation ($r= -0.019+/-0.006$, $n=41$, $p<0.01$ compared to zero). The anti-correlation was also significant ($p<0.05$) when we subtracted correlation coefficients from a 0.5 second spike-jittered cross-correlation, to remove slow correlations (see Figure 4E). In contrast, spike rate correlations between enhanced neurons (Enh-Enh) were positive ($r=0.20+/-0.01$, $n=64$, $p<10^{-21}$ comparing Enh-Supp and Enh-Enh).

We also tested whether Enh-Supp anti-correlations are time-locked to target stimulus onset. We verified that Enh-Supp pairs are significantly anti-correlated if we restrict our analyses to peri-target epochs ($r_{target}=-0.013+/-0.005$, $p<0.05$ compared to zero). Scrambling trial order (e.g., comparing neuron A trial 1 with neuron B trial 8, aligned to target onset) significantly reduced anti-correlation magnitude ($r_{scrambled}= -0.004+/-0.001$, $p<0.05$ comparing r_{target} to $r_{scrambled}$). These data indicate that Enh-Supp anti-correlations are not simply time-locked to stimulus onset, but reflect trial-by-trial variations in neural processing.

While the amplitudes of the pairwise correlation values are small, we believe that these values underestimate the antagonism between these populations during behavior. First, individual neurons have highly variable spike rates and timings; if we pool spiking across simultaneously recorded enhanced and suppressed neurons, we calculate correlation values of $-0.041+/-0.015$ (peri-target). Second, we find that correlation structure is highly state-dependent (Figures S4A–D), such that during strong low frequency LFP oscillations, the anti-correlated behavior between enhanced and suppressed populations breaks down and is replaced with positive correlations (Figure S4B). If we restrict our analyses to hit trials, either in anticipation or during licking, we observe stronger anti-correlations (Figure S4D). Third, weak pairwise correlation values, within the range presented here, may reflect quite strong correlations across a population of neurons (Schneidman et al., 2006).

Next, we conducted Granger analyses from each post-target enhanced and post-target suppressed pair. We found that 93% of Enh→Supp and 85% of Supp→Enh comparisons were Granger causal (n=41 pairs total; mouse 1: 10 of 12 E→S, 7 of 12 S→E; mouse 2: 20 of 21 E→S, 20 of 21 S→E; mouse 3: 8 of 8 E→S, 8 of 8 S→E; $p<0.001$ for each direction in each mouse). To determine the sign of forward prediction, we calculated the average of the regression coefficients of the VAR model. Enh→Supp and Supp→Enh coefficients were significantly negative (Enh→Supp: −0.0064+/−0.0007, $p<10^{-10}$ compared to zero; Supp→Enh: −0.0142+/−0.0017, $p<10^{-9}$ compared to zero). While 100% of Enh→Enh comparisons were Granger causal (n=128), the regression coefficients of these models were significantly positive (0.0251+/−0.0014, $p<10^{-25}$ compared to Enh→Supp or Supp→Enh) (Figure 4F). Anti-correlation and bidirectional Granger causality supports bidirectional inhibition (model in Figure 4D), and more generally supports models of competition between enhanced and suppressed neuronal ensembles (Figure S4E–G, see Discussion).

Relationships between spike patterns and motor gating

As described above, the correlation structure of our neural data suggest that the post-target enhanced and suppressed populations are interacting by competition via mutual inhibition. To assess how this circuit structure may influence task-related responses, we simulated the detection task conditions in the competitive ensemble model (Figures 5A–D and S5). We added a small bias current to ensemble 2 at the start of the simulation (I_2), and implemented the sensory stimulus as a transient input to ensemble 1 (I_1). As previously described (Wang, 2002; Wong and Wang, 2006), this architecture forms a winner-take-all circuit in which one ensemble is highly active and suppressing the other. Accordingly, we found that if the stimulus is small, ensemble 1 responds with a transient increase in spiking followed by decay back to rest (Figure 5C, black trace). If the stimulus is sufficiently large, ensemble 1 overcomes the lateral inhibition from ensemble 2 and maintains an increased spike rate, while suppressing ensemble 2 (Figures 5C and 5D, colored traces). Thus, the competitive ensemble model is able to convert a transient stimulus into a sustained response due to the intrinsic dynamics of the circuit.

As shown in Figure 5C–F, the transient and sustained responses of the model data share remarkable similarities with the neural activity on miss and hit trials in behaving mice, respectively. Miss trials are associated with transient increases in the enhanced population (Figure 5E, black trace), similar to the transient responses in the model (Figures 5C, black trace). Also, miss trials are associated with little change in the suppressed population (Figure 5F, black trace), similar to the model (Figure 5D, black trace). Hit trials are associated with sustained increases in the enhanced population (Figure 5E, red trace) and sustained decreases in the suppressed population (Figure 5F, blue trace), similar to sustained transitions observed in the model (Figures 5C and 5D, colored traces). In the model, the transient stimulus to ensemble 1 (‘sensory’, Figure 5B) shifts the network away from the stable pre-target state. At that point, the network either relaxes back to the pre-target state (‘miss’) or diverges from the pre-target state and follows the attractor basin to the state in which ensemble 1 is dominant over ensemble 2 (‘hit’, Figure 5B) (see also Figure S5). Notice that the ‘hit’ trajectory follows a curved path which causes a dip in ensemble 1 activity prior to diverging (Figure 5B, green asterisk). Remarkably, this feature is present in

both the model and neural data (Figures 5C and 5E, green asterisks). These analyses suggest that the neural responses associated with hit trials can be understood as state transitions, or switches, of a competitive ensemble circuit.

We next explore how sensory and motor parameters are represented within this competitive process. We had assumed above that the rapid post-target response (Figure 5E) was a sensory representation. Here we provide further evidence for this, showing that the rapid post-target response aligns to the stimulus onset irrespective of reaction time (Figures 6A–C). Of our post-target enhanced neurons, we identified 25 (of 38, 66%) units with rapid (20–130 ms, Figure S6D) post-target responses, which were present on both hit and miss trials or hit trials with fast and slow reaction times (see Methods).

To identify neural correlates of the motor command, we aligned spiking to the reaction time of each trial, and identified units with peak activity 50–150 ms prior to the first lick contact. We identified 16 (of 38, 42%) units with this profile (Figure S6E). Many of these units (10 of 16) also had robust sensory responses (as defined above), and therefore displayed both sensory and motor representations. The remaining units (6 of 16) did not have robust sensory responses. As seen in Figure 6D, these neurons displayed a small post-target response and an apparent ramp to peak activity. The ramp is absent or highly delayed on miss trials (Figure 6D). Furthermore, the peak activity aligns to the motor response (Figure 6F), rather than stimulus onset (Figure 6E). We hypothesize that these neurons trigger a motor command by accumulation-to-bound.

Overall, we hypothesize that the ‘sensory’ enhanced neurons are primarily engaged in competition with the suppressed neurons; following a successful network transition, tonic activity in the ‘sensory’ enhanced population can then drive the ‘ramping’ enhanced population to trigger movement (see Discussion). Relating these dynamics to the model (Figure 5B), we propose that the sensory-to-motor transformation unfolds as the network proceeds through the ‘hit’ trajectory, by sequentially activating partially-overlapping ensembles of neurons.

Mechanisms regulating task-related spiking

We next sought to determine the network and cellular processes regulating spiking during the task. We started with a characterization of cortical state, to determine the general temporal structure of network activity. Recent studies have argued that the slow, oscillatory state dominates cortical activity in awake, stationary mice (Crochet and Petersen, 2006; Poulet and Petersen, 2008). Interestingly, we find that during long periods (tens of minutes) of optimal task performance, M1 may be persistently activated, even in the absence of sensory stimuli or overt movements. The activated dynamics appeared as (extracellular) lack of slow, oscillatory LFP fluctuations, tonic multiunit spiking and (intracellular) sustained membrane potential (Vm) depolarization without transitions to the Down state (Figures 7B and 7E; n=9 LFP/MUA recordings and n=18 intracellular recordings). These sustained depolarizations were observed in all post-target enhanced and suppressed neurons, and therefore may reflect global modulation of M1 activity.

Alternatively, synchronized or slow, oscillatory states were present towards the end of behavioral sessions coinciding with miss trials, and infrequently expressed for short duration (0.5–2 s) during the inter-trial intervals. Slow, oscillatory states appeared as (extracellular) 3–8 Hz LFP fluctuations, multiunit spike phase locking to the negative phase of the LFP and (intracellular) large (15–20 mV) 3–8 Hz Vm fluctuations to the Down state (Figures 7A, 7C, 7D, 7F, 7I). Accordingly, we found that 3–8 Hz LFP power was significantly increased preceding miss compared to hit trials (52+/−14% increase on miss trials, $p<0.05$, n=9; Figures 7G and S7B). Currently we do not know why these slow rhythms impede behavior. However, we do observe that during these slow oscillations, enhanced and suppressed neurons are positively correlated, in stark contrast to the anti-correlated spiking observed during hit trials (Figure S4). Thus, there appears to be a disorganization of the task-relevant functional circuitry during the slow rhythms (Figure S6F). Importantly, we did observe sensory responses to target stimuli on miss trials (Figure S7A), indicating that misses are not simply due to lack of feedforward sensory inputs and consistent with our single unit data (Figures 3C and 6A).

