HHS Public Access

Author manuscript

Neuron. Author manuscript; available in PMC 2016 April 08.

Published in final edited form as:

Neuron. 2015 April 8; 86(1): 319–330. doi:10.1016/j.neuron.2015.02.043.

Visually-cued action timing in the primary visual cortex

Vijay Mohan K Namboodiri[1], **Marco Huertas**[2], **Kevin J Monk**[1], **Harel Z Shouval**[2], and **Marshall G Hussain Shuler**[1,*]

[1]Dept. of Neuroscience, Johns Hopkins University, Baltimore, MD

[2]Dept. of Neurobiology and Anatomy, University of Texas-Houston, Houston, TX

SUMMARY

Most behaviors are generated in three steps: sensing the external world, processing that information to instruct decision-making, and producing a motor action. Sensory areas, especially primary sensory cortices, have long been held to be involved only in the first step of this sequence. Here, we develop a visually-cued interval timing task that requires rats to decide when to perform an action following a brief visual stimulus. Using single-unit recordings and optogenetics in this task, we show that activity generated by the primary visual cortex (V1) embodies the target interval and may instruct the decision to time the action on a trial-by-trial basis. A spiking neuronal model of local recurrent connections in V1 produces neural responses that predict and drive the timing of future actions, rationalizing our observations. Our data demonstrate that the primary visual cortex may contribute to the instruction of visually-cued timed actions.

1 Introduction

The production of a behavior often requires an animal to sense the external world, make decisions based on that information and generate an appropriate motor response (Goldman-Rakic, 1988; Kandel et al., 2000; Miller and Cohen, 2001). The canonical view of brain organization is that these functions are performed hierarchically by sensory, association, and motor areas respectively (Felleman and Van Essen, 1991; Kandel et al., 2000; Miller and Cohen, 2001). The role of sensory areas—especially primary sensory areas—has long been regarded as providing a faithful representation of the external world (Felleman and Van Essen, 1991; Goldman-Rakic, 1988; Kandel et al., 2000; Miller and Cohen, 2001); several studies have shown that these areas convey sensory information (Ghazanfar and Schroeder, 2006; Hubel and Wiesel, 1962, 1968; Lemus et al., 2010; Liang et al., 2013), while others have shown causal roles in sensory perception (Glickfeld et al., 2013; Jaramillo and Zador, 2011; Sachidhanandam et al., 2013; Znamenskiy and Zador, 2013). However, this view has

*Correspondence and requests for material should be addressed to M.G.H.S at shuler@jhmi.edu.

Author Contributions: V.M.K.N. and M.G.H.S. designed the study. V.M.K.N. performed the experiments and analyzed the data. M.H. and H.S. created the spiking neuronal model. K.J.M. performed infection surgeries and histology. V.M.K.N. and M.G.H.S. wrote the manuscript with comments from the other authors. The authors have no competing financial interests to disclose.

recently been challenged by observations that sensory cortices represent not only stimulus features but also non-sensory information (Abolafia et al., 2011; Ayaz et al., 2013; Brosch et al., 2011; Fontanini and Katz, 2008; Gavornik and Bear, 2014; Jaramillo and Zador, 2011; Keller et al., 2012; Niell and Stryker, 2010; Niwa et al., 2012; Pantoja et al., 2007; Samuelsen et al., 2012; Serences, 2008; Shuler and Bear, 2006; Stanisor et al., 2013; Zelano et al., 2011). In the visual modality, it has been shown that V1 can predict the learned typical interval between a stimulus and a reward (Chubykin et al., 2013; Shuler and Bear, 2006) and that the ability to learn such intervals depends on cholinergic input from the basal forebrain (Chubykin et al., 2013). In fact, similar timing responses can be conditioned even within an isolated in-vitro preparation of V1 (Chubykin et al., 2013), advancing the site of learning as local to V1. However, whether such predictive signals (Brosch et al., 2011; Chubykin et al., 2013; Pantoja et al., 2007; Serences, 2008; Shuler and Bear, 2006; Stanisor et al., 2013) in primary sensory areas can be involved in instructing behavior is unclear.

Here, we developed a novel visually-cued interval timing task to address this question. Rats performing this task must decide when to lick on a spout to obtain the maximum water reward: licking at longer delays following the visual stimulus (up to a target interval) results in larger reward volumes. Delays longer than the target interval result in no reward (Figure 1). Hence, trained animals wait a stereotyped interval after the stimulus before deciding to lick. The design of our current task was motivated to address whether V1 activity that reflects the average delay between stimulus and reward could be used to directly instruct the lapse of a target interval in order to time an action. As in prior tasks, single unit recordings showed responses that represent the mean expected delay between the stimulus and the reward; but in addition, other neurons' responses correlate with the timed action on a trial-by-trial basis. Among this latter group, we found neurons that represent a target interval from the cue as well as others that report the expiry of the target interval, potentially informing the timing of the behavioral response. Crucially, we show that many neural responses correlate with the timing of the action only on trials in which the animal timed its behavioral action from the visual stimulus and that the action can be decoded on a trial-by-trial basis on these trials. In contrast, on trials in which the action was not visually timed, the firing of the neurons did not relate to the action, nor could the action be decoded from the firing rates, even though these trials contain the same visual stimulus and action. Further, even when the action was visually-timed, many neurons convey information about the timing of the action in an eye-specific manner. We further show that optogenetic perturbation of activity in V1 during the timed interval (but after cue offset) shifts timing behavior only on visually-timed (but not non-visually-timed) trials. Our results indicate that post-stimulus activity in V1 embodies the wait interval and may inform the timing of the behavioral response. We show that a recurrent network model of spiking neurons (Gavornik and Shouval, 2011; Gavornik et al., 2009) can produce the timing of future actions. As observed experimentally, single unit activity within this network shows trial-by-trial correlations with the action, and a perturbation of the network activity produces a shift in the timing of the action, confirming that a model of local recurrent connections within V1 can rationalize our observations.

2 Visually-cued timing behavior

To initiate a trial in our visually-cued timing task, animals enter a port ("nosepoke") containing a lick spout and remain so that after a random delay, a monocular full-field visual stimulus is presented. In order to ensure that the animal is not licking at the moment of visual stimulus delivery, a 0.2–1.2s lick-free pause prior to visual stimulus delivery was imposed. The reward is delivered immediately upon the first lick after the visual stimulus. Importantly, the amount of reward obtained has a ramp profile with respect to the *time waited by animals from the visual stimulus until the first lick* ("wait time") (Figure 1A). Early in training, the wait times of animals occur at very short lags after the visual stimulus and are likely independent of the stimulus (see Supplemental Introduction *Experimental Procedures* for details of behavioral shaping) (Figure 1C). Unlike early stages, as the animals advance to an intermediate stage of training, their wait times begin to be stereotyped. In other words, the cumulative distribution function (CDF) of their wait times acquires a sigmoidal shape—indicating that their lick behavior is increasingly being timed from the visual stimulus. In later stages of learning, the animals' behavior shows a tighter sigmoidal shape (Figure 1C, D) with a median greater than one second, and a low coefficient of variation (CV) (~0.25). To ascertain whether the asymptotic wait times are near-optimal (Balci et al., 2011), we calculated the mean wait time that leads to the maximum reward per trial for a given coefficient of variation (Figure 1D, red dotted line). Because timing behavior follows scalar variability (Buhusi and Meck, 2005; Buonomano, 2007; Matell and Meck, 2004; Merchant et al., 2013), longer wait times are associated with larger variances, such that a larger proportion of wait times exceed the target interval, receiving no reward (Figure S1). Hence, the optimal wait time (given ramping reward to 1.5s) for a CV close to 0.25 is approximately 1.1 seconds. The asymptotic wait times for the animals were thus near-optimal (Figure 1D).

If these stereotyped wait times truly reflected timing from the visual stimulus, they should have no dependence on other behavioral events (such as the nosepoke entry), as was largely observed (Figure 1E). However, a closer examination of the dependence between per-trial wait times and the corresponding delays between the nosepoke entry and visual stimulus revealed a small, but significant, negative correlation ($R^2 \sim 0.05$) (an example session is shown in Figure 1F). A simple explanation of this correlation is that a subset of trials was timed from the nosepoke entry rather than the visual stimulus. If trials were timed from the nosepoke, the longer the delay from nosepoke to visual stimulus presentation, the shorter the corresponding wait time would be from the visual stimulus. To detect such non-visually-timed trials, we analyzed the effect of removing individual trials from the correlation, so as to identify the trials that contributed the most to the original correlation (Figure S1). Such non-visually-timed trials (marked with red squares in Figure 1F, G) also include trials showing outlier wait times (wait times less than 300ms or greater than 3000ms; see Experimental Procedures). These "non-visually-timed" trials necessarily cannot be fully timed from the visual stimulus as they show correlations with the nosepoke, though they may yet be partly influenced by it. Nonetheless, aside from the outliers, the non-visually-timed trials show a consistent delay from nosepoke-entry independent of the visual stimulus, implying that they are likely timed from the nosepoke entry (Figure 1G). Separating out

timing signals from those that merely reflect the action used to indicate the expiry of the timed interval is a challenge associated with studying the neural genesis of timing (Brody et al., 2003; Narayanan and Laubach, 2009; Xu et al., 2014). Hence, we used the above method to separate trials containing nominally the same action (licks), based on whether or not the actions were timed from the visual stimulus. Thus, if activity in V1 was driven by the action itself, it would be present on both visually-timed and non-visually-timed trials. If, on the other hand, V1 instructed the timing of the action, the activity in V1 would correlate with the action only on visually-timed trials, and not on non-visually-timed trials.

