Subthreshold basis for reward-predictive persistent activity in mouse prefrontal cortex

Eunyoung Kim[1], Bilal A. Bari[1], Jeremiah Y. Cohen[1,2,*]

[1]The Solomon H. Snyder Department of Neuroscience, Brain Science Institute, Kavli Neuroscience Discovery Institute, The Johns Hopkins University School of Medicine, Baltimore, MD, USA

[2]Lead contact

SUMMARY

Nervous systems maintain information internally using persistent activity changes. The mechanisms by which this activity arises are incompletely understood. We study prefrontal cortex (PFC) in mice performing behaviors in which stimuli predicted rewards at different delays with different probabilities. We measure membrane potential (V_m) from pyramidal neurons across layers. Reward-predictive persistent firing increases arise due to sustained increases in mean and variance of V_m and are terminated by reward or via centrally generated mechanisms based on reward expectation. Other neurons show persistent decreases in firing rates, maintained by persistent hyperpolarization that is robust to intracellular perturbation. Persistent activity is layer (L)- and cell-type-specific. Neurons with persistent depolarization are primarily located in upper L5, whereas those with persistent hyperpolarization are mostly found in lower L5. L2/3 neurons do not show persistent activity. Thus, reward-predictive persistent activity in PFC is spatially organized and conveys information about internal state via synaptic mechanisms.

Graphical abstract
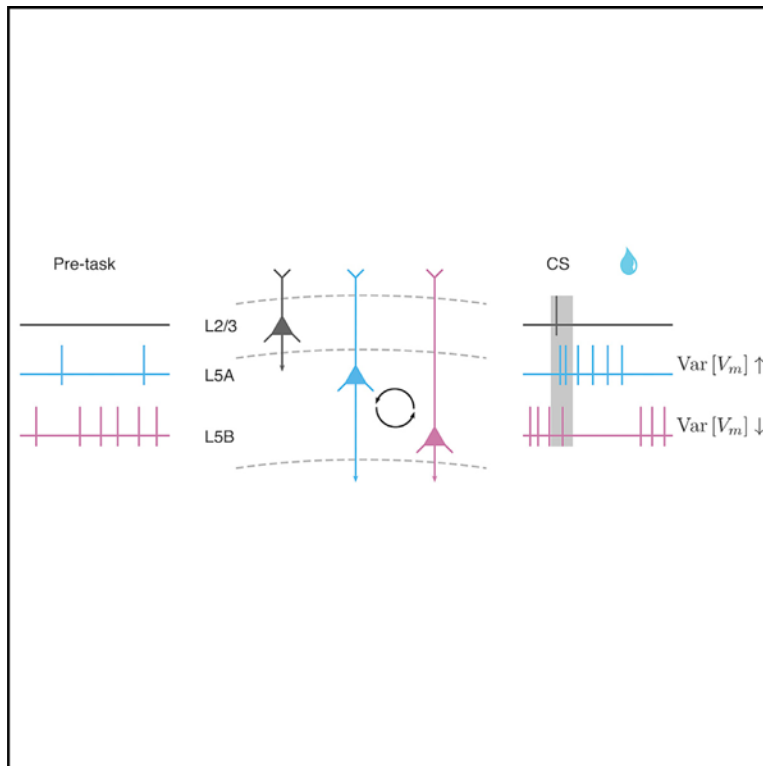
In brief

Kim et al. show sustained changes in membrane potential and firing rates in mouse frontal cortex leading up to an expected reward. These dynamics rely on underlying changes in mean and variance, directly testing prior theoretical studies. Neurons showing increased and decreased activity changes are located in different cortical layers.

INTRODUCTION

Predicting future reward is critical for successful adaptive behavior (Hull, 1943). Nervous systems anticipate the likely outcomes of stimuli in the environment through reinforcement learning (Bush and Mosteller, 1951; Rescorla and Wagner, 1972; Sutton and Barto, 1981). In the real world, many reward-predicting stimuli are followed by a time delay. This requires the nervous system to maintain activity in anticipation of the future reward.

Neurons are capable of maintaining changes in firing rates in the absence of external stimuli. This persistent activity was first observed in prefrontal cortex (PFC) (Fuster and Alexander, 1971; Kubota and Niki, 1971) and was traditionally viewed as a substrate of working memory (Fuster and Alexander, 1971, Funahashi et al., 1989). Subsequent work pointed to a more general strategy for the nervous system to bridge delays between events in the world, during decision making (Schall and Hanes, 1993; Kim and Shadlen, 1999), rule learning (Wallis et al., 2001), and anticipation of future reward (Watanabe, 1996; Leon and Shadlen, 1999; Shuler and Bear, 2006). In particular, persistent activity in PFC is critical for cognitive

functions that require integrating learned experience to predict future outcomes for flexible behavior (Miller and Cohen, 2001; Fuster, 2015).

Despite decades of theoretical work proposing how persistent activity may be generated, the mechanisms underlying its dynamics are still largely unknown. There are multiple challenges in understanding persistent activity. First, cortical neurons alone are biophysically incapable of maintaining information over behaviorally relevant timescales. Their intrinsic membrane time constants are on the order of tens of milliseconds, and postsynaptic potentials arising from synaptic input only last for hundreds of milliseconds (Koch, 1999). These are much shorter than the timescales of delay-related behaviors.

Second, persistent firing patterns in PFC neurons are highly irregular (Compte et al., 2003; Shafi et al., 2007). Early experimental and theoretical work suggested that persistent spiking rate changes during task delays are largely due to increased mean synaptic inputs, driving membrane potential (V_m) above spike threshold (Wang, 1999; Seung et al., 2000; Brunel and Wang, 2001). However, firing patterns in this regime are fairly regular, in contrast to experimentally observed irregular spike timing in delayed-response tasks (Compte et al., 2003; Shafi et al., 2007).

One solution to this puzzle lies in the dense synaptic connectivity between neurons. Cortical neurons receive extensive local and long-range synaptic inputs, and spikes are driven by integration of these inputs. Intracellular recordings in awake animals reveal that cortical neurons exhibit persistent depolarization (Steriade et al., 2001; Destexhe et al., 2003; Zagha and McCormick, 2014), and the observed irregular firing patterns are thought to be produced by fluctuations of their synaptic inputs (Softky and Koch, 1993; van Vreeswijk and Sompolinsky, 1996; Shadlen and Newsome, 1998). These characteristic features of cortical neurons in awake animals led to proposals that persistent firing rate changes with irregular timing arise not only by increasing the mean synaptic inputs but also through increased variance of synaptic inputs (Amit and Brunel, 1997; Renart et al., 2003, 2007).

Although these models are biologically plausible in explaining the underlying subthreshold dynamics of persistent activity, there is still a lack of experimental evidence supporting these models under conditions leading to persistent activity in PFC. Recently, intracellular recordings in the premotor cortex showed that increased spiking activity during motor preparation was correlated with V_m depolarization (Inagaki et al., 2019), but the direct relationship between V_m variability and the spike output irregularity remains to be tested. Moreover, in sensory and motor areas of mouse neocortex, large-amplitude V_m fluctuations observed during quiet resting states disappeared during movement, resulting in decreased V_m variability (Crochet and Petersen, 2006; Bennett et al., 2013; Polack et al., 2013; Schiemann et al., 2015), pointing to an apparent discrepancy between models and experimental data.

Another unexplained phenomenon—observed as soon as persistent activity was discovered in PFC (Fuster and Alexander, 1971, Funahashi et al., 1989)—is sustained decreases in firing rates during task delays. Recent intracellular recordings in sensory and motor cortical areas revealed layer- and cell-type-specific spike output and subthreshold dynamics

(Schiemann et al., 2015; Zhao et al., 2016). These observations suggest a possible laminar distribution of different subtypes of PFC neurons in mice (Douglas and Martin, 2004; Dembrow et al., 2010; Morishima et al., 2011).

Therefore, to study the subthreshold dynamics underlying reward-predictive persistent activity, we measured V_m in PFC neurons, while mice performed a delayed-reward task.

RESULTS

Persistent changes in V_m in anticipation of predicted rewards

To study persistent activity associated with reward anticipation, we trained thirsty, head-restrained mice on a classical trace-conditioning task (Figures 1A, 1B, and S1). Three olfactory cues, presented for 0.5 s, predicted one of three outcomes: no reward, a reward after a 1-s delay, or a reward after a 3-s delay. We measured V_m by making whole-cell patch clamp recordings (Figures 1C and 1D) (66 neurons from 39 mice) in the dorsomedial region of frontal cortex, previously characterized by its projections to mediodorsal thalamus, medial striatum, amygdala, ventral tegmental area, and dorsal raphe (Uylings and van Eden, 1990; Van De Werd et al., 2010). Injecting adeno-associated virus (AAV) into recording sites showed projections to each of these previously observed efferents (Figure S2). Axons did not target primary motor cortex, indicating that this area is distinct from neighboring secondary motor cortex (Hooks et al., 2013).

A subset of neurons (15 of 66, 24%) displayed significantly increased mean firing rates during delays to reward, relative to baseline (1 s prior to conditioned stimulus [CS]; pre-CS). This was the case for both 1-s (Figures 1E, 1F, and S3) (pre-CS: 1.23 ± 0.45; delay: 8.26 ± 1.22 spikes s^{-1} mean $\pm$ SEM; Wilcoxon signed rank test, $p < 0.01$) and 3-s delay trials (pre-CS: 0.87 ± 0.28; delay: 6.03 ± 2.32 spikes s^{-1}; Wilcoxon signed rank test, $p < 0.01$). After rewards, firing rates decreased significantly compared to delay periods (Figure S3) (1-s trials: 2.90 ± 1.26 spikes s^{-1}; Wilcoxon signed rank test, $p < 0.01$; 3-s trials: 3.15 ± 1.61 spikes s^{-1}; Wilcoxon signed rank test, $p < 0.01$). We found a similar proportion of neurons showing increased firing rates during the delay using extracellular recordings with tetrodes (Figure S4).

Subthreshold V_m, after removing action potentials from these neurons, showed a significant trial-type and time-window interaction (one-way repeated-measures ANOVA, $F_{2,84} = 11.18$, $p < 0.001$). There was significant depolarization in anticipation of reward, relative to pre-CS periods, during 1-s delays (pre-CS: -54.36 ± 1.36; delay: -48.04 ± 1.51 mV mean $\pm$ SEM; Wilcoxon signed rank test, $p < 0.01$), and 3-s delays (pre-CS: -54.28 ± 1.21; delay: -48.63 ± 1.38 mV; Wilcoxon signed rank test, $p < 0.01$). V_m significantly decreased after reward delivery compared to delay periods (Figure S3) (1-s trials: -51.37 ± 1.46 mV, Wilcoxon signed rank test, $p < 0.001$; 3-s trials: -51.56 ± 1.56 mV, $p < 0.01$). Depolarization was correlated with increases in firing rates (1-s: Pearson's $r = 0.541$, $p < 0.04$; 3-s: $r = 0.517$, $p < 0.05$) (Figure S3), indicating that firing rate increases in anticipation of reward were associated with V_m depolarization.

Next, we examined subthreshold V_m dynamics during the delay period. V_m rapidly depolarized following the onset of reward-predicting odors. Fitting logistic functions to the rise of V_m, we observed that the transition from baseline to a state of persistent depolarization was similar in 1-s (487 ± 45 ms after odor onset) and 3-s delay trials (684 ± 111 ms after odor onset, Wilcoxon signed rank test, $p > 0.1$) (Figure S3). After this transition, mean V_m remained in a sustained state of depolarization throughout the delay, showing no difference in V_m between the first 0.5 and last 0.5 s of the delay (1-s delay: first 0.5 s, -48.13 ± 1.47 mV, last 0.5 s, -47.96 ± 1.62 mV, Wilcoxon signed rank test, $p > 0.05$; 3-s delay: first 0.5 s, -48.33 ± 1.37 mV, last 0.5 s, -49.48 ± 1.46 mV, $p > 0.05$). These data indicate that reward-predicting cues evoked a persistent increase in V_m that appeared stable during delays to reward.

In tasks such as ours, reward anticipation and preparatory licking are correlated (Fiorillo et al., 2008; Cohen et al., 2012). As predicted, lick rates increased significantly in anticipation of reward compared to pre-CS (1-s: pre-CS, 0.50 ± 0.15, delay, 4.53 ± 0.39 licks s^{-1}; 3-s: pre-CS, 0.51 ± 0.06, delay, 4.17 ± 0.36 licks s^{-1}) (Figure S3) that was further maintained at an increased level during the consummatory period after the reward delivery (1-s: 5.91 ± 0.33; 3-s: 4.09 ± 0.36 licks s^{-1}). Previous studies have shown that a subset of neurons in the premotor cortex located adjacent to our area of study showed direct correlations between licking and neuronal activity in a lick/no-lick task (Komiyama et al., 2010), whereas another population showed ramping activity prior to lick onset only during motor preparation periods (Li et al., 2015; Inagaki et al., 2019). To test the temporal relationship between lick rates and V_m changes, we estimated cross-correlation coefficients for each neuron (Figure S3). We found that V_m and lick rates were positively correlated during reward-anticipation delays (Figure S3) (V_m-lick cross-correlation coefficient 1-s: Pearson's $r = 0.68 \pm 0.03$; 3-s: $r = 0.51 \pm 0.03$). However, the correlation decreased significantly during reward consumption (Figure S3) (1-s: Pearson's $r = 0.21 \pm 0.10$, Wilcoxon signed rank, $p < 0.001$; 3-s: $r = 0.23 \pm 0.07$, Wilcoxon signed rank, $p < 0.01$). After reward was delivered, licking persisted while mice harvested reward. At the same time, V_m depolarization terminated quickly, and decayed to baseline (Figures 2A–2C and S3) (1-s: 395 ± 118 ms after reward delivery; 3-s: 436 ± 71 ms; Wilcoxon signed rank test, $p > 0.1$). These results suggest that V_m depolarization was temporally correlated with licking during the reward anticipation, but the termination of depolarization was independent from ongoing consummatory licking.