We further explored the subthreshold dynamics leading to spike generation (Figure 8). Repetitive spiking could be driven by a sustained depolarization above spike threshold which is converted to repetitive spiking by intrinsic properties, or by transient depolarizations that briefly cross spike threshold and restrict spiking to specific temporal windows. We find that spiking is initiated by large (5–15 mV), transient depolarizations that emerge on top of a sustained Vm depolarization (Figures 8A, 8B, 8D). We observed these transient depolarizations in the spike-triggered average Vm for both enhanced and suppressed neurons (Enh: 8 mV depolarization in 60 ms rise time; Supp: 7 mV depolarization in 100 ms rise time; Figures 8E and 8H). Notably, these dynamics are unlike the phasic slow, oscillatory periods, which include Down states to −80 mV (Figure S8B).

Following the target stimulus we observed robust, ensemble-specific differences in subthreshold activity. The suppressed neurons (n=13) reduced spiking primarily due to elimination of the transient depolarizations (Figures 8A, 8B, 8F, 8G). This is reflected in reduced Vm variance during suppressed (post) compared to active (pre) periods, despite a similar mean Vm (Vm variance: pre: 7.7 +/− 0.9 mV^2, post: 2.9+/−0.4 mV^2, $p<0.0001$; Vm mean: pre: −66.1+/−1.4 mV, post: −66.9 +/− 1.4 mV, $p=0.14$). Reductions in Vm variance could result from shunting inhibition and/or suppression of pre-synaptic neurons leading to reduced synaptic inputs. For enhanced neurons (n=5), we observed a significant mean Vm depolarization during the active (post) periods (4.4+/−1.7 mV depolarization, $p<0.05$) (Figures 8C, 8D, 8I, 8J). This may be due to the summation of multiple transient inputs, tonic inputs related to movement, or both. Overall, we find that during task engagement there are at least two extrinsic components regulating Vm. There is a global, tonic component that maintains Vm out of the Down state and a transient (presumably ensemble-specific) component that initiates spiking in discrete intervals.

Discussion

In this study, we trained mice to selectively respond (lick) to somatosensory (whisker touch) stimuli. In expert mice, we recorded neural activity in layer 5 of M1, and characterized

neural dynamics at cellular, network and subthreshold resolution. We find that representations in M1 are highly redundant for this task, with most neurons displaying rapid, robust and sustained increases (enhancement) or decreases (suppression) of spike rate following the target stimulus. Amongst M1 neurons, we find that a transient sensory response is converted to a widespread and prolonged change in neural activity in anticipation of movement. Rather than independently encoding movement parameters, our analyses suggest that enhanced and suppressed neurons are engaged in a competition, and that the results of this competition gate behavior.

Classical studies of M1 activity during free whisking (Carvell et al., 1996) showed that most M1 neurons are highly positively correlated with whisking initiation (93% positively correlated, 3% [1 out of 30 units] negatively correlated). Since then, studies have shown strong positive correlations of M1 neurons with whisking and licking during active sensing (Huber et al., 2012; Petreanu et al., 2012), or modest changes in individual neurons, leading to no change in the population spike rate during free whisking (Hill et al., 2011). In contrast, we report robust and sustained changes in spiking prior to movement onset, including both enhancement and suppression. Lack of identification of a suppressed population in previous studies may be due to differences in recording method (electrical recordings versus calcium imaging) or focus on sustained movement rather than movement initiation. Alternatively, the suppressed population may reflect an adaptive response to our behavioral task. Unlike previous studies of rodent M1, our task encouraged strict suppression of motor output for most of the behavioral session. To optimize detection of transient whisker deflections (Ollerenshaw et al., 2012) and prevent aborting a trial without opportunity for reward, expert mice withheld whisking and licking throughout the pre-target intervals. Therefore it is possible that post-target suppression is a learned activity pattern, reinforced by prolonged training in our task.

Our suppression data suggest that M1 projections actively suppress motor commands initiated in other cortical or subcortical regions. This is supported by our finding that M1 suppression, either optogenetically or neurochemically, during the task resulted in an increase in non-productive licking (i.e., false alarms), rather than a decrease in licking. This is consistent with previous observations of motor cortical inhibition of movement (Stoltz et al., 1999). Currently, we do not know which M1 projection pathways contribute to this motor gating function. Possibilities include descending projections to brainstem/spinal cord (Ioffe, 1973), cortico-striatal projections to neurons of the indirect pathway (Hersch et al., 1995), or cortico-cortical projections to S1, which also projects to brainstem and striatum (Matyas et al., 2010; Miyashita et al., 1994; Porter and White, 1983). A consequence of this suppression may be to make movement dependent on M1-derived motor commands. It will be interesting to see if a gating function of rodent M1 generalizes across multiple behaviors, or if it is limited to stereotyped movements such as whisking and licking.

We find that licking following the target stimulus (hit rate) is similar during M1 suppression compared to control. We do not believe that these data invalidate a role for M1 in the normal sensory-motor transformation. While suppression experiments do reveal the consequences of removing a structure, in a highly non-linear system the function of a structure is not simply the opposite of what occurs when that structure is removed. For example, even

though primate frontal eye field (FEF) lesions do not abolish saccade generation in detection tasks (Schiller et al., 1987), this does not invalidate the decades of work demonstrating motor command signals in FEF and the importance of this activity in the normal performance of eye movements. Similarly, M1 is not only anatomically connected to structures involved in the performance of our task, it is highly active during the performance of the task, in a way that predicts the behavior of the animal.

In our unit recordings, amongst the post-target enhanced neurons, we observe a transformation from a sensory to a motor/decision representation. These representations involved partially overlapping neural ensembles. Interestingly, similar neural dynamics have been observed in primate frontal eye fields (FEF) and lateral intraparietal cortex (LIP) across diverse behavioral tasks (Bruce and Goldberg, 1985; Hanes et al., 1998; Roitman and Shadlen, 2002). Furthermore, ‘fixation-related’ neurons recorded in primate FEF (Bruce and Goldberg, 1985; Hanes et al., 1998) and ‘delay’ neurons in rodent M1 (Narayanan and Laubach, 2006, 2009) are similar to the post-target suppressed neurons described in our study. The similar neural dynamics across species and brain regions suggests a common functional organization of these cortical circuits.

Mechanisms underlying the decision process, however, are still heavily debated (Churchland et al., 2011; Gold and Shadlen, 2007; Lo et al., 2009; Purcell et al., 2010). We show that enhanced and suppressed neurons are anti-correlated and exhibit bidirectional Granger causality. This correlation structure is consistent with circuit models in which enhanced and suppressed neurons complete by lateral inhibition. Such models may display winner-take-all dynamics, which intrinsically convert graded, transient inputs to categorical outputs, as is needed in the conversion of a transient sensory stimulus to a motor response. We identify many similarities between the competitive ensemble model and our M1 neural data. The neural dynamics on miss trials are similar to a transient excursion of the competitive ensemble model. Furthermore, the anti-correlated and sustained neural dynamics on hit trials are similar to stable switches of the competitive ensemble model, including fine details of the switch process. Additionally, we provide a framework to explain how the hit trajectory drives sequential activation of neurons representing sensory followed by motor signals. These data support the conclusion that neural activity associated with hit trials may be understood as a state transition in a winner-take-all circuit.

In our neural data we observed a population of neurons that show apparent ramping activity late in the decision process, with peak activities aligning to motor initiation. We hypothesize that these neurons ultimately signal the motor command. These neurons are unlikely to drive the network transitions, due to their late activity, but rather appear to be driven by a successful transition. A major benefit of this circuit organization is that accumulation of activity in the motor command neurons is suppressed in the absence of a network transition. Previous modeling work has demonstrated the importance of such a filter or ‘gate’ for accumulator neurons to prevent stochastic triggering (Purcell et al., 2010), and here we demonstrate a plausible biophysical implementation of such a filter. While these late enhanced neurons precede and align with licking, it is currently unclear whether M1 drives whisking and/or licking directly, or whether the motor signals are first relayed to a different cortical area (Guo et al., 2014; Li et al., 2015).

Guided by the circuit model proposed in this study, future work may reveal the cellular substrates of the M1 circuit. Layer 5 of motor cortex contains both intratelencephalic and pyramidal tract projection neurons, which may have distinct circuit dynamics (Kiritani et al., 2012; Li et al., 2015). Furthermore, we do not know the identity or synaptic organization of the inhibitory neurons mediating lateral inhibition. Identifying these neurons will be important in confirming that competition occurs locally in M1, as we currently cannot exclude that a competition is occurring elsewhere in the brain and we are observing a read-out of that competition.

