Central cholinergic neurons are rapidly recruited by reinforcement feedback

Authors: Balázs Hangya[1,2,*], Sachin P. Ranade[1], Maja Lorenc[1] & Adam Kepecs[1,*]

Affiliations:

[1]Cold Spring Harbor Laboratory, Cold Spring Harbor, New York, 11724, United States.

[2]Institute of Experimental Medicine, Hungarian Academy of Sciences, Budapest, H-1083, Hungary.

* Correspondence to: Adam Kepecs: kepecs@cshl.edu and Balázs Hangya: hangya.balazs@koki.mta.hu

Summary

Basal forebrain cholinergic neurons constitute a major neuromodulatory system implicated in normal cognition and neurodegenerative dementias. Cholinergic projections densely innervate neocortex, releasing acetylcholine to regulate arousal, attention and learning. However, their precise behavioral function is poorly understood because identified cholinergic neurons have never been recorded during behavior. To determine which aspects of cognition their activity might support we recorded cholinergic neurons using optogenetic identification in mice performing an auditory detection task requiring sustained attention. We found that a non-cholinergic basal forebrain population — but not cholinergic neurons — were correlated with trial-to-trial measures of attention. Surprisingly, cholinergic neurons responded to reward and punishment with unusual speed and precision (18±3ms). Cholinergic responses were scaled by the unexpectedness of reinforcement and were highly similar across neurons and two nuclei innervating distinct cortical areas. These results reveal that the cholinergic system broadcasts a rapid and precisely timed reinforcement signal supporting fast cortical activation and plasticity.

Introduction

Neuromodulators are central to brain function, and have the ability to dramatically reconfigure circuits and change their dynamics (Bargmann and Marder, 2013). As the only classic neuromodulatory system with cell bodies located in the forebrain as opposed to the evolutionarily more ancient midbrain, it has been implicated in a range of cognitive functions from arousal and vigilance to attention and learning, and even consciousness (Everitt and Robbins, 1997; Hasselmo and Sarter, 2011). Cholinergic cell loss is a major feature of multiple diseases of cognition: the severity of cognitive impairment in Alzheimer's disease and Parkinson's dementia is correlated with the extent of deterioration of basal forebrain cholinergic neurons (Whitehouse et al., 1982). Notably, deep brain stimulation of the basal forebrain is being tested as a therapeutic option for dementia, and can improve the cognitive symptoms of some Alzheimer's and Parkinson-dementia patients (Freund et al., 2009; Kuhn et al., 2015). Thus, progressive degeneration of central cholinergic neurons is thought to play a key role in neurodegenerative dementias and age-related cognitive decline, lending acute pathophysiological significance to basal forebrain research.

It may not be surprising then, that perturbations of the central cholinergic system affect a wide range of behaviors. Rodents with selective lesions of cholinergic neurons, pharmacological blockade of acetylcholine receptors or optogenetic suppression of cholinergic activity show performance deficits in detecting and discriminating sensory stimuli (Everitt and Robbins, 1997; McGaughy et al., 2000, 2002;

Parikh et al., 2007; Pinto et al., 2013; Wrenn and Wiley, 1998), pointing to a causal role of the cholinergic system in these behaviors. However, how behavioral efficiency is modulated by higher level cognitive processes through the recruitment of the cholinergic system is largely unknown and there is a plethora of candidate behavioral functions that have been suggested to tap into cholinergic mechanisms.

One hypothesis is that cholinergic neurons are involved in the control of arousal (Buzsaki et al., 1988; Richardson and DeLong, 1991; Zhang et al., 2011), vigilance (Hassani et al., 2009; Lee et al., 2005) and attention (Everitt and Robbins, 1997; Sarter et al., 2009). Attention demanding tasks are accompanied by elevated cortical choline levels (Parikh et al., 2007; Sarter et al., 2009) and impaired by cholinergic blockers and lesions (Everitt and Robbins, 1997; McGaughy et al., 2002), suggesting that the cholinergic system may play a role in allocating attention at short temporal scales (Disney et al., 2007; Herrero et al., 2008). At the network level, cholinergic activation leads to rapid cortical activation and desynchronization in sensory cortices (Buzsaki et al., 1988; Eggermann et al., 2014; Kalmbach et al., 2012; Metherate et al., 1992; Pinto et al., 2013). These cholinergic effects are thought to be signatures of altered cortical operations that underlie the increased capacity for sensory detection and discrimination.

Another line of investigations has focused on the role of the cholinergic system in cortical plasticity and learning. Cholinergic lesions or pharmacological manipulations impair learning in spatial, working memory and other mnemonic tasks (Everitt and Robbins, 1997; McGaughy et al., 2000), pointing to a causal role of cholinergic neurons. Cholinergic activation is capable of changing the strength, sign and underlying molecular mechanisms of synaptic plasticity (Gu and Yakel, 2011; Gu et al., 2012; Seol et al., 2007), effects that likely underlie the widely observed cholinergic enhancement of receptive field plasticity in sensory cortices (Chubykin et al., 2013; Disney et al., 2007; Froemke et al., 2013; Kilgard and Merzenich, 1998). Through these mechanisms, the sensory cortex projecting cholinergic neurons may boost learning and thereby contribute to improvements in behavioral performance.

Nonetheless it remains unclear why behavioral performance decreases after loss of cholinergic tone, and the possible underlying mechanisms range from arousal to attention to learning processes. To gain insight into these processes, it is critical to first understand at what time scales the firing of cholinergic neurons vary with behavioral performance. For instance, fast modulation of cortical arousal might occur within a behavioral trial, leading to a trial-by-trial co-variation of cholinergic activity and behavioral performance. On the other hand, slow but steady decrease in vigilance throughout a behavioral session due to a concomitant diminution of cholinergic firing could lead to deterioration of behavioral performance. Importantly these possibilities would be expected to lead to similar changes in overall behavioral accuracy that are difficult to disentangle.

Therefore we reasoned that determining the conditions under which cholinergic neurons are

normally active is essential for revealing their behavioral functions across multiple time scales. Although there have been some recordings of unidentified neurons from the basal forebrain (Lin and Nicolelis, 2008; Richardson and DeLong, 1991; Wilson and Rolls, 1990; Zhang et al., 2011) there are no recordings of verified cholinergic neurons in behaving animals. The reasons for this are twofold. First, cholinergic neurons lie deep in the forebrain intermingled with other cell types including two cortically projecting populations, GABAergic and glutamatergic cells (Freund and Gulyás, 1991; Gritti et al., 1997, 2006). In addition, they lack distinguishing spike shape features or firing characteristics that could aid identification. Second, the cholinergic basal forebrain is comprised of a number of topographically projecting nuclei representing a high degree of anatomical complexity (Saper, 1984; Zaborszky et al., 2013), including the prefrontally projecting horizontal limb of the diagonal band (HDB) and the auditory/parietal cortex projecting caudal nucleus basalis (NB). Auditory projecting cholinergic neurons of the caudal nucleus basalis are present in a thin sheet on the lateral border of the internal capsule, making these experiments technically challenging even in the era of optogenetics (Lehmann et al., 1980; Saper, 1984; Zaborszky et al., 2013). Here we recorded identified cholinergic neurons from both the NB and the HDB during behavior for the first time. We report surprising dynamics of cholinergic firing including exceptionally fast and precise responses to innate reward (water) and punishment (air puff) – collectively referred to as primary reinforcers. The responses of cholinergic neurons were indistinguishable between the two nuclei despite their different projection targets, suggesting they constitute a unified broadcast system to cortex. Finally, we constructed a computational model to understand the variations in cholinergic responses and found they could be explained as responding to reinforcement surprise, showing stronger activation after unexpected reinforcement.