Internally generated termination of sustained V_m changes

Clearly, a reward-predictive stimulus initiates persistent changes in V_m. What terminates it? Is an external stimulus, such as reward delivery, required?

We observed that some neurons (6 of 66) showed significant hyperpolarization during no-reward trials. Remarkably, these hyperpolarized states were sustained for ~3 s after cue offset—precisely the time of the longest reward delay—and then terminated in the absence of any external stimulus (Figures 2D–2F) (3.54 ± 0.18 s from CS offset, not significantly different from termination time on 3-s rewarded trials, Wilcoxon rank-sum test, $p > 0.77$). This demonstrates that persistent changes in V_m do not require an external stimulus to terminate. They can be terminated purely internally.

To study sustained V_m termination further, we designed a behavioral task with uncertain reward. One stimulus predicted no reward, a second stimulus predicted reward after a 3-s delay, and a third stimulus predicted reward after a 3-s delay with probability 0.5 (Figure 3). This task is well-suited to address the question of whether sustained V_m changes can be terminated by purely internal mechanisms because following the third stimulus, reward expectation is fixed for 3 s, until the mouse does or does not receive reward. In the latter case, if activity terminates around the time of expected reward, it must be due to a purely internal process.

Consistent with data from the previous task, there was a significant interaction of trial type and time window for firing rate (one-way repeated-measures ANOVA, $F_{2,42} = 5.95$, p < 0.01) and V_m ($F_{2,42} = 10.54$, p < 0.001), and neurons showed significantly increased firing rates and sustained depolarizations in V_m during the delays of trials with reward probability of 1 (Figure 3) (6 neurons from 5 mice; pre-CS: 1.55 ± 0.40 spikes s^{-1}, -50.96 ± 1.62 mV; delay: 8.31 ± 1.95 spikes s^{-1}, -47.75 ± 1.35 mV; Wilcoxon signed rank tests, p < 0.05). We did not observe these differences during no-reward trials (pre-CS: 1.73 ± 0.48 spikes s^{-1}, -51.00 ± 1.50 mV; delay: 1.32 ± 0.27 spikes s^{-1}, -51.51 ± 1.79 mV; Wilcoxon signed rank test, p > 0.05).

Critically, the same neurons also showed sustained increases of firing rates and V_m during the delays of trials with reward probabilities of 0.5 on both rewarded (pre-CS: 2.06 ± 0.60 spikes s^{-1}, -50.51 ± 1.87 mV; delay: 9.66 ± 2.59 spikes s^{-1}, -46.85 ± 2.14 mV; Wilcoxon signed rank tests, p < 0.05) and unrewarded trials (pre-CS: 2.06 ± 0.38 spikes s^{-1}, -50.93 ± 1.62 mV; delay: 9.63 ± 2.78 spikes s^{-1}, -46.15 ± 2.13 mV; Wilcoxon signed rank tests, p < 0.05). There were no significant differences in either firing rates (Wilcoxon signed rank tests, p > 0.96) or V_m (Wilcoxon signed rank tests, p > 0.32) during the delays of these trials, compared to those with reward probabilities of 1, indicating similar dynamics while anticipating a possible reward at a fixed time.

As predicted, based on the previous experiment (Figure 2), sustained increases in firing rates and depolarization significantly decreased after reward delivery, compared to the delay period. This occurred during trials with reward probabilities of 1 or 0.5, when reward was delivered (Figure S5) (post-unconditioned stimulus [US], $P(R) = 1$: 2.59 ± 0.57 spikes s^{-1}, Wilcoxon signed rank test, p < 0.05, -50.31 ± 1.60 mV, p < 0.05; $P(R) = 0.5$: 3.87 ± 1.27 spikes s^{-1}, p < 0.05, -49.06 ± 1.97 mV, p < 0.05). Remarkably, even when reward was omitted during $P(R) = 0.5$ trials, V_m and firing activity also abruptly terminated around the time of expected reward and showed significantly decreased firing rates and V_m compared to the delay period (Figures 3F and S5) (2.98 ± 0.99 spikes s^{-1}, p < 0.05, -49.46 ± 1.88 mV, p < 0.05). We measured the times at which V_m changes terminated and found that transitions from depolarized states to baseline V_m were similar during trials with or without reward (Figure 3G) ($P(R) = 1$: 0.48 ± 0.63 s after reward; $P(R) = 0.5$: rewarded trials, 0.53 ± 0.32 s after reward, unrewarded trials, 0.19 ± 0.20 s after expected reward time; Wilcoxon signed rank tests, p > 0.1). Termination of V_m changes after expected reward on reward-omission trials was not solely a result of licking terminating at that time; mice continued licking even after the expected time of reward (Figure S5). In addition, V_m and licking rates showed weak temporal correlations during reward-anticipation delays ($P(R) = 1$, Pearson's $r = 0.13$

± 0.12; $P(R) = 0.5$ rewarded trials, $r = 0.14 \pm 0.11$; $P(R) = 0.5$ unrewarded trials, $r = 0.32 \pm 0.10$) and reward consumption ($P(R) = 1$, $r = 0.02 \pm 0.20$; $P(R) = 0.5$, $r = -0.01 \pm 0.19$) or during reward omission ($P(R) = 0.5$, $r = -0.01 \pm 0.19$).

Thus, based on results from both experiments, we conclude that persistent changes in V_m could be terminated in the absence of reward by a mechanism purely internal to the nervous system.

Persistent activity increases are primarily fluctuation-driven

We have observed that increased average V_m was associated with increased firing rates during reward-anticipation delay periods. Increased average V_m could arise by either one or a combination of two factors: tonic depolarization or a change in dynamics of V_m fluctuation resulting from changes in the patterns of presynaptic network activity (Hô and Destexhe, 2000; Chance et al., 2002; Shu et al., 2003). In order for a spike to occur, V_m must reach spike threshold. It has been proposed that persistent spiking arises from sustained V_m depolarization over threshold (mean driven) or by increasing the magnitude of fluctuation to enhance the probability that V_m exceeds spike threshold (fluctuation driven) (Amit and Brunel, 1997; Renart et al., 2003, 2007). To distinguish between these mechanisms during periods of persistent spiking, we combined neurons from the two behavioral tasks, and analyzed 3-s delay trials of $P(R) = 1$ with at least an average of 2 spikes s^{-1} during reward delays.

The example neurons in Figure 4A show V_m depolarization with large fluctuations, greater than 10 mV in magnitude, only during reward delays, suggesting that increased V_m mean (E(V_m)) and variance (Var(V_m)) underlie persistent activity during this interval. To quantify this, we plotted histograms of V_m in each time window (pre-CS, delay, and post-US) (Figure 4B) and measured the mean and variance of each V_m distribution (Figures 4C and S6). E[V_m] was significantly higher during delay periods (-47.53 ± 0.89 mV) relative to pre-CS (-52.49 ± 0.98 mV, Wilcoxon signed rank test, p < 0.001) (Figure 4C) or post-US (Figure S6) (-49.97 ± 1.16 mV, p < 0.01), reflecting the depolarized states during the delay. In addition, Var[V_m] was significantly larger during the delay than pre-CS or post-US (pre-CS: 17.17 ± 2.84 mV^2; delay: 23.73 ± 3.49 mV^2; post-US: 15.62 ± 1.67 mV^2; Wilcoxon signed rank test, p < 0.01) (Figures 4C and S6), which indicates that increased average V_m during the delay was due to a combination of V_m depolarization and increased V_m fluctuation. Var[V_m] was not significantly different between pre-CS and post-US periods (Wilcoxon signed rank test, p > 0.4), demonstrating that increased Var[V_m] was selective for the delay period. Despite weak correlations between mean V_m and lick rates (Figures S3 and S5), it is possible that trial-by-trial licking accounted for V_m changes. We compared trial-by-trial lick rates and V_m mean and variance during reward delays and post-US. We found no clear relationship between V_m and lick rates in either time interval Figures S6 and S7). We further compared trial-by-trial licking and V_m in the probabilistic reward task (Figure S6) and found a similar lack of relationship, indicating that V_m changes were dissociated from licking.

Changes in V_m variance affect spike timing irregularity (Softky and Koch, 1993; van Vreeswijk and Sompolinsky, 1996; Shadlen and Newsome, 1998). Theoretical studies of spiking neurons predicted that irregular spike timing during persistent activity arises from a

network operating in a fluctuation-driven regime, whereas regular spiking is due to mean-driven activity (Renart et al., 2003, 2007; Roxin et al., 2011; Petersen and Berg, 2016).

To determine how V_m variation related to spike irregularity, we measured spike timing variability by calculating the coefficient of variation (CV, SD/mean) of inter-spike interval distributions. Spike timing increased in irregularity during the delay, more so than during the pre-CS period (pre-CS: 0.7 ± 0.07; delay: 1.57 ± 0.14; Wilcoxon signed rank test, p < 0.001), consistent with previous results in primate PFC (Compte et al., 2003). Increased spike irregularity was correlated with increases in Var[V_m] (Figure 4D; $r = 0.58$, p < 0.02), suggesting that persistent activity during reward delays operated in a fluctuation-driven regime.

In fluctuation-driven networks, the mean inputs are subthreshold, whereas they are suprathreshold in mean-driven networks (Gerstner et al., 2014). When mean inputs are subthreshold, spikes are driven by fluctuations in V_m, resulting in increased probability of spikes within normally silent ranges of V_m (Hô and Destexhe, 2000; Miller and Troyer, 2002; Fellous et al., 2003; Roxin et al., 2011). For each neuron, we measured the probability of an action potential (P(AP)) in 1-mV intervals of V_m. When V_m was well below threshold, P(AP) increased monotonically with V_m during the delay, whereas there was no spike at the same voltage during pre-CS periods (Figure 4E). When V_m was subthreshold, the relationship between V_m and P(AP) during the delay was approximated by a power law (Figure 4E), confirming the contribution of V_m fluctuations in generating spikes during delay periods (Hô and Destexhe, 2000; Miller and Troyer, 2002). Across neurons, we fit a function of the form $a(V_m)^b$ separately for the 3-s delay ($R^2 = 0.98$) and for the pre-CS period ($R^2 = 0.97$), considering the monotonically increasing values of firing probability (Figure 4E). Values of b were smaller for the pre-CS period ($b = 2.85 \pm 0.29$) than the 3-s delay ($b = 3.55 \pm 0.27$), suggesting increased neuronal responsiveness during delay periods.

When neurons fired while V_m was above threshold, however, there was no clear relationship between V_m and P(AP), suggesting that neurons were no longer in a fluctuation-driven regime. Cumulative distributions of V_m showed that, during the delay, V_m between action potentials was mostly subthreshold ($84\% \pm 0.04\%$ of the total time) (Figure 4E), but spent more time above threshold than during the pre-CS period ($98.0\% \pm 0.01\%$ of the total subthreshold time, Wilcoxon signed rank test, p < 0.001) or the post-US period ($92.0\% \pm 0.03\%$, p < 0.002) (Figure S7). This suggests that, although delay-period activity was primarily fluctuation-driven, some periods of spiking may be more regular due to epochs of mean-driven activity over threshold. To test this prediction, we calculated instantaneous spike irregularity (CV_2), to measure the regularity of spiking over time (Holt et al., 1996). We found that, indeed, when mean V_m in the 25 ms preceding spikes were suprathreshold, spike irregularity was lower than when spikes were generated following subthreshold V_m (CV_2 suprathreshold: 0.57 ± 0.09; subthreshold: 0.75 ± 0.07; Wilcoxon rank-sum test, p < 0.02) (Figures 4G–4I). Interestingly, V_m remained above threshold more often at the beginning of the delay than during later delay periods, suggesting that strong synaptic inputs initiated persistent activity, to be maintained further by fluctuations of synaptic inputs (Figure 4H).