3 Neural activity conveys action timing

Neurons that embody an interval from a visual stimulus should reflect this information in their firing profile and can do so in different ways including sustained responses or linear ramps of a given duration (Figure S2A). It has been shown that neurons in V1 can represent both 1) temporal intervals (by sustained modulations of their firing rate for the duration of the mean delay between visual stimulus and reward (Chubykin et al., 2013; Shuler and Bear, 2006)), and, 2) the lapse of such intervals (by a peak in firing rate (Chubykin et al., 2013; Shuler and Bear, 2006)). If such representations are not used for the timing of an action, there will be no trial-by-trial correlation between the neural response and the action, as observed previously and termed “reward timing” (Chubykin et al., 2013; Shuler and Bear, 2006) since they expressed the typical delay to reward (Figure 2A–D shows a schematic). On the other hand, if the interval represented by such neurons is used to instruct the timing of an action, there would be a trial-by-trial correlation between the neural representation of the interval and the action. Conceptually, an “intervalkeeper” neuron that represents a target interval would persistently modulate its firing rate (by either an increase or a decrease) from the onset of the visual stimulus until the target interval expires. Since the animal’s indication of the lapse of the interval (lick) is informed by such a neuron, the lick would follow the moment at which the neuron indicates the expiry of the target interval. On trials in which the neural response lasts longer, the licks will be correspondingly delayed and vice-versa (Figure 2E, F). Neurons that decode the activity of such intervalkeepers to instruct the animal to lick would modulate their population firing rate immediately prior to the lick: the later they fire on a trial, the later the lick occurs (Figure 2G, H). V1 neurons recorded from well-trained animals with stereotyped timing behavior showed responses expressing these conceptualized forms, with strong trial-by-trial correlations with the action indicating the expiry of the timed interval (Figure 2I–L), in addition to other neurons showing reward timing (Figure S2B). Such units that show trial-by-trial correlations with the action in their response profile are labeled as “action units”.

We quantified the correlation between neural response and action by testing whether a neuron modulated its firing rate at fixed latencies with respect to the action (see Experimental Procedures and Figure S7D). To further test whether such an observed firing rate change truly results from timing the action and not, alternatively, reflecting merely the past presence of the corresponding visual stimulus (i.e. reward timing), we did a shuffle analysis that maintained the average relationship between the stimulus and the action, but shuffled the trial-by-trial relationship to the visual stimulus (see Experimental Procedures, Figure S7D and Figure S3). Of a total of 363 single units recorded from five trained animals,

we found significant correlations with the action (in excess of visually-driven correlations) in 122 (Figure 4A). These strong trial-by-trial correlations observed in the above "action" units can arise from two possible scenarios: either the neural response instructs the action ("action timing") or the action instructs the neural response ("action feedback").

Firing rate modulations that drive the visually-cued timing of the action should be absent when the visual stimulus is not being used to time the action—as in "non-visually-timed" trials. Therefore, we tested the activity of these 122 action-correlated neurons in trials with non-visually-timed licks. Additionally, such activity should be absent when the stimulus is not presented prior to the animal initiating licking in the nosepoke. To address this possibility, on some trials ("NoStim", see Experimental Procedures) we withheld the presentation of the visual stimulus, which, nevertheless, could result in licks, though not timed from a visual stimulus. Further, entries into the nosepoke during a mandatory intertrial interval sometimes resulted in licks (false entry licks) that also, then, were not timed from a visual stimulus (see Experimental Procedures). Such licks made in the absence of visual stimuli (NoStim licks and false entry licks), are referred to as "false first licks". These false first licks afford an additional opportunity to examine neural activity to first licks not timed from the visual stimulus. Note however, that the occurrence of all of these non-visually timed licks diminishes with increases in performance level.

Of the 122 "action" units, we found that 38 (labeled "action timing" units) showed a significant difference in activity between *both* visually-timed and non-visually-timed trials, *as well as* between visually-timed trials and false first licks (see Figure 3A, B for examples and Figure S3 for analysis). Hence, the activity of these units cannot be explained by the mere presence of the action. Seven units showed no significant differences between these trial types, but showed significant responses to *either* the false first licks *or* the non-visually-timed licks (see Figure S4A for an example). These units are labeled "action feedback" units as they contain information about the first lick, independent of whether or not it was timed from the visual stimulus. The remaining 77 units were classified as "action timing/feedback units" because they were from animals performing an insufficient number of false first licks or non-visually-timed trials to unambiguously classify as "action timing" or "action feedback" (see Figure S4B for an example). This is because as animals gain more and more experience, they perform fewer and fewer non-visually-timed or false first licks. Nonetheless, these units show significant trial-to-trial correlations with the action and are therefore likely to be predominantly "action timing" units, as 38 out of 45 (number of "action timing" + number of "action feedback") with sufficient statistical power are "action timing" (Figure S4B shows such a likely "action timing" unit). Also, note that as with any analysis, the classification of trials into visually-timed and non-visually-timed may contain false positives/negatives. Yet, the presence of such errors would but reduce our ability to distinguish neural activity between these trial types.

Hence, consistent with the hypothesis that neural activity drives behavior when animals time their action from the cue, we find that a significant portion of units show a difference between visually-timed trials and non-visually-timed trials. Additionally, if the action did in fact result from the firing pattern of these neurons, their activity should contain information about the action prior to its occurrence. In line with this prediction, we found that the earliest

latency at which “action timing” units carried information about the action (“neural report of time”; see Experimental Procedures) was prior to the action (median= −50ms, p=0.0013, one-tailed, W_{32}=104.5, z= −3.01, Wilcoxon signed-rank test). Further, the action feedback units showed a median neural report of time significantly later than the median report of the timing units (median= 50ms, p=0.005, one-tailed, U=242, z=2.55, Mann-Whitney U-test) (Figure 4C).

The prior results confirm that a linear decoder (ROC) can detect a change in firing rate locked to the action across all trials within a unit. To test whether the firing rate change on an individual trial can predict the moment of the action on that trial, we created a maximum likelihood decoder that detected the moment of firing rate transition (Figure 5). This decoder showed that the moment of action can be decoded on a trial-by-trial basis using the firing rate pattern on visually-timed trials for the population of “action timing units”. On the other hand, the firing pattern of “action timing units” in V1 could not be used to decode the action on non-visually-timed trials. These two observations, again, advance V1’s role in instructing the timed action.

To further examine the 77 units that defied classification as action timing or action feedback (due to the low number of false first licks and non-visually-timed trials performed by highly trained animals), we reasoned as follows: should these units simply reflect the action, there should be no difference in their responses to the action following stimulation of either eye. Contrary to this, we found that 29 out of these 77 showed significant differences (see Supplemental Experimental Procedures) at the time of the action (Figure 4D), according to which eye was stimulated in that trial. This is inconsistent with them merely reporting the presence of the action (see Figure S4 for an example). Further, differences in firing based on eye-of-origin are regarded as absent in areas outside of V1 (Burkhalter and Van Essen, 1986; Crick and Koch, 1995; Hubel and Livingstone, 1987; Maunsell and Van Essen, 1983; Shimojo and Nakayama, 1990). Thus, these eye-specific differences diminish the likelihood that the observed activity is merely relayed to V1 by “higher” brain regions.

We also found that out of 351 single units recorded from three naïve animals early in training (Experimental Procedures), only seven showed significant action correlations (Figure 4E), approximating the expected false positive rate of the test (see Figure S4C–I to see all of these responses). This observation indicates that trial-by-trial correlation exists in V1 only after animals learn the wait time-reward contingency to time their licking behavior, again confirming that the response is not simply driven by the action of licking.

4 Optogenetic perturbation consistently shifts timing

If visual stimuli evoke responses in V1 that not only convey the presence of the stimulus, but also instruct the timing of the action, a manipulation of neural activity in the interceding interval between the stimulus and the action should lead to a consistent shift in the timed behavior. In order to test this hypothesis, we measured the behavioral effect of a brief optogenetic perturbation of ongoing activity in V1 during the wait time (see Experimental Procedures and Figure 6A). This perturbation (lasting 200ms) was applied 300 ms after the visual stimulus offset so as to minimize interference with the ability of animals to sense the

stimulus (Figure 6A–D). We confirmed that the optogenetic perturbation was able to affect the firing properties of the network using single unit recordings in one animal (Figure S6). Should the optogenetic perturbation affect the timing of the action as hypothesized, the entire distribution of wait times must shift, as experimentally observed (Figure 6C). In order to quantify whether the *entire* CDF of wait times showed a shift, we measured the shift as the "normalized absolute median percentile shift" (defined in Experimental Procedures; also see Figure S6B). As a population, we found that perturbations over a range of intensities significantly (see Experimental Procedures) shifted the wait time (Figure 6E) compared to fluctuations expected by chance ($p<0.001$, two-tailed, bootstrapping, n=18; see Figure S6 and Experimental Procedures). To test whether this shift resulted from non-specific effects of the laser, we performed the same experiment in saline-injected animals and found no significant effect ($p>0.1$, two-tailed, bootstrapping, n=19). Additionally, the effect in virally-infected (mean = 0.111) animals was significantly higher than in saline-injected animals (mean = 0.061, $p=0.0036$, two-tailed, $t_{35}=3.22$, Welch's t-test).

