Mice can learn a stimulus-invariant orientation discrimination task

Dmitry R Lyamzin[1,*], Ryo Aoki[1], Mohammad Abdolrahmani[1], Andrea Benucci[1,2,3]

[1]RIKEN Center for Brain Science, 2-1 Hirosawa, Wako-shi, Saitama 351-0198, Japan

[2]University of Tokyo, Graduate School of Information Science and Technology, Department of Mathematical Informatics, 1-1-1 Yayoi, Bunkyo City, Tokyo 113-0032, Japan

[3]Lead Contact

Correspondence: andrea.benucci@riken.jp (A.B.); dmitry.lyamzin@riken.jp (D.R.L.)

DRL ORCID: 0000-0002-7082-2902

RA ORCID: 0000-0002-1097-0006

AB ORCID: 0000-0002-7698-6289

Summary

Understanding how the brain computes choice from sensory information is a central question of perceptual decision-making. Relevant behavioral tasks condition choice on abstract or invariant properties of the stimuli, thus decoupling stimulus-specific information from the decision variable. Among visual tasks, orientation discrimination is a gold standard; however, it is not clear if a mouse – a recently popular animal model in visual decision-making research – can learn an invariant orientation discrimination task and what choice strategies it would use.

Here we show that mice can solve a discrimination task where choices are decoupled from the orientation of individual stimuli, depending instead on a measure of relative orientation. Mice learned this task, reaching an upper bound for discrimination acuity of 6 degrees and relying on decision-making strategies that balanced cognitive resources with history-dependent biases.

We analyzed behavioral data from n=40 animals with the help of a novel probabilistic choice model that we used to interpret individual biases and behavioral strategies. The model explained variation

in performance with task difficulty and identified unreported dimensions of variation associated with the circularity of the stimulus space. Furthermore, it showed a larger effect of history biases on animals' choices during periods of lower engagement.

Our results demonstrate that mice can learn invariant perceptual representations by combining decision-relevant stimulus information decoupled from low-level visual features, with the computation of the decision variable dependent on the cognitive state.

Keywords

Vision, decision-making, mouse, behavioral model, probabilistic model

Introduction

Most of behaviorally relevant information in visual scenes is given by the objects and relationships between them rather than the low-level visual features. Accordingly, animals and humans can form high-level neural representations invariant to specific peripheral sensory activations (DiCarlo et al., 2012). These invariant perceptual representations can correspond to relative properties of objects, such as spatial arrangement (Krechevsky, 1938; Lashley, 1938), shape and color similarity (Martinho and Kacelnik, 2016), relative contrast (Burgess et al., 2017), and relative density or numerosity (Dakin et al., 2011).

In decision-making research, tasks that rely on such invariant relative categories are a valuable tool that allows researchers to separate the neural representation of the decision variable from sensory representations, which are often encoded by the same neural populations (Akrami et al., 2018; Pho et al., 2018; Romo et al., 1999; Steinmetz et al., 2019). In visual tasks of this type, information about the correct choice can be given by the difference between easy-to-parameterize features, such as the number of light flashes (Constantinople et al., 2019; Scott et al., 2015) or visual objects (Pinto et al., 2018), or contrast of oriented gratings (Burgess et al., 2017; Steinmetz et al., 2019), with the same value of decision variable given by many possible stimulus combinations.

Mice are a recently popular model animal for visual decision-making research (Carandini and Churchland, 2013); however, their ability to form invariances in an orientation discrimination task – a gold standard in vision research – and associated choice strategies have not been demonstrated. Although orientation discrimination is highly motivated by the efficiency of psychophysical (Campbell and Robson, 1968; Watson et al., 1983) and neural responses (Hubel and Wiesel, 1962), existing protocols for mice are based either on change detection (Glickfeld et al., 2013) or target-distractor discrimination paradigms (Reuter, 1987; You and Mysore, 2020). Thus, they should be interpreted as fundamentally detection tasks since a specific orientation is always either a target or a non-target stimulus. An invariant orientation discrimination task would instead rely on a relative orientation measure, permitting situations in which the same orientation is a target in one trial and a non-target in another.

Here we trained mice in an orientation discrimination task in which an animal was to indicate the more vertical orientation of the two simultaneously presented grating stimuli. Mice (n=40) learned the task, all exceeding the performance level of 65% correct, with up to 83% correct. Animals solved the task by adopting different strategies, which we captured with a novel behavioral model that quantified individual biases and the systematic variation of performance with both task difficulty and other factors. Furthermore, we explored the contribution of serial choice biases at different levels of engagement, and estimated orientation discrimination acuity that reached an upper bound of 6 degrees.

Finally, we discuss the advantages of our task for understanding the computation of decision in visual areas (DiCarlo and Cox, 2007), for linking neural and behavioral variability (Beck et al., 2012), and for studying suboptimal decision strategies and heuristics (Gardner, 2019). Given the availability of genetic tools, this work opens up a new level of questions about decision-making and the visual-to-cognitive link that can be pursued and explored in mice.

Results

Training and task

We trained n=40 transgenic mice in a 2AFC orientation discrimination task using an automated setup, in which the animal voluntarily fixed its head to initiate an experimental session (Aoki et al., 2017) (**Figure 1A,** top). Two oriented Gabor patches were simultaneously shown on the left and right sides of the screen, and in order to obtain water animals had to identify a *more vertically* oriented one (n=28; *more horizontally*, n=12) and move it to the center of the screen by rotating a wheel manipulator (Aoki et al., 2017; Burgess et al., 2017). Crucially, the target was not vertical on most trials, and animals had to compare the *verticality* of two orientations. The same physical stimulus could thus be a target or a non-target on different trials, making the task invariant relative to the orientation of individual stimuli (**Figure 1A,** middle). Orientations of both stimuli (θ_L, θ_R) were sampled at random from angles between -90° and 90° with a step size of 9° (3° for n=1 animal), with positive angles corresponding to clockwise and negative to counter-clockwise orientations relative to the vertical (**Figure 1A,** bottom). We used 9° spacing for most animals to sample a high number of responses for every angle condition, which was important for subsequent imaging experiments. Animals readily generalized across spatial frequencies and stimulus sizes (**Figure S1**). We analyzed a total of 1313355 trials, ranging from 4591 and 82065 per animal with an average of 32834±2962 trials (mean ± s.e.) (**Figure S1, Table 1**).

Mice reach high success rate in target-invariant task

We quantified task performance using a standard cumulative Gaussian psychometric function (Wichmann and Hill, 2001) of angular separation $\Delta\theta = |\theta_L| - |\theta_R|$ between the two orientations, where $|\cdot|$ denotes angular distance to the vertical (in short *verticality*), with small angular separation corresponding to difficult conditions and large angular separation corresponding to easy conditions. Angular separation $\Delta\theta = 0$ corresponds to two *equally vertical* orientations, which are not necessarily parallel. Conditions with $\Delta\theta < 0$ and $\Delta\theta > 0$ correspond to a more vertical orientation on the left and

right side respectively (example animal, **Figure 1B**; population, **Figure 1C**). Mice reached an average success rate of 0.75±0.01 with an average sensitivity of σ=42.93±1.18°.

The task disentangles any given probability of choice from specific orientations, and a fixed difficulty $\Delta\theta$ that corresponds to one point on the psychometric curve is given by many possible pairs of orientations (θ_L, θ_R). For example, $\Delta\theta$=30° corresponds to orientation pairs (30°,0°), (-60°,30°), and many others (**Figure 1D**). Conversely, no specific orientation was always rewarded, since for any orientation (except 0°) there was a possibility of the other orientation being more vertical.

As a consequence, this task design compels the animal to estimate the *verticality* of left and right orientations, $|\theta_L|$ and $|\theta_R|$, and compare their estimates, rather than detect a learned orientation.

Animals may not strictly adhere to this ideal strategy, so long as they get sufficient water reward in an experimental session. With choice variability taken into account, an animal looking at only one of the two stimuli will perform at 63.07±0.59% (Methods), exceeding the 50% chance level, but not reaching satiation. In the following we will introduce a model that will quantify how animals combine information from the two orientations while also capturing deviations from the ideal strategy.

Choice model

The psychometric curve quantifies the animal's behavior along a single dimension of difficulty $\Delta\theta$. However, given the task structure, the complete representation of the stimulus space is two-dimensional, with a unique stimulus condition corresponding to a pair of angles (θ_L, θ_R). In this space, a fixed $\Delta\theta$ is given by all stimulus conditions along the iso-difficulty lines (*branches*) that lie in four quadrants of the space and correspond to four different combinations of angle signs (**Figure 1D).** We therefore consider probability of right choice, P(R), for all stimulus conditions in this space.

To get a better insight into the factors that affect an animal's choices, we developed a psychometric model that provided a functional mapping from the two-dimensional stimulus space to the choice probability P(R).

We assume that in every trial a mouse makes noisy estimates (θ_L^*, θ_R^*) of both orientations (θ_L, θ_R), compares their verticalities $(|\theta_L^*|, |\theta_R^*|)$, and makes a choice (**Figure 2A**). The probability of a right choice P(R) in this procedure is expressed as an integral of the distribution of estimates $p(x, y)$ over the $|x| < |y|$ subspace, corresponding to $|\theta_R^*| < |\theta_L^*|$ (**Figure 2B**) (Methods: Eq. 1). The shape of the P(R) surface over the stimulus space (θ_L, θ_R) (**Figure 2C**, left) is therefore determined only by the parameters of the distribution $p(x, y)$.

We model these distributions using circular von Mises functions $p(x|\theta_R; \kappa_R)$ and $p(y|\theta_L; \kappa_L)$ centered at θ_R and θ_L equal to the true orientations, and with variability for each target quantified by the concentration parameters κ_R and κ_L. High concentrations correspond to low variability of angle estimates, and κ is thus qualitatively inverse to the standard deviation and can be interpreted as *certainty* of orientation estimates (Drugowitsch et al., 2016; Laquitaine and Gardner, 2018). For example, a distribution of estimates $p(x, y)$ is broader and shallower along the axis of lower concentration (**Figure 2D,** left column, top), making P(R) more independent of the respective stimulus (**Figure 2D,** left column, middle).

Estimates of each orientation can be systematically biased, with an animal consistently making choices as if the right or the left orientation were rotated more clockwise or counter-clockwise. These systematic errors are accounted for by *translational biases* b_R, b_L (**Figure 2D,** center column), which move $p(x, y)$ and consequently P(R) surface relative to the angle axes without changing their values.

Both the translational biases and certainty parameters change the slope of the psychometric curve but not its left-right choice bias (**Figure 2D,** bottom row), with this effect generally indistinguishable in the $\Delta\theta$ space as opposed to the complete stimulus space. Lower or higher certainty results in a shallower or steeper P(R) respectively, and a shallower or a steeper psychometric curve. On the other hand, a translational bias displaces the entire P(R) surface, overall decreasing performance for every $\Delta\theta$ in the space of the psychometric curve.

To model a *choice bias* towards right or left, we introduced a family of prior distribution functions or *choice priors* $p_b(x, y; \kappa_b)$ parameterized by prior concentrations κ_b (**Figure 2E**). Choice priors make an orientation on the right or on the left effectively appear more vertical – as opposed to more clockwise or counter-clockwise – or equivalently make an animal more certain about the verticality of that stimulus (**Figure 2D**, right column) (Methods, Eq. 2).

For example, choice prior for a rightward bias has a peak at (90°,0°) (**Figure 2E**, right, $\kappa_b > 0$) and increases probability of a right choice for any pair of orientations (**Figure 2D**, right column) by biasing $p(x, y)$ to the $|x| < |y|$ region (**Figure 2D**, right column, green arrows).

Concentrations, translational biases, and prior concentration $\{\kappa_R, \kappa_L, b_R, b_L, \kappa_b\}$ thus determine our model of choice, which allows for a more complete analysis of P(R) than the psychometric curve. The model predicted a previously unexplored property of P(R): its variation along the branches of a fixed $\Delta\theta$. A model with zero biases and an equal certainty for both orientations ($\kappa_R = \kappa_L$) predicted the decrease of P(R) whenever either orientation was close to 0° or 90°, and an increase when close to 45° (**Figure 2B-C**). We parameterized this variation using *reference orientation* $\theta_{ref} = \min(|\theta_L|, |\theta_R|)$, i.e. the orientation of the more vertical stimulus. The source of this variation is clear from the position of $p(x, y)$ relative to the category boundary $|x| = |y|$ when considered along one branch of a fixed $\Delta\theta$ (**Figure 2B**): the probability mass of orientation estimates that result in error judgments (e.g. $|\theta_R^*| > |\theta_L^*|$ when $|\theta_R| < |\theta_L|$) is higher around θ_{ref}=0° and θ_{ref}=90° than around θ_{ref}=45°. This effect arises from the variability of both orientation estimates and their interaction with the category boundary in the circular space, and cannot be replicated by the psychometric curve whose only input variable is $\Delta\theta$.