Reward-predictive persistent hyperpolarization

Previous studies using extracellular recordings in PFC found neurons with suppression of firing rates relative to baseline during task delays (Fuster and Alexander, 1971; Funahashi et al., 1989). We also observed a subpopulation of neurons in our first task (Figure 1A) that showed persistent decreases in firing rates during reward-anticipation delays (Figure 5) (no reward: pre-CS, 5.13 ± 0.98 spikes s^{-1}, delay, 6.37 ± 1.21; 1-s trials: pre-CS, 5.07 ± 0.96, delay, 1.93 ± 0.54, post-US, 2.93 ± 0.76; 3-s trials: pre-CS, 5.03 ± 1.01, delay, 1.68 ± 0.40, post-US, 2.98 ± 0.67). There was a strong interaction between trial type and time window for firing rates (one-way repeated-measures ANOVA, $F_{2,54} = 14.45$, p < 0.001), showing significant differences between pre-CS and delay intervals in both 1-s (Wilcoxon signed rank test, p < 0.01) and 3-s reward trials (p < 0.01), as well as between delay and post-US in 3-s reward trials (p < 0.05) (Figure S8). These decreases in firing rates were accompanied by hyperpolarized V_m relative to the pre-CS period (no reward: pre-CS, -48.06 ± 0.51 mV, delay, -47.35 ± 0.68; 1-s trials: pre-CS, -47.86 ± 0.55, delay, -49.62 ± 0.71, post-US, -48.86 ± 0.64; 3-s trials: pre-CS, -48.06 ± 0.55, delay, -49.91 ± 0.78, post-US, -48.81 ± 0.6) (Figures 5D and S8). There was also a significant interaction between trial type and time window (one-way repeated-measures ANOVA, $F_{2,54} = 14.32$, p < 0.001), with significant differences in V_m between pre-CS and delay (Wilcoxon signed rank test, 1-s, p < 0.01, 3-s, p < 0.01) and delay and post-US during 3-s trials (p < 0.05) (Figure S8).

Similar to depolarizing neurons, V_m hyperpolarization was maintained throughout the delay and terminated after reward (Figures 5F and 5G) (1-s delay: 880 ± 221 ms after reward; 3-s delay: 894 ± 164 ms). In contrast to delay-period depolarizing neurons, hyperpolarizing neurons showed significant decreases in Var[V_m] during reward delays compared to pre-CS (no reward: pre-CS, 6.48 ± 0.87 mV, delay, 7.56 ± 0.86, Wilcoxon signed rank test, p = 0.84; 1-s: pre-CS, 6.71 ± 0.85, delay, 5.14 ± 0.78, p = 0.004; 3-s: pre-CS, 6.70 ± 0.83, delay, 5.00 ± 0.52, p = 0.02) and post-US (1-s trials post-US, 7.85 ± 1.76, p = 0.01, 3-s trials post-US, 6.90 ± 0.60, p = 0.002). In addition, neuronal input-output transformations, comparing 3-s delay to pre-CS periods, reflected the decrease in Var[V_m] (Figure 5H). There was a marked decrease in P(AP) in each V_m interval during delay periods ($b = 4.17 \pm 1.25$, 95% confidence interval [CI], $R^2 = 0.94$) compared to pre-CS ($b = 7.59 \pm 0.88$, 95% CI, $R^2 = 0.99$).

As we observed in depolarizing neurons, lick rates did not clearly correlate with V_m in hyperpolarizing neurons. We compared trial-by-trial lick rates with V_m mean and variance during delay and post-US periods in 3-s delay trials and found no clear relationship (Figure S8), suggesting that reward anticipation may be represented by neurons with persistent decreases in activity as well as by those with increases.

Recent work showed that neurons with persistent activity during the task delay were robust to perturbation (Kopec et al., 2015; Li et al., 2016; Inagaki et al., 2019), demonstrating that persistent activity was maintained by attractor-like network dynamics (Hopfield, 1982; Aksay et al., 2001; Brody et al., 2003). Many of these experiments focused on neurons with increased activity; less is known about the robustness of persistent activity changes in neurons with decreased activity during the delay (Li et al., 2016). To test this, we expressed the light-gated ion channel, channelrhodopsin-2 (ChR2) in dorsomedial frontal cortex and

made whole-cell recordings in neurons expressing ChR2 (see STAR Methods) that showed hyperpolarization during the delay (5 neurons, V_m relative to pre-CS: no reward trials, 0.53 ± 0.46 mV, 1-s delay trials, −3.28 ± 1.68 mV, 3-s delay trials, −4.14 ± 2.14 mV). We directly excited these neurons and surrounding areas for 500 ms during the delays to reward. If the hyperpolarized state is not maintained by network activity, a brief excitation could induce a prolonged depolarized state driven by intrinsic mechanisms such as plateau potentials (Major and Tank, 2004; Milojkovic et al., 2005; Major et al., 2008). However, we found that V_m rapidly returned to its hyperpolarized state after the stimulation without showing any prolonged depolarization (Figures 5I–5K). These results suggest that reward anticipation could be represented by neurons with persistent decreases in activity that is actively maintained by network dynamics.

Persistent activity is layer-specific

Is there a circuit logic for this persistent activity? Neocortical pyramidal neurons are organized into layers, comprising subpopulations of neurons that send outputs to distinct targets (Thomson and Bannister, 2003; Douglas and Martin, 2004). Locally, L2/3 neurons provide prominent excitatory input to L5 neurons, whereas L5 neurons form reciprocal connections with each other (Douglas and Martin, 2004; Otsuka and Kawaguchi, 2008; Brown and Hestrin, 2009; Morishima et al., 2011). L5 pyramidal neurons are further subdivided into two major groups based on their projection targets. Pyramidal tract (PT) neurons send axons predominantly to midbrain and brainstem structures and have somata mainly in lower L5. Intratelencephalic neurons (IT) send axons to striatum and contralateral cortex and have somata mostly, but not exclusively, in upper L5. These two subpopulations of L5 neurons have different somatodendritic morphologies and biophysical properties (Hattox and Nelson, 2007; Dembrow et al., 2010; Avesar and Gulledge, 2012; Oswald et al., 2013; Kawaguchi, 2017; Anastasiades et al., 2018).

To test whether the distinct subsets of neurons found in different layers contributed distinct patterns of persistent activity, we first compared firing patterns and their recording depth. Somatic depth estimated from the brain surface during recordings revealed that neurons that showed persistent V_m during delays to reward were mostly located in L5 (Figures 6A and 6B). Of those, depolarizing neurons were mostly in upper L5, whereas hyperpolarizing neurons were almost exclusively found in lower L5. In addition, a subset of neurons visualized with GFP expression after recordings supported the correlation between depth of soma and V_m modulation (Figure S9).

In contrast, most PFC neurons recorded in superficial layers (14 neurons, <300 μm from the pial surface) did not fire action potentials, and only a few neurons fired action potentials briefly following odor cues (3 of 14) (Figures 6C and 6D). All three response types clustered at different somatic depths (Kruskal-Wallis $\chi^2_2 = 16.3$, p < 0.001).

To verify this result with a larger sample size, we recorded extracellularly from 167 neurons in 2 mice using 64-channel silicon probes (Figure 6E). These electrodes had a linear arrangement of contacts, which allowed us to record simultaneously from neurons with known relative depths. Qualitatively, data from individual sessions confirmed the laminar organization observed with whole-cell recordings. Neurons at depths of less than 300 μm

from the pial surface typically showed brief (100–200 ms) excitation after reward-predicting cues. Neurons at deeper locations had sustained firing rate changes, with predominantly excitation in more superficial L5, and predominantly inhibition in deeper L5. To quantify this, we calculated the average firing rate of each neuron during 3-s delay trials, and clustered them into three groups, using principal component analysis. These three groups had firing dynamics matching the shapes of the three firing patterns observed in the whole-cell recordings: phasic excitation, sustained excitation, and sustained inhibition (Figure 6E) (Shuler and Bear, 2006; Huertas et al., 2015). We compared the depths of each population of neurons (Figures 6F and 6G) and found that those showing phasic excitation (median depth, 320 μm) were more superficial than those showing sustained excitation (median depth, 620 μm, Wilcoxon rank-sum test, $p < 0.05$). Neurons showing sustained excitation were found to be more superficial than those showing sustained inhibition (median depth, 770 μm, Wilcoxon rank-sum test, $p < 0.05$).

To further demonstrate the physiologically distinct subpopulations of L5 neurons, we measured intrinsic properties of pyramidal neurons in different depths of L5 in anesthetized mice (12 neurons from 6 mice). There was no difference in firing patterns during current injection (frequency-adaptation ratio at +200 pA: upper L5, 0.80 ± 0.23; lower L5, 0.87 ± 0.24, $p > 0.05$) and input resistance (upper L5, 153 ± 7.52 MΩ; lower L5, 155 ± 5.75 MΩ, $p > 0.05$). However, there was a positive correlation between recording depth and ''voltage sag'' ratio ($r = 0.61$, $p < 0.05$): neurons in lower L5 had a greater voltage sag in response to hyperpolarizing current steps than those in upper L5 (Figures 7A and 7B) (upper L5, 1.04 ± 0.01; lower L5, 1.09 ± 0.01, $p < 0.01$).

The voltage sag is thought to be generated by hyperpolarization-activated cyclic nucleotide-gated channels that generate the h-current, which helps generate resting V_m (Biel et al., 2009). Given this biophysical difference between upper and lower L5 neurons, we predicted that resting V_m would be higher in lower L5 neurons. To test this, we examined V_m in a pre-task window, before the start of the first trial. Pre-task V_m was significantly higher in neurons that were hyperpolarized during delays to reward (Figures 7C–7E) (−61.4 ± 2.2 mV) than in neurons that were depolarized during those delays (−50.7 ± 0.88 mV; Wilcoxon rank-sum test, $p < 0.01$). In addition, firing rates before the task began were higher in neurons hyperpolarized (Figure 7F) (5.16 ± 1.13 spikes s^{-1}) than depolarized (0.70 ± 0.44 spikes s^{-1}; Wilcoxon rank-sum test, $p < 0.01$) during delays to reward. These results suggest that two different functionally defined populations in L5 sublayers have distinct intrinsic properties that set their baseline activity before the task and can determine how they behave during reward-delay periods.

DISCUSSION

In our experiments, mice learned associations between cues and their outcomes. Our data showed that these learned relationships elicited activity that was stable and depended on predicted future events.

Theoretical and experimental studies have proposed that attractor dynamics could produce stable, persistent firing rate changes (Hopfield, 1982; Seung, 1996; Amit and Brunel, 1997;

Renart et al., 2007; Lim and Goldman, 2013). In particular, a Hopfield network model (Hopfield, 1982) suggested that neurons can learn synaptic input patterns and store them as a set of synaptic weights that can be retrieved by learned inputs. Odor cues that predicted reward with probability 0.5 initiated persistent activity changes that terminated precisely at the time of the expected reward, even when reward was omitted. In a subset of our neurons, odor cues that predicted no reward also generated persistent activity that lasted for ~3 s, which was precisely the longest expected interval between cue and reward. This learned timing may have been stored as a reference interval (Gibbon et al., 1984), which could then be used to predict the expected time of reward or the delay on no-reward trials (Watanabe et al., 2002). Our findings indicate that persistent activity can represent an internal state of expectation about future events, which can be terminated without an external signal.

We found that neurons with persistent increases in firing rates exhibited both V_m depolarization and increased V_m variance. In addition, these neurons showed highly irregular firing patterns, suggesting that, during the delay period, they were in a fluctuation-driven regime. Although theoretical studies predicted increased V_m variance without changing mean V_m to underlie fluctuation-driven persistent activity maintenance, it has been shown that both depolarization and increased variance enhance neuronal gain and responsiveness (Hô and Destexhe, 2000; Chance et al., 2002; Fellous et al., 2003). In our data, input-output curves were approximated by a power law, characterized by a non-zero spike rate for mean V_m below spike threshold (Hô and Destexhe, 2000; Miller and Troyer, 2002; Roxin et al., 2011), which shifted to the left during delay periods. These results indicate that V_m depolarization during reward anticipation provided increased neuronal responsiveness, with increased V_m variance further enhancing spiking probability. These appear to be the ingredients to maintain persistent, irregular spike output for multiple seconds in our tasks.

What could be the source of V_m fluctuations? Previous studies proposed that temporally balanced excitatory and inhibitory synaptic inputs maintain neural activity in stable states, producing irregular spiking outputs (Amit and Brunel, 1997; Shadlen and Newsome, 1998; Renart et al., 2003, 2007). In these models, if V_m drifts above threshold, it produces regular spiking output. However, if excitation is balanced by inhibition, net input currents fluctuate and spike trains are irregular. Interestingly, we observed that although V_m was mostly subthreshold during the delay period, V_m tended to lie above threshold, accompanied by more regular spike patterns, in the beginning of the delay. In addition, there was no threshold-linear relationship when spikes were initiated above threshold, suggesting that total input was saturated during these periods. Strong, transient excitation, generated by the reward-predicting stimuli, could account for this early suprathreshold activity. This pulse of excitation could act as a command signal for the network to recruit local recurrent excitation, balanced by inhibition, to maintain persistent activity (Seung et al., 2000; Brunel and Wang, 2001). The transition from mean- to fluctuation-driven activity may be similar to previous reports of changes in firing dynamics over the course of delay periods (Suzuki and Gottlieb, 2013; Spaak et al., 2017).