It could potentially be argued that the shift in timing observed due to the optogenetic perturbation is a result of the animal treating the laser as a second visual stimulus, or of the animal becoming "confused" by the perturbation. We controlled for these possibilities in three ways. One, we observed that across all sessions from all experimental animals, the laser induced a shift only on visually-timed trials and not on non-visually-timed trials, even though both trial types contained the exact same visual stimulus, laser, and action (Figure 6F). The absence of a shift on non-visually-timed trials indicates that animals are not treating the laser as a visual stimulus or becoming "confused". Second, to further examine the hypothesis that animals might treat the laser as a visual stimulus, we assessed the probability of licking in "NoStim" trials in the presence and absence of laser for individual animals (at maximum intensity). We found no significant increase in licking induced by the laser in any animal ($p>0.05$ with Holm-Bonferroni correction for multiple comparisons, n=30 trials in each animal, one-tailed bootstrapping; Supplemental Experimental Procedures). Further, speeding up timing as shown in Figure 6C is inconsistent with animals resetting their timing upon receiving the optogenetic perturbation. Third, to address whether the animals become "confused", we measured the variability in timing behavior under perturbation. We found (Figure S7I) that the optogenetic perturbation in experimental animals did not affect the coefficient of variation of wait times ($p=0.85$, two-tailed Wilcoxon signed-rank test, $W_{18}=81$, $z = -0.1960$). This implies that the animals are still timing their licks (albeit shifted) as their actions remain consistent with scalar timing (Buhusi and Meck, 2005; Buonomano, 2007; Matell and Meck, 2004; Merchant et al., 2013), and, that the observed shift does not merely result from a non-specific disruption of behavior. In the next section, we show that a model of a *local* recurrent network within V1 is sufficient to rationalize the results of our optogenetic experiment. Taken together, these results indicate that optogenetically perturbing V1 activity in the interceding wait time between the visual stimulus and the action causally shifts the timing of visually-cued actions.

5 Spiking neuronal model & Reward responses

Based on these data, one can infer that V1 may be involved in instructing action timing. But what is the mechanism that gives rise to visually-cued action timing? Previous experimental

work has shown that the average delay between predictive cues and reward (i.e. reward timing) can be locally generated within V1 (Chubykin et al., 2013). Further, a computational model has been proposed for how such visually-cued temporal intervals may be learned and expressed (Gavornik and Shouval, 2011; Gavornik et al., 2009). Therefore, we postulated that action timing may arise from reward timing activity. To explore this possibility, we examined whether the same local recurrent connections in V1 that give rise to reward timing activity can also generate action timing. We reasoned that since a population of neurons can report an average temporal interval (reward timing), a subpopulation could be used to instruct the timing of an action, thereby expressing both reward and action timing activity in V1. If so, would individual neurons in this subpopulation show trial-by-trial correlations with the action despite the fact that only their aggregate activity instructs the action? To address these questions, we modified (Figure 7A) our previous model of visually-cued reward timing in V1 (Gavornik et al., 2009), so generating both action and reward timing responses (see Experimental Procedures for details). Individual neurons of the subpopulation reporting the expiry of the interval and driving the action showed responses remarkably similar to the experimental data, exhibiting significant trial-by-trial correlations with the action ($p<0.01$, n=60 trials, bootstrapping) (Figure 7B). The remaining neurons reflected only the average time between stimulus and reward (reward timing), and not the individual actions (Figure S7). Therefore, action timing can be derived from reward timing activity.

To test whether the above model (producing timed actions based on local recurrent connections within a network) can also explain the results of our optogenetic experiment, we introduced a perturbation in the network mimicking the optogenetic perturbation (see Experimental Procedures). In order to assess whether the experimental observations can be sufficiently explained by our model, we tested the following four questions: 1) does the perturbation produce a shift in the distribution of wait times, 2) if so, is this a consistent shift in the distribution, 3) is the magnitude of shift comparable to the experimental data (within reasonable ranges of parameters), and, 4) does the shifted distribution still abide by scalar timing? The answer to each question was in the affirmative, as experimentally observed (Figure 7C and Figure S7). This confirms that the experimentally observed shift in behavior can be rationalized by a model of V1 that *locally* generates and instructs the timing of the action.

In order for the model to learn to predict and drive the timing of future actions, it would need to receive a feedback signal that reflects the magnitude of reward received (Gavornik et al., 2009) (see Experimental Procedures). Such primary reward responses have been reported in auditory cortex (Brosch et al., 2011; FitzGerald et al., 2013) and somatosensory cortex (Pleger et al., 2008, 2009), but have not hitherto been reported within V1 (Weil et al., 2010). Because of the unique design of our task in which different wait times lead to differing magnitudes of reward, we were able to test whether the magnitude of received reward is conveyed by the firing pattern of neurons in V1. We found that 23 units (Figure 8) showed significant differences between rewarded and unrewarded trials in addition to showing a significant correlation with the amount of reward received (see Figure S8 and Supplemental Experimental Procedures for analysis). This confirms that primary reward

information reaches V1, potentially permitting the relationship between actions and rewards to be learned locally.

6 Discussion

The production of the stereotyped timing behavior observed in our visually-cued timing task requires that the brain sense the presence of a visual stimulus, generate the stereotyped waiting interval, and then produce the lick. The traditional view of primary sensory cortices would assume that V1 has a role only in sensing the stimulus. However, we show that the cue evoked responses in V1 evolve from those that merely reflect the presence of a stimulus to those consistent with driving the timing of delayed voluntary actions. By parsing the effect of the timed interval from the action indicating the lapse of that interval, we show that neurons in V1 correlate with the action only when it is visually-timed, ruling out the alternative that the correlation is merely due to the presence of the action. This is the first time that a neural signal generating a timed action by a persistent modulation of firing during the wait interval until action has been observed anywhere in the cortex (Buhusi and Meck, 2005; Buonomano and Laje, 2010; Johnson et al., 2010; Leon and Shadlen, 2003; Matell and Meck, 2000; Merchant et al., 2013). Perturbing activity in V1 during the interceding interval indicates that V1 neurons may contribute to instructing the decision of when to lick. Our spiking neuronal model confirms that such control of timed actions can be performed by a recurrent network of neurons. Interestingly, due to the unique design of our task (with a continuous gradient of reward depending on the time waited by animals), we also found primary reward responses in V1, confirming that it does indeed receive feedback regarding the outcome of recent actions. Our data indicates that “higher order” decisions, such as when to perform an action, may be instructed even by primary sensory areas. Note, however, that we do not claim that the timing of visually-cued actions in this task is *exclusively* controlled by V1, or that other regions do not play a role. Our optogenetic experiment was thus designed only to perturb the activity, rather than silence it, since intact timing behavior under silencing may simply mean that other regions also instruct timing or compensate for the absence of activity in V1. Perturbing (rather than silencing) activity in V1 to show a shift in timing behavior, however, directly addresses whether V1 itself has any role in the timing of the actions. Because of the metabolic cost incurred from using more neurons than strictly necessary to optimally perform a task, we posit that it may be sufficient for a primary sensory area to control highly stereotyped actions so as to efficiently balance the trade-off between redundancy of processing, metabolic demands, and reward maximization (Bullmore and Sporns, 2012; Namboodiri et al., 2014a, 2014b).

7 Experimental Procedures

7.1 Subjects

Animal subjects were wild-type adult male Long-Evans rats with weights between 300–400g and age between P60–100. All experimental procedures were approved by the Institutional Animal Care and Use Committee. The animals were singly-housed in a vivarium with a 12 hour light/dark cycle and the experiments were conducted during the day (light cycle).

7.2 Behavioral Task

The behavioral task was a visually-cued timing task as shown in Figure 1A, wherein the amount of reward obtained was dependent on the time waited until first lick post visual stimulus presentation (Supplemental Experimental Procedures).

7.3 Behavioral analysis

All analysis in this paper was performed by custom-written scripts in MATLAB (The Mathworks). In order to parse visually-timed trials from non-visually-timed trials, we first removed trials showing outlier wait times (shorter than 300ms or longer than 3000ms). Subsequently, the effect of removing individual trials on the R^2 of the correlation between the delay from nosepoke entry to visual stimulus presentation and the wait time was measured (whenever the original correlation was significant). If removal of a given trial produced a reduction in correlation that was 1.5 times the interquartile range from the median change in correlation, it was deemed to be a non-visually-timed trial (Figure S1). For more details, see Supplemental Experimental Procedures.

7.4 Neural recordings

See Supplemental Experimental Procedures for details of neural recording.

7.5 Analysis of neural response

To test whether there were any trial-by-trial correlations between the activity of a single unit and the first lick, we tested whether there was a significant change in firing rate locked to the first lick in visually-timed trials, separately for left-eye-stimulated and right-eye-stimulated trials. To this end, we checked whether a linear classifier could decode the firing rate change between a 100 ms bin before and after a given latency from the first lick. The latencies tested were −300 ms, −250 ms, −200 ms,…, 100 ms, and, 150 ms. We restricted the decoder analysis to only the two latencies that led to the maximum and minimum signed change in firing rate. Since we were testing only these two bins for significance, we corrected for multiple comparisons using Bonferroni correction. Trials with less than 1 Hz of firing in a window starting from 200 ms prior to the visual stimulus to 200 ms after the first lick were excluded.

Once the latencies with maximum and minimum firing rate changes were identified, units that fired less than 1 Hz in both of these bins were excluded from further analysis. Further, only trials in which the visual stimulus turned off at least 100 ms prior to each latency were included in the analysis for that latency. To test whether an ideal observer can measure the change in firing rate at a given latency, we quantified the area under an ROC curve (AuROC) formed by the spike count distributions for the 100 ms bins flanking the latency. Another possible method was to test whether the peri-stimulus time histogram showed a significant increase in firing rate locked to the first lick. In order to test significance, most standard methods assume Poisson statistics of firing. However, the units recorded in our task often violated this assumption (data not shown). Further, we were interested in quantifying whether a downstream decoding area could tell the difference in firing rate so as to instruct the action. For this purpose, a linear decoding strategy such as that assumed in ROC analysis was deemed more appropriate. Additionally, we were also interested in testing whether the

change in firing rate was consistently present across trials. This was implicitly tested by our resampling method (explained below) in which we resampled individual trials, as opposed to individual spike counts.

