



Visual Evidence Accumulation Guides Decision-Making in Unrestrained Mice

Onyekachi Odoemene,^{1,2}  Sashank Pisupati,^{1,2} Hien Nguyen,² and  Anne K. Churchland²

¹Watson School of Biological Sciences and ²Cold Spring Harbor Laboratory, Cold Spring Harbor, New York 11724

The ability to manipulate neural activity with precision is an asset in uncovering neural circuits for decision-making. Diverse tools for manipulating neurons are available for mice, but their feasibility remains unclear, especially when decisions require accumulating visual evidence. For example, whether mice' decisions reflect leaky accumulation is unknown, as are the relevant/irrelevant factors that influence decisions. Further, causal circuits for visual evidence accumulation are poorly understood. To address this, we measured decisions in mice judging the fluctuating rate of a flash sequence. An initial analysis (>500,000 trials, 29 male and female mice) demonstrated that information throughout the 1000 ms trial influenced choice, with early information most influential. This suggests that information persists in neural circuits for ~1000 ms with minimal accumulation leak. Next, in a subset of animals, we probed strategy more extensively and found that although animals were influenced by stimulus rate, they were unable to entirely suppress the influence of stimulus brightness. Finally, we identified anteromedial (AM) visual area via retinotopic mapping and optogenetically inhibited it using JAWS. Light activation biased choices in both injected and uninjected animals, demonstrating that light alone influences behavior. By varying stimulus–response contingency while holding stimulated hemisphere constant, we surmounted this obstacle to demonstrate that AM suppression biases decisions. By leveraging a large dataset to quantitatively characterize decision-making behavior, we establish mice as suitable for neural circuit manipulation studies. Further, by demonstrating that mice accumulate visual evidence, we demonstrate that this strategy for reducing uncertainty in decision-making is used by animals with diverse visual systems.

Key words: decision-making

Significance Statement

To connect behaviors to their underlying neural mechanism, a deep understanding of behavioral strategy is needed. This understanding is incomplete for mice. To surmount this, we measured the outcome of >500,000 decisions made by 29 mice trained to judge visual stimuli and performed behavioral/optogenetic manipulations in smaller subsets. Our analyses offer new insights into mice' decision-making strategies and compares them with those of other species. We then disrupted neural activity in a candidate neural structure and examined the effect on decisions. Our findings establish mice as suitable for visual accumulation of evidence decisions. Further, the results highlight similarities in decision-making strategies across very different species.

Introduction

Rodents have emerged as a powerful model organism for probing the neural circuits underlying decision-making (Carandini and

Churchland, 2013). Mice are an ideal model for studying neural circuits because of the tools for accessing and probing genetically defined cell types (Madisen et al., 2010, 2012, 2015; Taniguchi et al., 2011). Despite these advantages, other species are more commonly used in perceptual decision-making studies that involve temporal accumulation of sensory evidence, perhaps due to the assumption that such tasks are too difficult for mice. However, mice have been trained on numerous sensory perception tasks (Andermann et al., 2010; Busse et al., 2011; Sanders and Kepecs, 2012; Glickfeld et al., 2013; Z.V. Guo et al., 2014; Funamizu et al., 2016; Goard et al., 2016; Marbach and Zador, 2016; Burgess et al., 2017; Song et al., 2017). This suggests that they might be suitable for visual evidence accumulation tasks, and several studies report promising performance in mice on such tasks (Douglas et al., 2006; Morcos and Harvey, 2016; Stirman et al., 2016).

Received Dec. 5, 2017; revised Sept. 18, 2018; accepted Sept. 22, 2018.

Author contributions: O.O. and A.K.C. designed research; O.O. and H.N. performed research; O.O. and S.P. analyzed data; O.O., S.P., and A.K.C. wrote the paper.

This work was supported by NIH NRSA F31 (1F31-EY025164) to O.O., The Simons Collaboration on the Global Brain, The Klingenstein-Simons Foundation, and the Pew Charitable Trust to A.K.C. We thank Simon Musall for help with building wide-field epifluorescence microscope and providing the software for analysis, Rob Eifert and Barry Burbach for technical support, and Matthew Kaufman, Ashley Juavinett, and Lindsey Glickfeld for helpful discussions and feedback on the paper.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Anne Churchland, Cold Spring Harbor Laboratory, 1 Bungtown Road, Cold Spring Harbor, NY 11724. E-mail: churchland@cshl.edu.

<https://doi.org/10.1523/JNEUROSCI.3478-17.2018>

Copyright © 2018 the authors 0270-6474/18/3810143-13\$15.00/0

Two major gaps in our understanding of evidence accumulation decisions are apparent: a precise characterization of the time course of accumulation, and an understanding of how relevant and irrelevant factors jointly shape decisions. Recent work suggested that a sequence of visual cues collectively influence an eventual decision (Morcos and Harvey, 2016). These data constitute an essential first step in establishing mice as a suitable decision-making model, but outstanding questions remain. First, a large-scale study with many subjects and large trial counts is needed. This benefits a number of analyses, including those that evaluate the influence on the final decision of stimuli arriving at different times. Such analyses require many trials to achieve the level of precision that is required to fully characterize the evidence accumulation time course. Further, many animals are required to distinguish idiosyncratic strategies from the overall tendency of the species. For instance, in existing work (Morcos and Harvey, 2016), the use of five mice hinted that the most common strategy is to weight early evidence over late evidence, but the small animal number and the variability in strategy made a firm conclusion difficult. Closing these gaps in our understanding of accumulation decisions is essential, especially for interpreting causal manipulations (Krakauer et al., 2017).

Indeed, the causal circuits for visual evidence accumulation are not established in mice although inactivations have demonstrated a role for cortical and subcortical structures in other kinds of mouse decisions. For instance, suppressing the posterior parietal cortex (PPC) in mice impairs memory-guided decisions (Harvey et al., 2012; Funamizu et al., 2016; Goard et al., 2016). These studies suggest a role for cortical circuits in decision-making and visually guided behavior. However, the role of these structures in visual evidence accumulation is unknown. The importance of establishing a causal role for a putative neural structure is underscored by recent results that even areas strongly modulated during behavior might not be part of the causal circuit (Erlich et al., 2015; Katz et al., 2016).

Overall, mice have potential as an animal model for decision-making, but have been held back because of lack of detailed knowledge about behavior and little insight into the contribution of cortical circuitry. Here, we begin to close those gaps. First, we report that mice accumulate evidence for visual decisions. Our large dataset and use of stochastic stimuli allow precise characterization of how information presented at different times influences decisions. Next, we report that mice' decisions are jointly shaped by stimulus rate, brightness and the outcome of previous trials. Two independent experiments demonstrated that the influence of brightness was larger in rats compared with mice. Finally, we report that suppressing activity in the anteromedial visual area (AM) biases decisions, highlighting the role of cortex in evidence accumulation. Together, these findings establish mice as a suitable organism for visual accumulation-of-evidence decisions. Further, the results highlight similarities in decision-making strategies across very different species.

Materials and Methods

Animal subjects. The Cold Spring Harbor Laboratory Animal Care and Use Committee approved all animal procedures and experiments and all experimental procedures were in accordance with the National Institutes of Health's *Guide for the Care and Use of Laboratory Animals*. Experiments were conducted with female or male mice between the ages of 6–25 weeks. All mouse strains were of C57BL/6J background and purchased from The Jackson Laboratory. Ten GCaMP6f transgenic mice (Ai93/Emx1-cre/CaMKII α -tTA) of both sexes were used for retinotopic mapping and area AM photoinhibition experiments. Two male Long-Evans rats (6 weeks; Taconic) were also used for behavior experiments to

facilitate a comparison to other studies (Raposo et al., 2012; Brunton et al., 2013; Sheppard et al., 2013).

Behavioral training. Before behavioral training, mice were gradually water restricted over the course of 1 week. Mice were weighed daily and checked for signs of dehydration throughout training period (Z. Guo et al., 2014). Mice that weighed <80% of their original pretraining weight were supplemented with additional water. Behavioral training sessions lasted 1–2 h during which mice typically harvested at least 1 ml of water. Mice rested on the weekends. Mice who failed to harvest at least 0.4 ml on 2 consecutive days were supplemented with additional water.

Animal training took place in a sound isolation chamber containing a three-port setup described previously (Raposo et al., 2012). Mice poked their snouts into the center port to initiate trials and trigger sensory stimuli. Animals reported choices by moving to a left- or right-side port. In the first training stage, mice learned to wait for at least 1100 ms at the center port before reporting their decision. We shaped the behavior by rewarding the mice at the center port (0.5 μ l) and gradually increasing the minimum wait duration from 25 to 1100 ms over the course of 1–2 behavioral sessions. Without center reward, this stage typically took 10–12 sessions to learn.

During the first stage, mice were not rewarded for making the correct association between the stimulus and response port; rather, on each trial, a random port (left or right) was chosen as the reward port and a liquid reward (2–4 μ l) was delivered to the port. Trials in which the mouse waited the minimum required duration at the center port are referred to as completed trials.

In the second stage of training, mice learned to associate high-rate flash sequences (13–20 flashes/s) with the right port and low-rate flashes (1–11 flashes/s) with the left port. Trials with 12 flashes/s were randomly rewarded. For some mice, the contingency was reversed, such that high-rate flashes were rewarded at the left-hand port and low-rate flashes were rewarded at the right-hand port. Mice received a liquid reward for correct responses. They were punished for incorrect responses or for withdrawing early with a time-out period (2–4 s), during which they could not initiate a trial.

We used several anti-bias methods to correct the side bias, which often occurred when mice began stage two. Anti-bias strategies included: physically obstructing access to the biased port, changing the reward size, and modifying the proportion of left versus right trials.

Training was considered complete when mice waited at least 1100 ms at the center port, performed >80% correct on the easiest flash rates (Fig. 1B,C), and experienced at least eight or more flash rates. This required ~2 months, with one daily session 5 d per week.

Stimulus presentation, reward delivery, and data collection were performed through a MATLAB interface and Arduino-powered device (BPod, Sanworks).

Visual stimuli. Stimuli were sequences of 20 ms pulses of light from a LED panel (Ala Scientific). The interpulse intervals were randomly selected from a discrete exponential distribution (Brunton et al., 2013). This distribution was selected because it offers a major advantage over previous methods (Raposo et al., 2012; Sheppard et al., 2013), which is the independence of stimuli in each time bin. This greatly simplifies the process of computing psychophysical kernels. For the exponential interval stimulus, the minimum interpulse interval was 20 ms, and the number of flashes for a given stimulus determined the maximum interval. 4 to 20 flashes/s were presented on each trial, always over the course of 1000 ms. The stimulus was created using 25 fixed time bins each 20 ms in duration. A coin flip determined whether an event (flash) would occur in each bin. An empty 20 ms time bin followed each flash. This 1000 ms period was followed by a 100 ms delay, leading to a total time in the port of 1100 ms. The highest stimulus rates (20 Hz) have been shown to elicit reliable, steady-state flicker responses in retinal ERG (Krishna et al., 2002; Shirato et al., 2008; Tanimoto et al., 2015). Mice were presented with all rates from 4 to 20 Hz on each session in most cases. Some behavioral manipulations, such as the brightness manipulation and optogenetic stimulation, were especially challenging for the animals and tended to lower their overall reward rate. For those manipulations, we therefore omitted some of the more difficult stimuli to keep overall reward rate at a level that maintained the animals' motivation to work.