We found clear V_m modulation in the absence of licking (particularly on CS− trials), indicative of internally generated V_m fluctuations. However, it has also been shown that

activity across cortex covaries with movements (Musall et al., 2019; Stringer et al., 2019), forming a dynamic interaction between recurrent activity within cortex and inputs arising from movements. Although we found differences in the dynamics of V_m and licking, V_m fluctuations may be modulated by a combination of internally generated signals and the panoply of movements that cannot be summarized by a single variable.

Although most literature on persistent activity focused on neurons with increased firing rates (see Li et al., 2016), we also examined the subthreshold patterns of persistent activity in neurons that had decreased firing rates during reward-anticipation delays. These neurons showed persistent hyperpolarization and decreased V_m variance that was robust to brief local perturbation, indicating that decreased firing rates could be maintained by network activity. How, then, could the same network shared by depolarizing neurons maintain persistent hyperpolarization? One possibility is that strong excitatory input to the network also recruited local inhibition, resulting in hyperpolarization in a subpopulation of neurons. It has been reported that L5 cortical neurons form different local inhibitory connections, eliciting distinct patterns of responses in different subtypes of neurons (Lee et al., 2014; Morishima et al., 2017). Thus, the hyperpolarization we observed during reward delays in a subpopulation of neurons may have resulted from increased synaptic inhibition onto these neurons. On the other hand, it has been shown that increasing total synaptic conductance—without disrupting the balance between excitation and inhibition—leads to decreases in neural responsiveness via shunting inhibition (Hô and Destexhe, 2000; Chance et al., 2002). Indeed, we observed that input-output functions of hyperpolarizing neurons were shifted to the right, suggesting decreased neural responsiveness during delays. Thus, network-generated synaptic activity in these neurons could increase total synaptic conductance, including inhibitory synapses, further enhancing persistent hyperpolarization through shunting inhibition.

In the cortical area we studied, generation and maintenance of persistent activity was organized anatomically. It has been reported that neurons in different layers use different coding schemes: L2/3 neurons are sparsely active, whereas L5 neurons fire persistently on stimulation (de Kock et al., 2007; Niell and Stryker, 2008; Sakata and Harris, 2009; Schiemann et al., 2015). Similarly, we observed that neurons recorded in superficial layers were mostly silent, and only a few neurons showed brief excitation following reward-predicting stimuli. In contrast, although a large population of neurons was either silent or did not show task-relevant modulation during the delay, ~40% of neurons recorded in deeper layers showed persistent activity throughout the delay. The different connectivity of neurons in different layers may contribute to their coding schemes. Within a cortical column, L2/3 neurons often form interlaminar, feedforward synapses with L5 neurons, and L5 neurons form strong intralaminar recurrent connections (Douglas and Martin, 2004; Otsuka and Kawaguchi, 2008; Brown and Hestrin, 2009; Morishima et al., 2011). Brief excitation in L2/3 neurons could, therefore, act as a trigger to initiate persistent activity through recurrent synaptic networks in L5.

Within L5, we observed two distinct patterns of persistent activity. Upper L5 neurons showed increased firing rates with depolarized V_m, whereas lower L5 neurons showed decreased firing rates with hyperpolarized V_m. Rather than reflecting two tails of a

distribution of firing rates, we propose that differences in synaptic and intrinsic biophysical properties could explain these two opposing dynamics. Within L5, neurons are subdivided based on their projection targets and cluster into different sublayers (Morishima and Kawaguchi, 2006; Wang et al., 2006; Dembrow et al., 2010; Morishima et al., 2011; Lee et al., 2014). Furthermore, differences in intrinsic properties further differentiate their responses to synaptic inputs (Dembrow et al., 2010; Anastasiades et al., 2018). In agreement with these previous studies, we found that the intrinsic properties of neurons varied by response dynamics and cortical layer. The implication of opposing activity patterns in two sublayers of L5 is that their downstream targets receive distinct signals and form feedback loops that could maintain persistent activity. These efferents include thalamus (Schiemann et al., 2015; Bolkan et al., 2017; Guo et al., 2017), contralateral cortex (Li et al., 2016), and neuro-modulators such as norepinephrine (Wang et al., 2007; Dembrow et al., 2010; Schiemann et al., 2015; Breton-Provencher and Sur, 2019), acetylcholine (Egorov et al., 2002; Dembrow et al., 2010; Rahman and Berger, 2011; Baker et al., 2018), dopamine (Williams and Goldman-Rakic, 1995), and serotonin (Williams et al., 2002; Avesar and Gulledge, 2012; Stephens et al., 2014; Geddes et al., 2016; Zhou et al., 2017). Notably, the present laminar organization differs from that found in humans (Finn et al., 2019) and monkeys (Goldman-Rakic, 1995; Wang et al., 2013; Yang et al., 2013), likely reflecting differences across species (DeFelipe, 2011) and tasks.

Persistent activity is a general strategy for nervous systems to represent behaviorally relevant states over biophysically long timescales. As predicted from theoretical studies, the persistent activities elicited from learned association between reward-predicting cues and rewards primarily operated within a fluctuation-driven regime. These findings contrast with previous studies in sensory-motor cortex showing increased firing rates that were associated with decreased V_m variability. Although the reason for this discrepancy is unclear, one possibility is that because sensory-motor cortex is important for integrating sensory inputs and controlling motor output, they increase signal-to-noise ratio by reducing V_m variance, thereby enhancing signal detection. By contrast, because PFC has an important role in executive control, the increased V_m variance we observed may reflect the dynamics of multiple components of cognitive control, such as motivation, attention, and time estimation. Thus, it will be important for future studies to investigate the subthreshold mechanisms underlying various cognitive functions in PFC, which is critical for refining models of cognition.

STAR★METHODS

RESOURCE AVAILABILITY

Lead contact—Further information and requests for reagents should be directed to the Lead Contact, Jeremiah Y. Cohen (jeremiah.cohen@jhmi.edu).

Materials availability—This study did not generate new unique reagents.

Data and code availability—Data and computer code for experimental control and data analysis is available upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Animals—Six- to 12-week-old male C57BL/6J mice (The Jackson Laboratory, 000664) were used for all electrophysiological and behavioral experiments. All surgical and experimental procedures were in accordance with the *National Institutes of Health Guide for the Care and Use of Laboratory Animals* and approved by the Johns Hopkins University Animal Care and Use Committee.

METHOD DETAILS

Surgery—For whole-cell electrophysiological recordings, mice were surgically implanted with custom-made titanium head plates using dental adhesive (C&B-Metabond, Parkell) under isoflurane anesthesia (1.0%–1.5% in O_2). In a subset of mice, viruses were injected targeting dorsal medial PFC (3.0 mm anterior to bregma, 0.5 mm lateral from the midline). Following head plate implantation, the surface of the skull was covered with silicone elastomer (Kwik-Cast, WPI). For extracellular electrophysiological recordings, a custom-made microdrive containing 8–16 tetrodes made from nichrome wire (PX000004, Sandvik) positioned inside 39 ga polyimide guide tubes were implanted, targeted toward the same coordinates as above. Surgery was conducted under aseptic conditions and analgesia (ketoprofen, 5 mg kg^{-1} and buprenorphine, 0.05–0.1 mg kg^{-1}) was administered postoperatively. After at least one week of recovery, mice were water-restricted in their home cage with free access to food. Weight was monitored and maintained within 80% of their full body weight.

Behavioral task—Mice were head-restrained and positioned in a 38.1 mm acrylic tube in a sound-attenuated chamber. During each conditioning session, each trial began with the presentation of one of 3 different olfactory stimuli (A, B, and C), delivered for 0.5 s. Odor A was followed by an inter-trial interval (ITI). Odor B was followed by a 1 s trace delay, and then a reward (4 μL of 5% sucrose in water). Odor C was followed by a 3 s delay, and then a reward. ITIs were drawn from an exponential distribution with a rate parameter of 0.3, with a maximum cutoff of 5 s. For the task with reward probabilities of 0.5, reward was delivered on randomly chosen trials, but no more than 3 rewards were delivered consecutively. Odors were delivered with a custom-made olfactometer (Bari et al., 2019).

Each odor was dissolved in mineral oil at 1:10 dilution. Diluted odors (40 μl) were placed on filter-paper housing (Whatman, 2.7 μm pore size). Odors were p-cymene, (−)-carvone, (+)-limonene, and acetophenone, and differed across mice. Odorized air was further diluted with filtered air by 1:10 to produce a 1.0 L min^{-1} flow rate. Licks were detected by charging a capacitor (MPR121QR2, Freescale) or using a custom circuit. Task events were controlled with a microcontroller (ATmega16U2 or ATmega328). Mice were housed on a 12h dark/12h light cycle (dark from 08:00–20:00) and each performed behavioral tasks at the same time of day, between 09:00 and 18:00.

Extracellular recordings—Extracellular signals were recorded bilaterally from multiple neurons simultaneously at 30 kHz using a custom-built screw-driven microdrive with 8 tetrodes (32 channel total). All tetrodes were gold-plated to an impedance of 200–300 kΩ prior to implantation. Spikes were bandpass-filtered between 0.3–6 kHz and sorted online

and offline using Spikesort 3D (Neuralynx, Inc.) and custom software written in MATLAB. To measure isolation quality of individual units, we calculated the L-ratio (Schmitzer-Torbert et al., 2005) and fraction of inter-spike interval (ISI) violations within a 2 ms refractory period. All single units included in the dataset had an L-ratio less than 0.05 and fewer than 1% ISI violations. We collected data from 1,065 neurons from 3 mice in these experiments.

For silicon-probe recordings, we made acute penetrations with 64-channel probes (H3, Cambridge Neurotech) at 5-degree angles relative to the surface of cortex, at depths of 1 mm. Signals were acquired at 20 kHz, bandpass filtered between 0.1 and 6 kHz (Intan Technologies, RHD2164 headstage), and sorted offline using SpikeSort 3D. Depth estimates were corrected by 0.4 mm due to tissue compression during silicon probe penetrations. This value was drawn from post hoc reconstructions.

Patch-clamp recordings—For whole-cell recordings, mice were anesthetized with isoflurane (1%–1.5%) and a craniotomy was made over medial PFC (3.0 mm anterior to bregma, 0.5 mm lateral from the midline). Both hemispheres were sampled. A thin layer of Kwik-Cast (WPI) was applied over the skull, mice were returned to their home cage, and were given at least 2 hr to recover before being placed in the behavior apparatus. Glass electrodes (5–7 MΩ, fabricated using a PC-10 puller, Narishige) were filled with an internal solution composed of the following (in mM): 135 potassium gluconate, 4 potassium chloride, 10 sodium phosphocreatine, 4 ATP magnesium salt, 0.3 GTP sodium salt hydrate, 10 HEPES; pH was adjusted to 7.25 using KOH. In a subset of recordings, pCAG-GFP (50–100 ng μl^{-1}), was included in the internal solution for post hoc cell identification and reconstruction. pCAG-GFP (Matsuda and Cepko, 2004) was a gift from Connie Cepko (Addgene plasmid 11150).

Electrophysiological signals were low-pass filtered at 10 kHz (Multiclamp 700B, Molecular Devices) and acquired at 20 kHz on a PCIe-6323 (National Instruments) using Ephus (Vidrio Technologies, LLC). Standard blind patch methods were used to obtain whole cell recordings. Pipettes were lowered into the brain while high positive pressure (100 mmHg) was applied. Once in the brain, positive pressure was reduced (40 mmHg) and the pipette was advanced down slowly (approximately 2 μm s^{-1}) to search for neurons. If the pipette resistance increased abruptly by 10%–20%, positive pressure was released and whole-cell configuration was obtained when resistance was > 1 GΩ and stable. Series resistance was < 100 MΩ. After successful break in, the recording mode was switched to current clamp ($I = 0$), and the behavior session was initiated if the membrane potentials were stable over a 1 min period after break in. The recording was terminated if V_m became depolarized above −45 mV, or when the mouse stopped performing the task. After recordings, the patch pipette was slowly withdrawn, a thin layer of Kwik-Cast (WPI) applied again, and the animal returned to its home cage to recover. To measure the depth-dependent sag ratio and input resistance, a separate group of mice was anesthetized with a low level of isoflurane (< 1%), current step recording was performed, and post-recording procedures were followed.

Viral injections—To express ChR2 (500 nL for electrophysiological experiments), eGFP, or mCherry (30 nL each for anatomical experiments) in PFC neurons, we pressure-injected

each virus (bilaterally for ChR2) into PFC at a rate of approximately 1 nL s^{-1} (MMO-220A, Narishige). The injection pipette was left in place for > 5 min between each injection. The craniotomy was covered with silicone elastomer (Kwik-Cast, WPI). pAAV-CaMKIIa-hChR2(H134R)-EYFP (Lee et al., 2010) was a gift from Karl Deisseroth (Addgene viral prep 26969-AAV5; http://addgene.org/26969; RRID:Addgene_26969). pENN.AAV.CB7.CI.mCherry.WPRE.RBG was a gift from James M. Wilson (Addgene viral prep 105544-AAV1; http://addgene.org/105544; RRID:Addgene_105544). pENN.AAV.CB7.CI.eGFP.WPRE.rBG was a gift from James M. Wilson (Addgene 105542-AAV1; http://addgene.org/105542; RRID:Addgene_105542).