Network state is the temporal structure within which specific neural signaling occurs. Recently there have been significant efforts towards characterizing the diversity of network states in waking animals (Zagha and McCormick, 2014). In general, these studies have found that in stationary animals, cortex can be dominated by a low frequency, oscillatory state (Bennett et al., 2013; Crochet and Petersen, 2006; Hromadka et al., 2013; McGinley et al., 2015; Polack et al., 2013; Poulet and Petersen, 2008; Reimer et al., 2014; Tan et al., 2014; Vinck et al., 2015; Zagha et al., 2013). Here, we find that task engagement, even without movement, is sufficient to maintain M1 in the activated state. In this study and a recently published auditory detection task (McGinley et al., 2015), we find that slow, oscillatory network dynamics correlate with sub-optimal performance. The lack of such a relationship in a whisker detection task (Sachidhanandam et al., 2013) may be due to the use of stimuli far beyond the psychophysical detection threshold.

The precise mechanisms by which network states influence task performance are currently being investigated. In primary sensory cortices, cortical state influences the signal-to-noise ratio and reliability of sensory encoding (Bennett et al., 2013; Crochet and Petersen, 2006; Goard and Dan, 2009; Harris and Thiele, 2011; Marguet and Harris, 2011; McGinley et al., 2015; Vinck et al., 2015; Zagha et al., 2013; Zagha and McCormick, 2014). Computational studies suggest that in decision-related circuits, selective and stable ensemble activity are necessary for categorical representations of task and behavioral parameters (Brunel and Wang, 2001; Compte et al., 2000). Interestingly, we find that the anti-correlated activity of M1 ensembles is disrupted during the slow, oscillatory state. Further studies are needed to determine the drivers of cortical activation across cortical regions and across tasks, and determine the precise cellular mechanisms that promote competition amongst neural ensembles.

Experimental Procedures

Animal preparation and surgery

All protocols are in accordance with Yale University Institutional Animal Care and Use Committee. A light-weight metal head-holder with recording well was chronically implanted onto the skull of 2–3 month old male C57BL/6 wild type or PV-ChR2 (see below) mice under ketamine (90 mg/kg, i.p.) and xylazine (10 mg/kg, i.p.) anesthesia. A craniotomy (<0.5 mm in diameter) and durotomy were conducted on the day of recording, using stereotactic coordinates (from bregma [in mm], M1: 1 rostral, 1 lateral; S1: 1.5 caudal, 3.5 lateral). For M1 targeting, we first thinned the skull at the coordinates above and then prepared a small craniotomy either rostral or caudal the large surface vessel than typically

runs through M1. For recordings under anesthesia this was immediately after head-holder implantation, whereas for behaving mice this was after 2–3 weeks of training. Recordings from behaving mice began at least 1–2 hours after surgery, allowing recovery from anesthesia such that the animals appeared to be behaving normally in their own cages. For EMG recordings, fine tungsten wires (A-M Systems) were threaded into the whisker pad. For focal cortical suppression, 2 mM muscimol (Tocris) in normal saline was applied to the brain surface following head fixation and for the duration of the behavioral session. Experimenters were not blinded to pharmacological manipulation. Behavioral data was collected starting 15 minutes after muscimol application, which in control studies was sufficient for suppression of activity in layer 5 (data not shown).

Animal behavior

Mice were progressively trained in a sensory detection task. Water restriction was initiated 1–2 days after recovery from head-holder implantation, and continued throughout training. Mice were head-fixed during the training/recording sessions. Piezo-controlled (Physik Instrumente) paddles were placed in the whisker fields bilaterally, each contacting multiple (3–4) whiskers. Mice were trained in three stages. 1. Classical conditioning: we paired unilateral whisker deflection (either side, randomly selected) with fluid reward from a lickport (training period of 2–3 days, 1 session per day). 2. Operant conditioning: following unilateral whisker deflection, the mice must lick the lickport in order to initiate the fluid reward (lick detection window: 1.5 s) (training period of 3–4 days). 3. Impulse control training: mice are punished by aborting the trial and starting a new inter-trial interval for licking outside of the post-target period (training period of 7–14 days). Each trial started with a 10–15 second delay (inter-trial; chosen randomly from a negative exponential distribution). In addition to the target, in 80% of trials we first presented a tone after the inter-trial delay. Mice had to withhold licking to the tone (correct rejection) in order to be presented with the target stimulus 1–3 seconds later. Licking in the 1.5 second window following presentation of the target stimulus initiated fluid reward (hit). The tone-target delay had a negative exponential distribution to discourage timing. Moreover, in all expert mice, reaction times were tightly correlated to the target onset, as opposed to the tone, indicating that mice were not timing from the tone. Licking to the tone (false alarm) or spontaneous licking outside of the post-target window was punished by aborting the current trial and starting a new inter-trial delay (effectively a time out). By the third training stage, mice performed continuously for 40–70 minutes, completing 100–200 trials.

Electrophysiological recordings

All recordings were targeted to layer 5, from 500 to 1100 um from the brain surface. Recording methods are similar to a previously published study (Zagha et al., 2013). Details can be found in Supplemental Experimental Procedures.

Optogenetic inhibition

We generated mice expressing excitatory opsin channelrhodopsin-2 in PV-containing GABAergic interneurons by mating PV-Cre and Ai32 (Jackson) mice. We delivered 460 nm blue light to M1 bilaterally (1 mW each) through 200–300 um core diameter optical fibers

(ThorLabs), using LED-based light sources (Prizmatix, ThorLabs). The optical fibers were positioned at the dural surface above M1.

Data analyses

All analyses were conducted in Matlab. Juxtacellular single unit and tungsten electrode multiunit spike times were determined as threshold crossings well isolated (> 2X amplitude) from background noise. LFP was determined by low-pass filtering offline (100 Hz cutoff, 5^{th} order Bessel filter) and down-sampling to 200 Hz. Power spectral density was calculated using a multitaper method. Signals from EMG wires were high-pass filtered (100 Hz) and rectified. For multielecrode recordings, channels were segregated into four tetrodes. Putative spikes for each tetrode were determined by amplitude and template matching (Spike2), and clustered into units using MClust (Schmitzer-Torbert et al., 2005) and KlustaKwik (Kadir et al., 2014). Unit quality was assessed based on spike shape, refractory period, L_{ratio} (LR) and Isolation Distance (ID) (LR: 0.21+/−0.04, ID: 35+/−14). For calculation of membrane potential mean and variance, recordings were median filtered with a 10 ms sliding window to truncate spikes. For spike rasters, spikes are binned into 20 ms bins. Peri-event indices were calculated as $(FR_{post}-FR_{pre})/(FR_{post}+FR_{pre})$ [FR, firing rate]. Thus, suppression (enhancement) across an event would be negative (positive) using this statistic. 'Spontaneous' activity was selected during the inter-trial interval, at random latencies from the previous trial; 'tone' compared 1 second before/after tone onset for correct rejection trials; and 'target' compared 1 second before the target to 50–200 ms after the target for hit trials, which was before the onset of repetitive licking and whisker during the reward phase. To assess whether neural responses were aligned to the target, we compared hit versus miss trials. For units with insufficient miss trials collected, we compared 'fast' (<0.6 s) versus 'slow' (>0.6 s) reaction time trials. Both spike-spike cross-correlation and Granger causality analyses used spike rates in 20 ms bins across the entire recording length. Spike correlation amplitude is calculated as the zeroth lag of the normalized covariance function. Granger causality was tested using the MVGC multivariate Granger causality toolbox (Barnett and Seth, 2014) and maximum 10^{th} order VAR models. Data are presented as mean +/− standard error, unless otherwise specified. Fisher's exact test was used for comparing categorical distributions, including proportion of anti-correlated units or Granger causal connections; Student's t test was used for comparing distributions of continuous variables, including EMG amplitude, changes in HR and FAR, firing rates (pre vs post or RSU vs FSU), correlation magnitude, regression coefficients of the VAR model, LFP power, Vm mean and variance; Chi-square was used to compare the proportion of modulated neurons across various task periods.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank James Mazer, Alexandria Marino, David Salkoff, Daeyeol Lee, Bart Massi, Martin Vinck, Xiao-Jing Wang and Jah Chaisangmonkon for insightful discussions. We thank Daeyeol Lee, Mark Laubach and Ethan Mohns for helpful comments on a previous version of this manuscript. We thank Peter O'Brien for technical

solutions, and Liana Saussy and Ozge Yuzgec for assistance with mouse training. This work was supported by NIH NS026143, NS007224 and Kavli Institute for Neuroscience (DAM) and NIH F32NS077816 (EZ).