Next, we asked two questions (Supplemental Experimental Procedures) to ascertain whether there is a firing rate change locked to the action of the first lick (Figure S8D). First, we asked whether there was a statistically significant change in firing rate from that expected by chance. Second, we asked whether this change could be completely explained by visually-evoked activity. The latter was required to rule out changes locked to the action that did not truly result in or from the action. The significance testing for AuROC was done by the percentile method as explained in (Obuchowski and Lieber, 1998).

The next step in the action timing analysis was to determine whether the trial-by-trial correlations between the neural activity and the action could be explained by the mere presence of the action, or whether it required the trial to actually be timed from the visual stimulus. To this end, we tested whether the change in firing rate observed at the given latency was present also on non-visually-timed trials and false first licks. This was done by conducting a two-sample version of the significance test mentioned above (Obuchowski and Lieber, 1998) to generate sampling distributions of $AuROC_{visually\text{-}timed}$-$AuROC_{non\text{-}visually\text{-}timed}$ and $AuROC_{visually\text{-}timed}$-$AuROC_{falsefirstlicks}$, and determining the corresponding p values (labeled p_1 and p_2). An example is shown in Figure S3. Further, to test whether there were any significant trial-by-trial correlations of the activity with non-visually-timed licks and false first licks, we repeated the same analysis that was performed on visually-timed trials, i.e. it was tested whether there was any significant change in firing rate at the given latency on non-visually-timed licks (labeled p_3) and whether the observed firing rate change was explained by the visually-evoked activity (labeled p_4). It was also tested to see whether there was a significant change in firing rate at the given latency with respect to the false first licks (labeled p_5). For application of this general procedure and its requisite modification for the units recorded early in training as shown in Figure S4, see Supplemental Experimental Procedures.

The neural report of time for a unit was defined as the earliest latency at which it contained any information about the action (in either left-eye or right-eye visually-timed trials). Based on the p values obtained for the five comparisons (p_1-p_5) mentioned above, the unit was subsequently classified as “action timing” or “action feedback”, or as “action timing/feedback” if this classification could not definitively be made (owing to low number of non-visually-timed and false first licks). Action timing units were those that showed *both* p_1 and p_2 lower than 0.05, i.e. visually-timed licks were different from both non-visually-timed and false first licks. Action feedback units were those that were not action timing, but rather showed a significant change in firing rate locked *either* to non-visually-timed licks ($p_3<0.05$ and $p_4<0.05$) *or* false first licks ($p_5<0.05$). All remaining units were classified as “action timing/feedback”.

The trial-by-trial decoding analysis (see Supplemental Experimental Procedures for a full description and Figure S5 for a flowchart) was aimed at testing whether it was possible to decode the moment of first lick from the firing rate pattern of an individual trial. This is

different from the above analysis in that the previous analysis used a linear decoder across all trials, whereas the decoder in this analysis worked only on one trial at a time. The decoder was a maximum likelihood estimator of the moment of transition in firing rate on any given trial (West and Ogden, 1997), with the assumption that the firing rate before and after transition are known. The bin size was set to 100 ms. This decoder is trained using every other trial to learn the expected firing rates before and after transition. If there was no transition predicting the action on a trial-by-trial basis, the decoded moments of transition would be uniformly distributed with respect to the action time. Hence, the ability of the decoder was measured using the excess kurtosis over a uniform distribution of the distribution formed by the decoded moments of transition across all trials. The hypothesis testing for the entire population was done using bootstrapping as explained in Figure S5.

7.6 Optogenetics

Animals were injected with either the virus containing the channelrhodopsin (ChR2) construct (UPenn Vector Core: AAV9.CAG.ChR2-Venus.W.SV40 (Petreanu et al., 2009)) or saline randomly, using sterile surgical procedures (Supplemental Experimental Procedures).

After at least a month, animals underwent a second surgery in which optical fibers were placed at the surface of the brain bilaterally (1.5mm anterior and 4.2mm lateral to lambda) along with goggle posts. In one experimental animal infected with ChR2, custom-built bundles of sixteen electrodes were also implanted bilaterally at a depth of 1.0mm at the above mentioned coordinates, so as to verify that the optogenetic perturbation indeed affected neural activity (Figure S6D–K); Supplemental Experimental Procedures for the immunohistochemistry protocol. Prior to optogenetic testing, the dental cement headcap of each animal was painted black using nail polish so as to prevent the leakage of laser light (Cobolt, Blues 473nm DPSS laser).

For this experiment, the ramping profile of the reward was offset by 500ms (no reward for wait times less than 500ms; done to ensure that the bulk of first-licks fell after the laser offset), reaching the same maximum reward volume at 1500ms. There was a 50% probability that a given trial would result in a laser presentation lasting 200ms (300ms after visual stimulus offset). The laser and no-laser trials were computer-generated and randomly interleaved throughout a session of 360 trials. Laser was also presented on 50% of the NoStim trials to test whether the laser was treated as a visual stimulus, thereby triggering licking behavior. In each of the experimental animals (n=3), two sessions each were run at laser intensities of 0.21, 0.69 and 1.67 mW/mm^2 (measured at the tip of the optical fiber that is placed on the surface of the brain), thereby producing 18 data points. In one of the saline control animals (n=4), only the highest intensity was used since one implanted optical fiber failed prior to testing the lower intensities, and in another, one last session at 0.21 mW/mm^2 could not be performed for the same reason. This resulted in a total of 19 (6+6+2+5) data points from the saline group.

We defined the shift statistic as the absolute median percentile shift (see Supplemental Experimental Procedures for the rationale), calculated thusly: for the CDF of wait times, for each percentile (i.e. 0, 1, 2, …., 99, 100 percentile), the shift in wait times between the

"laser" and "no laser" trials was measured. The absolute value of the median of this list of one hundred data points (i.e. shifts at each percentile) was defined as the absolute median percentile shift (see Figure S6).

There were four different statistical tests performed on this dataset, corresponding to four different questions. In all cases, only first licks following the laser offset were analyzed. The first question was whether, in an individual session, laser trials had a significant shift in wait times. To this end, the observed median percentile shift was compared against that expected by chance. The null distribution was created by resampling with replacement from the wait time distribution in no-laser trials and measuring the absolute median percentile shift for each resampling. This was repeated 2000 times to calculate a two-tailed p value for the observed median percentile shift. This resulted in 8 statistically significant shifts in the experimental group out of a total of 18 sessions, along with 3 sessions with a p value between 0.05 and 0.06. There were no significant shifts (at $p>0.05$) observed in the 19 sessions of the control group.

The second test was to determine whether the population of all observed absolute median percentile shifts was significantly different from that expected by chance. As our aim was only to test whether there was an observed shift for the group as a whole (and not to quantify the shift for each individual laser intensity), we pooled the data from all sessions. Since the larger the variability in the wait times, the larger the shift expected by chance will be, the absolute median percentile shifts were normalized for each animal by dividing it by the inter-quartile range of the wait time distribution in no-laser trials. To test whether the population of normalized absolute median percentile shifts ($n=18$) was significantly different from that expected by chance, it was compared against the bootstrapped null distributions generated from the 18 individual sessions (using the procedure mentioned above after normalization) and was significant in all cases ($p<0.001$, one-tailed). For the purpose of Figure 6, the null distribution of one randomly selected session was used.

The above two questions only answer whether the laser presentation led to an observable shift in the wait time distribution in experimental animals. However, this observed shift could be the result of non-specific effects of laser presentation. To test against this possibility, we tested the population of normalized absolute median percentile shifts from the experimental animals (infected with ChR2) against the saline-control animals. Since we could not assume homoscedascity for these two populations, we performed the Welch's t test for the hypothesis that the shift in the experimental group is larger than the shift in the control group.

We also tested whether the significant effect of the laser in the aggregate as shown in Figure 6E was consistently present in all animals Supplemental Experimental Procedures).

The fourth test was to address whether or not the shift produced by the laser in experimental animals was present on both visually-timed and non-visually-timed trials. If so, one could argue that the effect is due to a non-specific effect of the laser and is not due to an effect on the timing of V1's responses. Since the statistical power for this test is limited within any given session (due to the low number of non-visually-timed trials), we performed it on data

pooled across all sessions from all experimental animals. In this dataset, there were about twice as many non-visually-timed trials with and without laser as there were any laser trials within any typical individual session. Hence, there was enough statistical power to observe a real effect. The test performed was similar to that performed in the first test mentioned above (but without taking the absolute values for the shift). To make the effect size on visually-timed and non-visually-timed trials comparable (due to the drastically different number of trials and hence, the expected effect size in each condition), we calculated the z-score of the normalized median percentile shift with respect to the null distribution (shown in Figure 6F).

7.7 Spiking neuronal model

The model (uploaded to ModelDB and available upon request) consisted of three populations of spiking neurons forming a network as illustrated in Figure 7A. Two populations, labeled E and P, contain excitatory neurons and the other, labeled I, is formed by inhibitory neurons. All synaptic weights are denoted by the form W_{ij}, where the indices i and j label post- and pre-synaptic populations. The magnitudes of these weights were learned as explained previously (Gavornik and Shouval, 2011; Gavornik et al., 2009) so as to represent a median interval of one second. Recurrent synaptic connections within E (W_{EE}) and P (W_{PP}) populations and projections between these (W_{PE}) had a spatial probability of connection that is Gaussian-distributed with a standard deviation equivalent to 12 neighboring neurons, while connections between excitatory and inhibitory populations (W_{IE}, W_{EI}, W_{IP}, W_{PI}) were sparse with a uniform 20% probability of connection.

To see how individual neurons were modeled, and for details of action generation in the model, see Supplemental Experimental Procedures.