Each 20 ms flash pulse was generated by a half-wave rectified sinusoidal signal thresholded at the peaks and with a base frequency of 200 Hz.

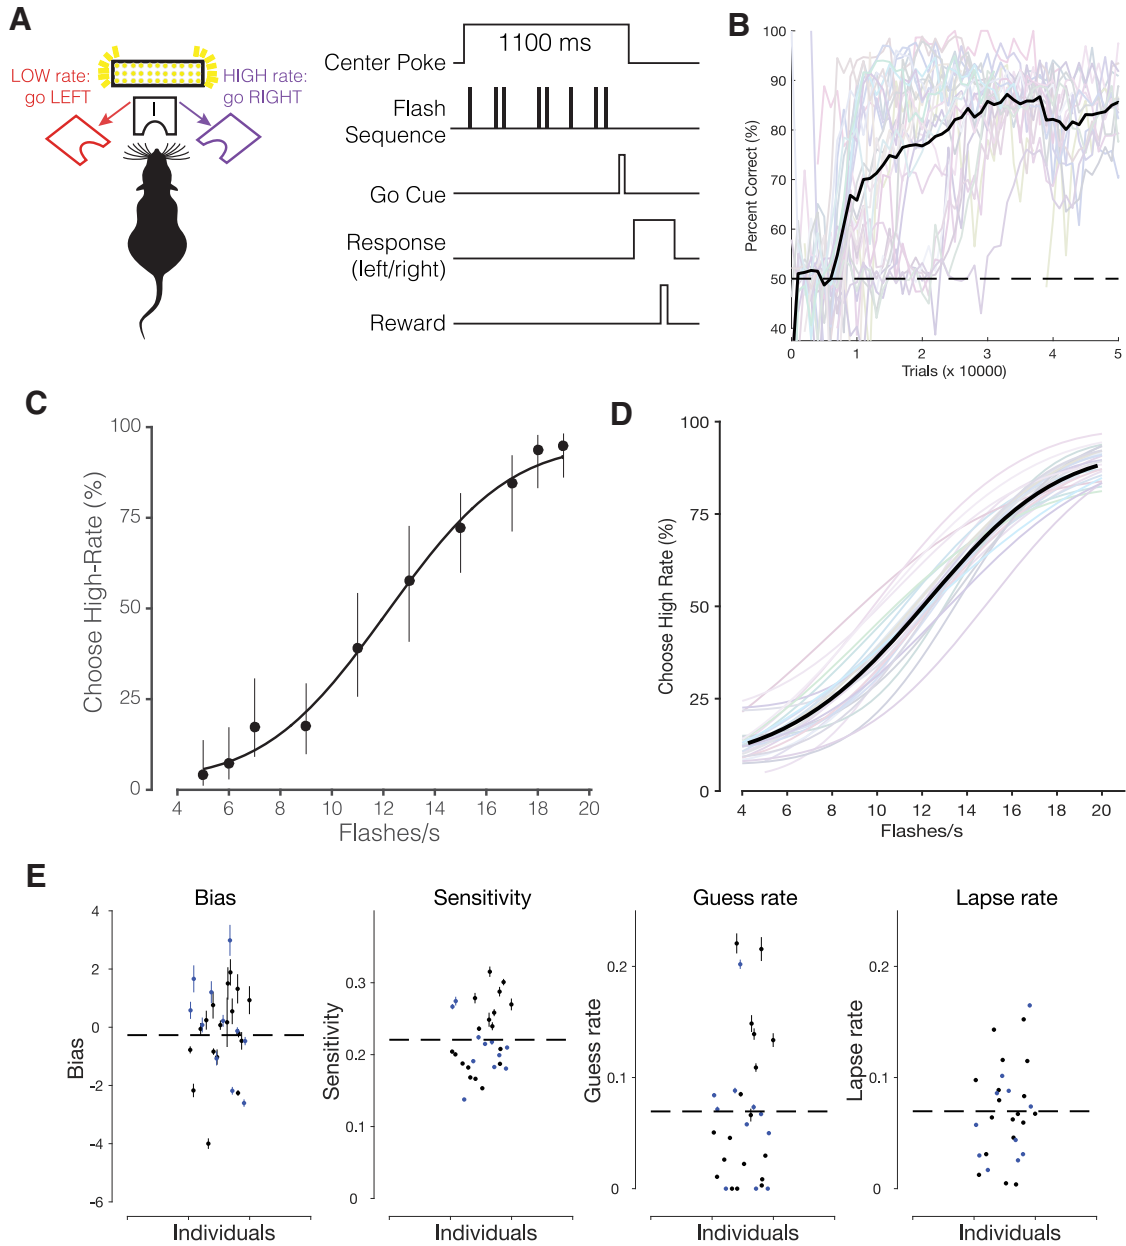


Figure 1. A dataset of over half a million trials demonstrates mice can be trained to make stable and reliable decisions about visual stimuli. **A**, Task schematic and trial structure of the three-port choice task. The mouse initiated trials and stimulus delivery by poking its nose into the center port. Mice reported whether stimuli were low-rate (left port) or high-rate (right port). Mice waited at the center port for at least 1100 ms, with the stimulus delivered after a variable delay (10–100 ms). At the end of the 1000 ms stimulus period, an auditory “Go” tone was played. Correct choices to the left or right were rewarded with a small drop of water (2 μ l), incorrect choices were followed by a 2–3 s timeout. **B**, Percentage correct on easiest stimulus conditions (4 and 20 flashes/s) plotted across total trials experienced by the mouse. Individual mice: gray traces and average: black trace, 29 mice. Colors arbitrarily selected to facilitate distinguishing subjects. **C**, Psychometric function fit for individual mouse from single session (494 trials; Error bars indicate Wilson binomial confidence intervals), and **(D)** data from 29 mice averaged across multiple sessions (537,288 trials). Individual mice, Colored traces; average, black trace. **E**, Parameter values from the model fit for each mouse in the cohort. Error bars indicate SEM calculated via bootstrapping (Palamedes Toolbox). Blue points indicate the 11 animals that were used for additional manipulations (11/11 were used for optogenetic or control experiments and 10/11 were used for the brightness manipulation).

This approach effectively controls the total LED on-time or the “density” of the 20 ms pulses. It is similar to pulse-width modulation technique used to control LED brightness. During normal sessions, the base frequency is multiplied by a brightness factor, which is kept constant across sessions.

Brightness manipulation. For the brightness manipulation experiments, the 20 ms flash duration was held constant, so that the subjects could not use the flash duration as a cue for the correct stimulus category. In the uniform brightness manipulation experiment, the normal brightness factor was either halved or doubled to produce the “dimmer” and “brighter” conditions. In the uncorrelated brightness manipulation experiment, we varied the LED on-time within the flash duration such that

the brightness factor was inversely scaled with the flash rate. Because the lowest number of flashes presented was 4 flashes/s, and we did not change the flash duration, we normalized all flash sequences such that the total LED on-time was equal to 4 flashes/s. All brightness manipulations were randomly introduced on 5% of all trials.

Head bar implantation and skull preparation. For retinotopic mapping experiments, mice were implanted with a custom titanium head bar. Mice were anesthetized with isoflurane (2%) mixed with oxygen and secured onto a stereotaxic apparatus. Body temperature was maintained at 37°C with a rectal temperature probe. The eyes were lubricated with Puralube ointment before the start of the surgery, followed by subcutaneous injection of analgesia (Meloxicam, 2 mg/kg) and antibiotic (Enro-

floxacin, 2 mg/kg). Fur on the scalp was removed with hair clippers and Nair (Sensitive Formula with Green Tea), followed by betadine (5%) swab. Lidocaine (100 μ l) was injected underneath the scalp before removing the scalp. The skull was cleaned with saline and allowed to dry. A generous amount of Vetbond tissue glue (3M) was then applied to seal the skull. Once the Vetbond was dry, the head bar was secured with Metabond (Parkwell) and dental acrylic. Postsurgical analgesia was applied. Mice were allowed 3 d to recover in their home cage before beginning the habituation process for retinotopic mapping.

Retinotopic mapping. To map visual area AM and surrounding cortical visual areas, we performed wide-field fluorescence retinotopic imaging in awake head-restrained mice (Garrett et al., 2014; Juavinett et al., 2017; Zhuang et al., 2017).

Mice were presented with a visual stimulus on one of two 27 inch monitors (Asus Design MX279H) rotated 90°. The monitor was placed in the visual hemifield contralateral to the imaging hemisphere and positioned at an angle of 77° (azimuth) from the midline of the mouse and at a distance of 15 cm from the eye. The mouse was positioned along the midline of the elevation axis. Retinotopic maps were generated by periodically sweeping a bar across the four cardinal directions of the monitor (Kalatsky and Stryker, 2003). The narrow bar (10°) contained a flickering black-and-white checkerboard pattern. The checkerboard squares alternated between black and white at 5 Hz. A trial repetition consisted of 11 sweeps of the bar during a 22 s period in one of the four cardinal directions. Each trial was repeated 15 times for each direction. The first cycle was discarded because it introduced stimulus onset transients. We used custom MATLAB and Psychtoolbox scripts for visual stimulus presentation and camera image acquisition during retinotopic mapping. During retinotopic mapping, mice were free to run on a custom 3D-printed wheel (4 inch diameter; Shapeways) sprayed with rubber coating (PlasticDip). Inclusion of a running wheel was selected because it tends to reduce stress. Running movements are known to modulate neural responses to visual stimuli, but likely not in a way that would systematically distort the borders between visual areas. Mice were habituated to head restraint on the wheel for at least four sessions across 2 d (15–30 min long) before retinotopic mapping sessions.

Wide-field imaging microscope. We used a custom wide-field epifluorescence microscope to acquire fluorescence images of the exposed skull during retinotopic mapping. Images were acquired at 20 frames/s with an sCMOS camera (pco.edge 5.5). A green emission filter (525/50 nm bandpass; Edmund Optics, 86-963) was positioned in front of the camera. Our microscope setup consisted of two macro lenses coupled face-to-face in a tandem lens configuration (Ratzlaff and Grinvald, 1991). The focal length of the first lens after the camera was 105 mm (B&H, NI1052DAF), the second lens had 85 mm (B&H, RO8514S), resulting in a magnification of 1.24 \times . Between both lenses was a dichroic mirror (Chroma T495LPXR), to couple blue light from a side-mounted LED (Thorlabs, M470L3) into the excitation light path. The light from the LED was passed through an excitation filter (470/40 nm bandpass, Chroma ET470/40 \times). Before retinotopic mapping, a focused image of the surface vasculature from the skull was acquired. The focal plane was subsequently moved \sim 400 μ m (relative to the pial surface) into cortex during retinotopic mapping, likely targeting responses in \sim layer 4. To increase optical clarity, a thin layer of mineral oil (Sigma-Aldrich) was applied as needed to the surface of the skull.