References

Barnett L, Seth AK. The MVGC multivariate Granger causality toolbox: a new approach to Granger-causal inference. Journal of neuroscience methods. 2014; 223:50–68. [PubMed: 24200508]

Beloozerova IN, Sirota MG, Swadlow HA. Activity of different classes of neurons of the motor cortex during locomotion. J Neurosci. 2003; 23:1087–1097. [PubMed: 12574439]

Bennett C, Arroyo S, Hestrin S. Subthreshold mechanisms underlying state-dependent modulation of visual responses. Neuron. 2013; 80:350–357. [PubMed: 24139040]

Bruce CJ, Goldberg ME. Primate frontal eye fields. I. Single neurons discharging before saccades. J Neurophysiol. 1985; 53:603–635. [PubMed: 3981231]

Brunel N, Wang XJ. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. Journal of computational neuroscience. 2001; 11:63–85. [PubMed: 11524578]

Carvell GE, Miller SA, Simons DJ. The relationship of vibrissal motor cortex unit activity to whisking in the awake rat. Somatosensory & motor research. 1996; 13:115–127. [PubMed: 8844960]

Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN. Variance as a signature of neural computations during decision making. Neuron. 2011; 69:818–831. [PubMed: 21338889]

Churchland MM, Afshar A, Shenoy KV. A central source of movement variability. Neuron. 2006; 52:1085–1096. [PubMed: 17178410]

Compte A, Brunel N, Goldman-Rakic PS, Wang XJ. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. Cereb Cortex. 2000; 10:910–923. [PubMed: 10982751]

Cook EP, Maunsell JH. Dynamics of neuronal responses in macaque MT and VIP during motion detection. Nat Neurosci. 2002; 5:985–994. [PubMed: 12244324]

Crochet S, Petersen CC. Correlating whisker behavior with membrane potential in barrel cortex of awake mice. Nat Neurosci. 2006; 9:608–610. [PubMed: 16617340]

de Lafuente V, Romo R. Neuronal correlates of subjective sensory experience. Nat Neurosci. 2005; 8:1698–1703. [PubMed: 16286929]

de Lafuente V, Romo R. Neural correlate of subjective sensory experience gradually builds up across cortical areas. Proc Natl Acad Sci U S A. 2006; 103:14266–14271. [PubMed: 16924098]

Erlich JC, Bialek M, Brody CD. A cortical substrate for memory-guided orienting in the rat. Neuron. 2011; 72:330–343. [PubMed: 22017991]

Gentet LJ, Avermann M, Matyas F, Staiger JF, Petersen CC. Membrane potential dynamics of GABAergic neurons in the barrel cortex of behaving mice. Neuron. 2010; 65:422–435. [PubMed: 20159454]

Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci. 1982; 2:1527–1537. [PubMed: 7143039]

Goard M, Dan Y. Basal forebrain activation enhances cortical coding of natural scenes. Nat Neurosci. 2009; 12:1444–1449. [PubMed: 19801988]

Gold JI, Shadlen MN. The neural basis of decision making. Annu Rev Neurosci. 2007; 30:535–574. [PubMed: 17600525]

Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K. Flow of cortical activity underlying a tactile decision in mice. Neuron. 2014; 81:179–194. [PubMed: 24361077]

Hanes DP, Patterson WF 2nd, Schall JD. Role of frontal eye fields in countermanding saccades: visual, movement, and fixation activity. J Neurophysiol. 1998; 79:817–834. [PubMed: 9463444]

Hanes DP, Schall JD. Neural control of voluntary movement initiation. Science. 1996; 274:427–430. [PubMed: 8832893]

Harris KD, Thiele A. Cortical state and attention. Nat Rev Neurosci. 2011; 12:509–523. [PubMed: 21829219]

Hersch SM, Ciliax BJ, Gutekunst CA, Rees HD, Heilman CJ, Yung KK, Bolam JP, Ince E, Yi H, Levey AI. Electron microscopic analysis of D1 and D2 dopamine receptor proteins in the dorsal striatum and their synaptic relationships with motor corticostriatal afferents. J Neurosci. 1995; 15:5222–5237. [PubMed: 7623147]

Hill DN, Curtis JC, Moore JD, Kleinfeld D. Primary motor cortex reports efferent control of vibrissa motion on multiple timescales. Neuron. 2011; 72:344–356. [PubMed: 22017992]

Hromadka T, Zador AM, DeWeese MR. Up states are rare in awake auditory cortex. J Neurophysiol. 2013; 109:1989–1995. [PubMed: 23343898]

Huber D, Gutnisky DA, Peron S, O'Connor DH, Wiegert JS, Tian L, Oertner TG, Looger LL, Svoboda K. Multiple dynamic representations in the motor cortex during sensorimotor learning. Nature. 2012; 484:473–478. [PubMed: 22538608]

Ioffe ME. Pyramidal influences in establishment of new motor coordinations in dogs. Physiology & behavior. 1973; 11:145–153. [PubMed: 4729155]

Kadir SN, Goodman DF, Harris KD. High-dimensional cluster analysis with the masked EM algorithm. Neural computation. 2014; 26:2379–2394. [PubMed: 25149694]

Kalaska JF, Cohen DA, Hyde ML, Prud'homme M. A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. J Neurosci. 1989; 9:2080–2102. [PubMed: 2723767]

Kiritani T, Wickersham IR, Seung HS, Shepherd GM. Hierarchical connectivity and connection-specific dynamics in the corticospinal-corticostriatal microcircuit in mouse motor cortex. J Neurosci. 2012; 32:4992–5001. [PubMed: 22492054]

Li N, Chen TW, Guo ZV, Gerfen CR, Svoboda K. A motor cortex circuit for motor planning and movement. Nature. 2015; 519:51–56. [PubMed: 25731172]

Lo CC, Boucher L, Pare M, Schall JD, Wang XJ. Proactive inhibitory control and attractor dynamics in countermanding action: a spiking neural circuit model. J Neurosci. 2009; 29:9059–9071. [PubMed: 19605643]

Marguet SL, Harris KD. State-dependent representation of amplitude-modulated noise stimuli in rat auditory cortex. J Neurosci. 2011; 31:6414–6420. [PubMed: 21525282]

Matyas F, Sreenivasan V, Marbach F, Wacongne C, Barsy B, Mateo C, Aronoff R, Petersen CC. Motor control by sensory cortex. Science. 2010; 330:1240–1243. [PubMed: 21109671]

McGinley MJ, David SV, McCormick DA. Cortical Membrane Potential Signature of Optimal States for Sensory Signal Detection. Neuron. 2015; 87:179–192. [PubMed: 26074005]

Miyashita E, Keller A, Asanuma H. Input-output organization of the rat vibrissal motor cortex. Experimental brain research. 1994; 99:223–232. [PubMed: 7523173]

Narayanan NS, Laubach M. Top-down control of motor cortex ensembles by dorsomedial prefrontal cortex. Neuron. 2006; 52:921–931. [PubMed: 17145511]

Narayanan NS, Laubach M. Delay activity in rodent frontal cortex during a simple reaction time task. J Neurophysiol. 2009; 101:2859–2871. [PubMed: 19339463]

Ollerenshaw DR, Bari BA, Millard DC, Orr LE, Wang Q, Stanley GB. Detection of tactile inputs in the rat vibrissa pathway. J Neurophysiol. 2012; 108:479–490. [PubMed: 22514290]

Petreanu L, Gutnisky DA, Huber D, Xu NL, O'Connor DH, Tian L, Looger L, Svoboda K. Activity in motor-sensory projections reveals distributed coding in somatosensation. Nature. 2012; 489:299–303. [PubMed: 22922646]

Polack PO, Friedman J, Golshani P. Cellular mechanisms of brain state-dependent gain modulation in visual cortex. Nat Neurosci. 2013; 16:1331–1339. [PubMed: 23872595]

Porter LL, White EL. Afferent and efferent pathways of the vibrissal region of primary motor cortex in the mouse. J Comp Neurol. 1983; 214:279–289. [PubMed: 6853758]

Poulet JF, Petersen CC. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. Nature. 2008; 454:881–885. [PubMed: 18633351]

Purcell BA, Heitz RP, Cohen JY, Schall JD, Logan GD, Palmeri TJ. Neurally constrained modeling of perceptual decision making. Psychological review. 2010; 117:1113–1143. [PubMed: 20822291]

Reimer J, Froudarakis E, Cadwell CR, Yatsenko D, Denfield GH, Tolias AS. Pupil fluctuations track fast switching of cortical states during quiet wakefulness. Neuron. 2014; 84:355–362. [PubMed: 25374359]

Robbins TW, Gillan CM, Smith DG, de Wit S, Ersche KD. Neurocognitive endophenotypes of impulsivity and compulsivity: towards dimensional psychiatry. Trends Cogn Sci. 2012; 16:81–91. [PubMed: 22155014]

Roitman JD, Shadlen MN. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci. 2002; 22:9475–9489. [PubMed: 12417672]

Sachidhanandam S, Sreenivasan V, Kyriakatos A, Kremer Y, Petersen CC. Membrane potential correlates of sensory perception in mouse barrel cortex. Nat Neurosci. 2013; 16:1671–1677. [PubMed: 24097038]

Schachar R, Tannock R, Marriott M, Logan G. Deficient inhibitory control in attention deficit hyperactivity disorder. Journal of abnormal child psychology. 1995; 23:411–437. [PubMed: 7560554]

Schall JD, Thompson KG. Neural selection and control of visually guided eye movements. Annu Rev Neurosci. 1999; 22:241–259. [PubMed: 10202539]

Schiller PH, Sandell JH, Maunsell JH. The effect of frontal eye field and superior colliculus lesions on saccadic latencies in the rhesus monkey. J Neurophysiol. 1987; 57:1033–1049. [PubMed: 3585453]

Schmitzer-Torbert N, Jackson J, Henze D, Harris K, Redish AD. Quantitative measures of cluster quality for use in extracellular recordings. Neuroscience. 2005; 131:1–11. [PubMed: 15680687]