To simulate the effect of optogenetic activation of neurons, individual neurons received an additional excitatory current 400 ms following the visual stimulus onset, lasting for 200ms (Supplemental Experimental Procedures). For the simulations in Figure 7C, 9% of the excitatory neurons and 20% of the inhibitory neurons were activated for 200 trials (see Supplemental Experimental Procedures for details). The activation ratios were systematically varied, showing perturbations similar to those observed experimentally (Figure S7).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Stefan Mihalas, Joshua Levy, Tanya Marton, Emma Roach and Camila Zold for helpful discussions throughout the work, and Simon Allard for comments on the manuscript. This work was funded by NIMH (R01 MH084911 and R01 MH093665) to M.G.H.S. The microscope used for visualizing histology was funded by NINDS (NS050274).

References

Abolafia JM, Martinez-Garcia M, Deco G, Sanchez-Vives MV. Slow Modulation of Ongoing Discharge in the Auditory Cortex during an Interval-Discrimination Task. Front Integr Neurosci. 2011; 5:60. [PubMed: 22022308]

Ayaz A, Saleem AB, Schölvinck ML, Carandini M. Locomotion controls spatial integration in mouse visual cortex. Curr Biol. 2013; 23:890–894. [PubMed: 23664971]

Balci F, Freestone D, Simen P, Desouza L, Cohen JD, Holmes P. Optimal temporal risk assessment. Front Integr Neurosci. 2011; 5:56. [PubMed: 21991250]

Brody CD, Hernandez A, Zainos A, Romo R. Timing and Neural Encoding of Somatosensory Parametric Working Memory in Macaque Prefrontal Cortex. Cereb Cortex. 2003; 13:1196–1207. [PubMed: 14576211]

Brosch M, Selezneva E, Scheich H. Representation of reward feedback in primate auditory cortex. Front Syst Neurosci. 2011; 5:5. [PubMed: 21369350]

Buhusi CV, Meck WH. What makes us tick? Functional and neural mechanisms of interval timing. Nat Rev Neurosci. 2005; 6:755–765. [PubMed: 16163383]

Bullmore E, Sporns O. The economy of brain network organization. Nat Rev Neurosci. 2012; 13:336–349. [PubMed: 22498897]

Buonomano DV. The biology of time across different scales. Nat Chem Biol. 2007; 3:594–597. [PubMed: 17876310]

Buonomano DV, Laje R. Population clocks: motor timing with neural dynamics. Trends Cogn Sci. 2010; 14:520–527. [PubMed: 20889368]

Burkhalter A, Van Essen DC. Processing of color, form and disparity information in visual areas VP and V2 of ventral extrastriate cortex in the macaque monkey. J Neurosci. 1986; 6:2327–2351. [PubMed: 3746412]

Chubykin AA, Roach EB, Bear MF, Shuler MGH. A cholinergic mechanism for reward timing within primary visual cortex. Neuron. 2013; 77:723–735. [PubMed: 23439124]

Crick F, Koch C. Are we aware of neural activity in primary visual cortex? Nature. 1995; 375:121–123. [PubMed: 7753166]

Felleman DJ, Van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. Cereb Cortex. 1991; 1:1–47. [PubMed: 1822724]

FitzGerald THB, Friston KJ, Dolan RJ. Characterising reward outcome signals in sensory cortex. Neuroimage. 2013; 83:329–334. [PubMed: 23811411]

Fontanini A, Katz DB. Behavioral states, network states, and sensory response variability. J Neurophysiol. 2008; 100:1160–1168. [PubMed: 18614753]

Gavornik JP, Bear MF. Learned spatiotemporal sequence recognition and prediction in primary visual cortex. Nat Neurosci. 2014; 2014

Gavornik JP, Shouval HZ. A network of spiking neurons that can represent interval timing: mean field analysis. J Comput Neurosci. 2011; 30:501–513. [PubMed: 20830512]

Gavornik JP, Shuler MGH, Loewenstein Y, Bear MF, Shouval HZ. Learning reward timing in cortex through reward dependent expression of synaptic plasticity. Proc Natl Acad Sci U S A. 2009; 106:6826–6831. [PubMed: 19346478]

Ghazanfar AA, Schroeder CE. Is neocortex essentially multisensory? Trends Cogn. Sci. 2006; 10:278–285.

Glickfeld LL, Histed MH, Maunsell JHR. Mouse primary visual cortex is used to detect both orientation and contrast changes. J Neurosci. 2013; 33:19416–19422. [PubMed: 24336708]

Goldman-Rakic PS. TOPOGRAPHY OF COGNITION: Parallel Distributed Networks in Primate Association Cortex. Annu Rev Neurosci. 1988; 11:137–156. [PubMed: 3284439]

Hubel DH, Livingstone MS. Segregation of Form, Color and Stereopsis in Primate Area 18. J Neurosci. 1987; 7:3378–3415. [PubMed: 2824714]

Hubel DH, Wiesel TN. RECEPTIVE FIELDS, BINOCULAR INTERACTION AND FUNCTIONAL ARCHITECTURE IN THE CAT'S VISUAL CORTEX. J Physiol. 1962; 160:106–154. [PubMed: 14449617]

Hubel DH, Wiesel TN. RECEPTIVE FIELDS AND FUNCTIONAL ARCHITECTURE OF MONKEY STRIATE CORTEX. J Physiol. 1968; 195:215–243. [PubMed: 4966457]

Jaramillo S, Zador AM. The auditory cortex mediates the perceptual effects of acoustic temporal expectation. Nat Neurosci. 2011; 14:246–251. [PubMed: 21170056]

Johnson HA, Goel A, Buonomano DV. Neural dynamics of in vitro cortical networks reflects experienced temporal patterns. Nat Neurosci. 2010; 13:917–919. [PubMed: 20543842]

Kandel, ER.; Schwartz, JH.; Jessel, TM. Principles of Neural Science. New York: McGraw-Hill; 2000.

Keller GB, Bonhoeffer T, Hübener M. Sensorimotor mismatch signals in primary visual cortex of the behaving mouse. Neuron. 2012; 74:809–815. [PubMed: 22681686]

Lemus L, Hernández A, Luna R, Zainos A, Romo R. Do sensory cortices process more than one sensory modality during perceptual judgments? Neuron. 2010; 67:335–348. [PubMed: 20670839]

Leon MI, Shadlen MN. Representation of time by neurons in the posterior parietal cortex of the macaque. Neuron. 2003; 38:317–327. [PubMed: 12718864]

Liang M, Mouraux A, Hu L, Iannetti GD. Primary sensory cortices contain distinguishable spatial patterns of activity for each sense. Nat Commun. 2013; 4:1979. [PubMed: 23752667]

Matell MS, Meck WH. Neuropsychological mechanisms of interval timing behavior. BioEssays. 2000; 22:94–103. [PubMed: 10649295]

Matell MS, Meck WH. Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes. Cogn Brain Res. 2004; 21:139–170.

Maunsell JH, Van Essen DC. Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. J Neurophysiol. 1983; 49:1127–1147. [PubMed: 6864242]

Merchant H, Harrington DL, Meck WH. Neural basis of the perception and estimation of time. Annu Rev Neurosci. 2013; 36:313–336. [PubMed: 23725000]

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

Namboodiri VMK, Mihalas S, Marton TM, Hussain Shuler MG. A general theory of intertemporal decision-making and the perception of time. Front Behav Neurosci. 2014a; 8:61. [PubMed: 24616677]

Namboodiri VMK, Mihalas S, Hussain Shuler MG. Rationalizing decision-making: understanding the cost and perception of time. Timing Time Percept Rev. 2014b; 1:1–40.

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

Niell CM, Stryker MP. Modulation of visual responses by behavioral state in mouse visual cortex. Neuron. 2010; 65:472–479. [PubMed: 20188652]

Niwa M, Johnson JS, O'Connor KN, Sutter ML. Activity related to perceptual judgment and action in primary auditory cortex. J Neurosci. 2012; 32:3193–3210. [PubMed: 22378891]

Obuchowski NA, Lieber ML. Confidence intervals for the receiver operating characteristic area in studies with small samples. Acad Radiol. 1998; 5:561–571. [PubMed: 9702267]

Pantoja J, Ribeiro S, Wiest M, Soares E, Gervasoni D, Lemos Na M, Nicolelis Ma L. Neuronal activity in the primary somatosensory thalamocortical loop is modulated by reward contingency during tactile discrimination. J Neurosci. 2007; 27:10608–10620. [PubMed: 17898232]

Petreanu L, Mao T, Sternson SM, Svoboda K. The subcellular organization of neocortical excitatory connections. Nature. 2009; 457:1142–1145. [PubMed: 19151697]

Pleger B, Blankenburg F, Ruff CC, Driver J, Dolan RJ. Reward facilitates tactile judgments and modulates hemodynamic responses in human primary somatosensory cortex. J Neurosci. 2008; 28:8161–8168. [PubMed: 18701678]

Pleger B, Ruff CC, Blankenburg F, Klöppel S, Driver J, Dolan RJ. Influence of dopaminergically mediated reward on somatosensory decision-making. PLoS Biol. 2009; 7:e1000164. [PubMed: 19636360]

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

Samuelsen CL, Gardner MPH, Fontanini A. Effects of cue-triggered expectation on cortical processing of taste. Neuron. 2012; 74:410–422. [PubMed: 22542192]

Serences JT. Value-based modulations in human visual cortex. Neuron. 2008; 60:1169–1181. [PubMed: 19109919]

Shimojo S, Nakayama K. Real World Occlusion Constraints and Binocular Rivalry. Vision Res. 1990; 30:69–80. [PubMed: 2321367]

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

Stanisor L, Togt C, Van Der Pennartz CMA, Roelfsema PR. A unified selection signal for attention and reward in primary visual cortex. Proc Natl Acad Sci. 2013; 110:9136–9141. [PubMed: 23676276]

Weil RS, Furl N, Ruff CC, Symmonds M, Flandin G, Dolan RJ, Driver J, Rees G. Rewarding feedback after correct visual discriminations has both general and specific influences on visual cortex. J Neurophysiol. 2010; 104:1746–1757. [PubMed: 20660419]

West BRW, Ogden RT. Continuous-time Estimation of a Change-point in a Poisson Process. J Stat Comput Simul. 1997; 56:293–302.