Results

Optogenetic identification of central cholinergic neurons

We sought to record identified basal forebrain cholinergic neurons to determine when and how they are recruited during behavior. We targeted two distinct nuclei of the cholinergic basal forebrain. First, we identified the auditory projecting portion of the nucleus basalis (NB) revealed by retrograde tracing (Fig. 1A, bottom). Second, we performed recordings from the prefrontally projecting horizontal limb of the diagonal band (HDB). These two nuclei are not only far apart (1.5 mm) but send non-overlapping cortical projections that are thought to underlie distinct functions (Nelson et al., 2005; Parikh et al., 2007). Virtually all cholinergic cells are projection neurons (Zaborszky et al., 2012, 2013), obviating the need for retrograde or antidromic identification of a projection subpopulation. However, both NB and HDB contain a diversity of cell-types, including GABAergic and glutamatergic projection neurons (Freund and Gulyás, 1991; Gritti

et al., 1997) that lack distinct electrophysiological signatures or pharmacological properties that could enable identification in extracellular recordings. Therefore we used optogenetic tagging to identify cholinergic neurons in extracellular recordings. We rendered cholinergic neurons light-sensitive using either viral transfection to deliver channelrhodopsin-2 in ChAT-Cre mice (Fig. 1B), or a ChAT-ChR2 mouse line (Fig. S1A-F; we observed no differences between the two lines, see Experimental Procedures). We recorded well-isolated single units and delivered brief (1 ms) blue light-pulses to elicit short-latency action potentials. Cholinergic neurons were identified by their significant short latency light-responses (n = 34 out of 1580 units, $p < 0.01$, SALT test for optogenetic identification; Fig. 1C-E, Fig. S1G-K). Note that only around 6% of the basal forebrain neurons are cholinergic (Gritti et al., 2006) and since our methods are designed to minimize false positives they might have missed some cholinergic cells because of insufficient viral infection or light access (see Experimental Procedures).

Slow fluctuations of cortical acetylcholine levels have been long hypothesized to mediate gradual changes of vigilance or arousal (Buzsaki et al., 1988). Therefore we first examined whether the baseline firing of cholinergic neurons was correlated with behavioral and brain states in freely moving mice. Video tracking data were used to differentiate segments of sleep (no motion, accompanied by delta band, 1-4 Hz oscillations in cortical local field potentials) and quiet wakefulness (characterized by head movements without locomotion) from freely moving epochs (Fig 1F, see Experimental Procedures). Cholinergic neurons showed the highest activity in freely moving mice (5.0 ± 1.4 Hz, median ± s.e.; n = 5; Fig. 1G), which decreased during quiet wakefulness (4.0 ± 1.7 Hz) and further in sleep (2.0 ± 1.1 Hz), in agreement with previous observations (Hassani et al., 2009; Lee et al., 2005).

Punishment promptly activates cholinergic neurons

Cholinergic lesions of the basal forebrain have been shown to impair sensory detection under attention demanding circumstances (Everitt and Robbins, 1997; McGaughy et al., 2002; Sarter et al., 2009). To investigate how the cholinergic system controls such cortical functions, we recorded cholinergic neurons in an auditory detection task that requires sustained attention (Fig. 2; Fig. S2). Head-fixed mice (n = 22) were trained to detect two pure tones, well separated at distinct frequencies, and respond to the 'go' tone with a lick while ignoring the 'no-go' tone. Responses to the 'go' tone were considered hits and resulted in the delivery of a drop of water reward, while responses to the 'no-go' tone constituted false alarms triggering a mild puff of air directed to the face as punishment (Fig. 2A,B). Thirsty mice learned to lick for water (Fig. 2C, left) and avoid air puffs (Fig. 2C, right; Fig.S2), thus demonstrating that water and air puff have positive and negative motivational value, respectively (Cohen et al., 2015). Mice consistently responded to air puffs by blinking, likely reflecting the aversive quality of the punishment (Fig. 2D). To

make the task attention-demanding, the stimulus was presented at unexpected moments following the trial start signal and tones of varying loudness were interleaved across trials in a white-noise background to create graded difficulty levels. Mice performed the task well and their accuracy and reaction time (RT) varied systematically as a function of signal-to-noise ratio (Fig. 2E,F; Fig. S2A).

Next we examined whether there are specific behavioral events that phasically recruit cholinergic neurons. Our major observation is that almost all cholinergic neurons showed short latency activation after the delivery of punishment, a brief, mild air puff ($n = 30/34$; $p < 0.05$, Mann-Whitney test, Fig. 3A-C). This was characteristic both to NB ($n = 19/22$) and HDB ($n = 11/12$) cholinergic neurons despite their anatomical separation and distinct projection targets.

Encouraged by the phasic nature of cholinergic activation after punishment, we further examined its temporal properties. The phasic activation of NB cholinergic neurons showed remarkably short latency (17.5 ± 0.6 ms, median ± s.e; range, 15-31 ms; Fig. 3D) and extremely high temporal precision (jitter, 3.2 ± 0.7 ms), unexpected for a neuromodulatory system. Cholinergic neurons either fired a single spike or a brief burst of action potentials in response to punishment, with high reliability ($76.9 \pm 6.2\%$; Fig. 3D; Fig. S3A). Within the narrow range of spike latencies, shorter response times were associated with higher spike counts (Fig. 3E; $p < 0.01$ for 8 out of 11 neurons firing at least three bursts; remaining 3/11 p-values, $p = 0.011$, $p = 0.06$, $p = 0.14$), consistent with stronger excitatory drive. Similar to NB, identified cholinergic neurons recorded from the HDB also exhibited fast response kinetics (median latency, 18.7 ± 2.3 ms; jitter, 3.8 ± 2.9 ms excluding two neurons showing atypical 220-230 ms activation with 15 ms onset; reliability, $75.5 \pm 8.4\%$). Such rapid punishment-elicited responses may be either related to cues associated with punishment (termination of the stimulus, touch of air on the face or click of the air valve) or the execution of a stereotypic motor program (mouth opening or licking). To dissociate these possibilities, we introduced a variable delay (200-400 ms Gaussian, s.d. = 30 ms) between the animal's motor response and the feedback (punishment or reward) delivery ($n = 16$ cholinergic neurons). We found that the phasic activation of cholinergic neurons was aligned to the timing of feedback and not the animals' motor response (Fig. 4A,B; Fig. S3A-B). This demonstrates that the rapid activation of cholinergic neurons was triggered by sensory cues associated with the behavioral feedback and not motor events.