Schneidman E, Berry MJ 2nd, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. Nature. 2006; 440:1007–1012. [PubMed: 16625187]

Sebastian A, Jung P, Krause-Utz A, Lieb K, Schmahl C, Tuscher O. Frontal dysfunctions of impulse control - a systematic review in borderline personality disorder and attention-deficit/hyperactivity disorder. Frontiers in human neuroscience. 2014; 8:698. [PubMed: 25232313]

Stoltz S, Humm JL, Schallert T. Cortical injury impairs contralateral forelimb immobility during swimming: a simple test for loss of inhibitory motor control. Behavioural brain research. 1999; 106:127–132. [PubMed: 10595428]

Tan AY, Chen Y, Scholl B, Seidemann E, Priebe NJ. Sensory stimulation shifts visual cortex from synchronous to asynchronous states. Nature. 2014; 509:226–229. [PubMed: 24695217]

Tanji J, Evarts EV. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. J Neurophysiol. 1976; 39:1062–1068. [PubMed: 824409]

Vinck M, Batista-Brito R, Knoblich U, Cardin JA. Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. Neuron. 2015; 86:740–754. [PubMed: 25892300]

Wang XJ. Probabilistic decision making by slow reverberation in cortical circuits. Neuron. 2002; 36:955–968. [PubMed: 12467598]

Wong KF, Wang XJ. A recurrent network mechanism of time integration in perceptual decisions. J Neurosci. 2006; 26:1314–1328. [PubMed: 16436619]

Zagha E, Casale AE, Sachdev RN, McGinley MJ, McCormick DA. Motor cortex feedback influences sensory processing by modulating network state. Neuron. 2013; 79:567–578. [PubMed: 23850595]

Zagha E, McCormick DA. Neural control of brain state. Curr Opin Neurobiol. 2014; 29:178–186. [PubMed: 25310628]

Highlights

- We implement a somatosensory detection task in which motor cortex gates behavior
- Motor cortex neurons display anti-correlated spiking in anticipation of movement
- These neurons may form a competitive ensemble circuit linking sensory/motor signals
- Both task performance and spiking correlation structure are state-dependent

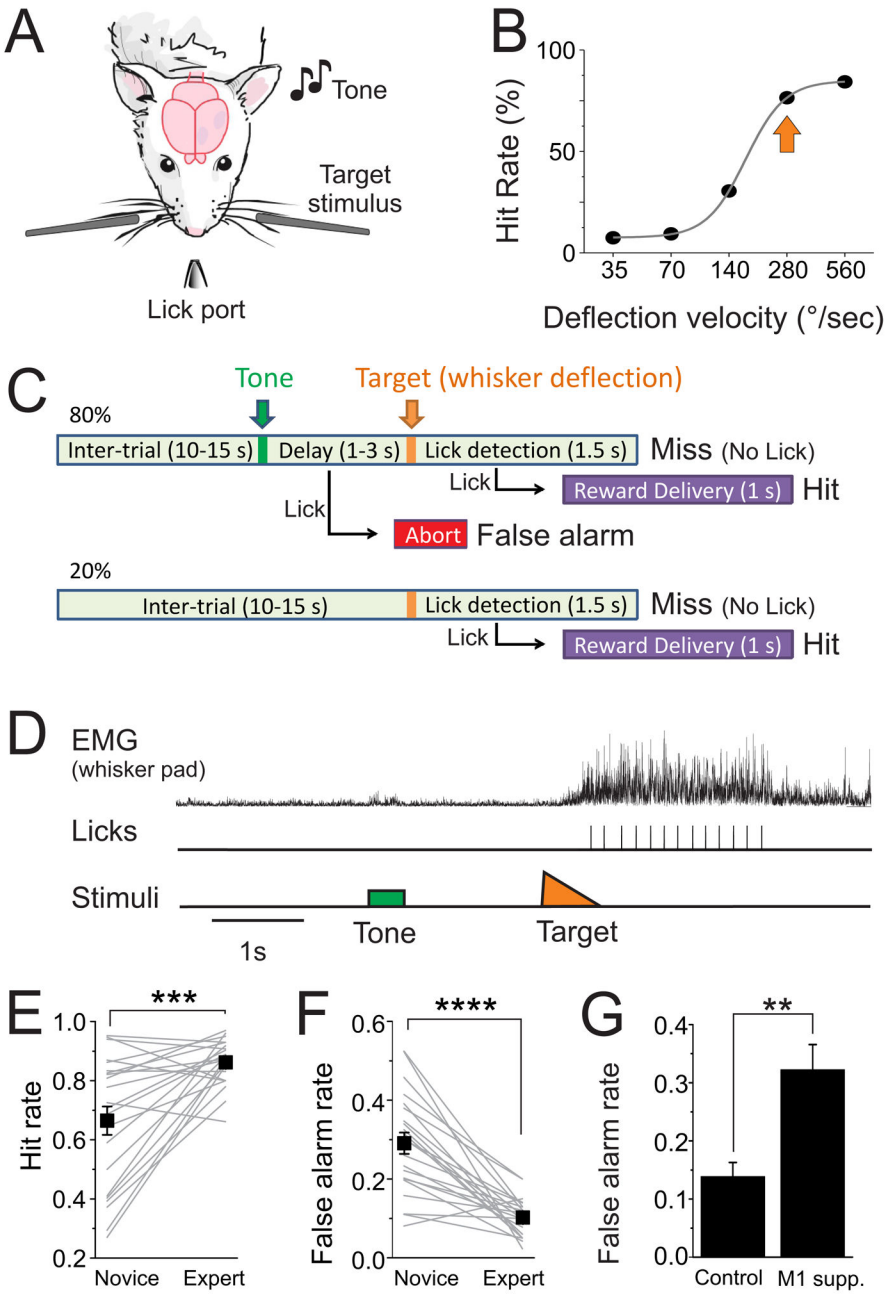

Figure 1. Somatosensory detection task

A) Schematic of the head-fixed detection task, with bilateral piezo-controlled paddles for delivering the target stimuli and lickport for behavior monitoring and reward delivery. B) Example psychophysical curve from one mouse. Arrow indicates the deflection velocity typically used during recordings. C) Trial structure. D) A hit trial, while monitoring whisker pad EMG (top) and licking of the lickport (middle). Note that EMG activity is at baseline prior to the target stimulus, which triggers repetitive whisking and licking. Hit (E) and false alarm (F) rates, comparing novice and expert performance. 'Novice' refers to the first session of impulse control training, after approximately 1 week of classical and operant

conditioning (see Methods). Boxes are population averages and gray lines are individual mice (n=22). Note that over the course of training mice significantly increase hit rate and decrease false alarms. G) In expert mice, bilateral M1 suppression (M1 supp.) caused an increase in false alarm rates (n=7 mice; n=5 muscimol, n=2 optogenetics, data combined). Data are represented as mean ± SEM. ** $p<0.01$; *** $p<0.001$; **** $p<0.00001$. See also Figure S1 and Table S1.