Wide-field imaging analysis. All imaging analysis was done in MATLAB (2014b, MathWorks). Analysis of retinotopic maps followed methods described in Zhuang et al. (2017). Briefly, the acquired image sequences are averaged across trials ($n = 15$) for each of the four stimulus directions. The first cycle in each trial was discarded to avoid stimulus onset transients. The fast Fourier transform was computed along the time axis of the average image sequence for each direction. To cancel the delay from the stimulus to response, the Fourier transforms corresponding to image sequences acquired from the same stimuli traveling in opposite directions (e.g., vertical bar traversing nasal-to-temporal and temporal-to-nasal) were divided, which is equivalent to pixel-wise subtraction of the image sequences in the time domain. To generate the azimuth and elevation phase maps, we used the MATLAB function *angle* to extract the phase of the Fourier transform at each pixel. The phase angles (in radi-

ans) were converted to degrees and normalized by the dimensions (in degrees) of the stimulus monitor. These angles determine the colors on azimuth and elevation maps. The azimuth and elevation phase maps were combined to compute visual field sign maps (Sereno et al., 1994, 1995; Garrett et al., 2014) as described in detail by Garrett et al. (2014). Briefly, we used the MATLAB function *gradient* to calculate the local gradients for the azimuth and elevation maps individually. To obtain the visual field sign map, we computed the sine of the difference of the angles between the local gradients of both maps. The magnitude of this quantity determines the color on visual field sign maps. Visual field sign maps were generated for each mouse individually. Visual area borders were determined from the visual field sign maps and drawn manually (Wekselblatt et al., 2016).

Viral injection and implant of optical fibers. For JAWS inhibition experiments, mice were injected with AAV8-CaMKII-JAWS-KGC-GFP-ER2 (UNC Vector Core) into area AM identified by retinotopic mapping. AM is a prominent candidate for decision-making as it appears to at least partially overlap with the previously defined location of posterior parietal cortex (Funamizu et al., 2016; Krumin et al., 2017). Indeed, recent work suggests that AM may be at the more posterior end of PPC (defined at 2 mm posterior and 1.7 mm lateral to bregma (Driscoll et al., 2017, their Fig. S1)). AM has projections to frontal and motor areas. Similar projection patterns have been observed in primate lateral intraparietal area LIP (Cavada and Goldman-Rakic, 1989a,b), an area routinely implicated in perceptual decision-making studies (Gold and Shadlen, 2007; Hanks and Summerfield, 2017).

To target AM, the visual field sign map was overlaid on the surface vasculature image captured before retinotopic imaging. Blood vessel landmarks located on or near area AM were used to guide the location of a small craniotomy (<1 mm diameter) and subsequent virus injection and optical fiber implantation. Viral injections were performed using Drummond Nanoject III, which enables automated delivery of small volumes of virus. To minimize virus spread, the Nanoject was programmed to inject slowly: six 30 nl boluses, 60 s apart, and each bolus delivered at 10nl/s. Approximately 180 nl of virus was injected at multiple depths (200 and 500 μ m) below the brain surface. Following the virus injection, 200 μ m fiber (metal ferrule, Thorlabs) was implanted above the injection site. The optical fiber was secured onto the skull with Vitrebond, Metabond, and dental acrylic. The animals were allowed at least 3 d to recover before behavioral training.

Optogenetic inactivation. A red 640 nm fiber-coupled laser (OptoEngine) was used for inactivation. Experiments were conducted with multiple laser power levels: 0.5, 1, and 2 mW (16, 32, and 64 mW/mm²). One power level was used per session. On inactivation sessions, laser light was externally triggered using a PulsePal (Sanworks) device. The laser stimulation pattern was a square pulse (1 s) followed by a linear ramp (0.25 s), which began at the offset of the stimulus. Stimulation occurred on 25% of trials during the weekday sessions. Animals were not run on weekends; a systematic difference on performance before and after the weekend was not apparent.

Psychometric function. We fitted a four-parameter psychometric function to the responses of subjects that performed the visual flashes categorization task. The general form of the psychometric function defines the probability (p_H) that the subject chooses the port associated with high flash rate as follows:

$$p_H = \gamma + (1 - \gamma - \lambda) F(x; \alpha, \beta), \quad (1)$$

where γ and λ are the lower and upper asymptote of the psychometric function, which parameterizes the guess rate and lapse rate, respectively; F is a sigmoidal function, in our case a cumulative normal; x is the event rate i.e., the number of flashes presented during the 1 s stimulus period; α parameterizes the horizontal shift or bias of the psychometric function, and β describes the slope or sensitivity. The psychometric function $F(x; \alpha, \beta)$ for a cumulative normal distribution is defined as follows:

$$F(x; \alpha, \beta) = \frac{\beta}{\sqrt{2\pi}} \int_{-\infty}^x \exp\left(-\frac{\beta^2(x - \alpha)^2}{2}\right). \quad (2)$$

The parameters of the psychometric function were estimated with the Palamedes Toolbox (Prins and Kingdom, 2018).

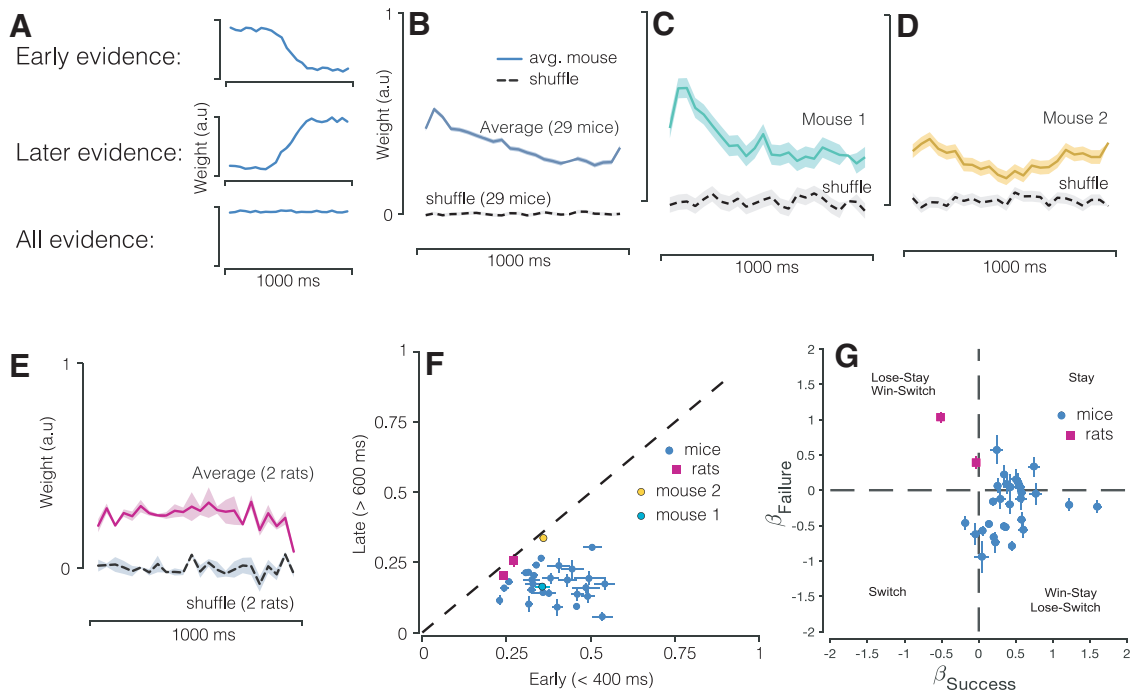


Figure 2. Decisions in mice and rats reflect evidence presented throughout the trial. **A**, Schematics indicating possible shapes of psychophysical kernels and the strategies they reveal. **B**, Psychophysical kernel based on pooled data from 29 mice (537,288 trials). **C, D**, Psychophysical kernels from example mice that typify early weighting and flat weighting. **E**, Psychophysical kernels from two rats (26,890 trials). Gray traces, Individual subjects; black trace, average. Values were >0 throughout the trial for almost all subjects, demonstrating that stimuli presented throughout the 1000 ms duration influence the animal’s eventual choice. **F**, Scatter plot relating the weight of evidence early versus late in the trial. Error bars indicate SEM. Dashed line: $x = y$. **G**, Effect of previous decision outcome (success/failure) on current choice.

Choice history. We implemented a probabilistic choice history model to evaluate the influence of prior choice(s) on the current choice of the subject. This approach assessed whether success or failure on the most recent trial influenced the performance on the current trial (Busse et al., 2011):

$$\ln\left(\frac{p_H}{1 - p_H}\right) = \beta_0 + \beta_E E(t) + \beta_S I_{\text{success}}(t - 1) + \beta_F I_{\text{failure}}(t - 1), \quad (3)$$

where t indicates the current trial and E is the signed stimulus evidence of the current trial. Evidence is computed as the difference between the flash rate of the trial and the category boundary (12 flashes/s). I_{success} was assigned a value of 1 or -1 to reflect a success on the right or left on the previous trial. I_{failure} indicated failures on the left or right. The coefficients (β_0 , β_E , β_S , and β_F) were estimated with MATLAB *glmfit*. The magnitude of evidence coefficient (β_E) can be interpreted in absolute terms in inverse units of stimulus magnitude, using signal detection theory (DeCarlo, 1998); a stimulus weight of 0.5 means that the noise in the animal’s internal estimate of the stimulus rate has a SD of ± 2 Hz. The magnitude of the success and failure coefficients (β_S and β_F) can be understood in relative terms to the stimulus weight; for instance, if the stimulus weight is 0.5 and the previous reward weight is 0.25, then a unit change in stimulus rate will affect the odds of going right twice as strongly as a previous reward. The inclusion of additional parameters and interaction terms did not reliably improve the fit (assessed by *stepwiseglm.m*). For example, interaction terms to evaluate whether the previous trial’s level of difficulty affected its impact on subsequent trials were rarely significant and had a much smaller influence on the choice compared with the current stimulus or previous outcome.