Optogenetic stimulation with recordings—Mice that were injected with AAV-CaMKII-ChR2 were used for optogenetic perturbation experiments. The optic fiber was inserted into the recording pipette, enabling direct light projection to the recorded neuron (Katz et al., 2013). After a whole-cell recording was obtained, a train (10 pulses, 3 ms at 10 Hz) of 473 nm light (Laserglow) stimuli was delivered using a shutter in series with the laser (Uniblitz) to induce action potentials to identify ChR2-expressing neurons. If the action potentials were elicited reliably with a short latency (< 3 ms), the light irradiance was lowered to a level at which the membrane potential crossed action potential threshold (except in one cell) to generate more than 1 action potential but not bursting during a long pulse (500 ms). Light stimulation during the delay was delivered in 30%–40% of trials, chosen randomly.

Histology—Seven to 10 days after recordings, mice were euthanized with an overdose of ketamine (100 mg kg^{-1}), exsanguinated with saline, and perfused with 4% paraformaldehyde. The brain was removed, post-fixed in the perfusion solution, and cut in 100-μm-thick sagittal sections. For immunostaining for GFP, rabbit anti-GFP (Invitrogen, 1:1000, 2 hr) primary antibody, followed by donkey Alexa 488 anti-rabbit (Invitrogen, 1:1000, overnight) secondary antibody was used. All confocal images were taken as tiled z stacks using a confocal microscope (Zeiss LSM 800, ZEN software) at 10X or 20X and reconstructions were done using ImageJ or Fiji (Schindelin et al., 2012).

QUANTIFICATION AND STATISTICAL ANALYSIS

All analyses were performed with MATLAB (Mathworks) and R (http://www.r-project.org/). All data are presented as mean ± SEM unless reported otherwise. All statistical tests were two-sided, and multiple-comparison (Bonferroni) corrections were used. For nonparametric tests, the Wilcoxon rank sum test was used, unless data were paired, in which case the Wilcoxon signed rank test was used.

For subthreshold V_m measurements, spikes were removed from raw traces by truncating data above spike threshold. Mean V_m was calculated by averaging spike-removed V_m traces in each time window. E[V_m] and Var[V_m] in Figure 4 were estimated from the probability distributions of V_m in each time window. Spike threshold was calculated as the value of V_m when $d^2 V_m/dt^2$ of each spike reached its maximum. Mean spike threshold of spontaneous action potentials was used to estimate the time V_m spent below or over threshold.

To determine the relationship between V_m and spike output, we first selected spikes that were not preceded by other spikes in a 30-ms window and calculated spike-triggered V_m by averaging V_m over 10 ms prior to the spike. Spike probability was estimated as a function of V_m by calculating the probability of spike-triggered V_m in 1-mV bins (Jahn et al., 2011; Petersen and Berg, 2016). Power-law fits were based on individual measurements of V_m and the estimated spike probability of each neuron, and fit over the range of V_m.

Time-dependent CV_2 was defined as $CV_2(i) = 2|ISI(i+1)|/(ISI(i)+ISI(i+1))$, where $ISI(i)$ is the i^{th} interspike interval (Holt et al., 1996). CV_2 above threshold was defined by intervals during which the mean V_m in the 25 ms preceding the spike was over the average spike threshold.

Sag ratio was measured by hyperpolarizing current steps (1 s, −200 pA, holding at −60 mV) and calculated as a ratio between the peak amplitude of the initial response (0–0.25 s) and the steady state response (0.75–1 s). Input resistance was obtained by calculating the slope from the current-voltage curve of the steady state response of hyperpolarizing current steps from −200 pA to 0 pA (100 pA increments). A frequency adaptation index was calculated as the ratio of the first ISI to the last ISI of the spike trains evoked by a depolarizing current injection.

Neurons were classified as showing persistent changes in V_m if it was significantly different from baseline using t tests. Neurons were classified as showing persistent changes in firing rates if area under receiver operating characteristic curves were > 0.65 or < 0.35 throughout the delay on 3 s delay trials. To measure persistent V_m onset and offset times, we fit sigmoidal curves to V_m traces in each window (Polack et al., 2013). The onset and offset times were defined as the time at which the sigmoids reached half of their maximum.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

REFERENCES

Aksay E, Gamkrelidze G, Seung HS, Baker R, and Tank DW (2001). In vivo intracellular recording and perturbation of persistent activity in a neural integrator. Nat. Neurosci 4, 184–193. [PubMed: 11175880]

Amit DJ, and Brunel N (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb. Cortex 7, 237–252. [PubMed: 9143444]

Anastasiades PG, Marlin JJ, and Carter AG (2018). Cell-type specificity of callosally evoked excitation and feedforward inhibition in the prefrontal cortex. Cell Rep 22, 679–692. [PubMed: 29346766]

Avesar D, and Gulledge AT (2012). Selective serotonergic excitation of callosal projection neurons. Front. Neural Circuits 6, 12. [PubMed: 22454619]

Baker AL, O'Toole RJ, and Gulledge AT (2018). Preferential cholinergic excitation of corticopontine neurons. J. Physiol 596, 1659–1679. [PubMed: 29330867]

Bari BA, Grossman CD, Lubin EE, Rajagopalan AE, Cressy JI, and Cohen JY (2019). Stable representations of decision variables for flexible behavior. Neuron 103, 922–933.e7. [PubMed: 31280924]

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

Biel M, Wahl-Schott C, Michalakis S, and Zong X (2009). Hyperpolarization-activated cation channels: from genes to function. Physiol. Rev 89, 847–885. [PubMed: 19584315]

Bolkan SS, Stujenske JM, Parnaudeau S, Spellman TJ, Rauffenbart C, Abbas AI, Harris AZ, Gordon JA, and Kellendonk C (2017). Thalamic projections sustain prefrontal activity during working memory maintenance. Nat. Neurosci 20, 987–996. [PubMed: 28481349]

Breton-Provencher V, and Sur M (2019). Active control of arousal by a locus coeruleus GABAergic circuit. Nat. Neurosci 22, 218–228. [PubMed: 30643295]

Brody CD, Romo R, and Kepecs A (2003). Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. Curr. Opin. Neurobiol 13, 204–211. [PubMed: 12744975]

Brown SP, and Hestrin S (2009). Intracortical circuits of pyramidal neurons reflect their long-range axonal targets. Nature 457, 1133–1136. [PubMed: 19151698]

Brunel N, and Wang XJ (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. J. Comput. Neurosci 11, 63–85. [PubMed: 11524578]

Bush RR, and Mosteller F (1951). A mathematical model for simple learning. Psychol. Rev 58, 313–323. [PubMed: 14883244]

Chance FS, Abbott LF, and Reyes AD (2002). Gain modulation from background synaptic input. Neuron 35, 773–782. [PubMed: 12194875]

Cohen JY, Haesler S, Vong L, Lowell BB, and Uchida N (2012). Neuron-type-specific signals for reward and punishment in the ventral tegmental area. Nature 482, 85–88. [PubMed: 22258508]

Compte A, Constantinidis C, Tegner J, Raghavachari S, Chafee MV, Goldman-Rakic PS, and Wang XJ (2003). Temporally irregular mnemonic persistent activity in prefrontal neurons of monkeys during a delayed response ktas. J. Neurophysiol 90, 3441–3454. [PubMed: 12773500]

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

de Kock CPJ, Bruno RM, Spors H, and Sakmann B (2007). Layer- and cell-type-specific suprathreshold stimulus representation in rat primary somatosensory cortex. J. Physiol. 581, 139–154. [PubMed: 17317752]

DeFelipe J (2011). The evolution of the brain, the human nature of cortical circuits, and intellectual creativity. Front. Neuroanat 5, 29. [PubMed: 21647212]

Dembrow NC, Chitwood RA, and Johnston D (2010). Projection-specific neuromodulation of medial prefrontal cortex neurons. J. Neurosci 30, 16922–16937. [PubMed: 21159963]

Destexhe A, Rudolph M, and Paré D (2003). The high-conductance state of neocortical neurons in vivo. Nat. Rev. Neurosci 4, 739–751. [PubMed: 12951566]

Douglas RJ, and Martin KAC (2004). Neuronal circuits of the neocortex. Annu. Rev. Neurosci 27, 419–451. [PubMed: 15217339]

Egorov AV, Hamam BN, Fransén E, Hasselmo ME, and Alonso AA (2002). Graded persistent activity in entorhinal cortex neurons. Nature 420, 173–178. [PubMed: 12432392]

Fellous JM, Rudolph M, Destexhe A, and Sejnowski TJ (2003). Synaptic background noise controls the input/output characteristics of single cells in an in vitro model of in vivo activity. Neuroscience 122, 811–829. [PubMed: 14622924]

Finn ES, Huber L, Jangraw DC, Molfese PJ, and Bandettini PA (2019). Layer-dependent activity in human prefrontal cortex during working memory. Nat. Neurosci 22, 1687–1695. [PubMed: 31551596]

Fiorillo CD, Newsome WT, and Schultz W (2008). The temporal precision of reward prediction in dopamine neurons. Nat. Neurosci 11, 966–973. [PubMed: 18660807]

Funahashi S, Bruce CJ, and Goldman-Rakic PS (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J. Neurophysiol 61, 331–349. [PubMed: 2918358]

Fuster J (2015). The Prefrontal Cortex (Academic Press).

Fuster JM, and Alexander GE (1971). Neuron activity related to short-term memory. Science 173, 652–654. [PubMed: 4998337]

Geddes SD, Assadzada S, Lemelin D, Sokolovski A, Bergeron R, Haj-Dahmane S, and Béïque JC (2016). Target-specific modulation of the descending prefrontal cortex inputs to the dorsal raphe nucleus by cannabinoids. Proc. Natl. Acad. Sci. USA 113, 5429–5434. [PubMed: 27114535]

Gerstner W, Kistler WM, Naud R, and Paninski L (2014). Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition (Cambridge University Press).

Gibbon J, Church RM, and Meck WH (1984). Scalar timing in memory. Ann. N Y Acad. Sci 423, 52–77. [PubMed: 6588812]

Goldman-Rakic PS (1995). Cellular basis of working memory. Neuron 14, 477–485. [PubMed: 7695894]

Guo ZV, Inagaki HK, Daie K, Druckmann S, Gerfen CR, and Svoboda K (2017). Maintenance of persistent activity in a frontal thalamocortical loop. Nature 545, 181–186. [PubMed: 28467817]

Hattox AM, and Nelson SB (2007). Layer V neurons in mouse cortex projecting to different targets have distinct physiological properties. J. Neurophysiol 98, 3330–3340. [PubMed: 17898147]

Hô N, and Destexhe A (2000). Synaptic background activity enhances the responsiveness of neocortical pyramidal neurons. J. Neurophysiol 84, 1488–1496. [PubMed: 10980021]

Holt GR, Softky WR, Koch C, and Douglas RJ (1996). Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. J. Neurophysiol 75, 1806–1814. [PubMed: 8734581]

Hooks BM, Mao T, Gutnisky DA, Yamawaki N, Svoboda K, and Shepherd GMG (2013). Organization of cortical and thalamic input to pyramidal neurons in mouse motor cortex. J. Neurosci 33, 748–760. [PubMed: 23303952]

Hopfield JJ (1982). Neural networks and physical systems with emergent collective computational abilities. Proc. Natl. Acad. Sci. USA 79, 2554–2558. [PubMed: 6953413]

Huertas MA, Hussain Shuler MG, and Shouval HZ (2015). A simple network architecture accounts for diverse reward time responses in primary visual cortex. J. Neurosci 35, 12659–12672. [PubMed: 26377457]

Hull CL (1943). Principles of Behavior: An Introduction to Behavior Theory (Appleton-Century-Crofts).

Inagaki HK, Fontolan L, Romani S, and Svoboda K (2019). Discrete attractor dynamics underlies persistent activity in the frontal cortex. Nature 566, 212–217. [PubMed: 30728503]

Jahn P, Berg RW, Hounsgaard J, and Ditlevsen S (2011). Motoneuron membrane potentials follow a time inhomogeneous jump diffusion process. J. Comput. Neurosci 31, 563–579. [PubMed: 21479618]

Katz Y, Yizhar O, Staiger J, and Lampl I (2013). Optopatcher–an electrode holder for simultaneous intracellular patch-clamp recording and optical manipulation. J. Neurosci. Methods 214, 113–117. [PubMed: 23370312]

Kawaguchi Y (2017). Pyramidal cell subtypes and their synaptic connections in layer 5 of rat frontal cortex. Cereb. Cortex 27, 5755–5771. [PubMed: 29028949]

Kim JN, and Shadlen MN (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci 2, 176–185. [PubMed: 10195203]

Koch C (1999). Biophysics of Computation: Information Processing in Single Neurons (Oxford University Press).