In summary, by combining information from two orientations, our model predicted a dependency of choice probability not only on difficulty but also on reference orientation. This latter variability necessarily follows from the circularity in the input stimulus space given a limited certainty in orientation estimates.

The model captures animals' choices

We next analyzed choice probabilities of mice in the two-dimensional stimulus space. For the population of animals, P(R) varied with difficulty $\Delta\theta$, as expected from the psychometric curves (**Figure 1B,C**), and with reference θ_{ref} (**Figure 2F**), as predicted by our model (**Figure 2B-C**). For a fixed $\Delta\theta > 0$, P(R) was higher (and choices were more often correct) when orientations were far from horizontal or vertical (**Figure 2F**); while for $\Delta\theta < 0$ P(R) was smaller (also more often correct) when the orientations were far from horizontal or vertical.

The model reproduced this performance variation for individual animals (**Figure 2G**). However due to individual biases P(R) curves for fixed $\Delta\theta$ were additionally distorted in comparison to the unbiased case (cf. **Figure 2C**, right). Counterintuitively, P(R) for the same $\Delta\theta$ in different quadrants of the stimulus space could represent on average opposite choices (**Figure 2G**, center, right), which our model accounted for thanks to translational biases. The model successfully captured animal-specific differences in choice probabilities (**Figure S3**), explained data significantly better than the psychometric curve (ΔAIC = 798.8±141.9; ΔAIC>0 for all animals), and explained significantly more deviance (Runyan et al., 2017) (ΔDE=9.03±1.48%, $p=1.07\cdot10^{-6}$, signed-rank test) (Methods).

Across the population of animals, average stimulus concentration values were high and positive – κ_R=2.22±0.69 ($p=3.73\cdot10^{-7}$, t-test) and κ_L=1.76±0.52 ($p=1.34\cdot10^{-7}$) (**Figure 2H**, left) – showing that animals used both targets for the decision. Bias concentration κ_b was small (κ_b=-0.06±0.05, p=0.01), pointing to the mixed bias across the population. Translational biases (b_R=0.14±0.05, $p=3.71\cdot10^{-6}$; b_L=0.19±0.05, $p=1.21\cdot10^{-7}$) were similarly small but significant.

Although the stimulus protocol, reward sizes, and session schedules were designed to motivate animals to use information of both orientations equally, we found that strategies of individual animals ranged from a balanced orientation comparison to relying on one target more than the other. We quantified this range of strategies with the ratio of log concentrations κ_R and κ_L, with ratios close to 1 for balanced strategies (**Figure 2H**, center). Right and left concentrations were significantly anti-

correlated (ρ=-0.57, p=4.45·10^{-4}, t-test α=5·10^{-3}, corrected), reflecting a trade-off in animals that preferentially used information from either stimulus (**Figure 2H**, right). Despite this trade-off, best-performing animals also had overall higher concentrations (p<0.05, ANCOVA, F-test of intercept with fixed slope), showing that while the task permitted a relative flexibility of choice strategies, a more accurate estimation of orientations was necessary for high success rates. Other parameters of the model did not significantly correlate between each other or with the concentration ratios.

In summary, our model accounted for biases in animals' behavior and explained performance variation with θ_{ref}. Individual animals weighted sensory information from two orientations differently, following strategies that were *sufficient* to obtain high amounts of reward, but were not perfectly aligned with the true stimulus-reward space. While left and right concentrations were anti-correlated across the population, high success rates required overall high certainty in orientation estimate.

Discrimination acuity

We used our model to estimate the minimal orientation difference that animals can reliably discriminate. A change in a pair of orientations that results in a significant change of P(R) is the smallest for conditions with the largest gradient of P(R). Since the numerical gradient directly computed from the data can be too noisy, we used our model to more accurately estimate maximum gradient conditions.

We compared choice probability in stimulus conditions with the highest gradient and in neighboring conditions (**Figure 3**). We found that a change of either orientation by 9° resulted in a significant change of P(R) in 62.5% (n=25) of animals, and a change of 27° resulted in a significant change in all (n=40) animals (**Figure 3A-E**). In the only animal tested with a 3° sampling of stimuli, we found that changing both orientations by 3° along or against the gradient – amounting to a total change of 6° – resulted in a significantly different P(R) (p<0.001, both cases) (**Figure 3F-I**).

In summary, our model allowed an in-depth analysis of discrimination acuity by utilizing a complete picture of the P(R) gradient and identifying stimulus conditions where the sensitivity to angle change

is the highest. We found that an angle change of 6° can be significantly detected based on the change of choice probability, thus establishing an upper bound for mouse orientation discrimination acuity.

Effects of trial history

Choice strategies are determined not only by preferential weighting of available sensory information, but also by trial history (Abrahamyan et al., 2016; Akrami et al., 2018; Busse et al., 2011; Corrado et al., 2005; Fründ et al., 2014; Urai et al., 2017; Yu and Cohen, 2008). To account for history-related biases, we included a history prior $p_h(x, y)$ parameterized with a concentration parameter κ_h and a term h that linearly depended on the previous trial's choice r and target orientation s through history weights, $h = sh_s + rh_r$ (Busse et al., 2011; Corrado et al., 2005; Fründ et al., 2014). A pair of weights (h_s, h_r) determined the choice strategy of an animal, such as "win-stay" (**Figure 4A**, model example) or "lose-stay" (**Figure 4A**, example animal) throughout all trials, and in combination with the choice and target of the previous trial (r, s) determined the history-dependent change of the P(R) (**Figure S4A-E**) and the psychometric curve (Fründ et al., 2014) (**Figure 4A**).

Through the flexible family of history priors, our model captured a variety of strategies besides win-stay (**Figure S4F**). Most of our mice showed a mild tendency for the "stay" strategy, followed by "win-stay", and rare cases of "lose-stay" (**Figure 4C**), largely in consistency with previous reports (Odoemene et al., 2018). The history-dependent model explained data significantly better than the history-independent model (ΔAIC = 211.8±40.1; ΔAIC>0 for all but n=4 animals), and explained significantly more deviance (ΔDE=5.07±0.98%, $p=8.1\cdot10^{-6}$, signed rank test).

We investigated whether animals rely on the history to a different extent during periods of relatively high and low engagement in the task, which we identified based on performance within a session (Methods). For each of the two engagement levels, we computed the difference between the trial-average log-likelihood of choices given a model with a history prior and without: ΔL_l for low-engagement trials, and ΔL_h for high-engagement. The increase in explanatory power was larger in the

low-engagement trials ($\Delta L_l > \Delta L_h$) (Wilcoxon test, p=2.12e-07) (**Figure 4E**), meaning that on these trials the choices were more strongly driven by the history biases.

Difficult stimulus conditions were more susceptible to the influence of history priors than easy conditions (**Figure 4F**, left), in a way that depended on the engagement state of the animal. During periods of high engagement, inclusion of history priors led to a substantial improvement in the model's performance only in the most difficult conditions (**Figure 4F**, center), while in the low engagement periods most stimulus conditions were affected (**Figure 4F**, right).

In summary, in expanding our model to capture history-dependent biases, we found that most prominent strategies were "win-stay" and "stay", and that choices were more affected by history biases during periods of lower engagement. Our observations demonstrate that choice heuristics can fluctuate together with the cognitive state of the subject (Whiteley and Sahani, 2012; Wyart and Koechlin, 2016).

Discussion

Using high throughput automated cages with voluntary head fixation, we trained a large cohort of mice (n=40; 1,313,355 trials) in a complex variant of a 2AFC orientation discrimination task. The task required mice to measure relative orientation, thereby decoupling choice from the orientation of individual stimuli. We quantified their behavior with a novel model of choice that accounted for the circularity of stimulus space and for individual choice biases and strategies. The model explained variation in choice probability not only with the task difficulty $\Delta\theta$, but also with the reference orientation θ_{ref}, an effect not reported previously. With the help of the model we found that the maximum acuity of orientation discrimination in expert animals can be as small as 6°. History biases, ubiquitous in human and animal psychophysical experiments (Abrahamyan et al., 2016; Akrami et al., 2018; Busse et al., 2011; Corrado et al., 2005; Fründ et al., 2014; Urai et al., 2017; Yu and Cohen, 2008), were modulated by animals' engagement, affecting choices more strongly and on a broader set of

stimulus conditions whenever the engagement was relatively low. Our work responds to the need for a visual task that depends on abstract choice categories and is invariant to specific visual stimuli, but can be learned by mice, relies on basic visual features, and allows straightforward quantification within the framework of signal detection theory. We argue that besides these advantages, our task can be useful in engaging higher visual areas in the computation of decision (DiCarlo and Cox, 2007), and can give valuable insight into the relationship between neural and behavioral variability (Beck et al., 2012; Britten et al., 1996; Brunton et al., 2013; Drugowitsch et al., 2016; Renart and Machens, 2014).

Object recognition and discrimination tasks can be used to introduce perceptual invariances into decision-making experiments (DiCarlo and Cox, 2007; Rust and Stocker, 2010). In rodents, such experiments have been successfully carried out in rats (Tafazoli et al., 2012; Zoccolan, 2015; Zoccolan et al., 2009), with more recent ongoing research into object recognition (Froudarakis et al., 2020) and its ethological relevance (Hoy et al., 2016) in mice. However, in discrimination tasks, where distinct objects (rather than transformations of the same object) can correspond to the same category (Cox, 2014; DiCarlo and Cox, 2007), quantification of the stimulus space and of subjective object similarity can be challenging, requiring the introduction of object-morphing axes (Tafazoli et al., 2012), and is yet to be demonstrated in mice. Our task, which can be learned by mice, is invariant to low-level visual activations and is straightforwardly parameterized with two angles, thus representing a useful alternative to object-based paradigms. Furthermore, the low dimensionality and periodicity of our stimuli allows exploration of the full stimulus space, while 3d shape-detection tasks are restricted to a set of possible objects (Cox, 2014). Our task builds upon orientation discrimination tasks in mice (Andermann et al., 2010; Goard et al., 2016; Long et al., 2015; Reuter, 1987; You and Mysore, 2020), in which a specific orientation is to be chosen over a distractor orientation (Long et al., 2015; Reuter, 1987; You and Mysore, 2020), or in which a change relative to a specific orientation is to be detected (Glickfeld et al., 2013). Because in these tasks a given orientation is always either a target or a non-target, they can be interpreted as fundamentally detection-type. In our task, in contrast, a successful

trial outcome is not guaranteed by detection of any specific orientation. Instead, the animal had to measure and compare the verticality of two orientations.

An adequate quantification of choice behavior differs accordingly between these two approaches. In detection-type tasks, models of one variable are appropriate, such as cumulative Gaussian or Weibull (Glickfeld et al., 2013) functions of angle difference (Reuter, 1987). Choices in our task could similarly be quantified with a one-dimensional function of $\Delta\theta$, but taking into account the full dimensionality of the stimulus space and using both orientations as inputs to the model allowed us to explain the data better than the standard psychometric function.

Our model helped estimate orientation discrimination acuity, which reached 6 degrees of angle difference. The orientation discrimination acuity of mice has been previously measured in detection-type tasks, such as 2AFC with a distractor (Reuter, 1987), and in change detection tasks (Glickfeld et al., 2013; Wang et al., 2020). Acuity measures have been reported as thresholds or just-noticeable differences (JND) and commonly rely on model-derived values, such as the model-based inverse of a certain success rate (Glickfeld et al., 2013), mean of the fitted Gaussian (Wang et al., 2020), or $\sqrt{2}$ times its standard deviation (Wang et al., 2020). We developed a new acuity estimation procedure suitable for our stimulus space, in which we identified stimulus conditions with the highest gradient of model-predicted P(R), and compared performance in these and neighboring conditions. Our approach took advantage of the complete stimulus space representation of P(R) instead of relying on a cruder psychometric model to compute a JND or threshold value.