Logistic regression reverse correlation. The logistic regression function was used to estimate the weights associated with each moment of the stimulus (the psychophysical kernel; Katz et al., 2016; Licata et al., 2017), noting that these weights can reflect both sensory and decision-making processes (Okazawa et al., 2018). This can be written as follows:

$$\ln\left(\frac{p_H}{1 - p_H}\right) = \beta_0 + x^T w, \quad (4)$$

where β_0 is a scalar bias term, x is a vector of the 25 successive time windows over the trial, and w is a vector of the weights for each time window. β_0 and weight vector w were estimated with the MATLAB function *glmfit*. To generate the average kernels across mice (Fig. 2B) and rats (Fig. 2E), we pooled data from all animals and generated a single psychophysical kernel using all the data. Pooling together multiple generative rates introduces an artefactual correlation between time bins. To remove this correlation, we ran the regression with a reparametrized version of Equation 4, which subtracted out the mean generative rate λ_k from each bin and added it as a separate regressor, scaled by the number of time bins T and its corresponding weight μ_w (Eq. 5):

$$\begin{aligned} \ln\left(\frac{p_H}{1 - p_H}\right) &= \beta_0 + \sum_{t=1}^T w_t x_t \\ &= \beta_0(1) + \sum_{t=1}^T w_t (x'_t + \lambda_k) \begin{cases} \lambda_k = \mathbb{E}_t[x_t | k] \\ x'_t = x_t - \lambda_k \\ \Rightarrow \sum_t x'_t = 0, \text{corr}(x'_t, x'_t) = 0 \end{cases} \\ &= \beta_0(1) + \sum_{t=1}^T (w'_t + \mu_w)(x'_t + \lambda_k) \begin{cases} \mu_w = \mathbb{E}_t[w_t] \\ w'_t = w_t - \mu_w \\ \Rightarrow \sum_t w'_t = 0 \end{cases} \\ &= \beta_0(1) + \sum_{t=1}^T w'_t x'_t + \mu_w \sum_{t=1}^T x'_t + \lambda_k \sum_{t=1}^T w'_t + \mu_w \lambda_k \sum_{t=1}^T 1 \end{aligned}$$

$$\ln\left(\frac{p_H}{1-p_H}\right) = \beta_0(1) + \mu_w(\lambda_k T) + \sum_{i=1}^T w'_i x'_i \quad (5)$$

Generalized linear mixed model. To statistically test whether there was a significant effect of photoinhibition of area AM on the population group level, we used a generalized linear mixed-model (GLMM). GLMMs are an extension of the generalized linear model, which can be used to model both fixed and random effects in categorical data. In psychophysics, GLMMs can be used to generalize results across multiple subjects and experimental conditions (Knoblauch and Maloney, 2012; Moscatelli et al., 2012; Erlich et al., 2015).

The GLMM model written in the Wilkinson notation:

$$r \sim 1 + evidence + opto + evidence : opto + (evidence | subject/opto). \quad (6)$$

Each term of the equation has a coefficient, β . The model specifies that the subject's response, r , is a function of the fixed effects: the *evidence*, defined as the difference between flash rate and the category boundary (its coefficient represents the slope of the psychometric function), the photoinhibition indicator variable *opto*, and the interaction between the *evidence* and *opto*. The interaction term *evidence:opto* evaluates whether photoinhibition alters the subject's sensitivity or the slope of the psychometric function. The model allows the four fixed effects parameters to vary for each individual subject (random effects). The model uses a probit linking function and was fit using a maximum likelihood procedure. The GLMM analysis was performed using the R package "lme4" as described by Erlich et al. (2015).

The effect of photoinhibition on the horizontal location of the psychometric function was quantified by the choice bias. The choice bias was defined as follows:

$$choice\ bias = \frac{\beta_{opto}}{\beta_{evidence} + \beta_{opto:evidence}}, \quad (7)$$

where β_{opto} , $\beta_{evidence}$, $\beta_{evidence:opto}$ are estimated coefficients from the GLMM equation above. The choice bias reflects the equivalent change in the stimulus that would recapitulate the observed effects of photoinhibition and is in units of flashes/second. Positive choice bias would indicate that photoinhibition caused the subject to be biased toward high-rate responses. Because the choice bias is computed from estimated parameters of the GLMM model, we computed the errors (95% confidence intervals) via error propagation.

Results

We trained mice to categorize a stochastic sequence of visual flashes (Fig. 1), similar to earlier studies with rats (Raposo et al., 2012; Brunton et al., 2013; Scott et al., 2015). Mice performed a three-port choice task (Uchida and Mainen, 2003), in which they judged whether the total number of full-field flashes presented during a 1000 ms period exceeded an experimenter-defined category boundary (12 flashes/s; Fig. 1A). Each flash in the sequence was 20 ms long and followed by an inter-flash interval drawn from a discrete exponential distribution.

Mice learned to categorize stochastic sequences of visual flashes

Mice performed hundreds of trials per session (median 767 trials) and reached high performance accuracy at the easiest level of the task (Fig. 1B). Behavioral performance was quantified by fitting a psychometric function (Eq. 1; Fig. 1C–E). Individual mice on single sessions (Fig. 1C) and across multiple sessions (Fig. 1D) made increasingly more high-rate choices as the flash rate increased, achieving psychometric performance comparable to rats trained on a similar task (Raposo et al., 2012).

Mice decisions were influenced most by flashes early in the sequence

To maximize accuracy, animals should count all the flashes presented during the fixed stimulus presentation period. Because all flashes in the sequence are equally informative about the overall count, subjects should apply an equal weight to all flashes (Fig. 2A, bottom). However, mice might instead make use of only a portion of the stimulus. Greater weighting of flashes early in the sequence is consistent with an impulsive strategy, which, in its extreme, would amount to making up one's mind right at the beginning (Fig. 2A, top). Greater weighting of flashes later in the sequence can indicate a forgetful/leaky strategy or, in its extreme, "procrastination" in the sense that only stimuli at the very end influence the decision (Fig. 2A, middle; Kiani et al., 2008).

To distinguish these strategies, we used the well established logistic regression approach to estimate the psychophysical kernel (Huk and Shadlen, 2005; Katz et al., 2016; Yates et al., 2017). The logistic regression-based reverse correlation approach reveals the time course of how incoming stimuli, on average, influence the subject's choice. Our use of stimuli that appear stochastically over time and our large dataset together enabled a continuous and precise characterization of this time course. Across mice (Fig. 2B), the entire sequence of flashes was informative, as indicated by non-zero regression weights throughout the trial. Interestingly, flashes presented earlier in the sequence informed the choice more strongly than flashes presented later in the sequence. This implies that mice tended to more heavily weight stimuli presented early in the trial, consistent with an impulsive integration strategy. Our inclusion of 29 mice also allowed us to compare individuals and potentially identify different strategies. The psychophysical kernel of many individual mice reflected the average (Fig. 2C) although there were exceptional animals for whom stimuli throughout the trial influenced decisions more evenly (Fig. 2D). The strategy of these rare mice was reminiscent of the behavior that has been observed previously in rats (Brunton et al., 2013; Sheppard et al., 2013), a finding we replicate here in a small cohort of rats (Fig. 2E).

Considering all animals together (Fig. 2F) clarifies that most mice assign a greater weight to early evidence (Fig. 2F, most circles are below the line), whereas the tendency to weight early and late evidence more evenly was rare in mice (Fig. 2F, yellow circle). Note that few animals ignored late evidence altogether. This is likely because we used a "go tone" at the end of the stimulus period. As in previous studies (Sheppard et al., 2013), inclusion of the go tone may eliminate the need for animals to estimate the duration of the stimulus period, preventing a tendency to sometimes misestimate it and thus stop accumulating evidence too soon.

Mice were influenced by performance on previous trial

Next, we used the same dataset to evaluate whether mice were influenced by performance on previous trials. Several studies have reported that human and animal subjects performing perceptual tasks are influenced by previous choices (Busse et al., 2011; Fründ et al., 2014; Scott et al., 2015; Abrahamyan et al., 2016; Hwang et al., 2017; Urai et al., 2017), even when the trials are independent. We used a quantitative model to assess whether the event-based, visual accumulation decisions used here were likewise influenced by choices made on previous trials.

This approach assessed whether success or failure on the most recent trial influenced the performance on the current trial (Materials and Methods; Busse et al., 2011). Figure 2G shows a scatter plot of the coefficients for previous success (β_s) and previous

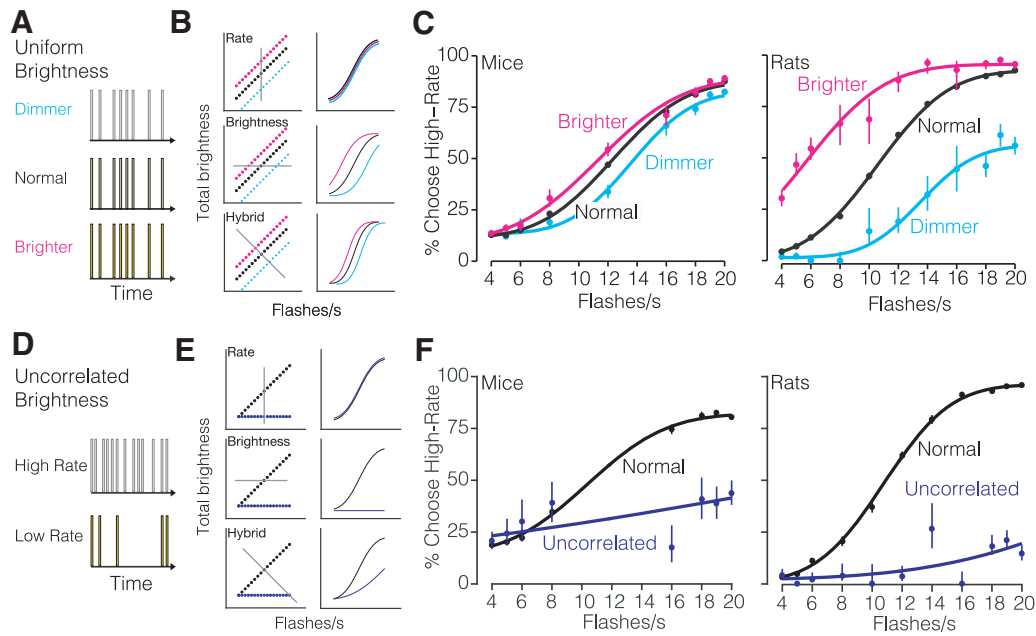


Figure 3. Stimulus brightness influences rate decisions. **A**, Schematic of the uniform brightness manipulation experiment. The intensity of individual flashes was varied such that all flashes were dimmer or brighter than normal on 5% of randomly selected trials. **B**, Left, Stimulus spaces and decision planes (gray lines). Right, Predicted psychometric functions. Each row reflects a candidate way in which the stimulus in each condition would influence decisions given the strategy indicated in the label. Top, Stimulus rate. Middle, Stimulus brightness. Bottom, Hybrid strategy in which both features are used. **C**, Measured psychometric functions. Left, Eight mice: 108,547 trials. Right, Two rats: 26,201 trials. Points, Subjects' responses; solid line, four-parameter cumulative normal psychometric function fit to the data. Error bars indicate Wilson binomial 95% confidence intervals. **D**, Schematic of uncorrelated brightness manipulation experiment. The intensity of individual flashes was scaled inversely with the flash rate on 5% of randomly selected trials. All sequences have the same cumulative brightness, independent of flash rate. **E**, Same as **B** but for the manipulation in **D**. **F**, Measured psychometric functions. Left, Two mice: 6326 trials. Right, Two rats: 9946 trials. Points, Subjects' responses; solid line, four-parameter cumulative normal psychometric function fit to the data. Error bars indicate Wilson binomial 95% confidence intervals.

failures (β_F). Nearly all the 29 mice had positive β_S coefficients, indicating that mice tended to repeat the same choice on the current trial if they were rewarded on the previous trial. Many of the mice also had positive β_F coefficients, meaning that they mice tended to repeat their choice following a failure (Fig. 2G, Stay quadrant), whereas others had negative β_F coefficients, indicating a tendency to switch choices following a failure (Fig. 2G, Win-Stay, Lose Switch quadrant). The overall observed trial history patterns were similar to that observed in human subjects performing a perceptual decision-making task (Abrahamyan et al., 2016). To assess the necessity of including history parameters in characterizing behavior, we compared a model with and without the choice history parameters using a likelihood ratio test. We quantified the magnitude of model improvement using BIC and AIC and tested whether the improvement was significant above and beyond the expected improvement in fit from adding new parameters (χ^2 test). The improvement was significant ($p < 10^{-5}$) for all 29 mice and for the two rats.