Xu M, Zhang S, Dan Y, Poo M. Representation of interval timing by temporally scalable firing patterns in rat prefrontal cortex. Proc Natl Acad Sci. 2014; 111:480–485. [PubMed: 24367075]

Zelano C, Mohanty A, Gottfried JA. Olfactory predictive codes and stimulus templates in piriform cortex. Neuron. 2011; 72:178–187. [PubMed: 21982378]

Znamenskiy P, Zador AM. Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. Nature. 2013; 497:482–485. [PubMed: 23636333]

HIGHLIGHTS

- V1 responses embody the intervals waited per trial in a visually-cued timing task.
- This trial-by-trial representation is absent in non-visually-timed trials.
- Optogenetically perturbing V1 during the timed interval shifts behavioral timing.
- A model of a local recurrent network in V1 reproduces and rationalizes our findings.

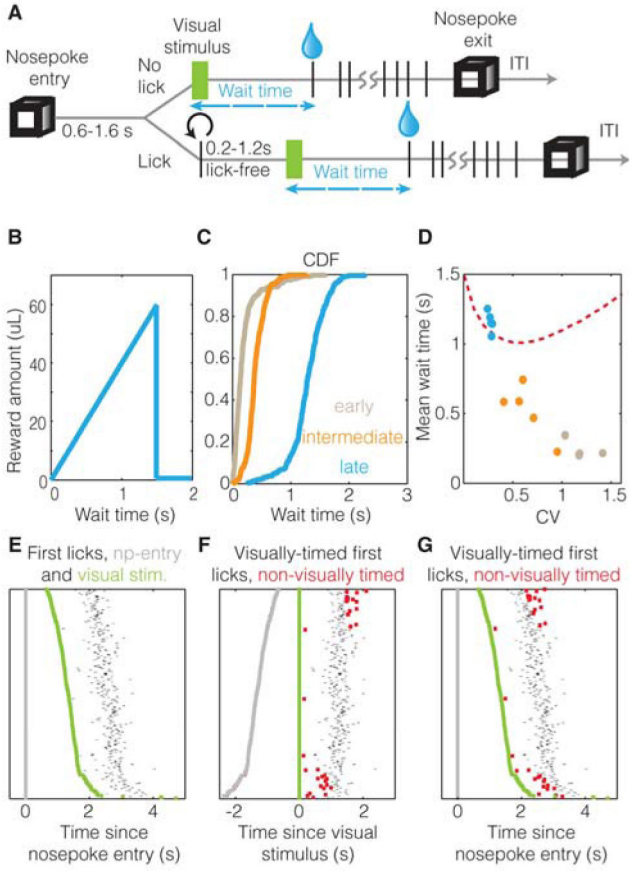

Figure 1. Visually-cued timing behavior

A & B. Visually-cued timing task, showing that the time waited by an animal from the visual stimulus onset until the first lick ("wait time") determines the reward obtained (see Experimental Procedures). **C & D.** Example performance of animals at early (brown), intermediate (orange) and late (cyan) stages of learning, shown using cumulative-distribution-functions (CDF) (**C**) and a population plot (**D**) of asymptotic coefficient of variation (CV) with mean wait time (5 animals each). Red dotted line shows optimal behavior (see Figure S1). **E–G**. Raster plots showing relation between delay from nosepoke-entry (grey) to visual stimulus (green), to the corresponding first licks (black dots), when aligned to nosepoke entry (**E, G**) and visual stimulus (**F**). Non-visually-timed trials (see text) are shown in red. This is the session in which the neuron showed in Figure 3A was recorded.

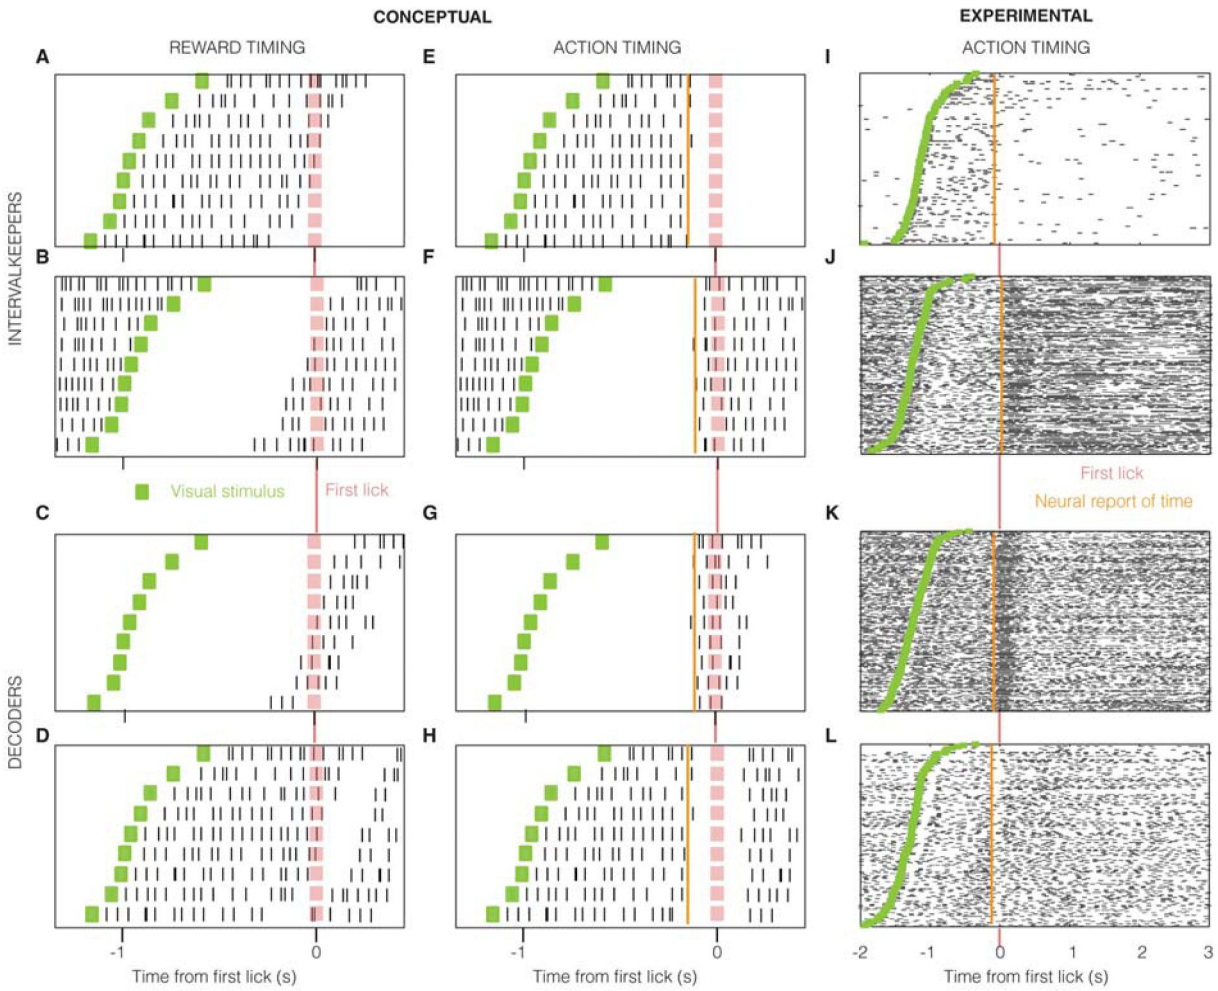

Figure 2. Raster plot showing conceptual and observed neural responses

Shown for example neurons that measure the passage of a target interval (**A, B, E, F, I, J**) from a visual stimulus (green) and those that indicate its expiry (**C, D, G, H, K, L**). The target interval measured by intervalkeepers (and thus reflected by decoders) shows trial-by-trial variability, which is consequently reflected by the action of first lick (pink). The neural report of time is defined as the earliest moment when neural activity conveys information about the action. The values for “Conceptual” neurons are merely illustrative, whereas those for the experimental neurons are quantified using the methods described in Experimental Procedures.