Komiyama T, Sato TR, O'Connor DH, Zhang YX, Huber D, Hooks BM, Gabitto M, and Svoboda K (2010). Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. Nature 464, 1182–1186. [PubMed: 20376005]

Kopec CD, Erlich JC, Brunton BW, Deisseroth K, and Brody CD (2015). Cortical and subcortical contributions to short-term memory for orienting movements. Neuron 88, 367–377. [PubMed: 26439529]

Kubota K, and Niki H (1971). Prefrontal cortical unit activity and delayed alternation performance in monkeys. J. Neurophysiol 34, 337–347. [PubMed: 4997822]

Lee JH, Durand R, Gradinaru V, Zhang F, Goshen I, Kim DS, Fenno LE, Ramakrishnan C, and Deisseroth K (2010). Global and local fMRI signals driven by neurons defined optogenetically by type and wiring. Nature 465, 788–792. [PubMed: 20473285]

Lee AT, Gee SM, Vogt D, Patel T, Rubenstein JL, and Sohal VS (2014). Pyramidal neurons in prefrontal cortex receive subtype-specific forms of excitation and inhibition. Neuron 81, 61–68. [PubMed: 24361076]

Leon MI, and Shadlen MN (1999). Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. Neuron 24, 415–425. [PubMed: 10571234]

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

Li N, Daie K, Svoboda K, and Druckmann S (2016). Robust neuronal dynamics in premotor cortex during motor planning. Nature 532, 459–464. [PubMed: 27074502]

Lim S, and Goldman MS (2013). Balanced cortical microcircuitry for maintaining information in working memory. Nat. Neurosci 16, 1306–1314. [PubMed: 23955560]

Major G, and Tank D (2004). Persistent neural activity: prevalence and mechanisms. Curr. Opin. Neurobiol 14, 675–684. [PubMed: 15582368]

Major G, Polsky A, Denk W, Schiller J, and Tank DW (2008). Spatiotemporally graded NMDA spike/plateau potentials in basal dendrites of neocortical pyramidal neurons. J. Neurophysiol 99, 2584–2601. [PubMed: 18337370]

Matsuda T, and Cepko CL (2004). Electroporation and RNA interference in the rodent retina in vivo and in vitro. Proc. Natl. Acad. Sci. USA 101, 16–22. [PubMed: 14603031]

Miller EK, and Cohen JD (2001). An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci 24, 167–202. [PubMed: 11283309]

Miller KD, and Troyer TW (2002). Neural noise can explain expansive, power-law nonlinearities in neural response functions. J. Neurophysiol 87, 653–659. [PubMed: 11826034]

Milojkovic BA, Radojicic MS, and Antic SD (2005). A strict correlation between dendritic and somatic plateau depolarizations in the rat prefrontal cortex pyramidal neurons. J. Neurosci 25, 3940–3951. [PubMed: 15829646]

Morishima M, and Kawaguchi Y (2006). Recurrent connection patterns of corticostriatal pyramidal cells in frontal cortex. J. Neurosci 26, 4394–4405. [PubMed: 16624959]

Morishima M, Morita K, Kubota Y, and Kawaguchi Y (2011). Highly differentiated projection-specific cortical subnetworks. J. Neurosci 31, 10380–10391. [PubMed: 21753015]

Morishima M, Kobayashi K, Kato S, Kobayashi K, and Kawaguchi Y (2017). Segregated excitatory-inhibitory recurrent subnetworks in layer 5 of the rat frontal cortex. Cereb. Cortex 27, 5846–5857. [PubMed: 29045559]

Musall S, Kaufman MT, Juavinett AL, Gluf S, and Churchland AK (2019). Single-trial neural dynamics are dominated by richly varied movements. Nat. Neurosci 22, 1677–1686. [PubMed: 31551604]

Niell CM, and Stryker MP (2008). Highly selective receptive fields in mouse visual cortex. J. Neurosci 28, 7520–7536. [PubMed: 18650330]

Oswald MJ, Tantirigama MLS, Sonntag I, Hughes SM, and Empson RM (2013). Diversity of layer 5 projection neurons in the mouse motor cortex. Front. Cell. Neurosci 7, 174. [PubMed: 24137110]

Otsuka T, and Kawaguchi Y (2008). Firing-pattern-dependent specificity of cortical excitatory feed-forward subnetworks. J. Neurosci 28, 11186–11195. [PubMed: 18971461]

Petersen PC, and Berg RW (2016). Lognormal firing rate distribution reveals prominent fluctuation-driven regime in spinal motor networks. eLife 5, e18805. [PubMed: 27782883]

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

Rahman J, and Berger T (2011). Persistent activity in layer 5 pyramidal neurons following cholinergic activation of mouse primary cortices. Eur. J. Neurosci 34, 22–30. [PubMed: 21645136]

Renart A, Song P, and Wang XJ (2003). Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. Neuron 38, 473–485. [PubMed: 12741993]

Renart A, Moreno-Bote R, Wang XJ, and Parga N (2007). Mean-driven and fluctuation-driven persistent activity in recurrent networks. Neural Comput 19, 1–46. [PubMed: 17134316]

Rescorla RA, and Wagner AR (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In Classical Conditioning II: Current Research and Theory, Black AH and Prokasy WF, eds. (Appleton-Century-Crofts), pp. 64–99.

Roxin A, Brunel N, Hansel D, Mongillo G, and van Vreeswijk C (2011). On the distribution of firing rates in networks of cortical neurons. J. Neurosci 31, 16217–16226. [PubMed: 22072673]

Sakata S, and Harris KD (2009). Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. Neuron 64, 404–418. [PubMed: 19914188]

Schall JD, and Hanes DP (1993). Neural basis of saccade target selection in frontal eye field during visual search. Nature 366, 467–469. [PubMed: 8247155]

Schiemann J, Puggioni P, Dacre J, Pelko M, Domanski A, van Rossum MCW, and Duguid I (2015). Cellular mechanisms underlying behavioral state-dependent bidirectional modulation of motor cortex output. Cell Rep 11, 1319–1330. [PubMed: 25981037]

Schindelin J, Arganda-Carreras I, Frise E, Kaynig V, Longair M, Pietzsch T, Preibisch S, Rueden C, Saalfeld S, Schmid B, et al. (2012). Fiji: an open-source platform for biological-image analysis. Nat. Methods 9, 676–682. [PubMed: 22743772]

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

Seung HS (1996). How the brain keeps the eyes still. Proc. Natl. Acad. Sci. USA 93, 13339–13344. [PubMed: 8917592]

Seung HS, Lee DD, Reis BY, and Tank DW (2000). Stability of the memory of eye position in a recurrent network of conductance-based model neurons. Neuron 26, 259–271. [PubMed: 10798409]

Shadlen MN, and Newsome WT (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J. Neurosci 18, 3870–3896. [PubMed: 9570816]

Shafi M, Zhou Y, Quintana J, Chow C, Fuster J, and Bodner M (2007). Variability in neuronal activity in primate cortex during working memory tasks. Neuroscience 146, 1082–1108. [PubMed: 17418956]

Shu Y, Hasenstaub A, and McCormick DA (2003). Turning on and off recurrent balanced cortical activity. Nature 423, 288–293. [PubMed: 12748642]

Shuler MG, and Bear MF (2006). Reward timing in the primary visual cortex. Science 311, 1606–1609. [PubMed: 16543459]

Softky WR, and Koch C (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. J. Neurosci 13, 334–350. [PubMed: 8423479]

Spaak E, Watanabe K, Funahashi S, and Stokes MG (2017). Stable and dynamic coding for working memory in primate prefrontal cortex. J. Neurosci 37, 6503–6516. [PubMed: 28559375]

Stephens EK, Avesar D, and Gulledge AT (2014). Activity-dependent serotonergic excitation of callosal projection neurons in the mouse prefrontal cortex. Front. Neural Circuits 8, 97. [PubMed: 25206322]

Steriade M, Timofeev I, and Grenier F (2001). Natural waking and sleep states: a view from inside neocortical neurons. J. Neurophysiol 85, 1969–1985. [PubMed: 11353014]

Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, and Harris KD (2019). Spontaneous behaviors drive multidimensional, brainwide activity. Science 364, 255. [PubMed: 31000656]

Sutton RS, and Barto AG (1981). Toward a modern theory of adaptive networks: expectation and prediction. Psychol. Rev 88, 135–170. [PubMed: 7291377]

Suzuki M, and Gottlieb J (2013). Distinct neural mechanisms of distractor suppression in the frontal and parietal lobe. Nat. Neurosci 16, 98–104. [PubMed: 23242309]

Thomson AM, and Bannister AP (2003). Interlaminar connections in the neocortex. Cereb. Cortex 13, 5–14. [PubMed: 12466210]

Uylings HB, and van Eden CG (1990). Qualitative and quantitative comparison of the prefrontal cortex in rat and in primates, including humans. Prog. Brain Res 85, 31–62. [PubMed: 2094901]

Van De Werd HJJM, Rajkowska G, Evers P, and Uylings HBM (2010). Cytoarchitectonic and chemoarchitectonic characterization of the prefrontal cortical areas in the mouse. Brain Struct. Funct 214, 339–353. [PubMed: 20221886]

van Vreeswijk C, and Sompolinsky H (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science 274, 1724–1726. [PubMed: 8939866]

Wallis JD, Anderson KC, and Miller EK (2001). Single neurons in prefrontal cortex encode abstract rules. Nature 411, 953–956. [PubMed: 11418860]

Wang XJ (1999). Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. J. Neurosci 19, 9587–9603. [PubMed: 10531461]

Wang Y, Markram H, Goodman PH, Berger TK, Ma J, and Goldman-Rakic PS (2006). Heterogeneity in the pyramidal network of the medial prefrontal cortex. Nat. Neurosci 9, 534–542. [PubMed: 16547512]

Wang M, Ramos BP, Paspalas CD, Shu Y, Simen A, Duque A, Vijayraghavan S, Brennan A, Dudley A, Nou E, et al. (2007). Alpha2A-adrenoceptors strengthen working memory networks by inhibiting cAMP-HCN channel signaling in prefrontal cortex. Cell 129, 397–410. [PubMed: 17448997]

Wang M, Yang Y, Wang C-J, Gamo NJ, Jin LE, Mazer JA, Morrison JH, Wang X-J, and Arnsten AFT (2013). NMDA receptors subserve persistent neuronal firing during working memory in dorsolateral prefrontal cortex. Neuron 77, 736–749. [PubMed: 23439125]

Watanabe M (1996). Reward expectancy in primate prefrontal neurons. Nature 382, 629–632. [PubMed: 8757133]

Watanabe M, Hikosaka K, Sakagami M, and Shirakawa S (2002). Coding and monitoring of motivational context in the primate prefrontal cortex. J. Neurosci 22, 2391–2400. [PubMed: 11896178]

Williams GV, and Goldman-Rakic PS (1995). Modulation of memory fields by dopamine D1 receptors in prefrontal cortex. Nature 376, 572–575. [PubMed: 7637804]

Williams GV, Rao SG, and Goldman-Rakic PS (2002). The physiological role of 5-HT2A receptors in working memory. J. Neurosci 22, 2843–2854. [PubMed: 11923449]

Yang Y, Paspalas CD, Jin LE, Picciotto MR, Arnsten AFT, and Wang M (2013). Nicotinic a7 receptors enhance NMDA cognitive circuits in dorsolateral prefrontal cortex. Proc. Natl. Acad. Sci. USA 110, 12078–12083. [PubMed: 23818597]

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

Zhao WJ, Kremkow J, and Poulet JFA (2016). Translaminar cortical membrane potential synchrony in behaving mice. Cell Rep 15, 2387–2399. [PubMed: 27264185]

Zhou L, Liu MZ, Li Q, Deng J, Mu D, and Sun YG (2017). Organization of functional long-range circuits controlling the activity of serotonergic neurons in the dorsal raphe nucleus. Cell Rep 18, 3018–3032. [PubMed: 28329692]

Highlights

- Persistent changes in membrane potential (V_m) occur in anticipation of upcoming reward