Because air puff punishment acts as an innate, primary reinforcer we hypothesized that rapid, unconditional neural responses to reinforcers should also occur outside of task performance. To test this we delivered air puffs at random, unsignaled moments to head fixed mice. All cholinergic neurons ($n = 6$ neurons from 5 mice; 2 from NB and 4 from HDB) showed fast, reliable activation after air puff delivery ($p < 0.05$, Mann-Whitney test; median latency, 19.8 ± 5.5 ms; jitter, 3.7 ± 2.7 ms; reliability, $70.6 \pm 14.5\%$; Fig. 4C,D). In addition, one NB cholinergic neuron also responded to mild foot shocks (latency, 9.5 ms;

jitter, 5.4 ms; reliability, 89%; Fig.S3C-D). Thus, primary punishment elicits rapid reliable cholinergic firing in naïve mice.

We wondered whether this phasic response to negative reinforcers is unique to cholinergic neurons. Therefore we selected all NB neurons significantly responding after punishment with either increased or suppressed firing (n = 717/1360; $p < 0.01$, Mann-Whitney test) and performed hierarchical clustering on several response features (Fig.S4A, Experimental Procedures). Identified cholinergic neurons clustered together with 22 unidentified cells that we labeled as putative cholinergic neurons (pChAT, probable false negatives, see Experimental Procedures). pChAT neurons were similar to identified cholinergic neurons in their responses to punishment (Fig. S4), while the rest of the population showed distinct response properties (Fig. S5A-C). Thus fast responses to punishment defined a separate, unique subpopulation of NB neurons.

Cholinergic responses are scaled by reward expectations

Cholinergic neurons were also activated after positive behavioral feedback, the water reward, albeit with greater heterogeneity. Some cholinergic neurons exhibited precise reward-associated responses similar to their responses to punishment (n = 8/22; $p < 0.05$, Mann-Whitney test; Fig. 5A,B). Other cholinergic neurons exhibited more delayed and less precise responses (n = 10/22) while a few neurons were entirely unresponsive to reward delivery (n = 4/22; Fig. 5A,B). NB neurons characterized as putative cholinergic (pChAT, see above) based on their punishment responses exhibited reward responses similar to identified cholinergic neurons (Fig. S4-5), while no other NB cells were found to exhibit such rapid activation by reward delivery. Identified cholinergic neurons recorded from the HDB were also similar in their reward responses (Fig. 5B). This diversity of response properties could arise from session-wise differences in behavior or variations in anatomical location. The long dorso-ventral axis of the NB (3.2 to 5 mm) allowed us to dissociate these hypotheses by correlating the ratio of reward to punishment responses with anatomical position, and variables parameterizing behavior and training history (number of previous sessions, trials performed, performance). The ratio of reward to punishment responses showed the strongest correlation with recording depth ($R = -0.75$, $p < 0.0001$; $|R| > 0.63$ in partial correlations controlling for training history, see Experimental Procedures), suggesting that the differences in response magnitude are related to anatomical location (Fig. 5C; Fig. S5D). These data point to a potential anatomical gradient of valence preference within the cholinergic NB.

Next we wondered whether these high fidelity responses are solely triggered by primary reinforcers or also modulated by behavioral expectancies. We compared trials with different levels of uncertainty, in which lower or higher signal-to-noise levels in the stimulus differentially predicted outcome probability.

Punishment invariably elicited strong responses independent of stimulus strength (NB, $p = 0.90$; HDB, $p = 0.76$; repeated measures ANOVA; Fig. 6A-C). In contrast, we found that responses to water reward were differentially modulated based on the preceding signal-to-noise ratio of the stimulus with strongest activation by the least expected reward (NB, $p < 0.0001$; HDB, $p = 0.0015$; repeated measures ANOVA; Fig. 6D-F). While the earliest responses were typically not influenced by expectancy (Fig. 6G,H), differential activation started as early as 20-30 ms after reward delivery for some cholinergic neurons ($p < 0.01$, receiver operator characteristic, ROC, analysis for quantifying the discriminability of the two distributions of firing rates) and was statistically significant from 50 to 70 ms post-reward on average ($p < 0.01$, $n = 17$, Wilcoxon signed rank test).

A computational model for reinforcement surprise

The graded cholinergic responses we observed led us to wonder whether these might represent reinforcement surprise, the deviation from outcome expectation. To test this hypothesis we formally defined 'reinforcement surprise' through a hidden Markov model (Dayan and Yu, 2006) for the auditory detection task (Fig. 6I, Fig. S6). This model accounted for psychometric detection performance (Fig. 6J) and also generated a measure of surprise for each reinforcement event. Since mice were well-trained by the time of the recordings, we assumed that the animals knew the task contingencies, i.e. learned a statistically veridical model of the task. In this model hidden states of the task are not directly observable to the decision maker, but generate probabilistic outputs ('observations') that allow Bayesian inference to produce an internal belief about the presence of tone stimuli (Fig. 6I).

We next considered how reinforcement surprise can be computed within this framework. Observations of the stimulus resolve the ambiguity about the hidden states and therefore make trial outcomes more expected and correspondingly less surprising. Thus, the cumulative number of observations provides a natural measure of reinforcement expectations, allowing us to test whether cholinergic responses to reinforcers match formally-defined reinforcement surprise. We found that theoretical reward surprise was graded by stimulus signal-to-noise ratios, whereas punishment surprise was uniformly high (Fig. 6K). The lack of modulation of punishment surprise reflects that false detections arise independent of the stimulus, thus punishment is always behaviorally unexpected in detection tasks. In summary, we found that reinforcement surprise in the model closely matched the amplitudes of cholinergic responses to reinforcers (Fig. 6A-F,K). These results point to the possibility that the graded responses of cholinergic neurons represent differential reinforcement surprise.

A non-cholinergic subpopulation of basal forebrain neurons shows trial-by-trial correlations with attention

Finally, we set out to test the long-standing hypothesis that the cholinergic system is involved in attentional functions (Everitt and Robbins, 1997; Sarter et al., 2009). The cholinergic system could theoretically control attention in two fundamentally different ways: either through slow modulation of vigilance (Fig. 1F,G) or by rapid control of the momentary state of attention. To dissociate between these possibilities, we next asked whether the activity of cholinergic neurons varies with and is predictive of behavior on a rapid, trial-to-trial basis. Mice in our task had to sustain attentional effort during the foreperiod from the start of the trial to stimulus delivery in order to respond to faint 'go' tones. In humans it is well established that the temporal focus of attention can be manipulated by varying the expected moments of stimulus presentation (temporal expectancy), which is reflected in faster reaction time for expected stimuli (Barnes and Jones, 2000; Coull and Nobre, 1998). To assess this in mice, we used a bimodal foreperiod distribution (Janssen and Shadlen, 2005) to manipulate temporal focus. We observed that RT was inversely correlated with temporal expectancy as characterized by the subjective hazard rate, the relative probability of the stimulus to be delivered at a given moment of time (Fig. 7A,B). Importantly, this variation was only observed for difficult to detect stimuli, revealing that temporal expectations aid signal detection in our task, a hallmark of sustained temporal attention.