Assumption of noisy orientation estimation was also sufficient to qualitatively explain the reference effect – the variation of the probability of choice with a varying reference orientation θ_{ref} and a fixed difficulty $\Delta\theta$. This variation arises from the interaction of noisy orientation estimates in a circular domain with the category boundary and is not due to animal-specific strategies. A similar effect could probably be demonstrated using our task in another periodic sensory domain, for example with comparisons of pure tones or chords.

Our model allowed for systematic errors in angle estimates via translational bias parameters, which shifted P(R) surfaces with respect to the angle axes in the 2d stimulus space, but were concealed in the psychometric curve representation, instead only changing its slope. These biases were relatively small across the population, and did not preclude high levels of performance, but were essential to explaining a counterintuitive difference between the P(R) curves for the same small $\Delta\theta$ in different quadrants of the stimulus space (**Figure 2G**). Furthermore, we accounted for choice bias – unlike translational biases – using a choice prior, which modelled an animal's tendency to choose predominantly the right or left stimulus. Unlike translational biases, choice bias increased or decreased all values of P(R) and corresponded to a shift of the psychometric curve.

With the described set of parameters $\{\kappa_R, \kappa_L, b_R, b_L, \kappa_b\}$, our model could capture variation of choice patterns across animals, assuming fixed concentrations for all stimulus conditions and a Bayes-optimal combination of sensory information with the choice prior (Knill and Pouget, 2004). Beyond the capabilities of the model, we find that at least in some animals variation of choice probability with θ_{ref} for a fixed $\Delta\theta$ can be larger than our prediction (**Figure S5**. This additional variation would be explained by a dependency of κ_R and κ_L on the proximity to the category boundary ($|\theta_R^*| = |\theta_L^*|$) (Jazayeri and Movshon, 2007). Furthermore, concentration values are likely decreased by non-sensory factors, such as noise in decision computation (Beck et al., 2012; Dosher and Lu, 1998; Drugowitsch et al., 2016), inherent priors (Girshick et al., 2011), and choice heuristics (Beck et al., 2012; Gardner, 2019).

One choice heuristic was evident in the trade-off of concentration values, with some animals unequally weighting stimulus information. Accuracy of orientation estimation was still necessary for high success rates, but even among the best-performing animals right and left concentrations were anti-correlated. This trade-off demonstrated that animals followed a range of "sufficiently good" strategies when solving the discrimination problem.

Such strategies can be interpreted as examples of suboptimal or approximate inference in an uncertain environment. Suboptimal inference is sometimes thought of as an adaptive phenomenon, a way for a subject to deal with the complexity of the task at hand by constructing and acting upon its approximate model (Beck et al., 2012). Adherence to a suboptimal strategy can therefore be linked to limited cognitive resources (Whiteley and Sahani, 2012; Wyart and Koechlin, 2016), which in our task fluctuate together with task engagement. Indeed, we find that history-dependent biases – another manifestation of suboptimal behavior – are stronger during periods of lower engagement. We demonstrate this by introducing history priors that increase the explanatory power of the model more in periods of lower engagement than in periods of higher engagement. Besides the state-dependent effect of history biases, the model allowed us to establish that average history-dependent choice strategies were mainly “stay” and “win-stay”, with rare examples of “lose-stay”.

Our task can be useful for linking neural activity and behavior. First, our task decouples the decision variable from specific sensory stimuli, which makes it easier to distinguish between decisions and sensory representations at the neural level. A similar approach relying on combinations of stimuli has been used extensively in the decision-making literature (Hernández et al., 1997), but has not been reported in mouse orientation discrimination experiments.

Second, our task can give a valuable insight into the relationship between neural and behavioral variability. Whether behavioral variability arises predominantly from sensory sources (Brunton et al., 2013), or from the deterministic or stochastic suboptimality of decision computation (Beck et al., 2012) is one of the central questions in the neuroscience of decision-making. One key advantage of our task is that it keeps the variability of purely sensory responses constant by fixing the contrast of the stimuli. If the neural variability is found to systematically change across the stimulus space, this variation can only be attributed to the noisy computation of decision (Beck et al., 2012). Neural variability across the stimulus space can thus be partitioned into the purely sensory variability

determined by the contrast, the non-sensory component independent of the stimuli, and a stimulus-dependent component which our task can help tease apart from the other two.

Third, our task is well-suited to isolating the contributions of visual cortical areas in the computation of decision. The importance of a particular visual area for decision-making depends on the type of task (Pinto et al., 2019): mice with a lesioned or silenced visual cortex can show above-chance performance in detection paradigms (Glickfeld et al., 2013; Prusky and Douglas, 2004), possibly reflecting a predominant role of the superior colliculus (Wang et al., 2020), while for orientation discrimination tasks with a distractor, the visual cortex seems to be necessary (Jurjut et al., 2017; Poort et al., 2015; Resulaj et al., 2018). However even in the latter cases the involvement of the visual cortex in the computation of decision rather than sensorimotor relay is unclear. In our task, both stimuli are predominantly encoded in V1 and their comparison can be done in V1 or possibly further downstream, since the visual cortical system is characterized by increasingly complex category boundaries that facilitate discrimination between high-level perceptual concepts (DiCarlo and Cox, 2007). Our task is therefore a useful addition to the common protocols for visual decision-making research in mice focused on cortical computations.

Acknowledgements

We thank Yuki Goya and Rie Nishiyama for their support with behavioral training. We thank O'Hara and CO., LTD., for their support with the equipment. This work was funded by RIKEN BSI and RIKEN CBS institutional funding, JSPS grants 26290011, 17H06037, C0219129 to AB, and Fujitsu collaborative grant.

Author contributions

A.B. and D.R.L. designed the study. D.R.L. analyzed the data and developed the model. R.A. developed the general framework for the behavioral paradigm. A.B. and D.R.L. wrote the manuscript.

Declaration of interests

The authors declare no competing interests

Tables

Table 1. Total number of trials per animal

ID	N trials	ID	N trials	ID	N trials	ID	N trials
1	82065	11	38624	21	30078	31	18637
2	76488	12	38583	22	29646	32	17624
3	66929	13	38263	23	28228	33	15681
4	63074	14	37961	24	27659	34	13006
5	58747	15	37189	25	27425	35	11893
6	56392	16	36422	26	25509	36	11885
7	48118	17	33938	27	22222	37	11465
8	46872	18	33003	28	20926	38	10069
9	45415	19	32946	29	20113	39	5602
10	39673	20	30988	30	19406	40	4591

"ID" columns show animal identification numbers as in Figures S2-3, "N Trials" columns show total number of trials of the corresponding mouse used in the analysis throughout the paper.

Methods

Experimental Model and Subject Details

All surgical and experimental procedures were approved by the Support Unit for Animal Resources Development of RIKEN CBS. We used n=40 transgenic mice: Thy1-GCaMP6f (n=37), Camk2-tTA TRE-GCaMP6s (n=2), Emx1-tTA TRE-GCaMP6s (n=1), with a total of 30 male and 10 female animals, aged 4 to 25 months. The triple transgenic strain Camk2-tTA TRE-GCaMP6s was established by cross-mating Camk2a-cre and Camk2a-tTA. The triple transgenic strain Emx1-tTA TRE-GCaMP6s was established by cross-mating Emx1-cre and Camk2a-tTA.

Animals were anesthetized with gas anesthesia (Isoflurane 1.5-2.5%; Pfizer) and injected with an antibiotic (Baytril®, 0.5ml, 2%; Bayer Yakuhin), a steroidal anti-inflammatory drug (Dexamethasone; Kyoritsu Seiyaku), an anti-edema agent (Glyceol®, 100µl, Chugai Pharmaceutical) to reduce swelling of the brain, and a painkiller (Lepetan®, Otsuka Pharmaceutical). The scalp and periosteum were retracted, exposing the skull, then a 4 mm-diameter trephination was made with a micro drill

(Meisinger LLC). A 4mm coverslip (120~170µm thickness) was positioned in the center of the craniotomy in direct contact with the brain, topped by a 6 mm diameter coverslip with the same thickness. When needed, Gelfoam® (Pfizer) was applied around the 4mm coverslip to stop any bleeding. The 6mm coverslip was fixed to the bone with cyanoacrylic glue (Aron Alpha®, Toagosei). A round metal chamber (6.1mm diameter) combined with a head-post was centered on the craniotomy and cemented to the bone with dental adhesive (Super-Bond C&B®, Sun Medical), mixed to a black dye for improved light absorbance during imaging.

Behavioral training

Animals were housed in individual cages connected to automated setups (Aoki et al., 2017) (O'Hara & CO., LTD., http://ohara-time.co.jp/) where two experimental sessions per animal per day were carried out. Sessions were initiated by animals themselves as they entered the setup and their head plate was automatically latched. Animals were trained in a 2AFC orientation discrimination task. Two oriented Gabor patches (20° visual angle static sinusoidal gratings, sf = 0.08 cpd, with randomized spatial phase, and windowed by a 2D Gaussian envelope with σ=0.25° v.a.) were shown on the left and right side of a screen positioned in front of the animal (LCD monitor, 25 cm distance from the animal, 33.6 cm × 59.8 cm [~58° × 100°dva], 1080 x 1920 pixels, PROLITE B2776HDS-B1, IIYAMA) at ±35° eccentricity relative to the body's midline. Mice reported which of the two stimuli was more vertical (more horizontal for n=12 animals; task details in "Phases of training") by rotating a rubber wheel with their front paws, which shifted the stimuli horizontally on the screen. For a response to be correct, the target stimulus had to be shifted to the center of the screen, upon which the animal was rewarded with 4µL of water (amount adjusted for a few animals with non-typical weight and age). Incorrect responses were discouraged with a prolonged (10s) inter-trial interval and a flickering checkerboard stimulus (2 Hz). If no response was made within 10 s (time-out trials), neither reward nor discouragement was given.

All trials consisted of an open-loop period (OL, 1.5s) during which the wheel manipulator did not move the stimuli on the screen, and a closed-loop period (CL: 0—10 s) during which the wheel controlled their position. Inter-trial interval was randomized (ITI: 3—5s). Stimuli appeared on the screen at the beginning of the OL.

Phases of training

Training in the automated behavioral setup went through three phases. First, the animal learned to rotate the wheel manipulator and was rewarded for consistent rotations to either side. During this phase no visual stimulus was presented. In the next phase, the animal was shown one vertical target (horizontal, n=12), on one side of the screen chosen at random, and was rewarded for moving it into the center of the screen. In the final phase, the animal was shown two orientations, and had to move the more vertical (horizontal) one into the center of the screen. Since both stimuli moved synchronously with wheel rotation, the non-target stimulus moved out of the screen. In this phase, we sampled both orientations at random from a range of angles between -90° and 90°, with $\theta > 0$ corresponding to clockwise and $\theta < 0$ – to counter-clockwise orientations relative to the vertical

(**Figure 1A**). Orientations were initially sampled with a minimal angular difference of 30°, i.e. with specific angles from the set {-90°,-60°,-30°,0°,30°,60°} (-90° and 90° are the same orientation). As the animal's performance reached 70% success rate on 5-10 consecutive days, we increased the difficulty by sampling angles at 15° angle difference, and later in the training – at 9°, with one animal's conditions eventually sampled at 3°.

Data selection

We analyzed trials from sessions in which the average success rate was at least 60%, and the proportion of time-out trials did not exceed 20%. We only used animals that had reached the orientation sampling step of 9°, and included the choice data from preceding sessions with all angle binning steps starting from 30°. We excluded the first trial of every session, all time-out trials and every trial that followed a time-out. The two dimensions of the stimulus space were flipped for horizontal-reporting animals when fitting our model. Same stimulus space transformation was done for all the population summaries where mice trained on horizontal targets were pooled together with mice trained on vertical targets.

Psychometric curve

We fitted the animal's probability of making a right choice P(R) as a function of task difficulty using a psychometric function $\psi(\Delta\theta;\ \alpha,\beta,\gamma,\lambda) = \gamma + (1-\gamma-\lambda)\ F(\Delta\theta;\alpha,\beta)$, where $F(x)$ is a Gaussian cumulative probability function, α and β are the mean and standard deviation, γ and λ are left and right (L/R) lapse rates, $\Delta\theta$ is the difference in the angular distance to the vertical, $\Delta\theta = |\theta_L| - |\theta_R|$. Confidence intervals were computed by bootstrapping (n = 999).