- V_m changes arise due to sustained increases in mean and variance of V_m

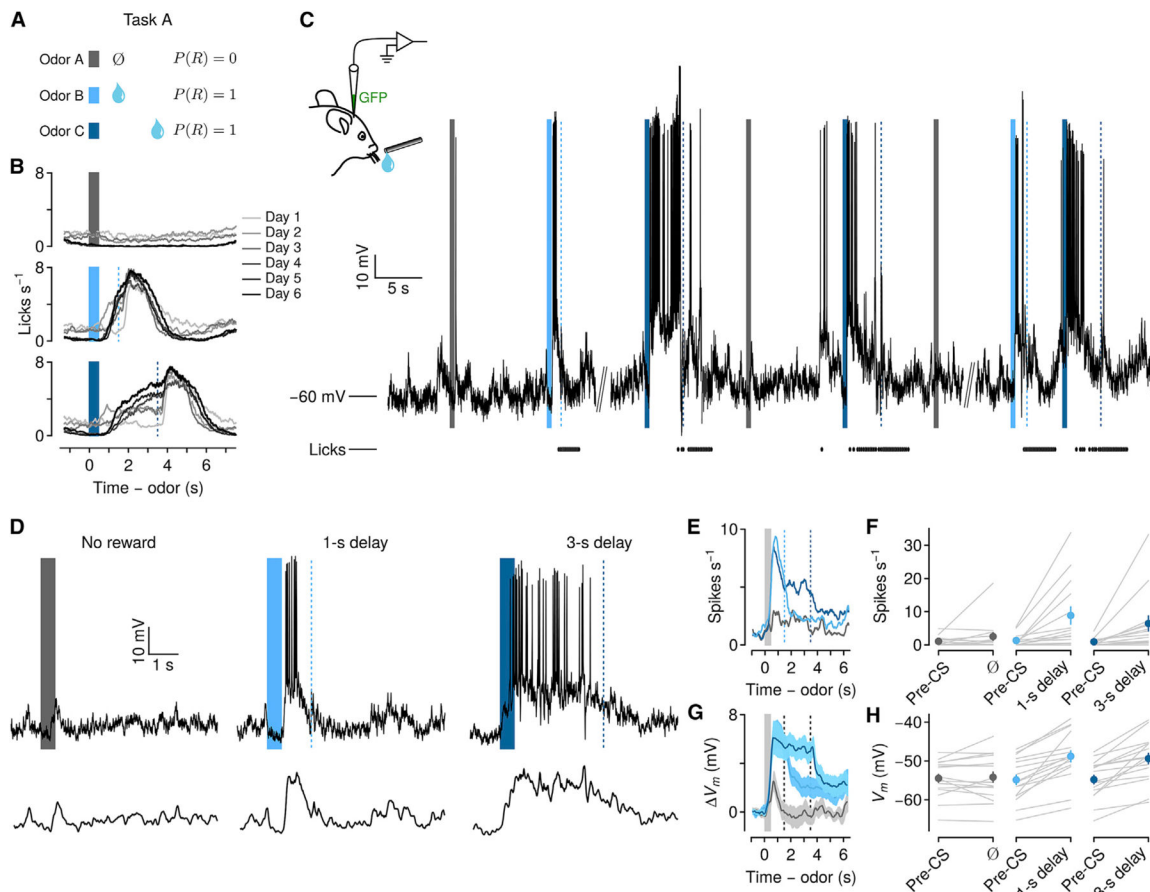- Persistent activity changes are layer-specific and robust to perturbation

Figure 1. Persistent firing rate and V_m changes in PFC during delays to expected reward

(A) Behavioral task in which odors predict no reward, reward following a short delay, or reward following a long delay.

(B) Behavioral learning curves show mean lick rates across days of exposure to the task in one mouse. Bars, odor cues; dashed lines, rewards.

(C) Left: schematic of whole-cell recording. GFP plasmid was included in the recording pipette to localize a subset of neurons. Right: V_m from an example neuron over several minutes. Ticks below V_m indicate licks.

(D) Top: example trials of each type from the neuron in (C). Note the persistent increase in spiking and sustained depolarization in the delay between cue and reward. Bottom: V_m with action potentials removed.

(E) Mean firing rates of 15 neurons showing persistent increases during delays to reward.

(F) Firing rates from the same neurons, comparing pre-CS period to delay (individual neurons in gray, points are mean ± SEM).

(G) Mean ± SEM V_m without action potentials from the same neurons.

(H) V_m from the same neurons, comparing pre-CS to delay (individual neurons in gray, points are mean ± SEM).
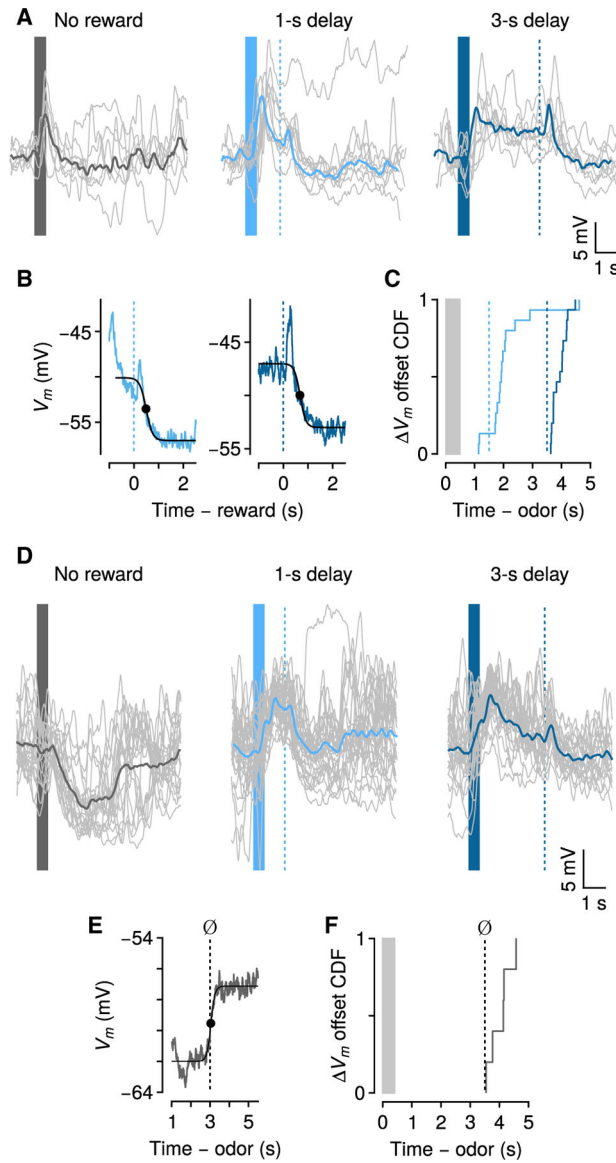
Figure 2. Persistent V_m changes can be terminated by reward or purely internally

(A) Action-potential-removed traces (individual trials in gray, averages in thick lines) of V_m changes relative to baseline from one neuron.

(B) Offset of persistent V_m changes for the same neuron, relative to reward times.

(C) Cumulative distribution function (CDF) of V_m offset times relative to odor and reward (n = 15).

(D) Example neuron showing hyperpolarization for ~3 s during no-reward trials (left), and depolarization during reward trials.

(E) Dynamics of V_m offset for the same example neuron in (D) during no-reward trials.

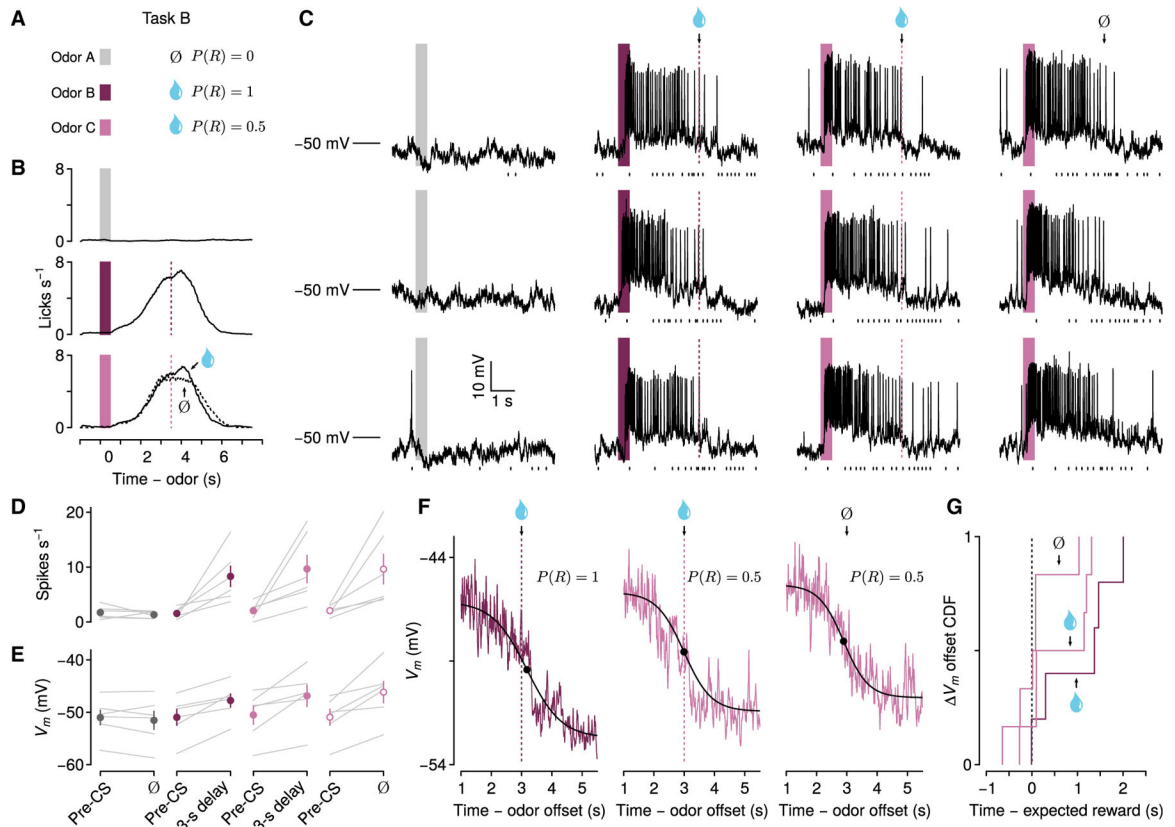(F) CDF of V_m offset times during no-reward trials (n = 6).

Figure 3. Persistent V_m changes in anticipation of probabilistic reward

(A) Behavioral task in which odors predict no reward, reward with probability 0.5 after a 3-s delay, or reward with probability 1 after a 3-s delay.

(B) Mean licking rates from one experiment on each trial type.

(C) V_m from one neuron during three example trials of each type. Ticks below V_m traces indicate lick times.

(D) Firing rates from 6 depolarizing neurons in this task, comparing pre-CS period to delay (individual neurons in gray, points are mean ± SEM).

(E) V_m without action potentials from the same neurons.

(F) Dynamics of sustained V_m termination across trial types from one neuron. Note the return of V_m to baseline even without reward.

(G) CDF of V_m termination times relative to expected reward times for rewarded and unrewarded trials (n = 6).
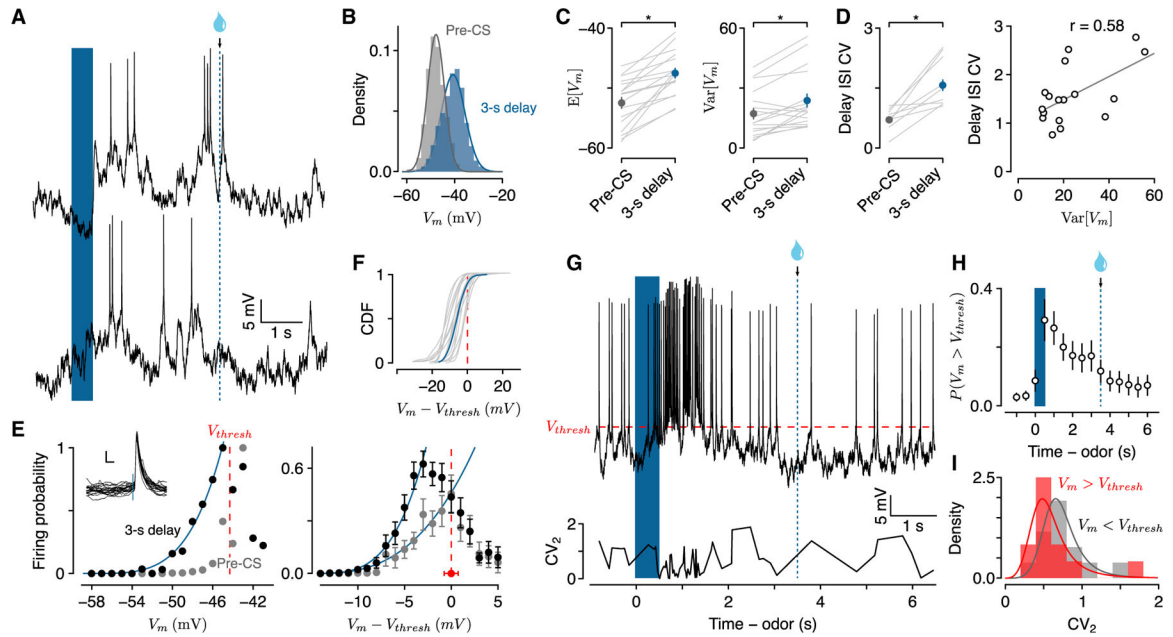
Figure 4. V_m changes are primarily fluctuation-driven during delays to reward

(A) V_m during two 3-s delay trials illustrating subthreshold fluctuations.

(B) Probability densities of spike-removed V_m during the pre-CS period and the 3-s delay for the same neuron.

(C) E[V_m] and Var[V_m] during pre-CS and 3-s delay periods (n = 17; individual neurons in gray, points are mean ± SEM).

(D) Left: CV of ISI during pre-CS and 3-s delay. Right: CV of ISI correlated with Var[V_m] during delays.

(E) Left: example transfer function of firing probability versus V_m during pre-CS period (gray) and 3-s delay (black). V_{thresh} indicates spike threshold. Inset: spike waveforms from the example neuron. Scale bars, 5 mV, 5 ms. Right: average transfer functions of 17 neurons (±SEM).

(F) CDF of the deviation of V_m from V_{thresh} during the delay for individual neurons (gray, n = 17) and average (blue).

(G) Example V_m from a 3-s delay trial showing the relationship between spikes and threshold and local interspike interval variability (CV_2).

(H) Probability of $V_m > V_{thresh}$ during 3-s delay trials across 17 neurons (mean ± SEM).