Sustained attention can wander from moment to moment and reaction time and performance are expected to correlate with the momentary level of attention at the time of the stimulus. Therefore, we operationalized attentional modulation as neural activity before stimulus onset that predicts either RT (i.e. shows significant negative correlation) or accuracy (i.e. shows positive correlations). Surprisingly, only two out of 34 cholinergic neurons (one in NB and one in HDB) showed activity that was predictive of RT and none predicted accuracy. In fact taken as a population the pre-stimulus firing of cholinergic neurons was slightly negatively correlated with behavioral performance ($p = 0.043$ in difficult trials, Wilcoxon signed rank test). On the other hand, a subpopulation of non-cholinergic neurons exhibited attention-related firing based on our operational definition. Figure 7C,D shows an example neuron with increased firing rate up to 1s before stimulus onset that is strongly correlated with short reaction times ($R = -0.36$, $p < 0.00001$). We found that a subset of NB neurons (96/1360 neurons, 7%, $p < 0.01$; only 2/220, 1% of HDB neurons) showed similar RT-predicting activity in the late foreperiod (Fig. 7E; Fig. S7A-C). We also found a population of NB neurons (68/1360 neurons, 5%; 8/220, 4% of HDB neurons) whose pre-stimulus firing predicted the animals' accuracy (Fig. 7F-H; Fig. S7D-F). Thus the behavioral task enabled us to identify attention-like responses that were predictive of future performance. These responses were present in a small non-cholinergic population but did not appear as significant features of cholinergic neurons, suggesting that

cholinergic neurons might contribute to attentional functions mostly through their slower modulation of brain states (Fig 1F,G).

Discussion

Here we recorded identified basal forebrain cholinergic neurons during behavior for the first time. We found that in addition to the behavioral state-dependent modulation of their tonic firing, cholinergic neurons were phasically activated with millisecond precision during behavior. Our experiments revealed that cholinergic neurons exhibit fast, precise and reliable responses to natural, primary reinforcers: water reward and air puff punishment. The response properties of cholinergic neurons were similar across two distinct nuclei, the prefrontally projecting HDB in the rostral forebrain and the auditory projecting NB located at the caudal end of the basal forebrain complex, despite the fact that these nuclei are often implicated in different functions. Cholinergic responses were graded by outcome expectancy and we could account for this with a quantitative model of reinforcement surprise.

A cholinergic broadcast signal to cortex

Cholinergic neurons responded most strongly and uniformly to punishment. This response was reliably elicited by reinforcement feedback (Fig. 3), unrelated to the signal-to-noise ratio of preceding stimuli (Fig. 6) and locked not to the motor event eliciting reinforcement feedback but to cues immediately preceding them (Fig. 4A-B). Primary negative reinforcers elicited similar responses outside the behavioral task (Fig. 4C). Therefore we suspect that the sensory cues triggering these responses must be related to the delivery of the reinforcers such as clicks of the valves controlling water or air flow, the touch of water or air on the face or the sound of air.

Neuromodulators are thought to broadcast signals widely to impact on-going processing across brain regions. However, whether cholinergic neurons across different basal forebrain nuclei respond in a sufficiently uniform manner to consider them a functionally single system has been unclear. We found that cholinergic responses were nearly identical in two distinct central cholinergic nuclei with non-overlapping projections thought to support different functions: the prefrontally projecting HDB mediating top-down attention, while NB implicated in bottom-up attention (Nelson et al., 2005). These results indicate that the cholinergic system is capable of reliably broadcasting of a unified signal to large areas of the brain.

Cholinergic neurons may signal reinforcement surprise

Reward-elicited responses showed greater diversity across cholinergic neurons (Fig. 5). We found that reward responses were scaled by the signal-to-noise ratio of auditory stimuli that usually occurred hundreds of milliseconds before reward delivery, suggesting that cholinergic activation was modulated by outcome expectations (Fig. 6A-G). These data indicate that the central cholinergic system does not simply relay primary reinforcements but can also convey cognitive information. To better understand the potential computational significance of this graded signal we constructed a hidden Markov model (Dayan and Yu, 2006) of the detection task that could reproduce behavioral performance (Fig. 6I-J). This model enabled us to show that a formally defined reinforcement surprise (unsigned inverse outcome expectation) could account for both the uniform response to punishment and the graded response to reward (Fig. 6K).

Mice interpret water reward and a puff of air to the face with opposing motivational valence: they express strong approach behavior to water while they avoid air puffs (Fig. 2C and Fig. S2D). This raises the interesting possibility that basal forebrain cholinergic firing is related to the motivational value of the outcomes. Alternatively, our model suggests that differences between reward and punishment responses can be to a large degree attributed to reinforcement surprise. Our definition of reinforcement surprise bears resemblance to reward prediction errors (RPE) represented by midbrain dopaminergic neurons (Schultz et al., 1997). Indeed, the response of cholinergic neurons is consistent with a representation of unsigned RPE, sometimes called 'salience'. Note, however, that RPE and reinforcement surprise are defined in two different behavioral contexts (sensory detection task vs. cued outcome task) and require distinct computations (trial-to-trial belief state inference vs. experience-dependent reinforcement learning). Therefore further experiments will be required to understand whether and how the signals represented by the dopaminergic and cholinergic systems are related.

The overall magnitude of the difference between punishment and reward responses was not fully captured by the model. Indeed, a correlation analysis revealed that there is an anatomical correlate of this difference, the scaling of reward responses along the unusually long dorso-ventral axis of the NB (Fig. 5C). This could be explained by a systematic difference in the excitability of NB cholinergic neurons or a systematic variation in the strength of bottom-up excitatory connectivity — which may constitute a gradient of surprise representation. Thus our findings resonate with previous theoretical accounts suggesting that acetylcholine signals different forms of uncertainty boosting learning and attention (Dayan et al., 2000; Doya, 2002; Yu and Dayan, 2005).

Cholinergic control of plasticity and learning

Lesions of cholinergic neurons and pharmacological studies have established a causal role of the cholinergic system in learning (Everitt and Robbins, 1997; McGaughy et al., 2000). For instance, stimulating auditory projecting NB neurons have been shown to reorganize receptive field maps in the auditory cortex (Froemke et al., 2013; Kilgard and Merzenich, 1998). However, there is a gap between the long-term impact of irreversible lesions or slow pharmacological manipulations and the cellular mechanisms of neuronal plasticity thought to underlie learning (Chubykin et al., 2013; Seol et al., 2007). Recent results revealed that at the synaptic level precisely timed acetylcholine can control the strength, sign and molecular rules of hippocampal plasticity with millisecond precision (Gu and Yakel, 2011; Gu et al., 2012). Our results demonstrate that the cholinergic system is indeed capable of such millisecond precision in behaving mice (Fig. 3-5). This may provide the missing link between the cellular mechanisms of cholinergic control over cortical plasticity and behavioral learning. Indeed, behavioral rewards can be replaced by optogenetic activation of basal forebrain input to visual cortex and thus sufficient to entrain reward timing activity in cortex (Liu et al., 2015). Taking our observations together with previous in vitro and theoretical studies on plasticity (Jimenez Rezende and Gerstner, 2014), we speculate that fast central cholinergic responses to reinforcers provide supervisory control over local unsupervised cortical plasticity and thereby support learning.