Model design

On each trial i the animal was shown a pair of stimuli $\{\theta_{Ri}, \theta_{Li}\}$, and made a right or a left choice r_i, which we set by convention to be $r_i = 1$ or $r_i = 0$ respectively. We denote response and correct target on the previous trial as r_{hi} and s_{hi} respectively, with $r_{hi} = -1$ or $r_{hi} = 1$ if the animal chose left or right respectively, and $s_{hi} = -1$ or $s_{hi} = 1$ if the correct answer was respectively left or right, and $s_{hi} = 0$ if targets had an equal distance to the vertical.

On every trial the animal made an estimate $\{\theta^*_{Ri}, \theta^*_{Li}\}$ of the presented stimulus orientations $\{\theta_{Ri}, \theta_{Li}\}$. We model θ^*_{Ri} and θ^*_{Li} as random variables whose probability densities are von Mises distributions centered at θ_{Ri} and θ_{Li}, with additional angle estimation biases (*translational biases*) b_R, b_L and with concentrations κ_R, κ_L (high concentration means smaller spread, with κ analogous to $1/\sigma$ of a normal distribution) **(Figure 2B,D)**:

$$\theta^*_{Ri} \sim p(x) = \mathrm{C}(\kappa_R) e^{\cos(\kappa_R(x - b_R - \theta_{Ri}))}$$
$$\theta^*_{Li} \sim p(y) = \mathrm{C}(\kappa_L) e^{\cos(\kappa_L(y - b_L - \theta_{Li}))} \quad (1)$$

where $\mathrm{C}(\kappa) = 1/2\pi I_0(\kappa)$, and I_0 is modified Bessel function of order 0. We reserve (x, y) notation for the domain of integration and $\{\theta^*_L, \theta^*_R\}$ – for angle estimates. We next assume that the joint probability density that underlies decision is a combination of a stimulus-based joint $p(x, y) =$

$p(x)p(y)$, a bias prior $p_b(x,y)$ that induces choice bias for right or left stimuli, and a history prior $p_h(x,y)$ that models choice dependency on previous choice and stimulus (r_{hi} and s_{hi}):

$$p_b(x,y) = C_b^2(\kappa_b)e^{\kappa_b(\cos(x)-\cos(y))} \quad (2)$$

$$p_h(x,y) = C_h^2(\kappa_h)e^{h_i\kappa_h(\cos(x)-\cos(y))} \quad (3)$$

Here, κ_b is a concentration parameter that regulates the strength and sign of choice bias, κ_h is a concentration parameter of history prior, $h_i = h_s s_{hi} + h_r r_{hi}$ determines the influence of the previous stimulus s_{hi} and choice r_{hi} with respective weights h_s and h_r fixed for a given animal, and $C_h = 1/2\pi I_0(\kappa_h h_i)$ and $C_b = 1/2\pi I_0(\kappa_b)$ are normalization constants.

Since by convention we set vertical orientation to zero, the angle with the smaller absolute value is the correct choice. Hence, the probability of choosing right on a given trial is given by:

$$P(R)_i = p(r_i = 1) = p(|\theta_{Ri}^*| < |\theta_{Li}^*|)$$
$$= \frac{\iint_{|x|<|y|} p(x,y)p_b(x,y)p_h(x,y)dxdy}{\iint p(x,y)p_b(x,y)p_h(x,y)dxdy} \quad (4)$$

Overall, the model has eight fitted parameters ($h_r, h_s, \kappa_R, \kappa_L, \kappa_h, b_R, b_L, \kappa_b$), or five parameters ($\kappa_R, \kappa_L, b_R, b_L, \kappa_b$) when we fit a history-free model. All angles were converted from (-90°,90°) range to (-180°,180°) to satisfy periodicity.

Optimization

To fit the model, we minimize the cross-entropy cost function

$$L = -\sum_{i=1\ldots N} r_i \log P(R)_i + (1-r_i)\log(1-P(R)_i) \quad (5)$$

using MATLAB built-in function fmincon. At every iteration of the optimizer we evaluated equation (4), first computing values of all probability densities on a grid of 300 by 300 points in the 2d domain $[-\pi,\pi]\times[-\pi,\pi]$, and integrating numerically using MATLAB function trapz over $|x|<|y|$ for the numerator and over the whole domain for the denominator. We ran these calculations on GPU (NVIDIA RTX 2080Ti) using MATLAB Parallel Computing Toolbox.

Success rate with a one-sided strategy

We estimated the success rate that animals could reach when taking into account only one stimulus by first computing P(R) for every trial using a model where one concentration was set to zero and the other one to $\sqrt{\kappa_R\kappa_L}$ of that animal. We sampled choices using the stimulus conditions as they appeared in the experimental dataset 1000 times and computed an average percent correct over repetitions and an average across animals.

Maximum perceptual acuity

By analogy with a 1d psychometric curve, we defined points of maximum perceptual acuity in the stimulus space as conditions (pairs of angles) where the change in $P(R)$ was the largest for a small fixed change in the stimuli. We found these conditions from the probability surface $P(R)$ given by the

full model by computing the squared norm of the gradient vector, $g(\theta_R, \theta_L) = \left(\frac{d}{d\theta_R}P(R)\right)^2 + \left(\frac{d}{d\theta_L}P(R)\right)^2$ and selecting $\{\theta_R, \theta_L\}$ conditions for which the values of g were in the top 5%. Among these conditions, we analyzed those with P(R)≈0.5 (0.48≤P(R)≤0.52), which we call *maximum gradient* conditions (**Figure 3C,H**, white) with a pooled right-choice probability of $P^{0.5}$. For n=28 animals this procedure gave at least 3 unique maximum gradient conditions. For n=12 animals, the initial criterion gave fewer than three maximum gradient conditions, and we expanded the allowed range to have at least 3: we set (0.47≤P(R)≤0.53) for n=7 animals, (0.46≤P(R)≤0.54) for n=1, (0.42≤P(R)≤0.58) for n=1, (0.40≤P(R)≤0.60) for n=1, (0.38≤P(R)≤0.62) for n=1, and (0.28≤P(R)≤0.72) for n=1.

We then determined the neighboring conditions by changing one orientation at a time by 9°, which resulted in an increase ("+") or decrease ("-") of P(R) relative to $P^{0.5}$ (**Figure 3C**). For example, P^{R-} corresponded to the probability of right choice pooled from all conditions in which θ_R changed relative to maximum gradient conditions in the direction of P(R) decrease. Here, the stimulus space was binned to a 9° grid. In a separate analysis, for an animal with 3° condition binning, we changed both orientations simultaneously by +-3°, "along" and "against" the gradient of P(R), and obtaining P^+ and P^+ respectively (**Figure 3H**).

We tested that probabilities in the neighboring conditions (P^{L+}, P^{R+}, P^{L-}, P^{R-} in case of 9°-binned conditions, and P^+, P^- in case of 3°-binned conditions) were significantly different from maximum gradient probabilities $P^{0.5}$ using a two-tailed χ^2 test with df=1. For a population summary (**Figure 3E**) we computed P^{L+}, P^{R+}, P^{L-}, P^{R-} with increasing angle increments of 9°, 18°, and 27° and reported the cumulative number of animals for which at least one of the four probabilities was significantly different from $P^{0.5}$, using a two-tailed χ^2 test with df=1 and a criterion α=0.05/4.

History biases during high and low engagement

We first identified periods of high and low engagement in every session. For a given session, we computed a running estimate of the success rate in a sliding window of 10 trials (average performance in the window was assigned to the last trial of that window). We centered the running estimate by subtracting the mean success rate of the session. All trials with the centered success rate estimate exceeding a fixed threshold of 10% were labeled as high engagement, and all trials in which the centered success rate estimate was lower than -10% were labeled as low engagement. We confirmed the stability of our results using threshold values of 5%, 15%, and 20% (data not shown). When identifying engagement epochs, time-out trials were counted as failures, but we discarded these trials for all the analysis that followed, consistently with the rest of this study.

Next, we computed the log-likelihood L of outcomes in high- and low-engagement trials (r^h and r^l respectively) given the probabilities predicted by the full model that accounted for trial history, and by a history-free model fitted separately (p_h and p_0 respectively) (see Methods: Model Design). For binary outcomes r and model-derived probabilities p, we computed trial-wise the log-likelihood using the formula $L(r,p) = r\log(p) + (1-r)\log(1-p)$ with stimulus conditions binned to a 9° grid.

Applying two different trial selections and two different models we obtained $L(r^h, p_h)$ for the likelihoods of high-engagement trial outcomes given the model with history, $L(r^l, p_h)$ for the likelihoods of low-engagement trial outcomes given the model with history, $L(r^h, p_0)$ for the likelihoods of high-engagement trials given the history-free model, and $L(r^l, p_0)$ for the likelihoods of low-engagement trials given the history-free model. We next computed the differences of likelihood averages between models with and without history terms, using high-engagement trials, $\Delta L_h = \langle L(r^h, p_h)\rangle - \langle L(r^h, p_0)\rangle$ and low-engagement trials, $\Delta L_l = \langle L(r^l, p_h)\rangle - \langle L(r^l, p_0)\rangle$, (**Figure 4D**).

Next, we computed the average of each of these likelihoods across all trials for every pair of orientations (θ_L, θ_R) thus obtaining maps of $\langle L(r^*, p_*)\rangle_\theta$ as a function of orientations (θ_L, θ_R). We discarded any stimulus conditions where the number of trials was <10. We computed the difference between history and history-free maps of $\langle L(r^*, p_*)\rangle_\theta$ separately for high- and low-performance trials, i.e. $\Delta L_{h\theta} = \langle L(r^h, p_h)\rangle_\theta - \langle L(r^h, p_0)\rangle_\theta$ and $\Delta L_{l\theta} = \langle L(r^l, p_h)\rangle_\theta - \langle L(r^l, p_0)\rangle_\theta$, and for all trials together, $\Delta L_\theta = \langle L(y, p_h)\rangle_\theta - \langle L(y, p_0)\rangle_\theta$. For a population summary (**Figure 4E**) of ΔL_θ, $\Delta L_{h\theta}$, and $\Delta L_{l\theta}$, we normalized ΔL_* maps of every animal by the standard deviation across all stimulus conditions, and averaged the resulting maps across animals.

Model comparison

AIC

We compared the cumulative Gaussian psychometric model to our history-free model, and the history-free model to the model with history priors, using the Akaike Information Criterion (AIC) defined as $AIC = -2L + 2k$ where k is the number of parameters (4 for Gaussian model, 5 for the history-free model, 8 for model with history) and L is the log-likelihood value of the best fit. We computed L using the binomial likelihood formula

$$L = \sum_i y_i n_i \log(p_i) + n_i(1 - y_i)\log(1 - p_i) + \log\binom{n_i}{y_i n_i}$$

where i corresponds to a 9°-binned unique stimulus condition defined by (θ_L, θ_R) for the history-free to Gaussian model comparison and $(\theta_L, \theta_R, r_h, s_h)$ for the history-free to the history-dependent model comparison, y_i is the proportion of successes, n_i is the total number of trials, and p_i is the success rate given by either one of three models. We computed and reported $\Delta AIC = AIC(Gauss) - AIC(HistFree)$, and $\Delta AIC = AIC(HistFree) - AIC(HistDependent)$ for the final quantification. Given this measure, the model whose AIC is the minuend in the expression for ΔAIC is $\exp(-\Delta AIC/2)$ times more probable than the other model to minimize the information loss.

Percent of explained deviance

To estimate how much explanatory power is gained by fitting the history-free model in comparison to the Gaussian psychometric model, and by the history-dependent model in comparison to the history-free model, we computed the fraction of explained deviance. Deviance is defined as two times the log of the ratio of the saturated model likelihood $l(\theta_{max}; y)$ to optimal model likelihood $l(\hat{\theta}; y)$

$$D = 2\log\left(\frac{l(\theta_{max}; y)}{l(\hat{\theta}; y)}\right) \quad (6)$$

where y are observations, $\hat{\theta}$ are estimated parameters, and θ_{max} are parameters of the saturated model.