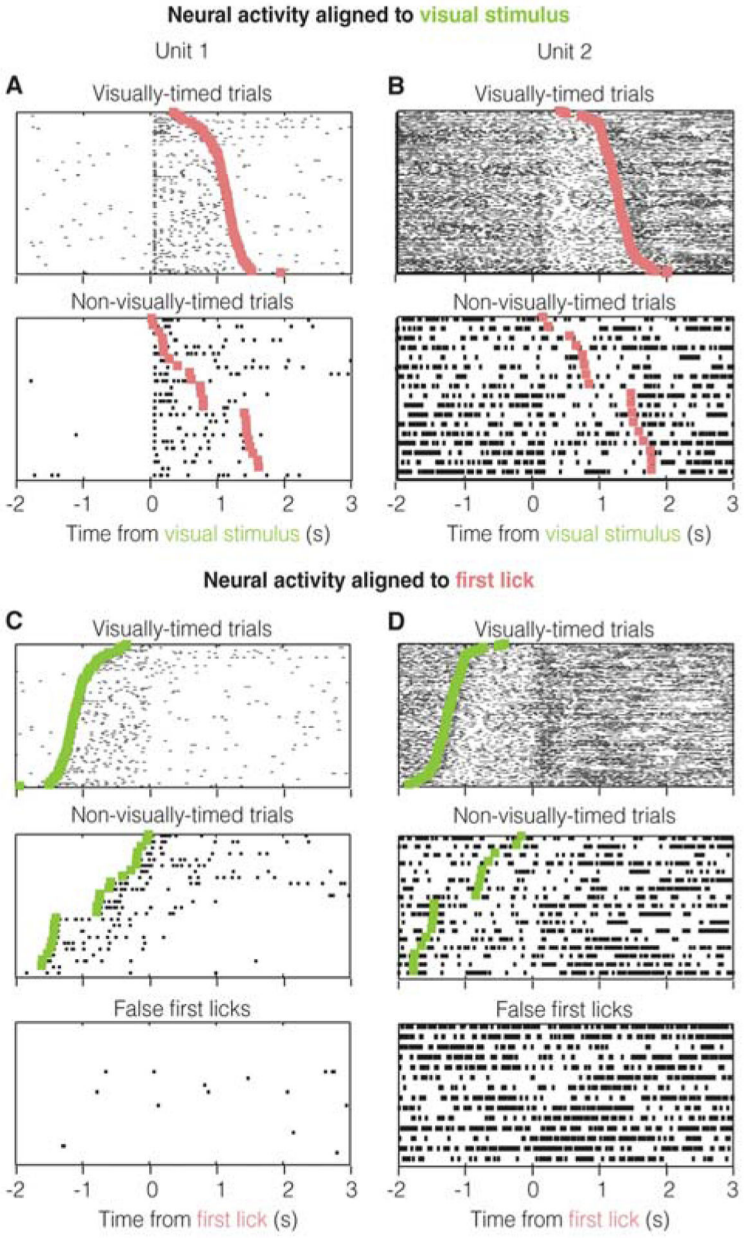

Figure 3. Trial-by-trial correlations reflect timing and not the first lick itself

A & B. Neural activity from example units in Figure 2**I, J** plotted with respect to the visual stimulus on both visually-timed and non-visually-timed trials. The trials are arranged in increasing order of wait times and not chronologically. **C & D.** Activity aligned to first lick, showing that trial-by-trial correlations with first licks are absent when they are not timed from the visual stimulus (see text).

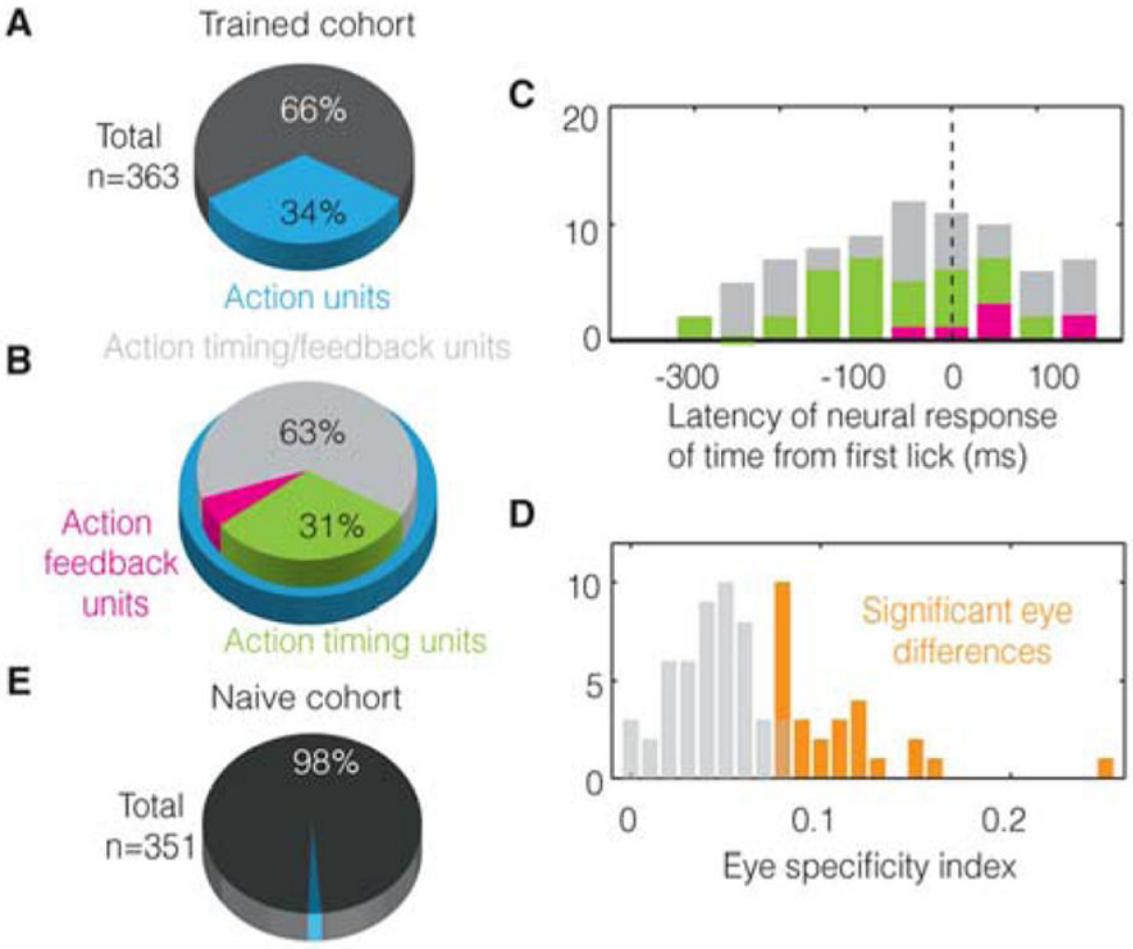

Figure 4. Population analysis of single unit data

A. 122 out of 363 units recorded from 5 trained animals show trial-by-trial correlations with action and are labeled as “action units” (see Experimental Procedures). **B.** Of these, 31% can be definitively classified as action timing and 6% as action feedback (see text). The others are likely action timing units but cannot be definitively classified due to the low number of non-visually timed licks performed by trained animals (see text). **C.** Histogram showing the earliest moment at which units contain information about the action. **D.** 29 out of 77 action timing/feedback units show a significant difference in the action response for both eyes (orange) (see text). **E.** Only 2% units show significant correlations with action early in training, reflecting the false positive rate of our statistical test.

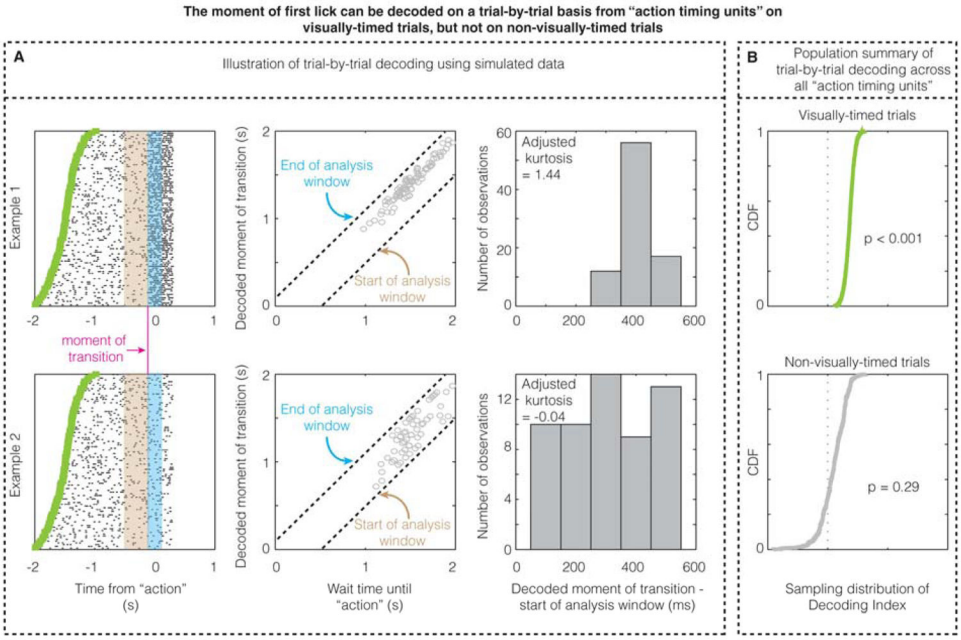

Figure 5. The moment of first lick can be decoded on a trial-by-trial basis from "action timing units" on visually-timed trials, but not on non-visually-timed trials

A. The three columns illustrate the single-trial decoding method using simulated data from two example units (shown in each row), one of which actually contains a transition in firing rate that predicts the moment of the action (100 ms prior to the action; "moment of transition"), while the other contains no transition in firing rate. The decoder calculates the maximum likelihood estimate of transition within an analysis window (Figure S5) for a given trial after being trained on the other trials to calculate the expected firing rates before (in the mauve bin) and after (in the cyan bin) transition (Figure S5). When there is a real transition from which the action time can be decoded, the decoded moments of transition across trials form a peaked, non-uniform distribution (example 1), whereas when there is no real transition from which the action time can be decoded (null hypothesis), the decoded moment of transition forms a uniform distribution within the analysis window. **B.** The CDF of the sampling distribution of Decoding Index (see Experimental Procedures and Figure S5) across all "action timing units" for visually-timed and non-visually-timed licks. These CDFs show that the action time can be decoded on a trial-by-trial basis on visually-timed trials, but not on non-visually-timed trials.