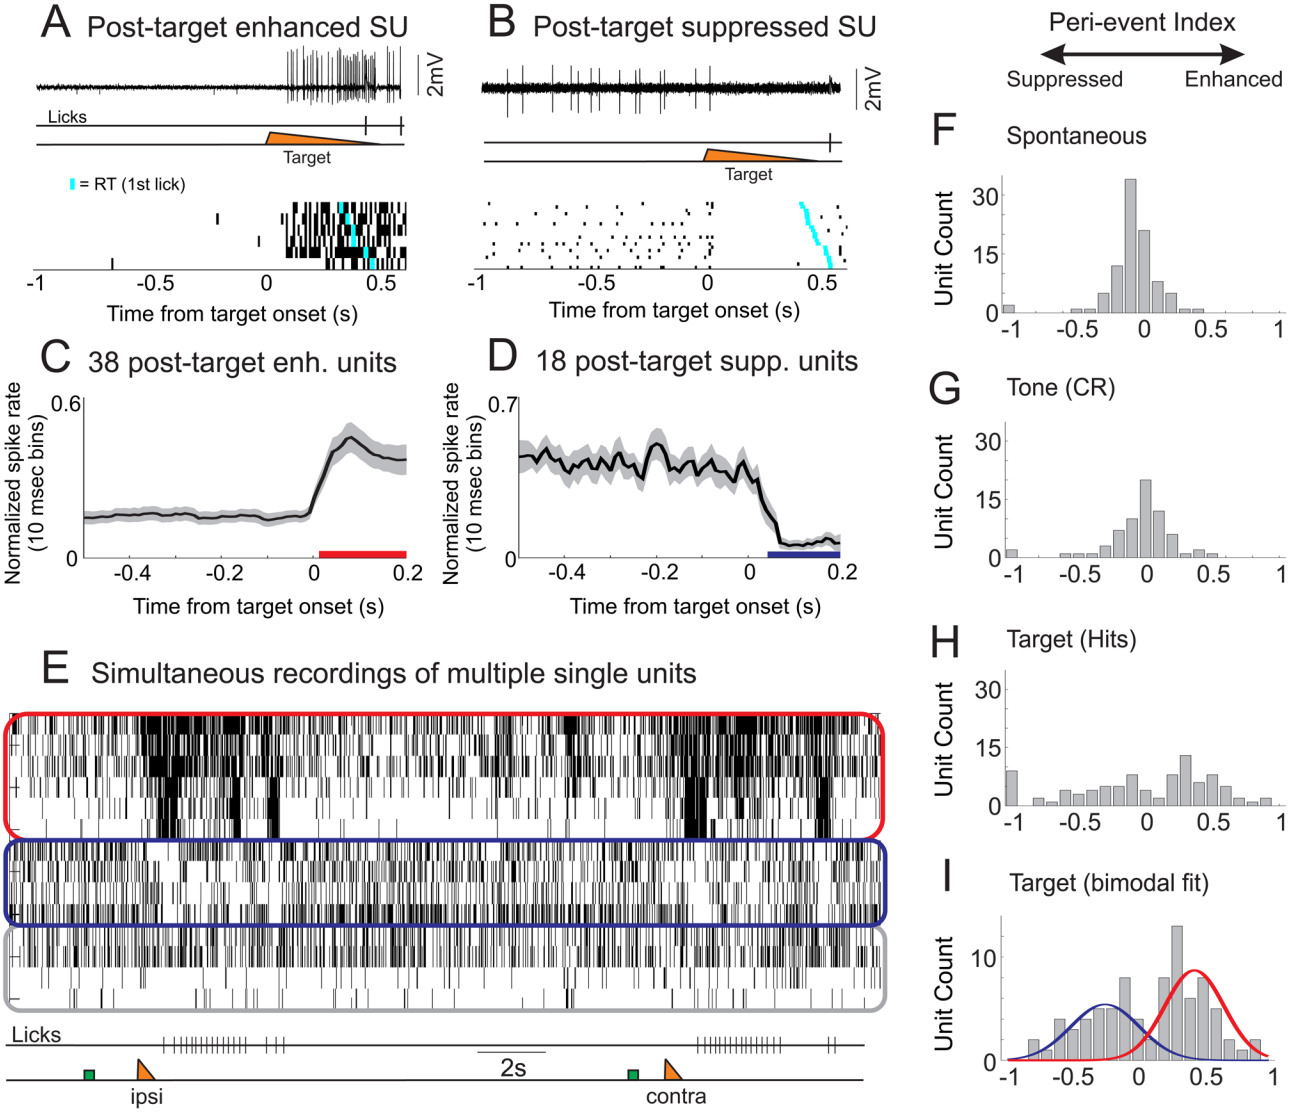

Figure 2. Populations of neurons exhibited enhanced or suppressed activity post target

A) Top, single trial example of a neuron that increased its spiking activity following whisker deflection (target) and prior to licking. We refer to these cells as post-target enhanced neurons. This cell was recorded in the loose-patch juxtacellular configuration. Bottom, raster plot of 6 hit trials, the cyan tick mark indicating the first lick. B) Top, single trial example of a neuron that decreased its activity following whisker stimulation and prior to licking (post-target suppressed neuron). Bottom, raster plot of 20 hit trials. (C, D) Average, normalized spike rate of 38 post-target enhanced (C) and 18 post-target suppressed (D) units. Colored bars beneath the traces indicate significant changes from baseline (pre-target) firing rates ($p<0.01$). E) Spike raster showing simultaneous recordings of 14 units using multi-electrode arrays. Units are clustered into post-target enhanced (red outline), post-target suppressed (blue outline) or unclassified (gray outline). Note that the population activity is relatively stable during the pre-target interval. The target stimulus produces rapid and robust changes in spiking of the enhanced and suppressed units prior to the onset of licking. These activity

patterns occurred for both ipsilateral (ipsi) and contralateral (contra) target stimuli. Green box, tone; orange triangle, target. (F–I) To determine the distribution of activity profiles across the population (n=92 units), we calculated the peri-event indices (PI) across various task periods [firing rate: (Post-Pre)/(Post+Pre)]. For spontaneous activity (F, during the inter-trial interval) and across the tone (G), the PI distribution was unimodal centered at zero. H) Across the target the PI distribution consisted of multiple non-zero peaks. I) Fit of the peri-target data as the sum of two Gaussian curves. Note, PI of −1 or 1 reflects a spike rate of 0 Hz in the post or pre interval, respectively. As they reflect a floor effect of spiking data, they were excluded from the Gaussian fits. See also Figure S2.

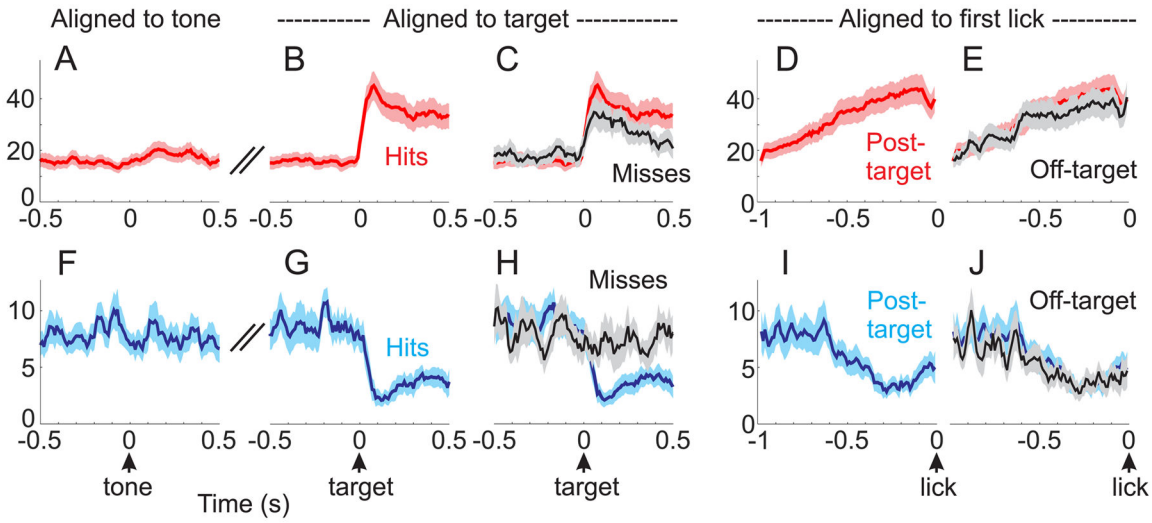

Figure 3. Coordinated enhancement and suppression anticipate motor responses

Average spike rates for significantly target-modulated units, segregated into enhanced (A–E, n=38 units) and suppressed (F–J, n=18 units) populations. A) Alignment of spike rates to the onset of the tone for correct rejection trials. B, C) Alignment to the onset of the target stimulus, showing hit trials [red in B and C] and miss trials [black in C]. D, E) Alignment to the first lick in a lick bout for post-target responses [red in D and E] and off-target responses [black in E]. (F–J), same alignments as (A–E), except for the suppressed population as indicated in blue. Note that the suppressed population responds selectively preceding the lick response. The enhanced population also responds preceding the lick response, but additionally responds to the target on miss trials. See also Figure S3.