Decisions are influenced by cumulative brightness

To more deeply probe the animal's strategy we conducted additional experiments in smaller subsets of animals. Specifically, we wished to evaluate an alternate strategy to accumulating sensory events: to base the decision on the overall brightness experienced over the course of the stimulus. This is a feasible strategy given that the flash event rate is directly proportional to the total LED on-time and therefore the total photons emitted in a sequence.

To test whether mice were influenced by brightness, we performed two brightness manipulation experiments on a subset of animals (Fig. 3). First, the intensity of all flashes in a given sequence was randomly increased or decreased on 5% of all trials

(Fig. 3A). If subjects are influenced by rate alone, their decisions will be the same regardless of brightness (Fig. 3B, top row). By contrast if subjects are influenced by brightness, they will report more high-rate choices on brighter trials, and more low-rate choices on dimmer trials (Fig. 3B, middle and bottom rows). This is what we observed (Fig. 3C, left): brighter stimuli drove a high-rate bias (shift of 1.0 ± 0.5 flashes/s), while dimmer stimuli drove a low-rate bias (shift of 0.8 ± 0.5 flashes/s). When we tested rats on the same manipulation, the changes were even larger (Fig. 3C, right): we observed a high-rate bias of 4.8 ± 0.9 flashes/s for brighter stimuli and a low-rate bias of 3.1 ± 0.8 flashes/s for dimmer stimuli.

Second, we removed the correlation between brightness and flash rate by adjusting the flash intensity in each sequence to the flash rate on 5% of trials. As a result, the total brightness over time was the same across all flash rates (Fig. 3D). If subjects are influenced by rate alone, their decisions will be unaffected by brightness (Fig. 3E, top row). By contrast if subjects are influenced by brightness, they will have a low-rate bias on uncorrelated trials because the brightness level used was that of the lowest rate stimulus (Fig. 3E, middle and bottom rows). This is what we observed: both mice and rats had a low-rate bias (shift of 3 ± 8 flashes/s for mice, 21 ± 7 for rats; Fig. 3F). Importantly, the dependence of decisions on stimulus rate was reduced but still present (Fig. 3F, blue lines not completely flat). For example, the sensitivity parameter for mice was 0.053, indicating that the animals' estimates of rate were very noisy (e.g., the SD on their estimate was ± 19 Hz; see Materials and Methods). However, this parameter was nonetheless highly significant ($p = 4 \times 10^{-10}$, likelihood ratio test, corrected for parameter on boundary), indicating that rate did have a significant impact on choice. Similarly, rats' sensitivity was 0.077; small but highly significant ($p = 6 \times 10^{-15}$). Together,

these findings argue in favor of a hybrid strategy (Fig. 3E, bottom row). If animals had used brightness alone, the psychometric functions would have been flat.

Inactivation of secondary visual area AM biases perceptual decisions

To test whether the evidence accumulation paradigm engaged cortical circuitry, we sought to reversibly silence secondary visual area AM in mice during decision-making. To target AM, we performed wide-field retinotopic mapping in each mouse (Fig. 4). Briefly, we imaged visually evoked activity in awake transgenic mice expressing GCaMP6 in excitatory neurons (Ai93; CamkII α -tTA; Emx-cre) in response to a vertical or horizontal bar that periodically drifted across the screen in the four cardinal directions (Kalatsky and Stryker, 2003; Garrett et al., 2014; Zhuang et al., 2017). This procedure enabled generation of phase maps for altitude and azimuth visual space (Fig. 4A,B) and subsequently visual field sign maps (Fig. 4C), which were used to estimate the borders between cortical visual areas and reliably identify AM (Fig. 4D). The maps we observed were similar to those previously reported using either calcium imaging or intrinsic signal imaging (Garrett et al., 2014; Zhuang et al., 2017). Note that the location of area AM is very close to the stereotaxic coordinates used to target PPC (Harvey et al., 2012; Goard et al., 2016; Hwang et al., 2017). Future studies are needed to establish whether AM is a separate cortical area from PPC or whether there is overlap (partial or complete) between the two.

To reversibly silence AM, we used cruxhalorhodopsin JAWS (Halo57), a red light-driven chloride ion pump capable of powerful optogenetic inhibition (Chuong et al., 2014; Acker et al., 2016). Optogenetic stimulation was randomly interleaved on 25% of trials within a session. The optogenetic stimulus pattern consisted of a 1 s long square wave followed by a 0.25 s long linear downward ramp to reduce the effect of rebound excitation that may occur after strong inhibition (Fig. 5A, red line; Chuong et al., 2014; Guo et al., 2014a).

A potential confound when using JAWS for optogenetic inhibition is that the presence of red light alone may influence behavior. While it was long assumed that rodents are unable to perceive red light, a recent study showed that red-light delivery in the brain can activate the retina and influence behavior (Danskin et al., 2015). Having demonstrated here that decisions can be influenced by brightness (Fig. 3), we feared that red-light stimulation might induce a perceived increase in brightness and thus a behavioral bias toward high-rate decisions on photoinhibition trials. To test whether decisions were affected by the presence of red light in the absence of JAWS, we implanted and trained mice injected with a sham virus (AAV-GFP) in AM.

In vivo red light stimulation of sham-injected mice resulted in a high-rate bias (Fig. 5B). This confirms the hypothesis that the red light alone influences decisions. The bias is likely because the red lights increased the perceived brightness of the stimulus, driving the animal to make more high-rate choices. Importantly,

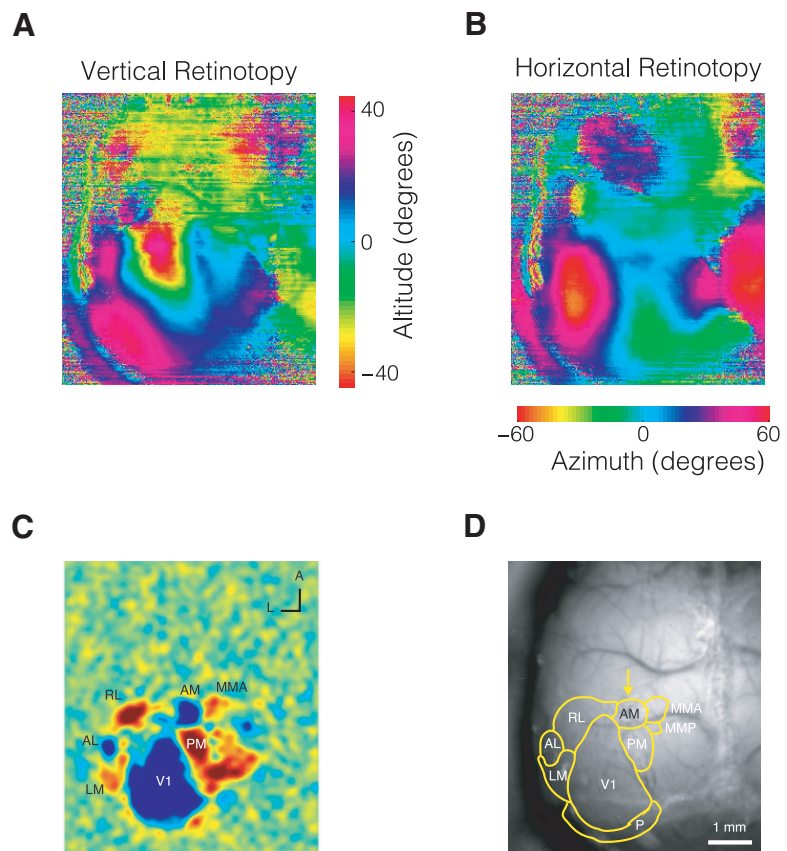


Figure 4. Retinotopic Mapping allows precise localization of visual areas for subsequent manipulation. **A**, Altitude and **B** azimuth phase maps. **C**, Visual field sign map with labeled visual areas. **D**, Visual area borders overlaid on photograph of skull.

choices were biased away from the site of the implant, arguing against phosphenes that drew the animal's attention toward the stimulation side. To counter the red-light bias, we installed additional red lights in the behavior booth to adapt long-wavelength-sensitive photoreceptors (Danskin et al., 2015). These external "house lights" strongly reduced the effect of the laser stimulation on behavior (Fig. 5C,D).

Although the house lights reduced the red-light bias in uninjected animals, the presence of a residual red-light bias in any individual injected animal is difficult to ascertain. If the house lights were incompletely effective in masking the red light, the bias could diminish or possibly enhance the effects of direct neural manipulation, depending on how the red-light bias and neural manipulation interact. To surmount this problem, we developed an experimental design in which the stimulus–response contingency was varied, whereas the stimulated hemisphere was held constant. Specifically, we trained two groups of mice on opposing behavioral contingencies: Group A was trained on the contingency: High-Rate, go LEFT; Group B was trained on the reverse: High-Rate, go RIGHT (Fig. 6A). Both groups were implanted on the left.

One scenario would be difficult to interpret: specifically, JAWS suppression could bias the animal's estimate of a sensory parameter: rate or perceived brightness for example. If so, the (e.g., high) rate bias from the neural manipulation and the high-rate bias from the red light would be combined similarly in both groups of animals, simply changing the magnitude of the effect in both groups (Fig. 6B, left, solid and dashed lines are similar in both top and bottom). Distinguishing the effect of the neural manipulation from that of the red light would be difficult.




Figure 5. Long wavelength laser stimulation biases decisions in control animals. **A**, Experimental configuration. Mice were injected with AAV-GFP and implanted with a fiber in right hemisphere area AM. **B**, Psychometric function without masking red light (2 mice; 2011 Laser-off trials, 610 Laser-on trials). Irradiance was 32 mW/mm². **C**, Psychometric performance with masking red light (2 mice) with easiest flash rate conditions (2699 Laser-off trials, 823 Laser-on trials). **D**, Same as **C** but for sessions including multiple flash rates (2866 Laser-off trials, 903 Laser-on trials). Irradiance was 64 mW/mm².

Fortunately, in an alternate scenario, JAWS suppression could drive a bias to the ipsilateral side that is independent of the rate associated with that side. In this scenario, the effects will differ for the two groups because the ipsilateral side is associated with high-rate for Group A and low-rate for Group B (Fig. 6B, top right). The difference across groups will be present even if a red-light bias persists (Fig. 6B, solid and dashed lines differ in both top and bottom). This because the red-light bias will increase high rate choices and thus simply shift both curves leftward, leaving the difference between them unchanged. For Group B animals, a contralateral (high rate) bias from the red light, combined with an ipsilateral (low rate) bias from the JAWS suppression may potentially cancel each other out (Fig. 6B, bottom right), leading to an interesting scenario in which the effect of JAWS suppression is only apparent in one group. A comparison of the two groups is therefore essential in interpreting the dual effects of the red-light bias and the JAWS suppression.