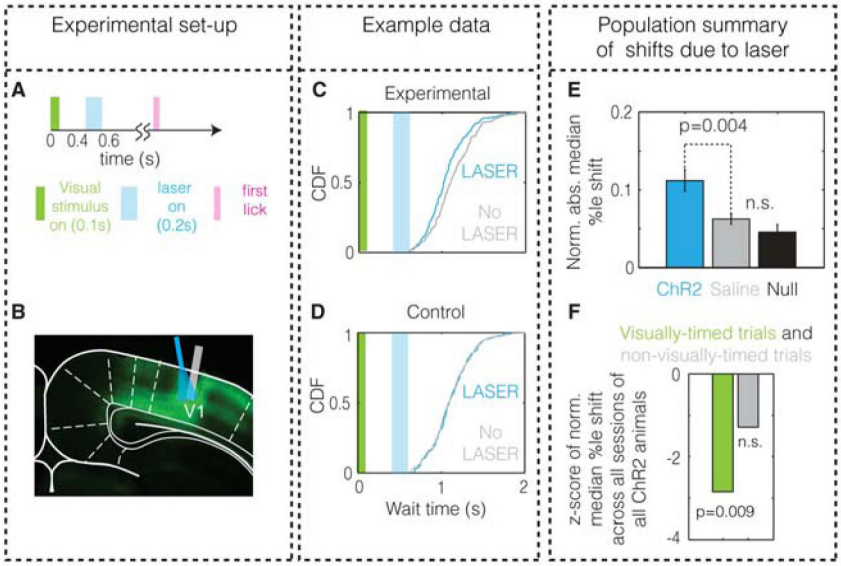

Figure 6. Optogenetic perturbations of V1 cause a significant shift in wait times

A. Schematic of a trial containing the laser presentation. **B.** Histology showing expression of ChR2 and the placement of optical fiber (cyan). Grey bar indicates the placement of electrodes in one animal used to verify the neural perturbation. **C & D.** Example wait time distributions on laser and no-laser trials from an animal infected with ChR2 (**C**) and another infected with saline (**D**). **E**. The quantified shift in wait time, the normalized absolute median percentile shift (this is *not* the ratio of shift to the median wait time, see Experimental Procedures for definition), for the experimental population shows significant differences from both control and chance levels (null). Error bars denote s.e.m. **F.** The pooled shift induced by the laser across all sessions of all experimental animals (infected with ChR2) is significant only on visually-timed trials (p=0.009, bootstrapping, two-tailed), but not on non-visually-timed trials (p=0.18, bootstrapping, two-tailed) (see Experimental Procedures).

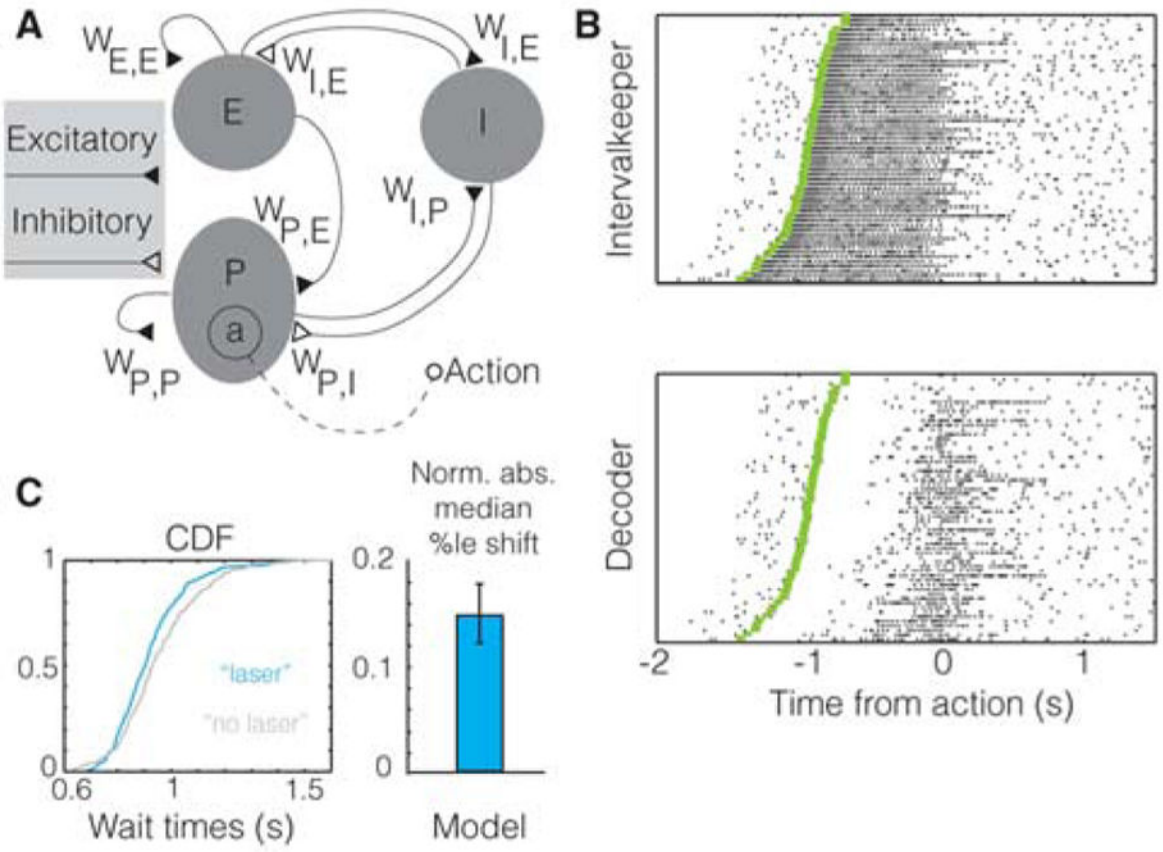

Figure 7. Spiking neuronal model

A. Schematic of the model showing the network architecture ("a" is the subpopulation driving the action; see Experimental Procedures). **B.** Raster plot showing significant trial-by-trial correlations with action, of example intervalkeeper (E population; $p<0.001$) and decoder (P population; $p=0.010$) units. Each trial began 500 ms prior to visual stimulation. **C.** The effect of the introduction of a "laser" perturbation in the model is shown using a CDF plot and is quantified using its normalized median percentile shift (see Experimental Procedures). The resultant significant shift in wait times ($p=0.039$, bootstrapping) is similar to the population data shown in Figure 6D. Error bar is the s.e.m, as obtained using bootstrapping. This result holds for a range of parameters in the model (Figure S7H).

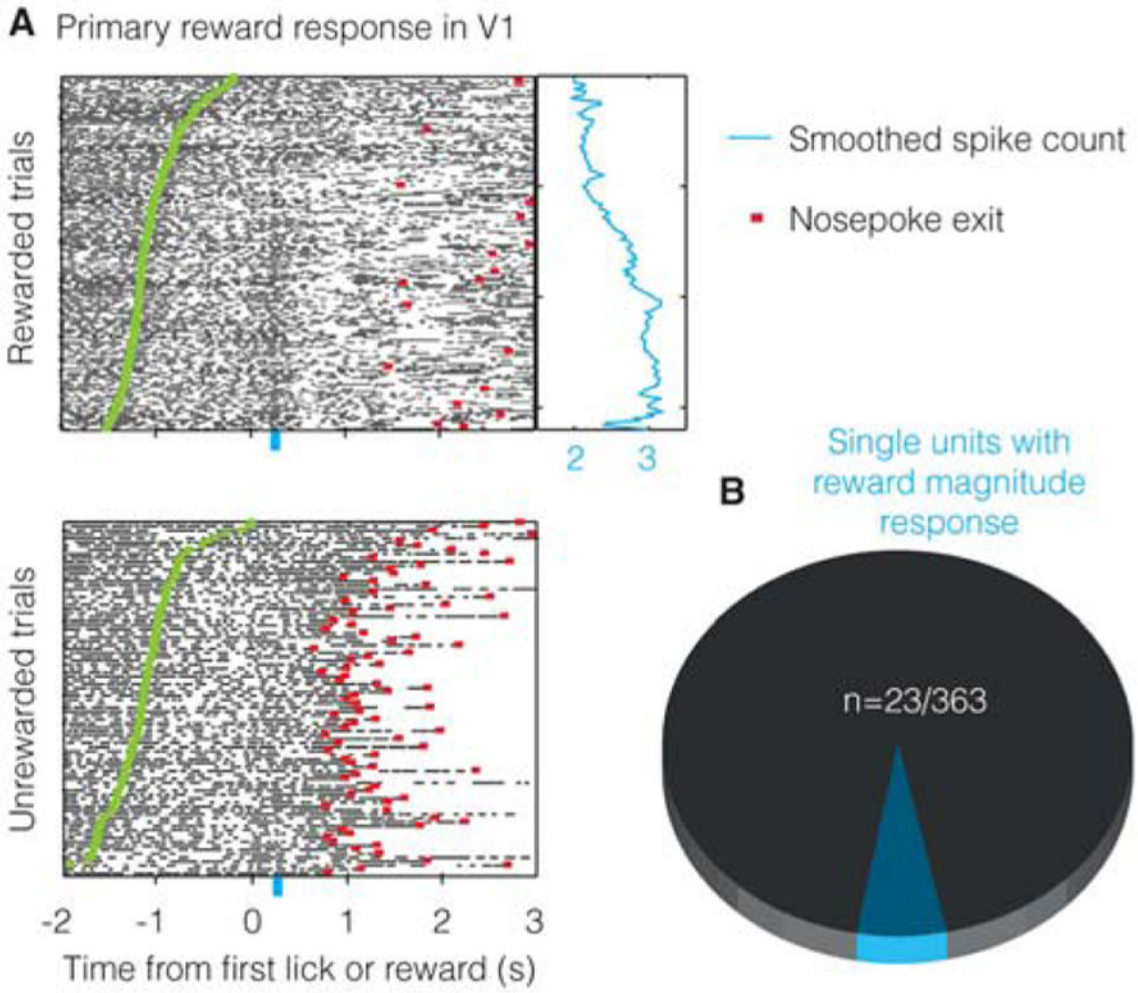

Figure 8. Reward magnitude response observed in V1

A. Example unit showing a response indicating the receipt of reward of a given magnitude (larger the wait time, larger the reward obtained). Blue curve shows smoothed spike counts (during 200–300ms following first lick and reward; shown by the blue bar) for individual trials. **B.** 23/363 units show reward magnitude coding.