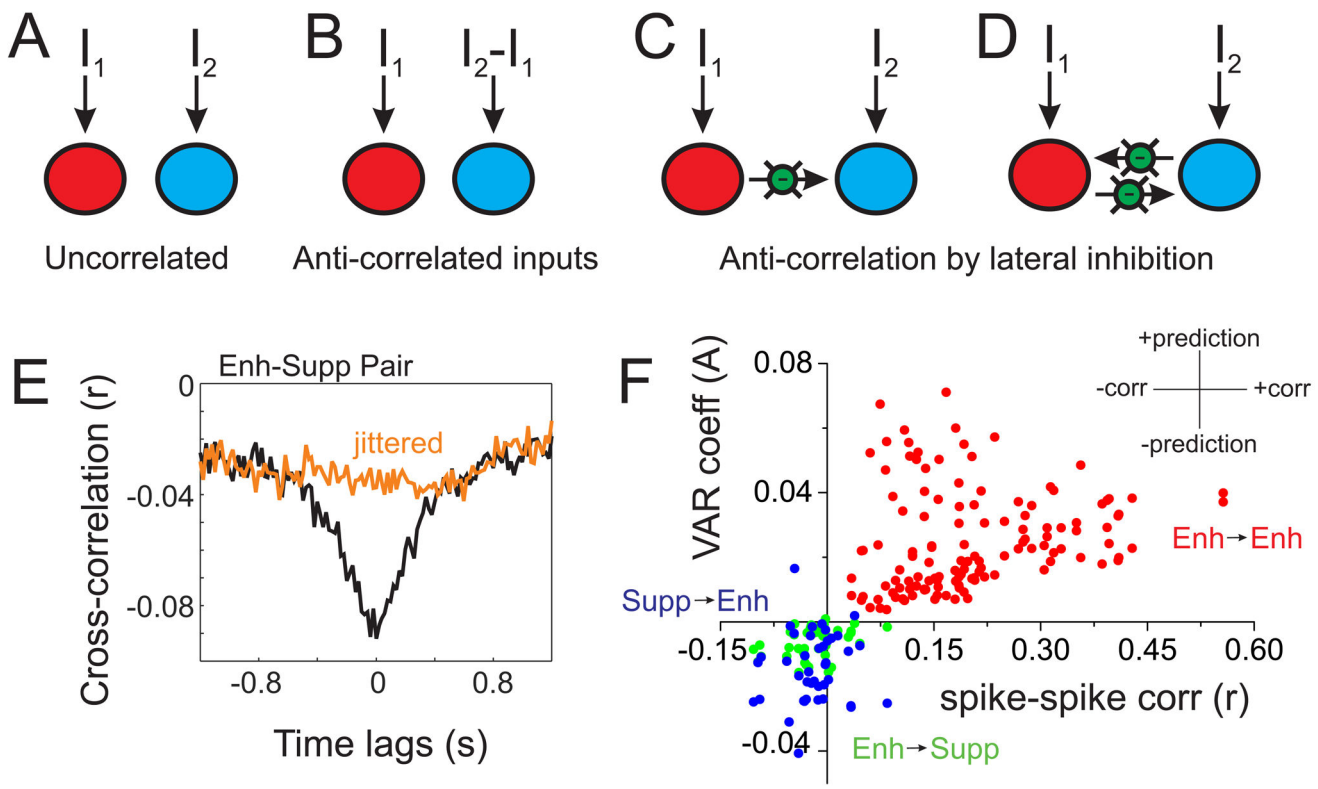

Figure 4. Circuit structure predicted by correlation analyses

A–D) Two ensemble circuit models used for comparison to the neural data. Each large circle (red or blue) represents an ensemble of neurons with recurrent excitation (see Figure S5), and may receive external inputs (I_1 and I_2) and lateral inhibition (green circles). E) Example cross-correlation from one Enh-Supp pair. The robust anti-correlation centered near zero (black trace) is disrupted by spike jittering (orange). F) Data from all pairwise comparisons (Enh→Supp (green), Supp→Enh (blue) and Enh→Enh (red); we did not have sufficient data for Supp→Supp comparisons). For each pair we plotted the spike rate correlation (r, x-axis) and average of the regression coefficients of the Granger model (A, y-axis). As shown in the inset, the x-axis reflects positive or negative correlations, whereas the y-axis reflects positive or negative weights of forward prediction. Enh→Supp and Supp→Enh pairs cluster in the lower left quadrant, indicating negative correlation and bidirectional, negative forward prediction. This profile is consistent with competition between ensembles, as illustrated in [D]. Enh→Enh pairs, in contrast, cluster in the upper right quadrant, consistent with recurrent excitation within an ensemble. See also Table S2 and Figure S4.

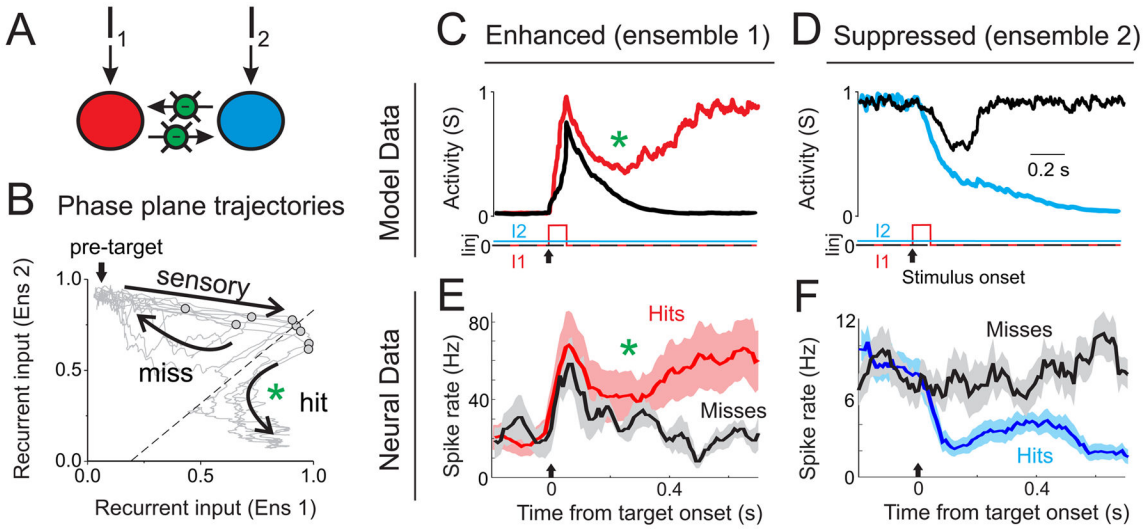

Figure 5. Similar dynamics in competitive network simulations and neural data

A) Schematic of the competitive ensemble model used for simulations. B) Phase plane plot of the recurrent input to each ensemble. The 'pre-target' state as indicated by the arrow is maintained by providing a small bias current to ensemble 2. Light grey lines show individual trial simulations, with the gray circles indicating the state of the network at the offset of the transient stimulus. The dashed line reflects the boundary between two attractors (see Figure S5 for more details). In 'miss' trials the network returns to the pre-target node, whereas in 'hit' trials the network diverges to a different stable node (at the bottom right of this plot). C, D) Activity profile for two example simulations of a near-threshold stimulus; external inputs I_1 and I_2 are shown below traces. C) Recurrent activity (S, which is proportional to firing rate) of ensemble 1 in which the transient stimulus did (red) or did not (black) lead to a stable transition. D) Same as (C), except for ensemble 2, with the successful transition shown in blue. E, F) Activity profile for hit and miss trials of enhanced (E, average of n=3 units) and suppressed (F, average of n=18 units) neurons. Note that sustained transitions in ensemble 1 and hit trials in the enhanced population are both associated with a curved activity trajectory, indicated by the green asterisks (in B, C and E).

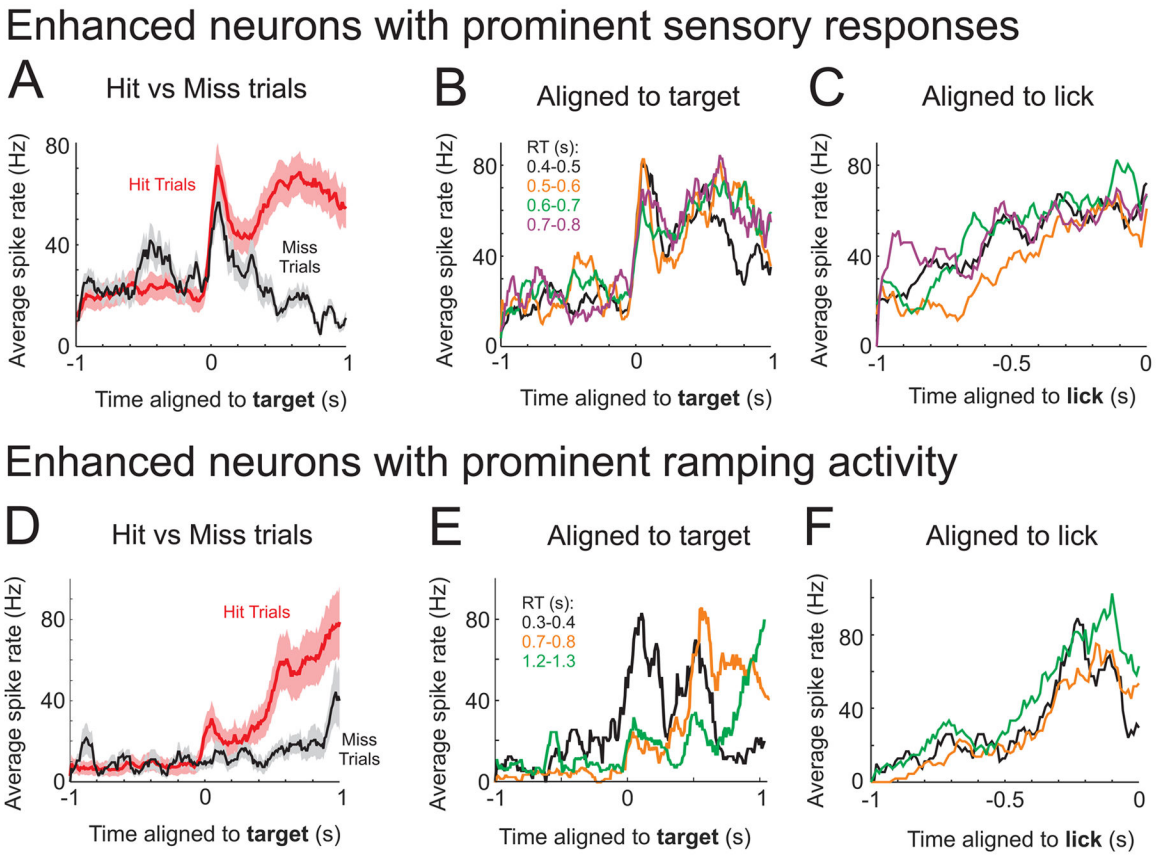

Figure 6. Correlates of the sensory-motor transformation in enhanced neurons

A–C) Spike patterns of three simultaneously recorded post-target enhanced neurons with strong onset and sustained activity. These are the same neurons as shown in Fig. 5E. A) Average spike rates for hit (red) and miss (black) trials. Note that the onset peak is present in both conditions, whereas the sustained activity is only present on hit trials. B, C) Trials were binned according to reaction time and aligned either to the target (B) or first lick (C). The onset peak aligns to the target across all reaction times, whereas there is a slow ramp up of activity to the first lick. D–F) Spike patterns of three simultaneously recorded post-target enhanced neurons with strong ramping activity. D) Average spike rates for hit (red) and miss (black) trials. The ramping activity is strongly reduced on miss trials. Across different reaction times, peak activity aligns to the first lick (F) rather than the target onset (E). See also Figure S6.