Another possibility is that cholinergic neurons drive learning by activating disinhibitory circuits in cortex and thereby gate plasticity. This is in agreement with a recent finding that some auditory cortical interneurons receive cholinergic input elicited by punishment during fear conditioning (Letzkus et al., 2011). Indeed, cortical inhibitory interneurons express both ionotropic and metabotropic cholinergic receptors (Alitto and Dan, 2012; Demars and Morishita, 2014; Disney et al., 2007). Thus cholinergic neurons could also drive reinforcement responses observed in cortical VIP^+ (Pi et al., 2013) and hippocampal SOM^+ interneurons (Kaifosh et al., 2013)

Role of cholinergic and non-cholinergic basal forebrain neurons in arousal and attention

A long-standing hypothesis is that cholinergic neurons are involved in the control of arousal and forms of attention (Everitt and Robbins, 1997; Hasselmo and Sarter, 2011). In agreement with previous observations we found that the tonic firing rates of cholinergic cells varies as a function of the sleep-wake cycle and arousal (Fig. 1F,G) (Duque et al., 2000; Hassani et al., 2009; Lee et al., 2005). These slower changes may underlie attention-like effects associated with the cholinergic system (Disney et al., 2007; Everitt and Robbins, 1997; Herrero et al., 2008), including recent results that optogenetic manipulations of cholinergic neurons can lead to performance changes in a visual discrimination task (Pinto et al., 2013).

On the other hand, whether the cholinergic system modulates attention at rapid time scales has not been previously tested. We probed two central cholinergic nuclei that are considered good candidates for such attentional effects: the NB projecting to primary auditory as well as other sensory cortices, capable of influencing sensory detection and input processing functions (Froemke et al., 2013), and the HDB, sending prefrontal projections thought to underlie top-down attentional modulation (Hasselmo and Sarter, 2011; Nelson et al., 2005; Parikh et al., 2007). Surprisingly, not the cholinergic but a subpopulation of unidentified neurons' activity predicted behavioral variables classically associated with attention, such as reaction time and performance accuracy (Fig. 7). This supports the idea that the basal forebrain also has attentional functions, albeit served by non-cholinergic neurons. This is also consistent with previous recordings of unidentified basal forebrain neurons showing a diversity of responses that were likely sampled from the more numerous and fast firing non-cholinergic populations (Lin and Nicolelis, 2008; Richardson and DeLong, 1991; Wilson and Rolls, 1990), some correlated with reaction time (Avila and Lin, 2014).

Conclusions

Our results support previous computational theories proposing that acetylcholine conveys a global reinforcement signal that enables the brain to associate prior events with behavioral outcomes (Doya, 2002; Hasselmo and Sarter, 2011; Jimenez Rezende and Gerstner, 2014; Yu and Dayan, 2005). Cholinergic responses were remarkably fast, 30-50 milliseconds faster than midbrain dopamine neurons (Cohen et al., 2012), raising important questions about how acetylcholine might impact processing. First, cholinergic cells may recruit disinhibitory circuitry via nicotinic receptors (Letzkus et al., 2011; Pi et al., 2013) leading to rapid dynamic modulation of cortical arousal (Buzsaki et al., 1988; Richardson and DeLong, 1991; Zhang et al., 2011). Second, the fast and precisely timed cholinergic responses can provide a powerful computational mechanism for global modulation of timing-dependent synaptic plasticity across cortex (Frémaux et al., 2010; Gu and Yakel, 2011). Thus we propose that the rapid phasic responses of basal forebrain cholinergic neurons represent reinforcement surprise and their broadcast serves as an alert signal capable of triggering rapid reconfiguration of cortical state and plasticity.

Experimental Procedures

Adult (over 2 months old) ChAT-Cre (n = 15), ChAT-ChR2 (n = 5) and PV-Cre (n = 4) mice were used for behavioral recording experiments and nine additional mice for behavior-only experiments under a protocol approved by Cold Spring Harbor Laboratory Institutional Animal Care and Use Committee in

accordance with National Institutes of Health standards. See Extended Experimental Procedures for details.

Microdrive construction, injection and microdrive implantation

Custom-built light-weight (2.2g) microdrives (Fig. S1G) were constructed for deep brain recording and optogenetic stimulation. A moveable shuttle held an optic fiber and 7-8 tetrodes for unit recordings. Two stereotrodes were also connected for cortical local field potential recordings. Standard surgical techniques were employed for virus injection and microdrive implantation.

Behavior, recording and optogenetics

Mice were trained on an auditory detection attention task in a head-fixed go/no-go detection paradigm using a custom-built apparatus. Extracellular recordings were performed using a DigitalLynx data acquisition system (Neuralynx). A blue laser (473 nm; 100 mW; Lasermate Group Inc.) was triggered through a data acquisition board (National Instruments) controlled by custom-built MATLAB programs (MathWorks) for optogenetic stimulation.

Histology and track reconstruction

To identify the recording sites electrolytic lesions were made under deep anesthesia. After perfusion, brains were post-fixed and sections were imaged by fluorescence (Olympus MVX10) and confocal microscopes (Zeiss 710LSM); then images were aligned to an atlas to accurately reconstruct the recording locations.

Data analysis

Data analyses were carried out using built-in and custom-built software in Matlab (MathWorks). Action potentials were sorted into clusters (MClust software, A. D. Redish). Significant light-activation was assessed by the Stimulus-Associated spike Latency Test (SALT; http://kepecslab.cshl.edu/salt.m). Peri-event firing rates were estimated by using an adaptive spike density function (SDF) approach. We implemented a hidden Markov model (HMM) of the auditory go/no-go detection task to test whether cholinergic neurons signal reinforcement surprise.

Author Contributions

B.H. and A.K. designed the experiments, data analyses and the model, B.H. performed the experiments, analyzed the data, performed model simulations and prepared the figures, M.L. performed additional experiments during the revision, S.P.R. and B.H. established the behavior, S.P.R. performed retrograde tracing and video analysis, B.H. and A.K. wrote the paper.

Acknowledgements

This work was supported by grants from the John Merck and McKnight Foundations and the National Institute of Neurological Disorders and Stroke (National Institutes of Health) grant R01NS075531 to A.K. B.H. received support from the Swartz Foundation and Marie Curie International Outgoing Fellowship within the EU Seventh Framework Programme for Research and Technological Development. We are grateful to Barry Burbach and Rob Eifert for invaluable technical assistance, Goncalo Lopez for help with Bonsai, Sanchari Ghosh for assisting in recording cholinergic neurons outside the behavioral task and to Hyun-Jae Pi, Duda Kvitsiani, Joshua I. Sanders, Michael Long, Stephen Shea and Jessica Tollkuhn for helpful comments and discussions.

References

Alitto, H.J., and Dan, Y. (2012). Cell-type-specific modulation of neocortical activity by basal forebrain input. Front. Syst. Neurosci. *6*, 79.

Avila, I., and Lin, S.-C. (2014). Motivational Salience Signal in the Basal Forebrain Is Coupled with Faster and More Precise Decision Speed. PLoS Biol. *12*, e1001811.

Bargmann, C.I., and Marder, E. (2013). From the connectome to brain function. Nat. Methods *10*, 483–490.

Barnes, R., and Jones, M.R. (2000). Expectancy, attention, and time. Cogn. Psychol. *41*, 254–311.