A comparison of effects in Group A and Group B revealed a striking difference. Specifically, Group A (High-Rate, go LEFT) had a large, high-rate bias [Fig. 6C; choice bias = 10.09 flashes/s CI: (2.84 18.72)] accompanied by a change in the slope of the psychometric functions [$\beta_{\text{evidence:opto}} = -0.045$ CI: (-0.064 -0.026) $p = 3.4e-06$, GLMM Test]. This cohort also had a low-rate bias on control trials (Fig. 6C, right, black points above red points). Biases on control trials in the opposite direction as the stimulation-induced bias have been reported previously (Salzman et al., 1992; Carello and Krauzlis, 2004; Fetsch et al., 2014). The bias on control trials suggests that animals cannot detect stimulation on individual trials, but can detect the lower likelihood of reward on the biased side across all trials, and adap-

tively shift their subjective category boundary to maximize this likelihood. If the animal only experienced stimulation trials (or could detect stimulation on individual trials), one would expect a complete compensatory shift, eventually leading to unbiased performance even in the presence of stimulation. However, because the animal only experiences stimulation on 25% of trials, the expected effect on the likelihood is small, leading to a larger bias on stimulation trials than control trials, as we observed.

By contrast, Group B (High-Rate, go RIGHT) showed no effect [Fig. 6D; $\beta_{\text{evidence:opto}} = 0.008$ CI: (-0.006 0.022), $p = 0.27$; choice bias = 0.43 flashes/s CI: (-0.75 1.62)]. The group difference was present despite matched laser power, injection volume, and percentage of disruption trials. These data most resemble Figure 6B, bottom right, in which the two effects drive the psychometric function in opposing directions: the red light drives a high-rate bias, and the JAWS suppression drives an ipsilateral bias, largely canceling each other out. In keeping with the hypothesis that AM suppression mainly drives a change in bias, and only a modest change in sensitivity, the psychophysical kernels were quite similar in both groups of animals (Fig. 6E,F) although some interesting subtle differences are apparent.

Discussion

We describe a quantitative behavioral paradigm for studying visual evidence accumulation decisions in freely behaving mice. Mice trained on our paradigm performed hundreds of trials per session and maintained stable performance across sessions. A dataset of over half a million trials allowed us to characterize the time course of accumulation with precision. Mice were influenced by visual evidence presented throughout the trial, but they

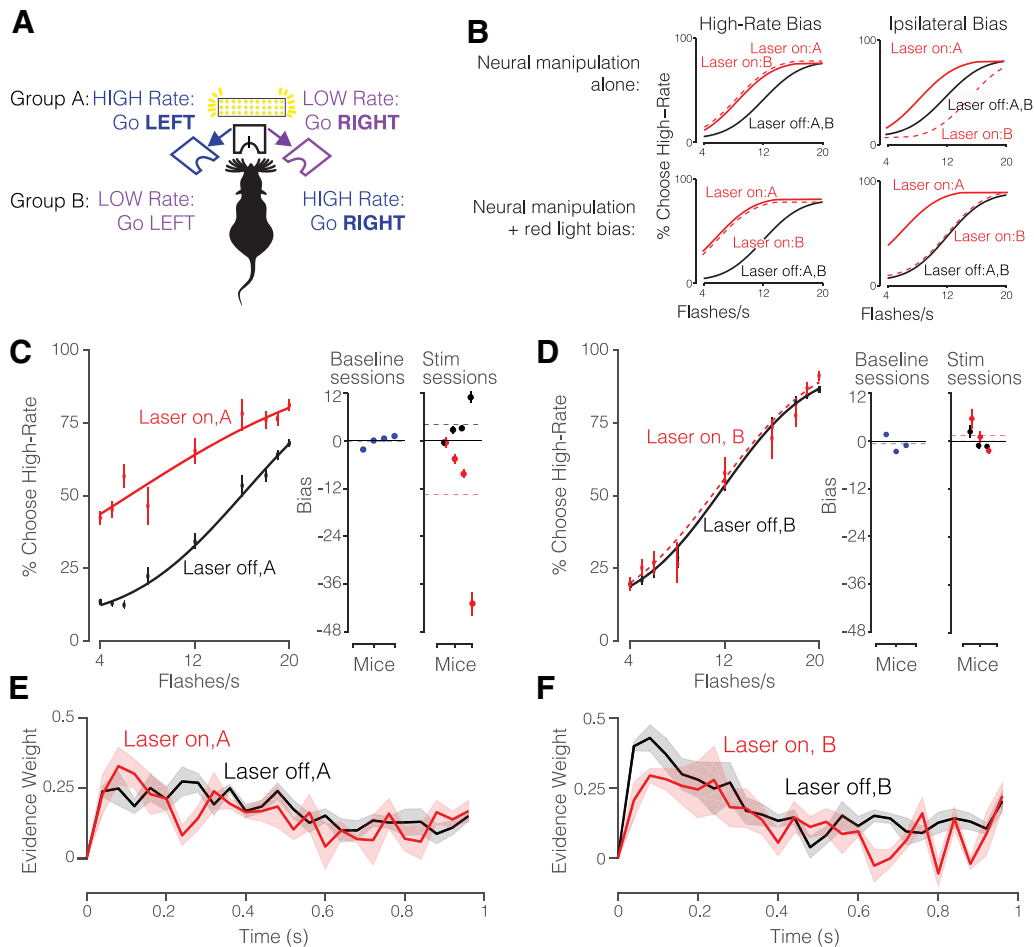


Figure 6. JAWS Photoinhibition of visual area AM. **A**, Schematic of experimental configuration of AM photoinhibition experiment. Group A mice trained on the contingency: High-Rate, go LEFT and Group B mice trained on the reverse contingency: High-Rate, go RIGHT. Both groups of mice were injected with JAWS virus and implanted with an optical fiber on the left hemisphere. Photoinhibition occurred on 25% of trials during the stimulus period. **B**, Predicted behavioral outcomes of AM photoinhibition. Top, Predictions for neural manipulation alone. Left, If AM inhibition drives a high-rate bias, the outcomes would be similar for the two groups. Right, If AM inhibition drives a bias toward the ipsilateral side, Groups A and B would show biases in opposite directions because the high-rate side differs for Groups A and B. Bottom, Predictions for neural manipulation alongside a bias driven by the presence of visible light (as in Fig. 5). Left, If AM inhibition drives a high-rate bias, both groups would again exhibit the same bias. Right, If AM inhibition drives an ipsilateral bias, Groups A and B would again show biases in opposite directions; potentially with a very weak effect for Group B because the red light and neural manipulations are in opposition. **C**, Left, Decisions for Group A (left, 4 mice; $n = 5722$ Laser OFF trials and $n = 1958$ Laser ON trials) and at laser power irradiance of 64 mW/mm^2 . Circles represent the subject's behavioral response during laser OFF (black) and laser ON (red) trials. Solid line represents the psychometric function fit to cumulative normal. Error bars represent Wilson binomial (95%) confidence intervals. Right, Bias values for individual animals during baseline sessions (blue points) as well as inactivation sessions which included stimulation (red) and control (black) trials. Black solid lines indicate unbiased performance; dashed lines indicate the mean bias across subjects for the corresponding color. **D**, Same as **C** but for Group B (right, 3 mice; $n = 4404$ Laser OFF trials and $n = 1381$ Laser ON trials). **E**, Psychophysical kernels for Group A (37,025 Laser off trials; 12,952 Laser on trials). **F**, Psychophysical kernels for Group B (32,936 Laser off trials; 12,423 Laser on trials). Shading indicates SEM.

assigned more weight on average to flashes presented earlier in the sequence, similar to monkeys and unlike rats. Further, despite overall high accuracy, decisions were nonetheless influenced by additional information, such as stimulus brightness and previous reward and choice history. Finally, we demonstrated that area AM plays a causal role in visual decisions. Our experimental design was key in allowing this conclusion because control experiments demonstrated that the red stimulation light biases mice even in the absence of JAWS. Together, these results fill a gap in our understanding of accumulation of evidence behavior in mice and begin to define the causal circuit that supports this behavior.

Our observations of early weighting of information (Fig. 2) are consistent with results from the evidence accumulation paradigm from Morcos and Harvey (2016) in head-fixed mice. Our large cohort of animals, high trial count and use of stochastic stimulus arrival times allowed us to more fully characterize this function and extend it to a new paradigm. Interestingly, the

shapes of the psychophysical kernels we observed in the mice are qualitatively similar to those observed in nonhuman primates (Katz et al., 2016; Yates et al., 2017). The shape of the kernel differed from that observed in rats trained on the same task (Fig. 2E) and previously reported by other evidence accumulation paradigms (Raposo et al., 2012; Brunton et al., 2013; Sheppard et al., 2013; Scott et al., 2015). The difference across species is intriguing because it suggests that although different species achieve comparable levels of performance, their internal behavioral strategies may differ. This underscores the importance of using stochastic stimuli, which make it possible to uncover the animal's strategy (Churchland and Kiani, 2016).

Our brightness manipulations revealed that decisions were not based solely on rate. For both mice and rats, the cumulative brightness of the flash sequence also influenced decisions (Fig. 3C,F). Incorporating brightness in decision-making reflects a clever strategy because in almost all trials, brightness and rate

provide evidence in favor of the same decision. In fact, combining these two sources of information is the optimal strategy (Fig. 3*B,E*, bottom row), in the same way that combining auditory and visual information is the optimal strategy in multisensory experiments (Raposo et al., 2012; Sheppard et al., 2013). It would be surprising if animals elected to marginalize brightness when it is such a useful source of information.

The influence of brightness that we observed here contrasts findings from a recent study, which reported that rats performing a visual evidence accumulation task counted individual flashes rather than cumulative flash on-time (Scott et al., 2015). Two experimental design features may explain the difference in results. First, the inter-flash interval in Scott et al. was more than an order of magnitude longer than that used here (250 vs 20 ms) which may discourage a cumulative brightness strategy. Although the highest stimulus rates used here are below the flicker fusion rate (Krishna et al., 2002; Shirato et al., 2008; Tanimoto et al., 2015), it is possible that the rates were nonetheless sufficiently high to make a brightness strategy more appealing than it would be at lower rates. Second, Scott et al. (2015) did not manipulate brightness on catch trials, as we did here, but instead conducted separate experiments in which all trials contained jittered light on-times. This raises the interesting possibility that in Scott et al. (2015), the jittered on-times increased the uncertainty of cumulative on-time information. The increased uncertainty would change the optimal strategy, leading animals to weight cumulative on-time less than they would in other experiments in which brightness and event number are correlated.