For binomial data, deviance is

$$D = 2\sum_i y_i n_i \log\left(\frac{y_i}{p_i}\right) - (1 - y_i) n_i \log\left(\frac{1 - y_i}{1 - p_i}\right) \quad (7)$$

where $y_i n_i$is the number of successes for stimulus condition i, n_i is the number of trials, and p_i is the probability of success in condition i given by the fitted model with parameters $\hat{\theta}$. For the cumulative Gaussian psychometric function $\psi(\Delta\theta;\ \alpha, \beta, \gamma, \lambda)$ a stimulus condition is defined by a pair of angles $\{\theta_R, \theta_L\}$ in a history-free model, and a pair of angles with trial history $\{\theta_R, \theta_L, s_h, r_h\}$ in a model with history.

We first computed the deviance of the null model, with the same P(R)=p_{null} rate for all conditions (computed as a grand average P(R) across trials). We then used the formula for deviance D (7), with $p_i = p_{null}$ when computing null deviance D_{null}, $p_i = p_i(HF)$ as predicted by history-free model when computing history-free deviance D_{HF}, $p_i = p_i(HD)$ as predicted by the history-dependent model when computing history-dependent deviance D_{HD}, and $p_i = p_i(Gauss)$ as predicted by the Gaussian model when computing Gaussian deviance D_{Gauss}. Here a condition i corresponded to a unique pair of orientations (θ_L, θ_R) when comparing the Gaussian model with the history-free model, and to a pair of orientations together with history inputs $(\theta_L, \theta_R, s_h, r_h)$ when comparing the history-free model and the history-dependent model; the fraction of right choices y_i and the total number of trials per condition n_i changed accordingly. We computed deviance explained (DE) for the three models as $DE_{HF} = 100\% \cdot (D_{null} - D_{HF})/D_{null}$, $DE_{HD} = 100\% \cdot (D_{null} - D_{HD})/D_{null}$ and $DE_{Gauss} = 100\% \cdot (D_{null} - D_{Gauss})/D_{null}$, and finally we computed difference in deviance explained as $\Delta DE = DE_{HF} - DE_{Gauss}$ or $\Delta DE = DE_{HD} - DE_{HF}$. For this analysis, we trained each model on 50% randomly sampled trials and computed deviances from the other 50% of trials. We tested the significance of $\Delta DE > 0$ for a population of animals using the Wilcoxon signed rank test.

References

Abrahamyan, A., Silva, L.L., Dakin, S.C., Carandini, M., and Gardner, J.L. (2016). Adaptable history biases in human perceptual decisions. Proc. Natl. Acad. Sci. U. S. A. *113*, E3548-57.

Akrami, A., Kopec, C.D., Diamond, M.E., and Brody, C.D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. Nature *554*, 368–372.

Andermann, M., Kerlin, A., and Reid, C. (2010). Chronic cellular imaging of mouse visual cortex during operant behavior and passive viewing. Front. Cell. Neurosci. *4*, 3.

Aoki, R., Tsubota, T., Goya, Y., and Benucci, A. (2017). An automated platform for high-throughput mouse behavior and physiology with voluntary head-fixation. Nat. Commun. *8*, 1196.

Beck, J.M., Ma, W.J., Pitkow, X., Latham, P.E., and Pouget, A. (2012). Not noisy, just wrong: the role of suboptimal inference in behavioral variability. Neuron *74*, 30–39.

Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis. Neurosci. *13*, 87–100.

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. Science (80-.). *340*, 95 LP – 98.

Burgess, C.P., Lak, A., Steinmetz, N.A., Zatka-Haas, P., Bai Reddy, C., Jacobs, E.A.K., Linden, J.F., Paton, J.J., Ranson, A., Schröder, S., et al. (2017). High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice. Cell Rep. *20*, 2513–2524.

Busse, L., Ayaz, A., Dhruv, N.T., Katzner, S., Saleem, A.B., Schölvinck, M.L., Zaharia, A.D., and Carandini, M. (2011). The detection of visual contrast in the behaving mouse. J. Neurosci. *31*, 11351–11361.

Campbell, F.W., and Robson, J.G. (1968). Application of Fourier analysis to the visibility of gratings. J. Physiol. *197*, 551–566.

Carandini, M., and Churchland, A.K. (2013). Probing perceptual decisions in rodents. Nat. Neurosci. *16*, 824–831.

Constantinople, C.M., Piet, A.T., and Brody, C.D. (2019). An Analysis of Decision under Risk in Rats. Curr. Biol. *29*, 2066-2074.e5.

Corrado, G.S., Sugrue, L.P., Seung, H.S., and Newsome, W.T. (2005). Linear-Nonlinear-Poisson models of primate choice dynamics. J. Exp. Anal. Behav. *84*, 581–617.

Cox, D.D. (2014). Do we understand high-level vision? Curr. Opin. Neurobiol. *25*, 187–193.

Dakin, S.C., Tibber, M.S., Greenwood, J.A., Kingdom, F.A.A., and Morgan, M.J. (2011). A common visual metric for approximate number and density. Proc. Natl. Acad. Sci. *108*, 19552 LP – 19557.

DiCarlo, J.J., and Cox, D.D. (2007). Untangling invariant object recognition. Trends Cogn. Sci. *11*, 333–341.

DiCarlo, J.J., Zoccolan, D., and Rust, N.C. (2012). How does the brain solve visual object recognition? Neuron *73*, 415–434.

Dosher, B.A., and Lu, Z.L. (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. Proc. Natl. Acad. Sci. U. S. A. *95*, 13988–13993.

Drugowitsch, J., Wyart, V., Devauchelle, A.-D., and Koechlin, E. (2016). Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality. Neuron *92*, 1398–1411.

Froudarakis, E., Cohen, U., Diamantaki, M., Walker, E.Y., Reimer, J., Berens, P., Sompolinsky, H., and Tolias, A.S. (2020). Object manifold geometry across the mouse cortical visual hierarchy. BioRxiv 2020.08.20.258798.

Fründ, I., Wichmann, F.A., and Macke, J.H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. J. Vis. *14*.

Gardner, J.L. (2019). Optimality and heuristics in perceptual neuroscience. Nat. Neurosci. *22*, 514–523.

Girshick, A.R., Landy, M.S., and Simoncelli, E.P. (2011). Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. Nat. Neurosci. *14*, 926–932.

Glickfeld, L.L., Histed, M.H., and Maunsell, J.H.R. (2013). Mouse Primary Visual Cortex Is Used to Detect Both Orientation and Contrast Changes. *33*, 19416–19422.

Goard, M.J., Pho, G.N., Woodson, J., and Sur, M. (2016). Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. Elife *5*, 1–30.

Hernández, A., Salinas, E., García, R., and Romo, R. (1997). Discrimination in the Sense of Flutter: New Psychophysical Measurements in Monkeys. J. Neurosci. *17*, 6391 LP – 6400.

Hoy, J.L., Yavorska, I., Wehr, M., and Niell, C.M. (2016). Vision Drives Accurate Approach Behavior during Prey Capture in Laboratory Mice. Curr. Biol. *26*, 3046–3052.

Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J. Physiol. *160*, 106–154.

Iwasato, T., Datwani, A., Wolf, A.M., Nishiyama, H., Taguchi, Y., Tonegawa, S., Knöpfel, T., Erzurumlu, R.S., and Itohara, S. (2000). Cortex-restricted disruption of NMDAR1 impairs neuronal patterns in the barrel cortex. Nature *406*, 726–731.

Jazayeri, M., and Movshon, J.A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. Nature *446*, 912–915.

Jurjut, O., Georgieva, P., Busse, L., and Katzner, S. (2017). Learning Enhances Sensory Processing in Mouse V1 before Improving Behavior. J. Neurosci. *37*, 6460–6474.

Knill, D.C., and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. Trends Neurosci. *27*, 712–719.

Krechevsky, I. (1938). An experimental investigation of the principle of proximity in the visual perception of the rat. J. Exp. Psychol. *22*, 497–523.

Laquitaine, S., and Gardner, J.L. (2018). A Switching Observer for Human Perceptual Estimation. Neuron *97*, 462-474.e6.

Lashley, K.S. (1938). The Mechanism of Vision: XV. Preliminary Studies of the Rat's Capacity for Detail Vision. J. Gen. Psychol. *18*, 123–193.

Long, M., Jiang, W., Liu, D., and Yao, H. (2015). Contrast-dependent orientation discrimination in the mouse. Sci. Rep. *5*, 15830.

Martinho, A. 3rd, and Kacelnik, A. (2016). Ducklings imprint on the relational concept of "same or different". Science *353*, 286–288.

Odoemene, O., Pisupati, S., Nguyen, H., and Churchland, A.K. (2018). Visual Evidence Accumulation Guides Decision-Making in Unrestrained Mice. J. Neurosci. *38*, 10143 LP – 10155.

Pho, G.N., Goard, M.J., Woodson, J., Crawford, B., and Sur, M. (2018). Task-dependent representations

of stimulus and choice in mouse parietal cortex. Nat. Commun. *9*, 2596.

Pinto, L., Koay, S.A., Engelhard, B., Yoon, A.M., Deverett, B., Thiberge, S.Y., Witten, I.B., Tank, D.W., and Brody, C.D. (2018). An Accumulation-of-Evidence Task Using Visual Pulses for Mice Navigating in Virtual Reality. Front. Behav. Neurosci. *12*, 1–19.

Pinto, L., Rajan, K., DePasquale, B., Thiberge, S.Y., Tank, D.W., and Brody, C.D. (2019). Task-Dependent Changes in the Large-Scale Dynamics and Necessity of Cortical Regions. Neuron *104*, 810-824.e9.

Poort, J., Khan, A.G., Pachitariu, M., Nemri, A., Orsolic, I., Krupic, J., Bauza, M., Sahani, M., Keller, G.B., Mrsic-Flogel, T.D., et al. (2015). Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. Neuron *86*, 1478–1490.

Prusky, G.T., and Douglas, R.M. (2004). Characterization of mouse cortical spatial vision. Vision Res. *44*, 3411–3418.

Renart, A., and Machens, C.K. (2014). Variability in neural activity and behavior. Curr. Opin. Neurobiol. *25*, 211–220.

Resulaj, A., Ruediger, S., Olsen, S.R., and Scanziani, M. (2018). First spikes in visual cortex enable perceptual discrimination. Elife *7*, e34044.

Reuter, J.H. (1987). Tilt discrimination in the mouse. Behav. Brain Res. *24*, 81–84.

Romo, R., Brody, C.D., Hernández, a, and Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. Nature *399*, 470–473.

Runyan, C.A., Piasini, E., Panzeri, S., and Harvey, C.D. (2017). Distinct timescales of population coding across cortex. Nature *548*, 92–96.

Rust, N.C., and Stocker, A.A. (2010). Ambiguity and invariance: two fundamental challenges for visual processing. Curr. Opin. Neurobiol. *20*, 382–388.

Scott, B.B., Constantinople, C.M., Erlich, J.C., Tank, D.W., and Brody, C.D. (2015). Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. Elife *4*, e11308.

Steinmetz, N.A., Zatka-haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of choice , action and engagement across the mouse brain. Nature.

Tafazoli, S., Di Filippo, A., and Zoccolan, D. (2012). Transformation-Tolerant Object Recognition in Rats Revealed by Visual Priming. J. Neurosci. *32*, 21 LP – 34.

Urai, A.E., Braun, A., and Donner, T.H. (2017). Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. Nat. Commun. *8*, 14637.

Wang, L., McAlonan, K., Goldstein, S., Gerfen, C.R., and Krauzlis, R.J. (2020). A Causal Role for Mouse Superior Colliculus in Visual Perceptual Decision-Making. J. Neurosci. *40*, 3768 LP – 3782.

Watson, A.B., Barlow, H.B., and Robson, J.G. (1983). What does the eye see best? Nature *302*, 419–422.

Whiteley, L., and Sahani, M. (2012). Attention in a Bayesian Framework. Front. Hum. Neurosci. *6*, 100.

Wichmann, F.A., and Hill, N.J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. Percept. Psychophys. *63*, 1293–1313.

Wyart, V., and Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. Curr. Opin. Behav. Sci. *11*, 109–115.

You, W.-K., and Mysore, S.P. (2020). Endogenous and exogenous control of visuospatial selective

attention in freely behaving mice. Nat. Commun. *11*, 1986.