Buzsaki, G., Bickford, R.G., Ponomareff, G., Thal, L.J., Mandel, R., and Gage, F.H. (1988). Nucleus basalis and thalamic control of neocortical activity in the freely moving rat. J. Neurosci. *8*, 4007–4026.

Chubykin, A.A., Roach, E.B., Bear, M.F., and Shuler, M.G.H. (2013). A cholinergic mechanism for reward timing within primary visual cortex. Neuron *77*, 723–735.

Cohen, J., Amoroso, M., and Uchida, N. (2015). Serotonergic neurons signal reward and punishment on multiple timescales. Elife 06346.

Cohen, J.Y., Haesler, S., Vong, L., Lowell, B.B., and Uchida, N. (2012). Neuron-type-specific signals for reward and punishment in the ventral tegmental area. Nature *482*, 85–88.

Coull, J.T., and Nobre, a C. (1998). Where and when to pay attention: the neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. J. Neurosci. *18*, 7426–7435.

Dayan, P., and Yu, A.J. (2006). Phasic norepinephrine: a neural interrupt signal for unexpected events. Network *17*, 335–350.

Dayan, P., Kakade, S., and Montague, P.R. (2000). Learning and selective attention. Nat. Neurosci. *3 Suppl*, 1218–1223.

Demars, M.P., and Morishita, H. (2014). Cortical parvalbumin and somatostatin GABA neurons express distinct endogenous modulators of nicotinic acetylcholine receptors. Mol. Brain *7*, 75.

Disney, A.A., Aoki, C., and Hawken, M.J. (2007). Gain modulation by nicotine in macaque v1. Neuron *56*, 701–713.

Doya, K. (2002). Metalearning and neuromodulation. Neural Netw. *15*, 495–506.

Duque, A., Balatoni, B., Detari, L., and Zaborszky, L. (2000). EEG correlation of the discharge properties of identified neurons in the basal forebrain. J. Neurophysiol. *84*, 1627–1635.

Eggermann, E., Kremer, Y., Crochet, S., and Petersen, C.C.H. (2014). Cholinergic signals in mouse barrel cortex during active whisker sensing. Cell Rep. *9*, 1654–1660.

Everitt, B.J., and Robbins, T.W. (1997). Central cholinergic systems and cognition. Annu. Rev. Psychol. *48*, 649–684.

Frémaux, N., Sprekeler, H., and Gerstner, W. (2010). Functional requirements for reward-modulated spike-timing-dependent plasticity. J. Neurosci. *30*, 13326–13337.

Freund, T.F., and Gulyás, A.I. (1991). GABAergic interneurons containing calbindin D28K or somatostatin are major targets of GABAergic basal forebrain afferents in the rat neocortex. J. Comp. Neurol. *314*, 187–199.

Freund, H.-J., Kuhn, J., Lenartz, D., Mai, J.K., Schnell, T., Klosterkoetter, J., and Sturm, V. (2009). Cognitive functions in a patient with Parkinson-dementia syndrome undergoing deep brain stimulation. Arch. Neurol. *66*, 781–785.

Froemke, R.C., Carcea, I., Barker, A.J., Yuan, K., Seybold, B. a, Martins, A.R.O., Zaika, N., Bernstein, H., Wachs, M., Levis, P. a, et al. (2013). Long-term modification of cortical synapses improves sensory perception. Nat. Neurosci. *16*, 79–88.

Gritti, I., Mainville, L., Mancia, M., and Jones, B.E. (1997). GABAergic and other noncholinergic basal forebrain neurons, together with cholinergic neurons, project to the mesocortex and isocortex in the rat. J. Comp. Neurol. *383*, 163–177.

Gritti, I., Henny, P., Galloni, F., Mainville, L., Mariotti, M., and Jones, B.E. (2006). Stereological estimates of the basal forebrain cell population in the rat, including neurons containing choline acetyltransferase, glutamic acid decarboxylase or phosphate-activated glutaminase and colocalizing vesicular glutamate transporters. Neuroscience *143*, 1051–1064.

Gu, Z., and Yakel, J.L. (2011). Timing-dependent septal cholinergic induction of dynamic hippocampal synaptic plasticity. Neuron *71*, 155–165.

Gu, Z., Lamb, P.W., and Yakel, J.L. (2012). Cholinergic coordination of presynaptic and postsynaptic activity induces timing-dependent hippocampal synaptic plasticity. J. Neurosci. *32*, 12337–12348.

Hassani, O.K., Lee, M.G., Henny, P., and Jones, B.E. (2009). Discharge profiles of identified GABAergic in comparison to cholinergic and putative glutamatergic basal forebrain neurons across the sleep-wake cycle. J. Neurosci. *29*, 11828–11840.

Hasselmo, M.E., and Sarter, M. (2011). Modes and models of forebrain cholinergic neuromodulation of cognition. Neuropsychopharmacology *36*, 52–73.

Herrero, J.L., Roberts, M.J., Delicato, L.S., Gieselmann, M. a, Dayan, P., and Thiele, a (2008). Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. Nature *454*, 1110–1114.

Janssen, P., and Shadlen, M.N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. Nat. Neurosci. *8*, 234–241.

Jimenez Rezende, D., and Gerstner, W. (2014). Stochastic variational learning in recurrent spiking networks. Front. Comput. Neurosci. *8*, 38.

Kaifosh, P., Lovett-Barron, M., Turi, G.F., Reardon, T.R., and Losonczy, A. (2013). Septo-hippocampal GABAergic signaling across multiple modalities in awake mice. Nat. Neurosci. *16*, 1182–1184.

Kalmbach, A., Hedrick, T., and Waters, J. (2012). Selective optogenetic stimulation of cholinergic axons in neocortex. J. Neurophysiol. *107*, 2008–2019.

Kilgard, M.P., and Merzenich, M.M. (1998). Cortical map reorganization enabled by nucleus basalis activity. Science *279*, 1714–1718.

Kuhn, J., Hardenacke, K., Lenartz, D., Gruendler, T., Ullsperger, M., Bartsch, C., Mai, J.K., Zilles, K., Bauer, A., Matusch, A., et al. (2015). Deep brain stimulation of the nucleus basalis of Meynert in Alzheimer's dementia. Mol. Psychiatry *20*, 353–360.

Lee, M.G., Hassani, O.K., Alonso, A., and Jones, B.E. (2005). Cholinergic basal forebrain neurons burst with theta during waking and paradoxical sleep. J. Neurosci. *25*, 4365–4369.

Lehmann, J., Nagy, J.I., Atmadia, S., and Fibiger, H.C. (1980). The nucleus basalis magnocellularis: the origin of a cholinergic projection to the neocortex of the rat. Neuroscience *5*, 1161–1174.

Letzkus, J.J., Wolff, S.B.E., Meyer, E.M.M., Tovote, P., Courtin, J., Herry, C., and Lüthi, A. (2011). A disinhibitory microcircuit for associative fear learning in the auditory cortex. Nature *480*, 331–335.

Lin, S.-C., and Nicolelis, M. a L. (2008). Neuronal ensemble bursting in the basal forebrain encodes salience irrespective of valence. Neuron *59*, 138–149.