To understand the neural mechanisms that enable perceptual decision-making, we tested the causal role of secondary visual area AM. Our results, an ipsilateral bias on inactivation trials, suggest that AM normally drives contralateral choices. This is consistent with anatomical projections of AM to motor areas (Wang et al., 2012; Allen Brain Atlas,) and the recently proposed role for mouse parietal cortex in navigation (Krumin et al., 2017). The results are also reminiscent of some perturbation studies of rat PPC, which may have some overlap with area AM. When neural activity in rat PPC was disrupted using an activating perturbation on a very similar task, an ipsilateral bias was observed in some sessions. However, the bias was a less consistent effect than a loss of sensitivity (Licata et al., 2017), which we observed in the current study to a small extent (Fig. 6*C*). It is unclear whether effects seen in bias versus sensitivity in mice versus rats reflect a species difference, or rather that inactivation of this region, while consistently demonstrating a role in visually guided behavior, has variable effects (Raposo et al., 2014; Erlich et al., 2015; Licata et al., 2017). This variability may arise partly from the fact that rat PPC is identified via stereotaxic coordinates (Reep and Corwin, 2009), a method that cannot take into account individual differences in neuroanatomy. The ability to more precisely localize areas using retinotopic mapping, as we have done here, is a strength of conducting these studies in mice. Future studies that target visual areas using this approach may define with more precision the contribution of regions in this area to the sensory processing and action-planning aspects of decision-making.

Nevertheless, additional studies are needed to more definitively establish the role of AM in mouse visual decision-making. Because we observed that red light alone biases choices (Fig. 5*B*), we inferred that the red light and the optogenetic suppression may, in some configurations, bias decisions in opposite directions and cancel each other out (Fig. 6*C*, right). The use of non-optical suppression methods such as muscimol (Raposo et al.,

2014; Erlich et al., 2015) or the use of other wavelengths of light to suppress activity (Lien and Scanziani, 2013) could provide additional evidence about the role of area AM in decision-making. The latter is appealing in part because shorter wavelengths activate more restricted regions of cortex compared with the red light used here, making it possible to target a region with more precision.

Our results make clear the need to carefully control for light-induced artifacts, both by adapting the animal and by experimental design that disentangles light-induced artifacts from true behavioral changes. The artifact we observed is most likely caused by red light propagating from the stimulation site through the brain and directly activating the retina. Danskin et al. (2015) measured retinal activation during *in vivo* red-light stimulation and found the largest activation ipsilateral to the implanted stimulation fiber. By contrast, the bias that we observed here was contralateral to the implanted stimulation fiber. Notably, in this cohort, the contralateral side was associated with high-rate choices (Fig. 5*A*). A likely explanation is that light from the fiber increased overall brightness; as we demonstrated in separate experiments (Fig. 3), increased brightness can be interpreted as evidence for high rate choices. Additional experiments that systematically vary the stimulus response contingency in sham-injected animals could confirm this hypothesis.

The proposed ipsilateral bias caused by AM photoinhibition is consistent with spatial hemineglect observed in visual parietal lesions. Spatial hemineglect, also referred to as contralateral neglect, is a phenomenon that occurs when subjects ignore the contralateral hemifield as a result of lesion to the parietal cortex. Although hemineglect has been reported in humans (Stone et al., 1991; Kerkhoff, 2001) and rats (Crowne et al., 1986; Reep and Corwin, 2009), we could not find a report on mice. The presence of hemispacial neglect would suggest that the mice are neglecting the tendency to go toward the affected (contralateral) visual hemifield. A related interpretation of the ipsilateral bias due to suppression of AM activity is that neurons in AM are active in advance of contralateral choices. In this scenario, the two hemispheres of AM would represent competing movement intentions (Andersen and Buneo, 2002), such that inactivation of one hemisphere leads to movement in the opposing direction. This prediction could be verified in future studies in by temporally restricting the inactivation to different epochs within the trial.

This behavior is suitable for the freely moving mice, a setup that can be preferable to head restraint because it more closely approximates the animal's natural state (Juavinett et al., 2018). Options for measuring neural activity in freely moving animals are proliferating, and include electrophysiology (Raposo et al., 2014; Nikbakht et al., 2018), fiber photometry (Adelsberger et al., 2005; Gunaydin et al., 2014), or an imaging miniscope (Jung et al., 2004; Cai et al., 2016). Further, characterizing the behavior in freely moving mice, as we have done here, provides a baseline to which head-fixed studies (using standard two-photon imaging or wide-field imaging) can be compared.

Together, these observations begin to address two major gaps in our understanding of accumulation of evidence behavior: we have precisely characterized the time course of evidence accumulation, and have uncovered that rate, brightness, and trial history jointly shape decisions. Finally, our observations provide reassurance that mice are free from sensory and cognitive limitations that would preclude their ability to accumulate visual evidence to guide decisions.

References

- Abrahamyan A, Silva LL, Dakin SC, Carandini M, Gardner JL (2016) Adaptable history biases in human perceptual decisions. *Proc Natl Acad Sci U S A* 113:E3548–E3557. [CrossRef Medline](#)
- Acker L, Pino EN, Boyden ES, Desimone R (2016) FEF inactivation with improved optogenetic methods. *Proc Natl Acad Sci U S A* 113:E7297–E7306. [CrossRef Medline](#)
- Adelsberger H, Garaschuk O, Konnerth A (2005) Cortical calcium waves in resting newborn mice. *Nat Neurosci* 8:988–990. [CrossRef Medline](#)
- Andermann ML, Kerlin AM, Reid RC (2010) Chronic cellular imaging of mouse visual cortex during operant behavior and passive viewing. *Front Cell Neurosci* 4:3. [CrossRef Medline](#)
- Andersen RA, Buneo CA (2002) Intentional maps in posterior parietal cortex. *Annu Rev Neurosci* 25:189–220. [CrossRef Medline](#)
- Brunton BW, Botvinick MM, Brody CD (2013) Rats and humans can optimally accumulate evidence for decision-making. *Science* 340:95–98. [CrossRef Medline](#)
- Burgess CP, Lak A, Steinmetz NA, Zatzka-Haas P, Bai Reddy C, Jacobs EAK, Linden JF, Paton JJ, Ranson A, Schröder S, Soares S, Wells MJ, Wool LE, Harris KD, Carandini M (2017) High-yield methods for accurate two-alternative visual psychophysics in head-fixed mice. *Cell Rep* 20:2513–2524. [CrossRef Medline](#)
- Busse L, Ayaz A, Dhruv NT, Katzner S, Saleem AB, Schölvinck ML, Zaharia AD, Carandini M (2011) The detection of visual contrast in the behaving mouse. *J Neurosci* 31:11351–11361. [CrossRef Medline](#)
- Cai DJ, Aharoni D, Shuman T, Shobe J, Biane J, Song W, Wei B, Veshkini M, La-Vu M, Lou J, Flores SE, Kim I, Sano Y, Zhou M, Baumgaertel K, Lavi A, Kamata M, Tuszynski M, Mayford M, Golshani P, et al. (2016) A shared neural ensemble links distinct contextual memories encoded close in time. *Nature* 534:115–118. [CrossRef Medline](#)
- Carandini M, Churchland AK (2013) Probing perceptual decisions in rodents. *Nat Neurosci* 16:824–831. [CrossRef Medline](#)
- Carello CD, Krauzlis RJ (2004) Manipulating intent: evidence for a causal role of the superior colliculus in target selection. *Neuron* 43:575–583. [CrossRef Medline](#)
- Cavada C, Goldman-Rakic PS (1989a) Posterior parietal cortex in rhesus monkey: I. Parcellation of areas based on distinctive limbic and sensory corticocortical connections. *J Comp Neurol* 287:393–421. [CrossRef Medline](#)
- Cavada C, Goldman-Rakic PS (1989b) Posterior parietal cortex in rhesus monkey: II. Evidence for segregated corticocortical networks linking sensory and limbic areas with the frontal lobe. *J Comp Neurol* 287:422–445. [CrossRef Medline](#)
- Chuong AS, Miri ML, Busskamp V, Matthews G A C, Acker LC, Sørensen AT, Young A, Klapoetke NC, Henninger MA, Kodandaramaiah SB, Ogawa M, Ramanlal SB, Bandler RC, Allen BD, Forest CR, Chow BY, Han X, Lin Y, Tye KM, Roska B, Cardin JA, et al. (2014) Noninvasive optical inhibition with a red-shifted microbial rhodopsin. *Nat Neurosci* 17:1123–1129. [CrossRef Medline](#)
- Churchland AK, Kiani R (2016) Three challenges for connecting model to mechanism in decision-making. *Curr Opin Behav Sci* 11:74–80. [CrossRef Medline](#)
- Crowne DP, Richardson CM, Dawson KA (1986) Parietal and frontal eye field neglect in the rat. *Behav Brain Res* 22:227–231. [CrossRef Medline](#)
- Danskin B, Denman D, Valley M, Ollerenshaw D, Williams D, Groblewski P, Reid C, Olsen S, Blanche T, Waters J (2015) Optogenetics in mice performing a visual discrimination task: measurement and suppression of retinal activation and the resulting behavioral artifact. *PLoS One* 10:e0144760. [CrossRef Medline](#)
- DeCarlo LT (1998) Signal detection theory and generalized linear models. *Psychol Methods* 3:186–205. [CrossRef](#)
- Douglas RM, Neve A, Quittenbaum JP, Alam NM, Prusky GT (2006) Perception of visual motion coherence by rats and mice. *Vision Res* 46:2842–2847. [CrossRef Medline](#)
- Driscoll LN, Pettit NL, Minderer M, Chettih SN, Harvey CD (2017) Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell* 170:986–999.e16. [CrossRef Medline](#)
- Erlich JC, Brunton BW, Duan CA, Hanks TD, Brody CD (2015) Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *eLife* 4:25869470. [CrossRef Medline](#)
- Fetsch CR, Kiani R, Newsome WT, Shadlen MN (2014) Effects of cortical microstimulation on confidence in a perceptual decision. *Neuron* 83:797–804. [CrossRef Medline](#)
- Fründ I, Wichmann FA, Macke JH (2014) Quantifying the effect of intertrial dependence on perceptual decisions. *J Vis* 14(7):9 1–16. [CrossRef](#)
- Funamizu A, Kuhn B, Doya K (2016) Neural substrate of dynamic bayesian inference in the cerebral cortex. *Nat Neurosci* 19:1682–1689. [CrossRef Medline](#)
- Garrett ME, Nauhaus I, Marshel JH, Callaway EM (2014) Topography and areal organization of mouse visual cortex. *J Neurosci* 34:12587–12600. [CrossRef Medline](#)
- Glickfeld LL, Histed MH, Maunsell JH (2013) Mouse primary visual cortex is used to detect both orientation and contrast changes. *J Neurosci* 33:19416–19422. [CrossRef Medline](#)
- Goard MJ, Pho GN, Woodson J, Sur M (2016) Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. *eLife* 5:e13764. [CrossRef Medline](#)
- Gold JL, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574. [CrossRef Medline](#)
- Gunaydin LA, Grosenick L, Finkelstein JC, Kauvar IV, Fenno LE, Adhikari A, Lammel S, Mirzabekov JJ, Airan RD, Zalocusky KA, Tye KM, Anikeeva P, Malenka RC, Deisseroth K (2014) Natural neural projection dynamics underlying social behavior. *Cell* 157:1535–1551. [CrossRef Medline](#)
- Guo ZV, Hires SA, Li N, O'Connor DH, Komiyama T, Ophir E, Huber D, Bonardi C, Morandell K, Gutnisky D, Peron S, Xu NL, Cox J, Svoboda K (2014) Procedures for behavioral experiments in head-fixed mice. *PLoS One* 9:e88678. [CrossRef Medline](#)
- Guo Z, Li N, Huber D, Ophir E, Gutnisky D, Ting J, Feng G, Svoboda K (2014) Flow of cortical activity underlying a tactile decision in mice. *Neuron* 81:179–194. [CrossRef Medline](#)
- Hanks TD, Summerfield C (2017) Perceptual Decision Making in Rodents, Monkeys, and Humans. *Neuron* 93:15–31. [CrossRef Medline](#)
- Harvey CD, Coen P, Tank DW (2012) Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484:62–68. [CrossRef Medline](#)
- Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci* 25:10420–10436. [CrossRef Medline](#)
- Hwang EJ, Dahlen JE, Mukundan M, Komiyama T (2017) History-based action selection bias in posterior parietal cortex. *Nat Commun* 8:1242. [CrossRef Medline](#)
- Song YH, Kim JH, Jeong HW, Choi I, Jeong D, Kim K, Lee SH (2017) A neural circuit for auditory dominance over visual perception. *Neuron* 93:940–954. [CrossRef Medline](#)
- Juavinett AL, Nauhaus I, Garrett ME, Zhuang J, Callaway EM (2017) Automated identification of mouse visual areas with intrinsic signal imaging. *Nat Protoc* 12:32–43. [CrossRef Medline](#)
- Juavinett AL, Erlich JC, Churchland AK (2018) Decision-making behaviors: weighing ethology, complexity, and sensorimotor compatibility. *Curr Opin Neurobiol* 49:42–50. [CrossRef Medline](#)
- Jung JC, Mehta AD, Aksay E, Stepnoski R, Schnitzer MJ (2004) *In vivo* mammalian brain imaging using one- and two-photon fluorescence microendoscopy. *J Neurophysiol* 92:3121–3133. [CrossRef Medline](#)
- Kalatsky VA, Stryker MP (2003) New paradigm for optical imaging: temporally encoded maps of intrinsic signal. *Neuron* 38:529–545. [CrossRef Medline](#)
- Katz LN, Yates JL, Jonathan W, Pillow JW, Huk AC (2016) Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature* 535:285–288. [CrossRef Medline](#)
- Kerkhoff G (2001) Spatial hemineglect in humans. *Prog Neurobiol* 63:1–27. [CrossRef Medline](#)
- Kiani R, Hanks TD, Shadlen MN (2008) Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci* 28:3017–3029. [CrossRef Medline](#)
- Knoblauch K, Maloney LT (2012) Modeling psychophysical data in R. New York: Springer.
- Krakauer JW, Ghazanfar AA, Gomez-Marín A, MacIver MA, Poeppel D (2017) Neuroscience needs behavior: correcting a reductionist bias. *Neuron* 93:480–490. [CrossRef Medline](#)
- Krishna VR, Alexander KR, Peachey NS (2002) Temporal properties of the mouse cone electroretinogram. *J Neurophysiol* 87:42–48. [CrossRef Medline](#)