Yu, A.J., and Cohen, J.D. (2008). Sequential effects: Superstition or rational behavior? In Advances in Neural Information Processing Systems, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, eds. (Curran Associates, Inc.), pp. 1873–1880.

Zoccolan, D. (2015). Invariant visual object recognition and shape processing in rats. Behav. Brain Res. *285*, 10–33.

Zoccolan, D., Oertelt, N., DiCarlo, J.J., and Cox, D.D. (2009). A rodent model for the study of invariant visual object recognition. Proc. Natl. Acad. Sci. U. S. A. *106*, 8748–8753.

Figures

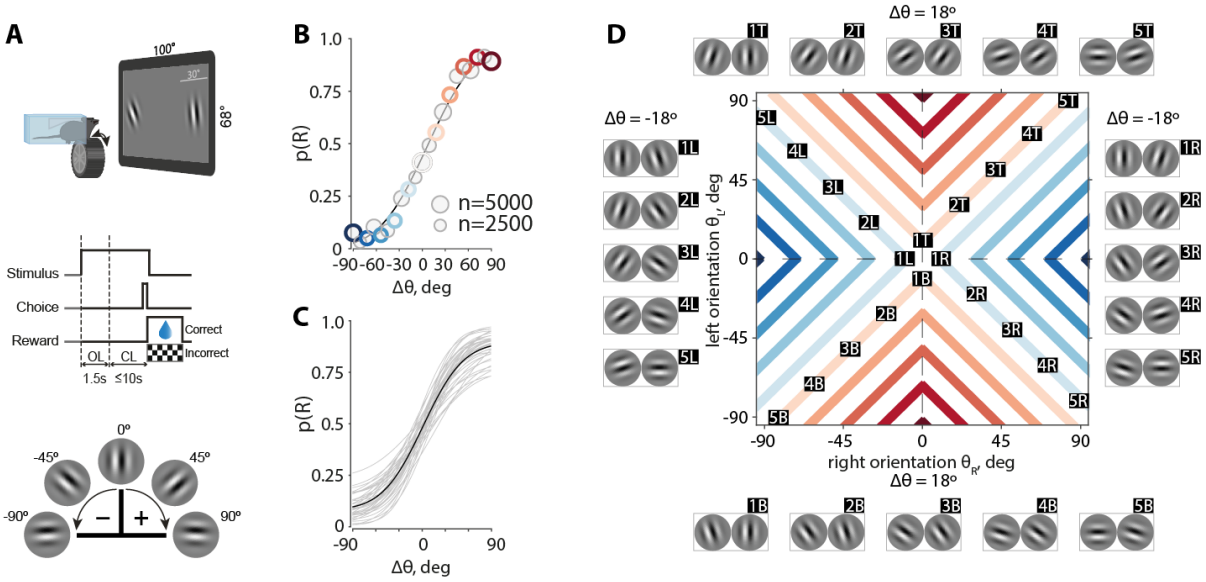

Figure 1. Mice successfully learned a novel invariant orientation discrimination task. **A.** Top. Schematic of a mouse during an experimental session. Middle. Epochs of one trial. Bottom. Convention for the angle signs. **B.** Psychometric curve of an example animal. Solid line – best fit of the cumulative Gaussian psychometric function, circles – data points, circle sizes represent numbers of trials, colors correspond to colors in D, gray circles are data points not explicitly marked in D. **C.** Psychometric curves for all animals in the study, solid black line – population average. **D.** Many orientation pairs give the same task-relevant information quantified by angular separation or difficulty ($\Delta\theta$). Conditions with a fixed $\Delta\theta$ in the 2d stimulus space (colored lines) correspond to $\Delta\theta$ conditions (circles) of the same color in B. Example stimuli for four branches of $\Delta\theta$=const (two branches for 18°, and two for -18°) are displayed along the sides of the stimulus space map. Labels next to the images of orientation pairs correspond to labels on the stimulus space map.

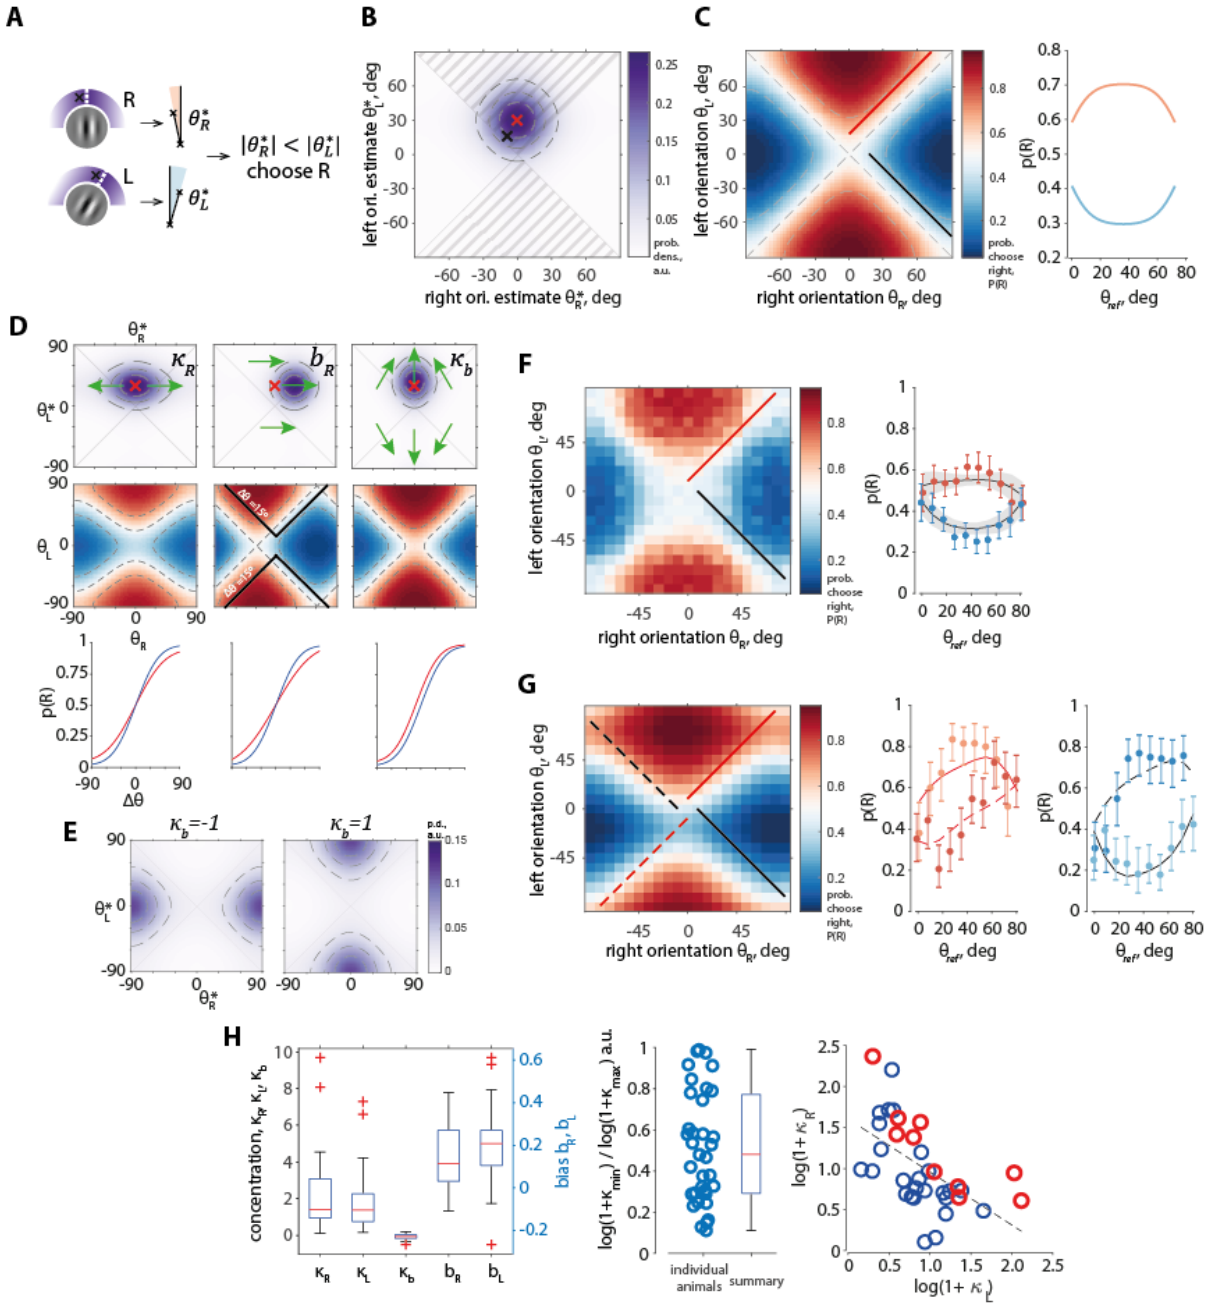

Figure 2. The choice model characterizes individual biases and strategies, and predicts variation of performance with reference orientation, as found in the data. **A.** Choice model schematic. The angles of two oriented Gabor patches (white dashed lines, left column) are estimated as samples from circular distributions (density – in purple, estimates – black crosses), and compared between each other (middle column), which generates a choice. **B**. Distribution $p(x, y)$ of orientation estimates as in a, in 2d space, for $(\theta_L, \theta_R) = (30°, 0°)$ (red cross) in an unbiased model with $(\kappa_R, \kappa_L) = (2,2)$, and a sample from this distribution (black cross). Probability mass inside the shaded areas ($|x| < |y|$) is equal to the probability of right choice P(R). Dashed lines – distribution quartiles. **C**. Left. P(R) of model in B evaluated at all stimulus pairs (θ_L, θ_R). Red and black lines – one example branch of $\Delta\theta = 18°$ and $\Delta\theta = -18°$ respectively. Right. P(R) along the branches of $\Delta\theta = const$ marked on the left panel with a red and black lines. **D**. Effect of model parameters on the distributions of estimates (top row), P(R) surface (middle row), and the corresponding psychometric curves (bottom row); red crosses – distribution means before parameter manipulation, green arrows – transformation of the

distributions with parameter change. Left column: reduction of the concentration κ_R leads to a shallower P(R) along the x axis, and a shallower psychometric curve (blue curve – before κ_R change, red – after). Middle column: non-zero bias b_R results in a translation of the estimate distribution and of P(R), and to a shallower psychometric curve since pooling for the psychometric curve is determined by $\Delta\theta$. Center panel: P(R) values are displaced relative to $\Delta\theta$ isolines (solid black: example isoline for $\Delta\theta$=15°), resulting in stimulus pairs with predominantly left choice (P(R)<0.5, blue) to be pooled alongside P(R)>0.5 values. Right column: choice bias $\kappa_b > 0$ leads to an increased (decreased for $\kappa_b < 0$) probability mass in the $|x| < |y|$ region, an expanded P(R)>0.5 area of the P(R) surface, and a psychometric curve shifted to the left, in favor of the right choices. **E.** Choice priors $p_b(x, y)$ with κ_b equal to -1 (left panel; left choice bias) and 1 (right panel; right choice bias). **F**. Left. population average P(R) (n=40 mice) with one example branch of $\Delta\theta = 9°$ and $\Delta\theta = -9°$ marked with red and black lines. Right: P(R) values as on the left panel (dots with error bars, mean ± c.i.), and average of model predictions (black lines with shaded areas, mean ± c.i.) across all animals. See Figure S3A for P(R) of every animal. **G**. Example mouse, left: P(R) of the fitted model, middle: P(R) along the red dashed and solid lines on the left panel predicted by the model (lines) and computed from the data (dots with error bars, darker dots correspond to the dashed line), right: P(R) along the black dashed and solid lines on the left panel, as predicted by the model (lines) and computed from the data (dots with error bars, darker dots correspond to the dashed line). See Figure S4D for the model of P(R) for every animal. **H**. Left: population summary of model parameters fitted to all mice (n=35; n=5 animals with κ_R or κ_L estimated on the edge of the allowed range of values are excluded). Middle: ratio of $\log(\kappa_R + 1)$ and $\log(\kappa_L + 1)$ with the smaller of the two values divided by the larger value for each mouse (n=35). Circles – individual animals. Box plot – population summary: red line – median value, box borders – 25th and 75th percentiles, whiskers are up to most extreme parameter values, red crosses – outliers. Animals with ratios close to one use right and left orientation information in their choices equally. Right: $\log(\kappa_R + 1)$ and $\log(\kappa_L + 1)$ across the population are significantly anti-correlated. Linear regression line for all animals together (R^2=0.32,p=4.5·10^{-4}, significant at α=0.005). Red circles – 10 mice with best performance; the best performing mice have significantly higher concentrations than the rest of the population (p<0.05, ANCOVA, F-test of intercept with fixed slope).