Liu, C.-H., Coleman, J.E., Davoudi, H., Zhang, K., and Shuler, M.G.H. (2015). Selective activation of a putative reinforcement signal conditions cued interval timing in primary visual cortex. Curr. Biol.

McGaughy, J., Everitt, B.J., Robbins, T.W., and Sarter, M. (2000). The role of cortical cholinergic afferent projections in cognition: impact of new selective immunotoxins. Behav. Brain Res. *115*, 251–263.

McGaughy, J., Dalley, J.W., Morrison, C.H., Everitt, B.J., and Robbins, T.W. (2002). Selective behavioral and neurochemical effects of cholinergic lesions produced by intrabasalis infusions of 192 IgG-saporin on attentional performance in a five-choice serial reaction time task. J. Neurosci. *22*, 1905–1913.

Metherate, R., Cox, C.L., and Ashe, J.H. (1992). Cellular bases of neocortical activation: modulation of neural oscillations by the nucleus basalis and endogenous acetylcholine. J. Neurosci. *12*, 4701–4711.

Nelson, C.L., Sarter, M., and Bruno, J.P. (2005). Prefrontal cortical modulation of acetylcholine release in posterior parietal cortex. Neuroscience *132*, 347–359.

Parikh, V., Kozak, R., Martinez, V., and Sarter, M. (2007). Prefrontal acetylcholine release controls cue detection on multiple timescales. Neuron *56*, 141–154.

Pi, H.-J., Hangya, B., Kvitsiani, D., Sanders, J.I., Huang, Z.J., and Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. Nature *503*, 521–524.

Pinto, L., Goard, M.J., Estandian, D., Xu, M., Kwan, A.C., Lee, S.-H., Harrison, T.C., Feng, G., and Dan, Y. (2013). Fast modulation of visual perception by basal forebrain cholinergic neurons. Nat. Neurosci. *16*, 1857–1863.

Richardson, R.T., and DeLong, M.R. (1991). Electrophysiological studies of the functions of the nucleus basalis in primates. Adv. Exp. Med. Biol. *295*, 233–252.

Saper, C.B. (1984). Organization of cerebral cortical afferent systems in the rat. II. Magnocellular basal nucleus. J. Comp. Neurol. *222*, 313–342.

Sarter, M., Parikh, V., and Howe, W.M. (2009). Phasic acetylcholine release and the volume transmission hypothesis: time to move on. Nat. Rev. Neurosci. *10*, 383–390.

Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. Science *275*, 1593–1599.

Seol, G.H., Ziburkus, J., Huang, S., Song, L., Kim, I.T., Takamiya, K., Huganir, R.L., Lee, H.K., and Kirkwood, A. (2007). Neuromodulators control the polarity of spike-timing-dependent synaptic plasticity. Neuron *55*, 919–929.

Whitehouse, P.J., Price, D.L., Struble, R.G., Clark, A.W., Coyle, J.T., and Delon, M.R. (1982). Alzheimer's disease and senile dementia: loss of neurons in the basal forebrain. Science *215*, 1237–1239.

Wilson, F.A., and Rolls, E.T. (1990). Learning and memory is reflected in the responses of reinforcement-related neurons in the primate basal forebrain. J. Neurosci. *10*, 1254–1267.

Wrenn, C.C., and Wiley, R.G. (1998). The behavioral functions of the cholinergic basal forebrain: lessons from 192 IgG-saporin. Int. J. Dev. Neurosci. *16*, 595–602.

Yu, A.J., and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. Neuron *46*, 681–692.

Zaborszky, L., Pol, A. Van Den, and Gyengesi, E. (2012). The Basal Forebrain Cholinergic Projection System in Mice. In The Mouse Nervous System, C. Watson, G. Paxinos, and L. Puelles, eds. (Amsterdam: Elsevier), pp. 684–718.

Zaborszky, L., Csordas, A., Mosca, K., Kim, J., Gielow, M.R., Vadasz, C., and Nadasdy, Z. (2013). Neurons in the Basal Forebrain Project to the Cortex in a Complex Topographic Organization that Reflects Corticocortical Connectivity Patterns: An Experimental Study Based on Retrograde Tracing and 3D Reconstruction. Cereb. Cortex.

Zhang, H., Lin, S.-C., and Nicolelis, M.A.L. (2011). A distinctive subpopulation of medial septal slow-firing neurons promote hippocampal activation and theta oscillations. J. Neurophysiol. *106*, 2749–2763.

Figure legends

Figure 1 | Optogenetic tagging of central cholinergic neurons

A, Auditory projecting cholinergic neurons are in the caudal nucleus basalis (NB; Fig. S1), which includes the ventromedial globus pallidus (GP) and the caudal substantia innominata (SI). Top, coronal section with increasing magnification. ChAT-Cre mouse; green, neurons infected with AAV-flex-GFP; red, ChAT staining; white arrowhead, location of neurons enlarged on the right. Scale bars; left, 1 mm; right, 50 μm. Bottom, retrograde labeling from the auditory cortex. Red, cholinergic neurons; green, retrograde Lumaflour beads; yellow, double labeled neurons. Scale bar, 150 μm. CPu, caudate putamen; int, internal capsule; Rt, reticular thalamic nucleus. **B,** Left, coronal sections showing expression of virally transfected ChR2-eYFP in the caudal NB (top) and horizontal limb of the diagonal band (HDB; bottom). Scale bar, 1 mm. Middle, enlarged images of the marked areas. Right, reconstructed location of identified cholinergic neurons projected onto two coronal planes (top, NB; bottom, HDB; numbers, antero-posterior distance from bregma). Different symbols indicate individual mice. VP, ventral pallidum. **C,** Left, spike raster of an identified cholinergic neuron aligned to light stimulation (blue line). Right, peri-stimulus time histograms aligned to photostimulation onset (blue line) for all identified cholinergic neurons (normalized by peak value, sorted by peak latency; all pulses of the most efficient stimulation frequency were used; colors from black to white correspond to higher firing rates). **D,** Cumulative histograms of light-evoked spike latency (left) and jitter (right) for all identified cholinergic neurons. **E,** SALT (stimulus-associated spike latency test) for optical tagging showed strongly bimodal p-value distribution (blue, $p<0.01$). **F,** Left, example recording of a cholinergic neuron in an awake freely moving mouse. Top, spike times; middle, auditory local field potential (LFP); bottom, wavelet spectrogram of the LFP. Right, example recording of the same cell during sleep. Note the lower firing rate and delta oscillations in the auditory LFP. **G,** Median firing rate of cholinergic neurons was highest in awake freely moving epochs, lower in quiet wakefulness and lowest

during sleep. Black lines, individual cells; solid lines, significantly different firing rate ($p<0.01$, Mann-Whitney test).