- Krumin M, Harris KD, Carandini M (2017) Decision and navigation in mouse parietal cortex. *bioRxiv* 166413. [CrossRef](#)
- Licata AM, Kaufman MT, Raposo D, Ryan MB, Sheppard JP, Churchland AK (2017) Posterior parietal cortex guides visual decisions in rats. *J Neurosci* 37:4954–4966. [CrossRef](#) [Medline](#)
- Lien AD, Scanziani M (2013) Tuned thalamic excitation is amplified by visual cortical circuits. *Nat Neurosci* 16:1315–1323. [CrossRef](#) [Medline](#)
- Madisen L, Zwingman TA, Sunkin SM, Oh SW, Zariwala H a, Gu H, Ng LL, Palmiter RD, Hawrylycz MJ, Jones AR, Lein ES, Zeng H (2010) A robust and high-throughput cre reporting and characterization system for the whole mouse brain. *Nat Neurosci* 13:133–140. [CrossRef](#) [Medline](#)
- Madisen L, Mao T, Koch H, Zhuo JM, Berenyi A, Fujisawa S, Hsu YW, Garcia AJ 3rd, Gu X, Zanella S, Kidney J, Gu H, Mao Y, Hooks BM, Boyden ES, Buzsáki G, Ramirez JM, Jones AR, Svoboda K, Han X, et al. (2012) A toolbox of cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nat Neurosci* 15:793–802. [CrossRef](#) [Medline](#)
- Madisen L, Garner AR, Shimaoka D, Chuong AS, Klapoetke NC, Li L, van der Bourg A, Niino Y, Ego L, Monetti C, Gu H, Mills M, Cheng A, Tasic B, Nguyen TN, Sunkin SM, Benucci A, Nagy A, Miyawaki A, Helmchen F, et al. (2015) Transgenic mice for intersectional targeting of neural sensors and effectors with high specificity and performance. *Neuron* 85:942–958. [CrossRef](#) [Medline](#)
- Marbach F, Zador AM (2016) A self-initiated two-alternative forced choice paradigm for head-fixed mice. *bioRxiv* 073783. [CrossRef](#)
- Morcos AS, Harvey CD (2016) History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat Neurosci* 19:1672–1681. [CrossRef](#) [Medline](#)
- Moscattelli A, Mezzetti M, Lacquaniti F (2012) Modeling psychophysical data at the population-level: the generalized linear mixed model. *J Vis* 12(11):261–17. [CrossRef](#) [Medline](#)
- Nikbakht N, Tafreshiha A, Zoccolan D, Diamond ME (2018) Supralinear and supramodal integration of visual and tactile signals in rats: psychophysics and neuronal mechanisms. *Neuron* 97:626–639.e8. [CrossRef](#) [Medline](#)
- Okazawa G, Sha L, Purcell BA, Kiani R (2018) Psychophysical reverse correlation reflects both sensory and decision-making processes. *Nat Commun* 9:3479. [CrossRef](#) [Medline](#)
- Prins N, Kingdom FAA (2018) Applying the model-comparison approach to test specific research hypotheses in psychophysical research using the Palamedes Toolbox. *Front Psychol* 9:1250. [CrossRef](#) [Medline](#)
- Raposo D, Sheppard JP, Schrater PR, Churchland AK (2012) Multisensory decision-making in rats and humans. *J Neurosci* 32:3726–3735. [CrossRef](#) [Medline](#)
- Raposo D, Kaufman MT, Churchland AK (2014) A category-free neural population supports evolving demands during decision-making. *Nat Neurosci* 17:1784–1792. [CrossRef](#) [Medline](#)
- Ratzlaff EH, Grinvald A (1991) A tandem-lens epifluorescence microscope: hundred-fold brightness advantage for wide-field imaging. *J Neurosci Methods* 36:127–137. [CrossRef](#) [Medline](#)
- Reep RL, Corwin JV (2009) Posterior parietal cortex as part of a neural network for directed attention in rats. *Neurobiol Learn Mem* 91:104–113. [CrossRef](#) [Medline](#)
- Salzman CD, Murasugi CM, Britten KH, Newsome WT (1992) Microstimulation in visual area MT: effects on direction discrimination performance. *J Neurosci* 12:2331–2355. [CrossRef](#) [Medline](#)
- Sanders JI, Kepecs A (2012) Choice ball: a response interface for two-choice psychometric discrimination in head-fixed mice. *J Neurophysiol* 108:3416–3423. [CrossRef](#) [Medline](#)
- Scott BB, Constantinople CM, Erlich JC, Tank DW, Brody CD (2015) Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. *eLife* 4:e11308. [CrossRef](#) [Medline](#)
- Sereno MI, McDonald CT, Allman JM (1994) Analysis of retinotopic maps in extrastriate cortex. *Cereb Cortex* 4:601–620. [CrossRef](#) [Medline](#)
- Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889–893. [CrossRef](#) [Medline](#)
- Sheppard JP, Raposo D, Churchland AK (2013) Dynamic weighting of multisensory stimuli shapes decision-making in rats and humans. *J Vis* 13(6):41–19. [CrossRef](#) [Medline](#)
- Shirato S, Maeda H, Miura G, Frishman LJ (2008) Postreceptor contributions to the light-adapted ERG of mice lacking b-waves. *Exp Eye Res* 86:914–928. [CrossRef](#) [Medline](#)
- Stirman JN, Townsend LB, Smith SL (2016) A touchscreen based global motion perception task for mice. *Vision Res* 127:74–83. [CrossRef](#) [Medline](#)
- Stone SP, Wilson B, Wroot A, Halligan PW, Lange LS, Marshall JC, Greenwood RJ (1991) The assessment of visuo-spatial neglect after acute stroke. *J Neurol Neurosurg Psychiatry* 54:345–350. [CrossRef](#) [Medline](#)
- Taniguchi H, He M, Wu P, Kim S, Paik R, Sugino K, Kvitsani D, Fu Y, Lu J, Lin Y, Shima Y, Fishell G, Nelson SB, Huang ZJ (2011) A resource of cre driver lines for genetic targeting of GABAergic neurons in cereb cortex. *Neuron* 71:995–1013. [CrossRef](#) [Medline](#)
- Tanimoto N, Sothilingam V, Kondo M, Biel M, Humphries P, Seeliger MW (2015) Electoretinographic assessment of rod- and cone-mediated bipolar cell pathways using flicker stimuli in mice. *Sci Rep* 5:10731. [CrossRef](#) [Medline](#)
- Uchida N, Mainen ZF (2003) Speed and accuracy of olfactory discrimination in the rat. *Nat Neurosci* 6:1224–1229. [CrossRef](#) [Medline](#)
- Urai AE, Braun A, Donner TH (2017) Uncertainty and alters serial choice bias. *Nat Commun* 8:14637. [CrossRef](#) [Medline](#)
- Wang Q, Sporns O, Burkhalter A (2012) Network analysis of corticocortical connections reveals ventral and dorsal processing streams in mouse visual cortex. *J Neurosci* 32:4386–4399. [CrossRef](#) [Medline](#)
- Wekselblatt JB, Flister ED, Piscopo DM, Niell CM (2016) Large-scale imaging of cortical dynamics during sensory perception and behavior. *J Neurophysiol* 115:2852–2866. [CrossRef](#) [Medline](#)
- Yates JL, Park IM, Katz LN, Pillow JW, Huk AC (2017) Functional dissection of signal and noise in MT and LIP during decision-making. *Nat Neurosci* 20:1285–1292. [CrossRef](#) [Medline](#)
- Zhuang J, Ng L, Williams D, Valley M, Li Y, Garrett M, Waters J (2017) An extended retinotopic map of mouse cortex. *eLife* 6:e18372. [CrossRef](#) [Medline](#)