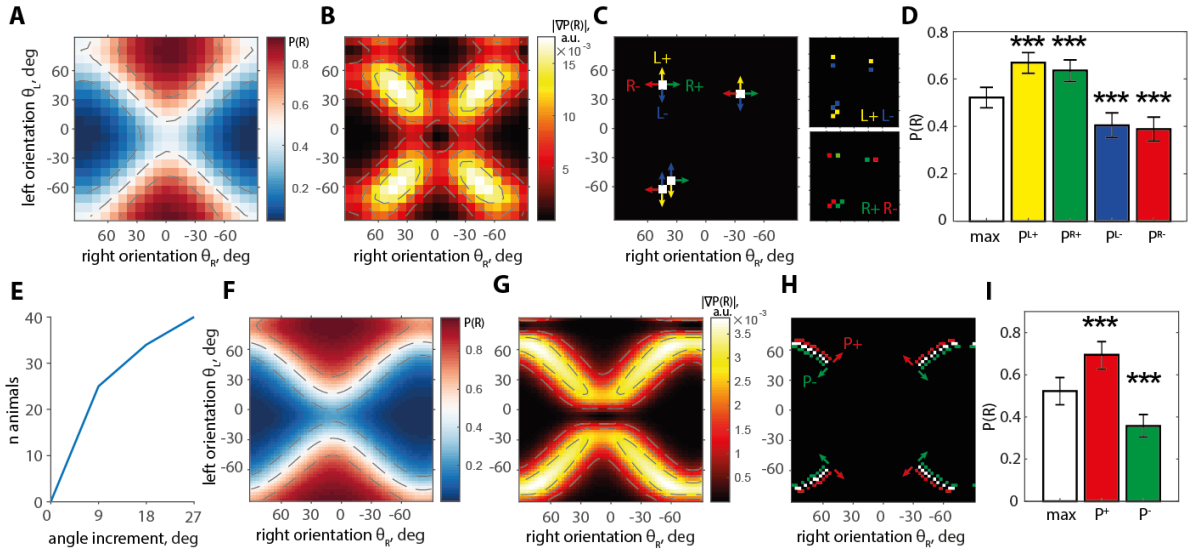

Figure 3. Mice can reach an orientation acuity of 6°. **A**. P(R) model surface of an example mouse. **B**. |∇P(R)| absolute value of the gradient of P(R) surface in A. **C**. Four stimulus conditions (white) with |∇P(R)| in the top 5% of values and P(R) close to 0.5 (0.48<P(R)<0.52). Arrows of the same color show angle change yielding neighboring conditions: P^{L+} (yellow; "L+" for left stim. change that increases P(R)), P^{R+} (green), P^{L-} (red), P^{R-} (blue). Insets (right) show these conditions. **D**. Pooled P(R) in maximum gradient conditions (white), $P^{0.5}$=0.52±0.04, differed from P(R) in the four neighboring conditions P^{L+}=0.67±0.04, $p<2.5\cdot10^{-4}$ (yellow), P^{R+}=0.63±0.05, $p<2.5\cdot10^{-3}$ (green), P^{L-}=0.40±0.05, $p<2.5\cdot10^{-3}$ (red), P^{R-}=0.39±0.05, $p<2.5\cdot10^{-4}$ (blue) (binomial confidence intervals, χ^2 test with n=4 comparisons). **E.** Cumulative number of animals for which at least one direction of angle change gives a P(R) significantly different from $P^{0.5}$, as a function of angle change. **F,G**. Similar to A,B for an animal trained with 3° angle binning. **H.** Maximum gradient conditions (white, same criteria as in C), and neighboring conditions obtained by changing both angles by ±3° in the direction of P(R) increase (P^+, yellow), and decrease (P^-, blue). **I.** Pooled P(R) in three groups highlighted in H: $P^{0.5}$=0.52±0.04, P^{L+}=0.69±0.07 (yellow), P^-=0.36±0.05 (blue), both different from $P^{0.5}$ with p<0.0005 (n=2 comparisons).

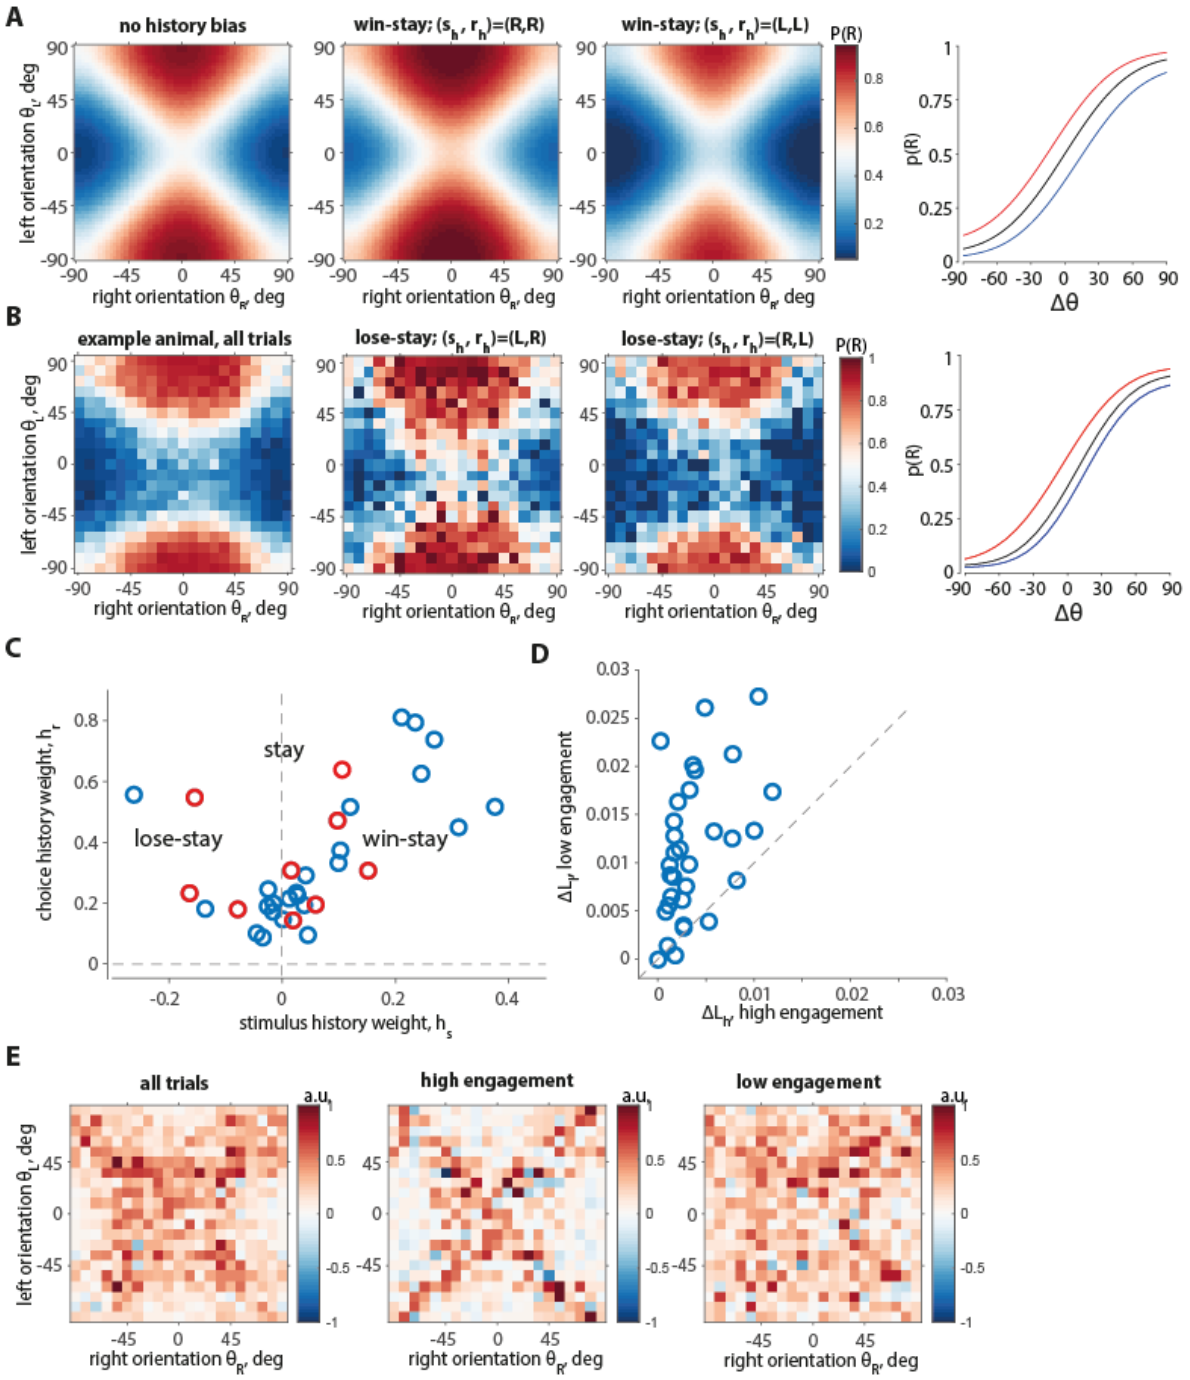

Figure 4. The probability of choice is affected by the choice and reward on the previous trial, with a larger effect during periods of lower task engagement. **A**. P(R) of an unbiased example model $(\kappa_R, \kappa_L, \kappa_b, b_R, b_L) = (1.5, 1.5, 0, 0, 0)$ with a win-stay strategy $(h_s, h_r, \kappa_p) = (0.4, 0.4, 0.5)$; from left to right: [1] without history bias (after a 'neutral' trial, $(s_h, r_h) = (0,0)$), [2] after a successful right choice $(s_h, r_h) = (1,1)$ with P(R) biased to the right choices as a result, [3] after a successful left choice $(s_h, r_h) = (-1,-1)$ with P(R) biased to the left choices, [4] psychometric curves corresponding to [1-3]: without a history effect (black), after a correct right choice (red), after a correct left choice (blue). **B**. P(R) of an example animal with large history biases $(h_s, h_r) = (-0.26, 0.55)$ (lose-stay strategy); from left to right: [1] average P(R) on all trials, [2] P(R) after an unsuccessful right choice $(s_h, r_h) = (-1,1)$ biased to right choices, [3] P(R) after an unsuccessful left choice $(s_h, r_h) = (1,-1)$ biased to

left choices; P(R) of all conditions in [2] and [3] with fewer than N=5 trials was set to the average P(R) of its neighbors; [2] and [3] are conditioned on preceding errors and thus had a relatively low number of trials since the performance of all mice exceeded a 65% success rate, [4] psychometric curves of the corresponding surfaces, colors as in A. **C**. History weights (h_s, h_r) and corresponding strategies of all animals (4 outliers not shown); "stay" and "win-stay" strategies are predominant, with few examples of "lose-stay". Blue circles – animals trained to detect a more vertical orientation, red circles - more horizontal orientation. **D**. Increase in the likelihood due to inclusion of history terms is larger for low-engagement trials than high-engagement trials (Wilcoxon test, p=2.12e-07). Abscissa – difference (ΔL_h) between average log-likelihood under the model with history (p_h) and without history (p_0) of high-engagement trial outcomes (r^h), $\Delta L_h = \langle L(r^h, p_h)\rangle - \langle L(r^h, p_0)\rangle$, ordinate – difference (ΔL_l) between average log-likelihood under the model with and without history of the low-engagement trial outcomes (y_l), $\Delta L_l = \langle L(r^l, p_h)\rangle - \langle L(r^l, p_0)\rangle$. **E**. Increase in the likelihood due to inclusion of history terms changes with stimulus conditions and engagement modes; left to right: [1] ΔL_θ – average likelihood difference for every stimulus condition, $\Delta L_\theta = \langle L(r, p_h)\rangle_\theta - \langle L(r, p_0)\rangle_\theta$, all trials are taken, maps are Z-scored and averaged across animals, conditions with fewer than 10 trials are excluded, [2] $\Delta L_{h\theta}$ – same value computed for high-engagement trials only, [3] $\Delta L_{l\theta}$ – same value computed for low-engagement trials only.