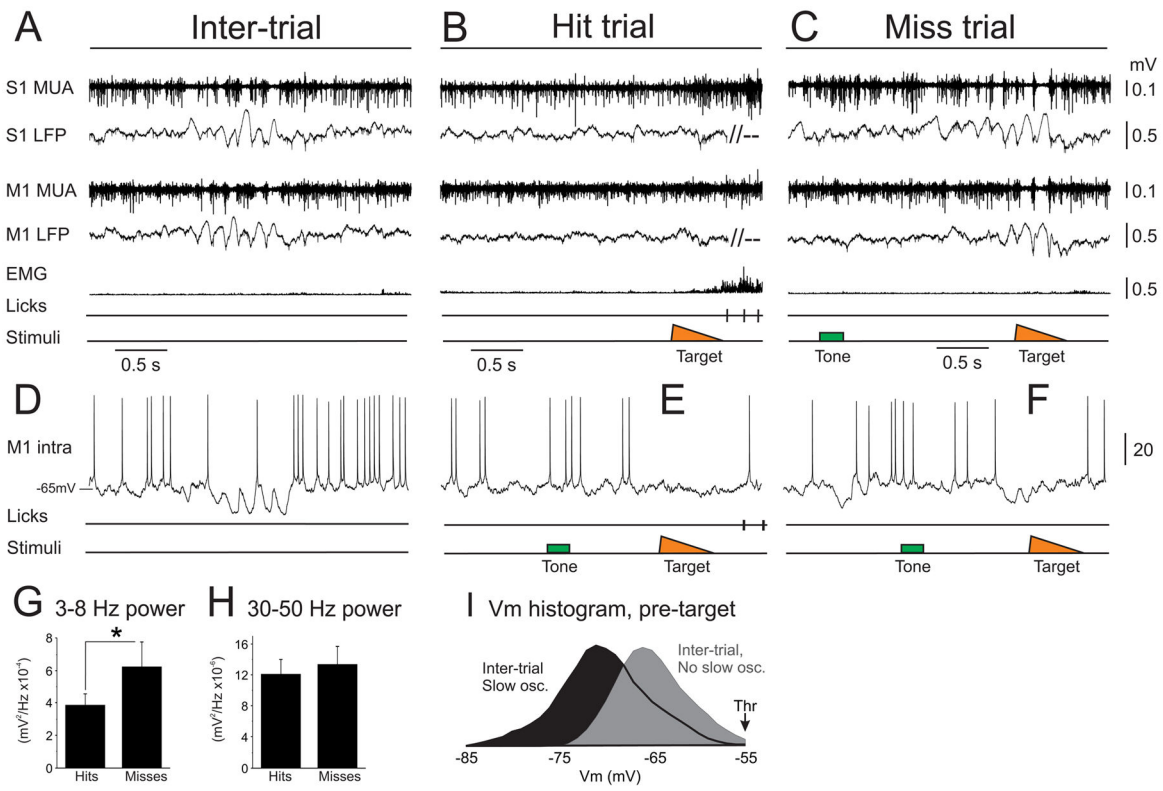

Figure 7. Cortical state correlates with task engagement

(A–C) Simultaneous extracellular recordings from S1 (top) and M1 (middle) and whisker pad EMG (bottom) during the detection task. Extracellular signals were filtered into multiunit activity (MUA, above) and local field potential (LFP, below). B) During periods of optimal performance, these cortical circuits were maintained in the activated state, even in the absence of whisker activity or sensory stimuli. Brief, slow oscillatory epochs were observed rarely during the inter-trial intervals (A), and commonly towards the end of a behavioral session when miss trials were more frequent (C). (D–F) Intracellular recording from an M1 neuron during the detection task, showing the intracellular correlates of the extracellular signals presented above (not simultaneously recorded). Note that activation was associated with sustained depolarization out of the Down state (E), whereas slow, oscillatory periods included large Vm fluctuations to the Down state (D, F). (G, H) LFP power across the 2 seconds prior to target onset for hit and miss trials. 3–8 Hz power was lower for hit trials, whereas there was no difference in power at 30–50 Hz. I) Vm histrograms during the inter-trial interval, parsed into epochs of slow, oscillatory fluctuations (black) and activation (gray). Note that the Vm histogram containing slow, oscillatory fluctuations has a long tail to −85 mV, reflecting brief transitions into the Down state. *, $p<0.05$. See also Figure S7.

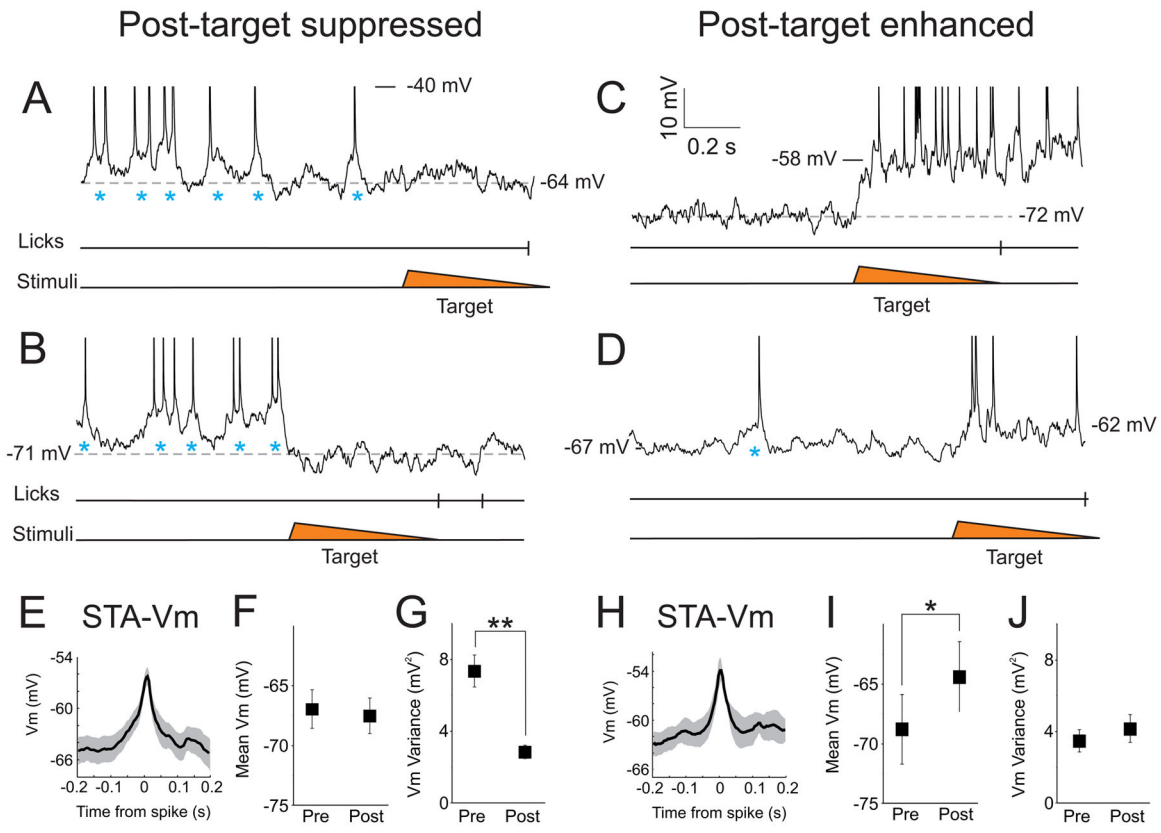

Figure 8. Spiking during active periods is initiated by large, transient depolarizations

(AD) Intracellular recordings from different M1 neurons on hit trials. Spikes are truncated at −40 mV. Blue asterisks reference transient depolarizations from baseline that initiate spikes. (A, B) Post-target suppression is associated with elimination of the large depolarizations, revealing a relatively stable Vm which is between the Down state and spike threshold. (C, D) Post-target enhancement is associated with a rapid depolarization and repetitive firing. (E, H) Spike-triggered average of the Vm (STA-Vm) for suppressed (E) and enhanced (H) neurons. For neurons of both spike patterns, spiking was associated with large (7–8 mV) and transient (100–200 ms) depolarizations. We calculated the mean Vm (F, I) and Vm variance (G, J) for neurons of both spike patterns, comparing one second pre-target (pre) to 300 ms post-target (post). For suppressed neurons, the post-target interval is associated with a strong reduction in Vm variance (G) whereas enhanced neurons experience a moderate depolarization (I). *, $p<0.05$; **, $p<0.0001$. See also Figure S8.