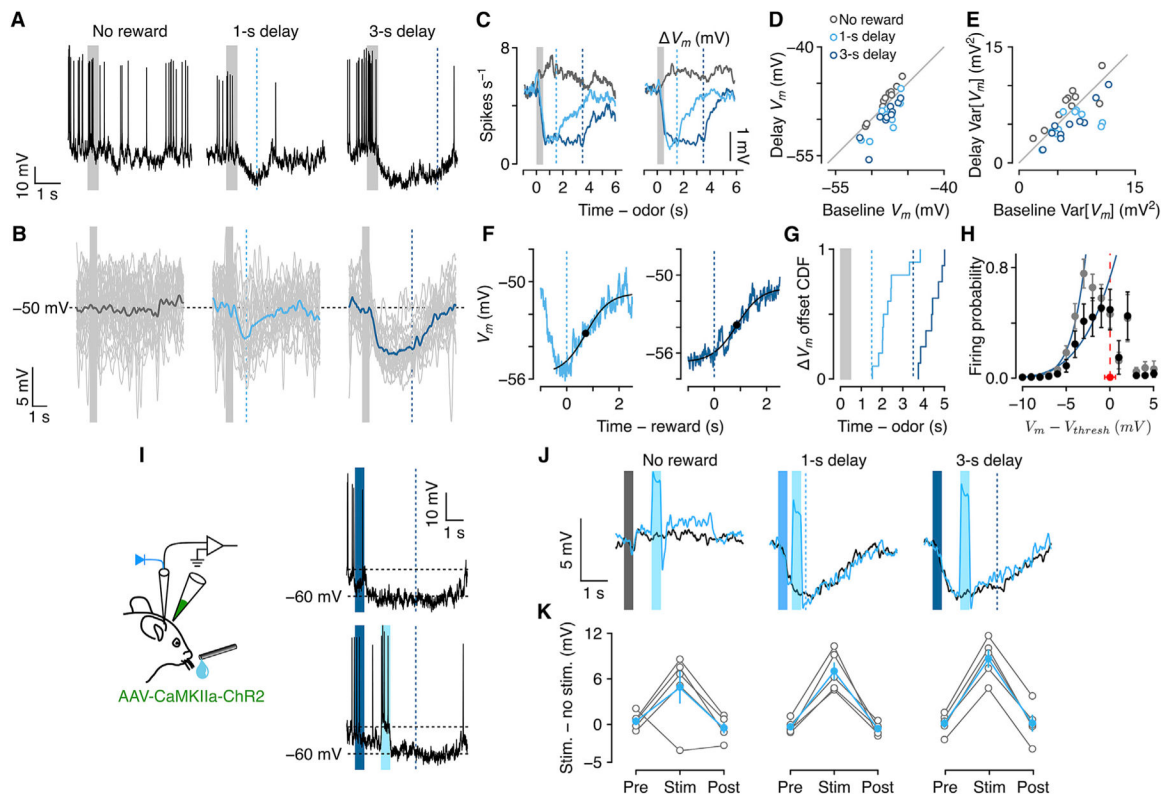(I) Probability densities of CV_2 when $V_m > V_{thresh}$ (red) or $V_m < V_{thresh}$ (gray).

Figure 5. Hyperpolarizing persistent V_m changes

(A) V_m from an example neuron showing hyperpolarization during delays to reward.

(B) V_m from the same neuron with action potentials removed (individual trials in gray).

(C) Mean firing rates and V_m from each trial type across 10 hyperpolarizing neurons.

(D) V_m during delay versus pre-CS periods.

(E) Var[V_m] during delay versus pre-CS periods.

(F) Dynamics of termination of hyperpolarization after reward from an example neuron.

(G) CDFs of hyperpolarization offset times for 1-s and 3-s delay trials.

(H) Average transfer functions (±SEM). Black, 3-s delay; gray, pre-CS.

(I) Experimental schema and example trials from one neuron, showing hyperpolarization between CS and delayed reward and a similar response despite a burst of stimulation-induced spikes during the delay (cyan period).

(J) Mean V_m from an example neuron during unstimulated versus stimulated (cyan) trials, for trials with no reward, 1-s delay, and 3-s delay.

(K) Mean ± SEM (in cyan; individual neurons in gray, n = 5) difference between stimulated (''stim'') and unstimulated (''no stim'') trials. ''Pre'' indicates the interval 500 ms before stimulation onset. ''Stim'' is during stimulation (or the corresponding period during unstimulated trials). ''Post'' is 500 ms after stimulation offset.
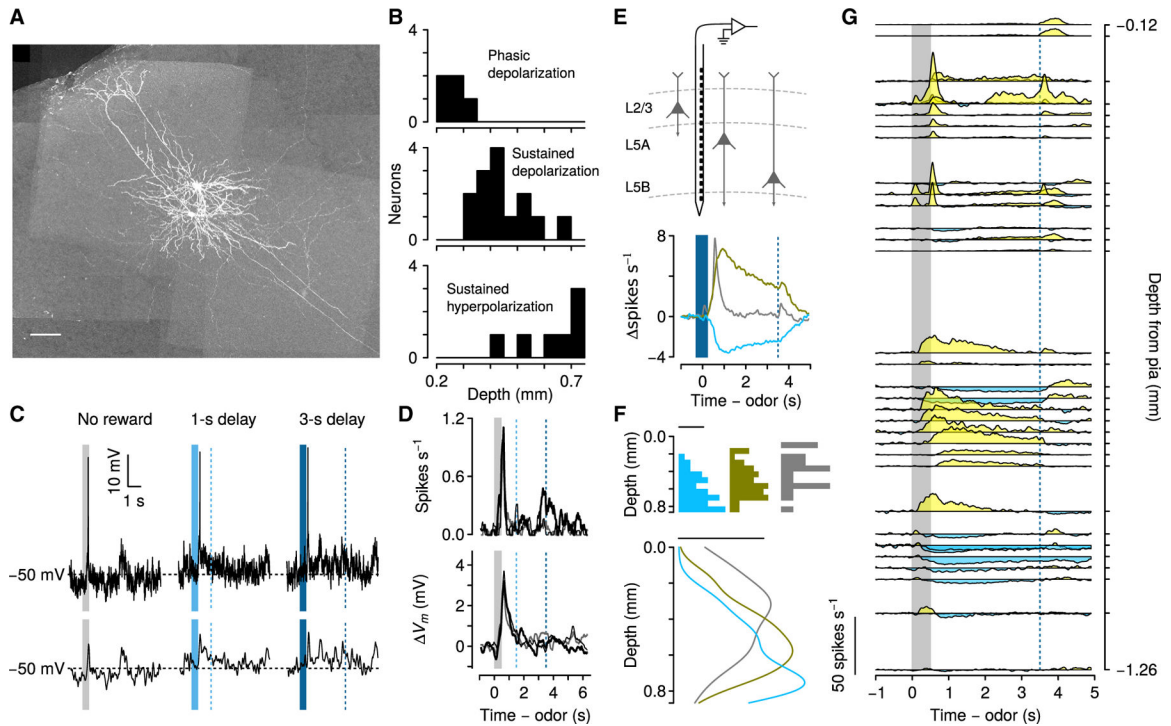
Figure 6. Persistent activity is layer-specific

(A) Maximum-intensity projection of GFP expression in two neurons recorded from one mouse. One had a soma in upper L5 and showed persistent depolarization during the delay to reward. The other had a soma in lower L5 and showed persistent hyperpolarization during the delay to reward. Scale bar, 100 μm.

(B) Recording depth of the three populations of neurons.

(C) Example V_m from a L2/3 neuron on each trial type. Below are spike-removed traces.

(D) Mean firing rates and V_m from 14 L2/3 neurons. Gray, no-reward trials; thin black, 1-s delay trials; thick black, 3-s delay trials.

(E) Schema of extracellular recordings with silicon probe contacts spanning cortical layers (top) and average firing rates on 3-s delay trials for each neuronal response type (bottom).

(F) Histograms (top) and estimates of probability densities (bottom) for neurons. Scale bar, 0.001 density.

(G) Right: average firing rates of 39 simultaneously recorded neurons plotted by depth from pial surface during 3-s delay trials.
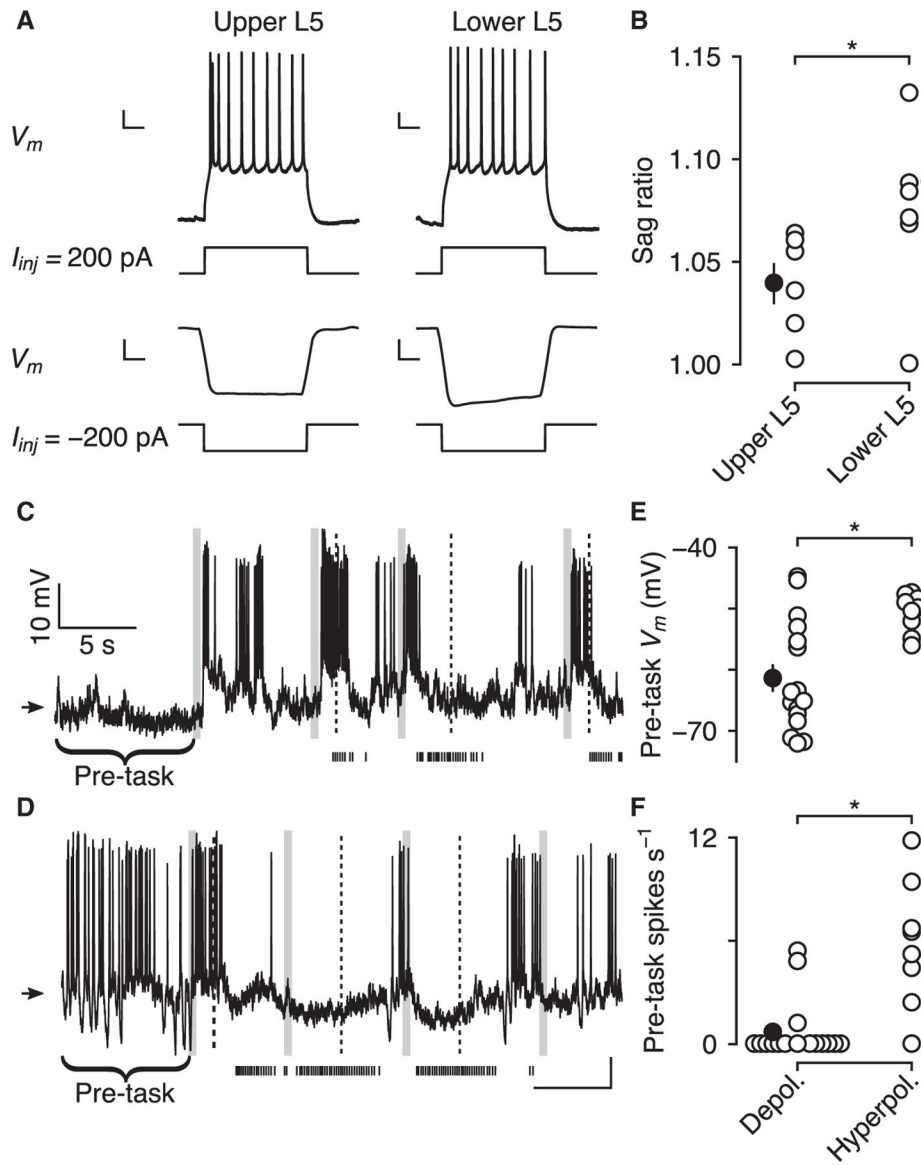
Figure 7. Layer-specific biophysical properties

(A) Biophysical properties of upper and lower L5 pyramidal neurons in anesthetized mice. Example upper (left) and lower (right) L5 neurons in response to current injections. Note the pronounced sag in the hyperpolarizing response to negative current injection in the lower L5 neuron. Scale bars, 10 mV, 200 ms.

(B) Sag ratios are larger in lower versus upper L5 neurons (black points, mean ± SEM).

(C) Example depolarized neuron exhibiting a sharp increase in V_m at the beginning of the task. Arrow indicates −50 mV.

(D) Example hyperpolarized neuron exhibiting a sharp decrease in V_m at the beginning of the task.

(E and F) Neurons depolarized during delays to reward had lower pre-task V_m and firing rates than those hyperpolarized during delays to reward (black points, mean ± SEM).

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and virus strains		
pAAV-CaMKIIa-hChR2(H134R)-EYFP	Addgene; Lee et al., 2010	Addgene viral prep 26969-AAV5; http://addgene.org/26969; RRID:Addgene_26969
pENN.AAV.CB7.CI.mCherry.WPRE.RBG	Addgene	Addgene viral prep 105544-AAV1; http://addgene.org/105544; RRID:Addgene_105544
pENN.AAV.CB7.CI.eGFP.WPRE.rBG	Addgene	Addgene 105542-AAV1; http://addgene.org/105542; RRID:Addgene_105542
Experimental models: organisms/strains		
C57BL/6J mice	The Jackson Laboratory	000664
Recombinant DNA		
pCAG-GFP	Addgene; Matsuda and Cepko, 2004	Addgene plasmid 11150