Figure 2 | Auditory detection task

A, Schematic of the head fixed auditory detection task. **B,** Structure of a trial and possible outcomes. Trial start was signaled by turning off an LED. After a variable delay, pure tones of well separated pitch but varying intensity signaled water reward or air puff punishment upon licking. **C,** Left, thirsty (water restricted) mice learned to lick for water, showing a median lick probability close to 1 after training. The same mice did not lick for water when water was available ad libitum in their home cage (free water condition; $n = 4$; $p < 0.0001$ in all animals; chi-square test). Right, mice licked significantly more for water than the same amount of water combined with air puff ($n = 5$; $p < 0.0001$ on the population level; $p < 0.05$ in 4/5 individual mice; chi-square test). **D,** Left, average eye blink response after air puff (red) and water (green) delivery, quantified by normalized pixel density based on video analysis (34 sessions from 7 mice). Right, zoomed in to the first 50 ms after reinforcement delivery. **E,** Left, performance in a single session: lick probability in 'go' (green) and 'no-go' (red) trials (labelled 'Hit rate' and 'False alarm rate', respectively) as a function of stimulus difficulty (psychometric function; SPL, sound pressure level of the cue). Right, average performance for individual mice (light; mice contributing at least 3 sessions are shown) and grand average (dark). Norm. SPL, intermediate SPLs were pooled to allow averaging across sessions. **F,** Left, cumulative reaction time (RT) histograms and median RT as a function of stimulus difficulty in the same session as C. Right, average for individual mice (light) and grand average (dark). Error bars, s.e.m. See also Fig. S2.

Figure 3 | Punishment uniformly activates cholinergic neurons

A, Spike rasters (top) and peri-event time histograms (PETHs, bottom) of two identified cholinergic neurons aligned to air puff punishment (orange line). Trials were sorted by RT (grey ticks, stimulus onset). Cholinergic neurons showed precisely timed short latency response to air puff. **B,** Top, individual PETHs (color coded from black to white) of all identified cholinergic neurons revealed homogeneous phasic responses to punishment (left, NB; right, HDB). Cells are sorted by response latency. Arrows indicate the example neurons in panel A (black, left neuron). Bottom, average PETH. **C,** The area in the purple rectangle in panel A is magnified to reveal the low latency and jitter. **D,** Cumulative histogram of punishment response peak latency, first spike jitter, reliability and number of spikes in response to punishment (spike count). **E,** Spike latency showed negative correlation with spike count. Error bars, s.e.m. See also Fig. S3.

Figure 4 | Cholinergic neurons respond to primary reinforcers

A, Raster plot aligned to air puff (left) and realigned to the animal's response (right; same cell as in Fig. 3A, right). **B,** Average PETH aligned to punishment (left) and realigned to the animals motor response (right). Shading, s.e.m. **C,** Spike raster and PETH aligned to air puff of a cholinergic neuron outside the detection task. **D,** Top, individual PETHs of cholinergic neurons recorded outside the task (sorted by response latency) revealed homogeneous phasic responses to air puffs. Bottom, average PETH. See also Fig. S3-4.

Figure 5 | Cholinergic neurons are activated by water reward

A, Spike rasters (top) and PETHs (bottom) of two identified cholinergic neurons (same as in Fig. 3A) aligned to water reward (orange line). Trials were sorted by RT (grey ticks, stimulus onset). The cholinergic neuron on left showed precisely timed short latency response to water, while the neuron on right exhibited a weaker and less precise reward response. **B,** Top, individual PETHs of all identified cholinergic neurons revealed heterogeneous responses to water reward (left, NB; right, HDB). Arrows indicate the example neurons in panel A (black, left neuron). The order of neurons corresponds to that of Fig. 3B. Bottom, average PETHs. **C,** Identified cholinergic neurons (purple) showed a valence preference towards negative reinforcement with increasing depth (Fig. S5). Putative cholinergic neurons are overlaid in grey (regression statistics were calculated from identified neurons). Different symbols indicate individual mice.

Figure 6 | Cholinergic responses are scaled by reinforcement surprise

A, Spike rasters and PETHs aligned to punishment separated according to stimulus signal-to-noise ratio of an identified NB cholinergic neuron. **B,** Top, average PETH across identified NB cholinergic neurons (n = 17; neurons for which the application of four different stimulus intensities allowed this analysis were included). Bottom, bar graph showing no modulation of punishment-evoked cholinergic activation by the strength of the preceding stimulus. **C,** Average PETH and bar graph for HDB cholinergic neurons (two neurons that did not show any response to water reward were excluded). **D,** Spike rasters and PETHs aligned to water reward corresponding to the same neuron as in A. **E,** Average PETH across identified NB cholinergic neurons. Bottom, bar graph showing significant modulation of reward-evoked cholinergic activation by the strength of the preceding stimulus ($p < 0.0001$). **E,** Average PETH and bar graph for HDB cholinergic neurons show significant modulation of reward responses ($p = 0.0015$). **G,** ROC analysis quantifying the firing rate difference after faint (20-30dB) or loud (40-50dB) tones for an example cholinergic neuron. Firing rates were significantly different after water reward during the period marked by

the green bar ($p < 0.05$). AUC, area under the ROC curve; a measure of discrimination between two distributions. **H,** Average AUC for identified NB cholinergic neurons. AUC was significantly positive after water reward during the period marked by the green bar ($p < 0.01$). **I,** HMM model of the auditory detection task (see Results and Experimental Procedures). **J,** The HMM model successfully reproduced psychometric functions (left) and RT (right) of the animals (compare with Fig. 2). **K,** The HMM model provided a measure of reinforcement surprise, which closely matched the firing responses of cholinergic neurons (Fig. S6). Punishm., punishment; rew., reward.

Figure 7 | Attentional correlates are characteristic of a population of non-cholinergic neurons

A, Left, subjective hazard rate corresponding to the bimodal foreperiod distribution (overlaid in grey). Middle, median RT as a function of the foreperiod from a single mouse, smoothed with a moving average of order 3 (only 20 and 30 dB trials included). Right, RT as a function of the foreperiod separated by signal-to-noise ratio of stimuli. RT for difficult stimuli was modulated by the foreperiod. Shading, s.e. **B,** As a control, reaction times (right) were constant for constant subjective hazard rates (left) corresponding to exponential foreperiod distributions (overlaid in grey). **C,** Top left, raster plot of an unidentified NB neuron during the foreperiod. Trials are aligned to stimulus onset (brown line) and sorted by RT (orange ticks). High firing rate in the foreperiod predicted fast RT. Top right, RT and pre-stimulus firing rate showed negative trial-to-trial correlation. Bottom, RT (orange) tracked firing rate changes (black) at multiple time scales. **D,** Top, PETH for the same neuron separated by slow and fast RT (median split). Bottom, ROC analysis quantifying the difference between firing rates for slow and fast RT trials. Shading, s.e.m. AUC, area under the curve, quantifying predictive value. **E,** Top, ROC analysis for all unidentified basal forebrain neurons (NB and HDB combined) that showed significant RT predictive firing. Middle, ROC analysis for identified cholinergic neurons. Bottom, average ROC for the two populations. **F,** Performance was plotted separately for trials with high (solid lines) and low firing rate (dashed lines) in the foreperiod (median-split) for two NB neurons. Pre-stimulus firing rate predicted performance accuracy. **G,** Difference in discrimination performance between high and low firing rate trials (quantifying predictive value) as a function of stimulus intensity for the neurons in F. Shading, bootstap s.e.m. **H,** Average discrimination difference for all performance predictive unidentified basal forebrain cells (black; NB and HDB combined) and all cholinergic neurons (purple). Shading, s.e.m. FR, firing rate. See also Fig. S7.