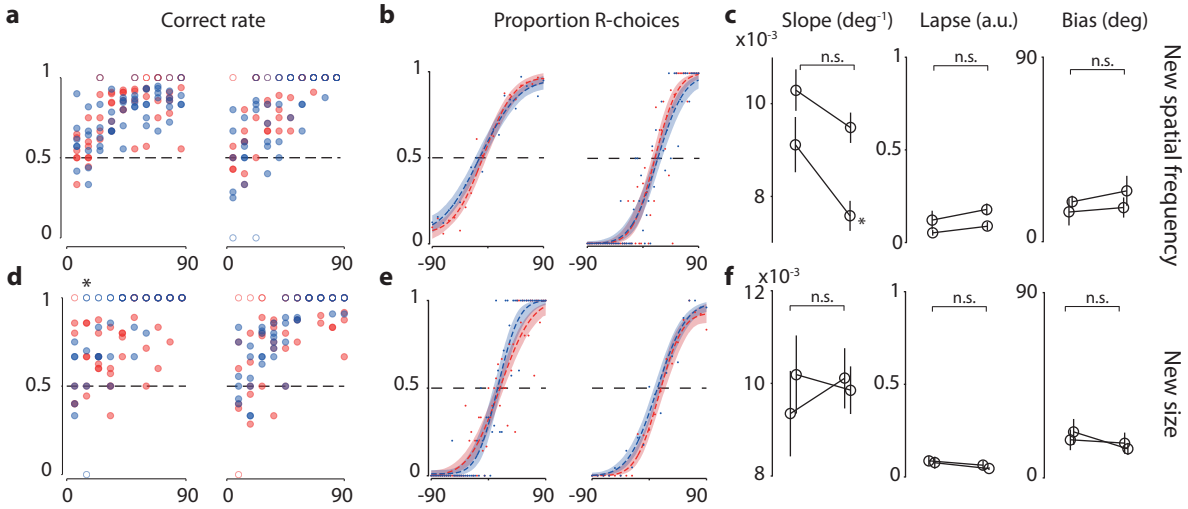

Figure S1. Performance is invariant relative to various stimulus transformations. **a.** Correct rate as a function of trial easiness |Δθ|, comparing 3 sessions (divided into 6 groups of trials, with ~75 trials/group) before a change in spatial frequency of the gratings (red dots) and after the change (blue dots) (spatial frequency, SF=0.008 -> 0.016 cpd for mouse ID: 39, left panel, and SF=0.0016 -> 0.032 cpd for mouse ID:15098, right panel). Open circles for correct rates {0, 1}. Data for the right panel (mouse ID: 15098, angle sampling step 3°) has been grouped into 9° bins to improve visualization. For statistical comparison, we compared binned data (non-overlapping 18° bins) from before vs after conditions and found no significant difference ($p>0.05$, Wilcoxon rank-sum test). **b.** Psychometric curves form 3 sessions before (red) and after (blue) changing the spatial frequency of stimuli. Filled dots for the data; dotted-line for the fits; colored bands for bootstrap confidence intervals. **c.** Comparison of fitting parameters: slope, lapse rate, and bias, before and after changing stimuli spatial frequency (mean ± s.e.m., n=2 mice, n.s. for $p>0.05$, and '*' for $p<0.05$, unpaired t-test for individual animals, paired t-test for comparison across animals). **d-f.** Same as a-c, but for changes in stimulus size (20° -> 25° visual angle, n = 2 mice, ID: 15098, 15100). Data sampled at 3° angle difference has been grouped into 9° bins to improve visualization. Left panels: statistical difference for |Δθ| bin = 18° ('*' for $p<0.05$) reflects an improvement in the performance after changing stimulus size.

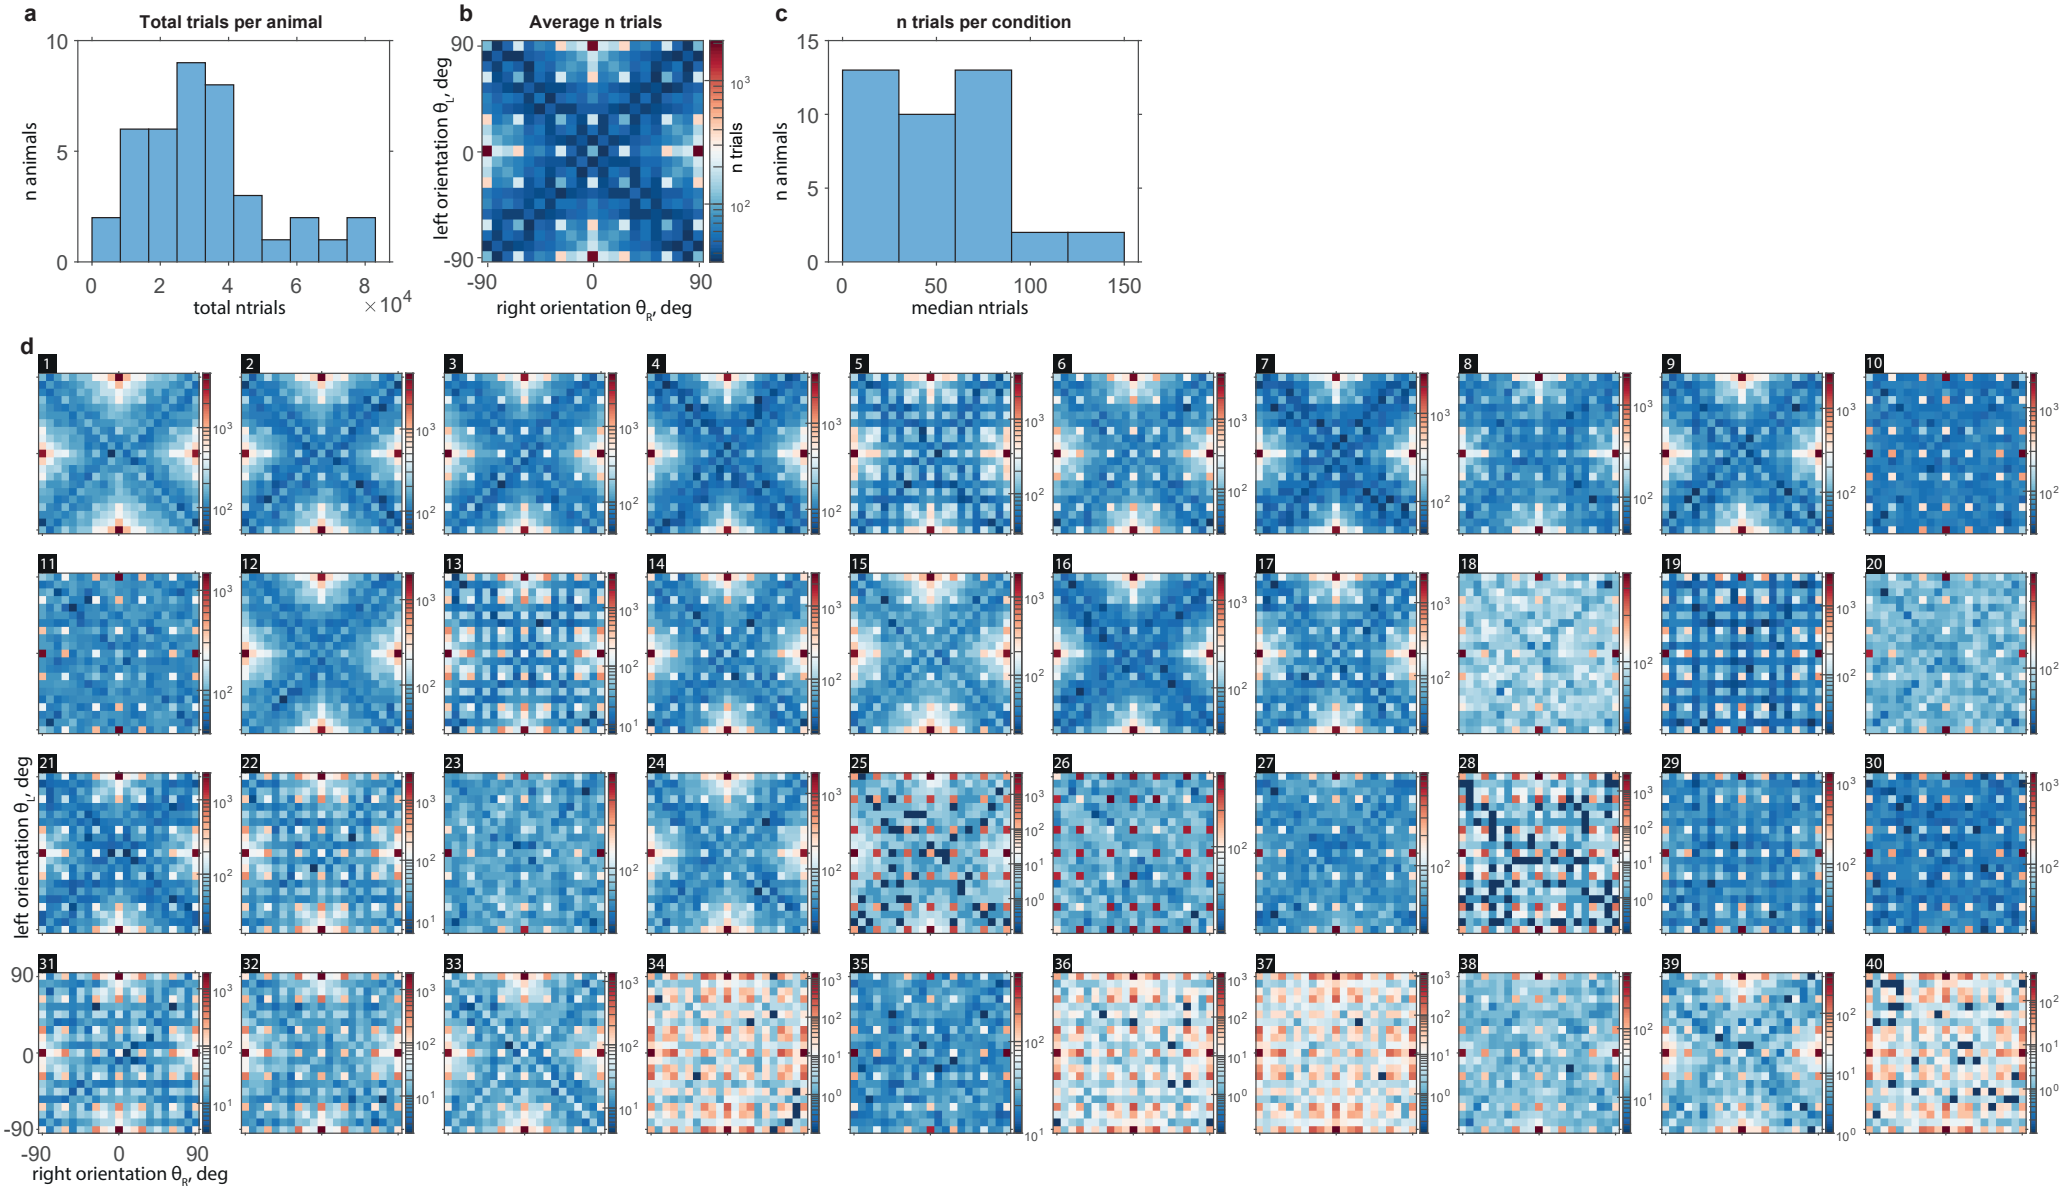

Figure S2. A high number of trials was collected from n=40 animals. **a.** Total number of trials collected for all animals. **b.** Population-average number of trials for every stimulus condition (pair of angles); color bar – number of trials, log scale. **c.** Median number of trials across conditions for every animal. **d.** Number of trials for every stimulus condition and every animal, axes as in b; number in black square - animal ID, the same as in Supplementary Figure 3 and Table 1.

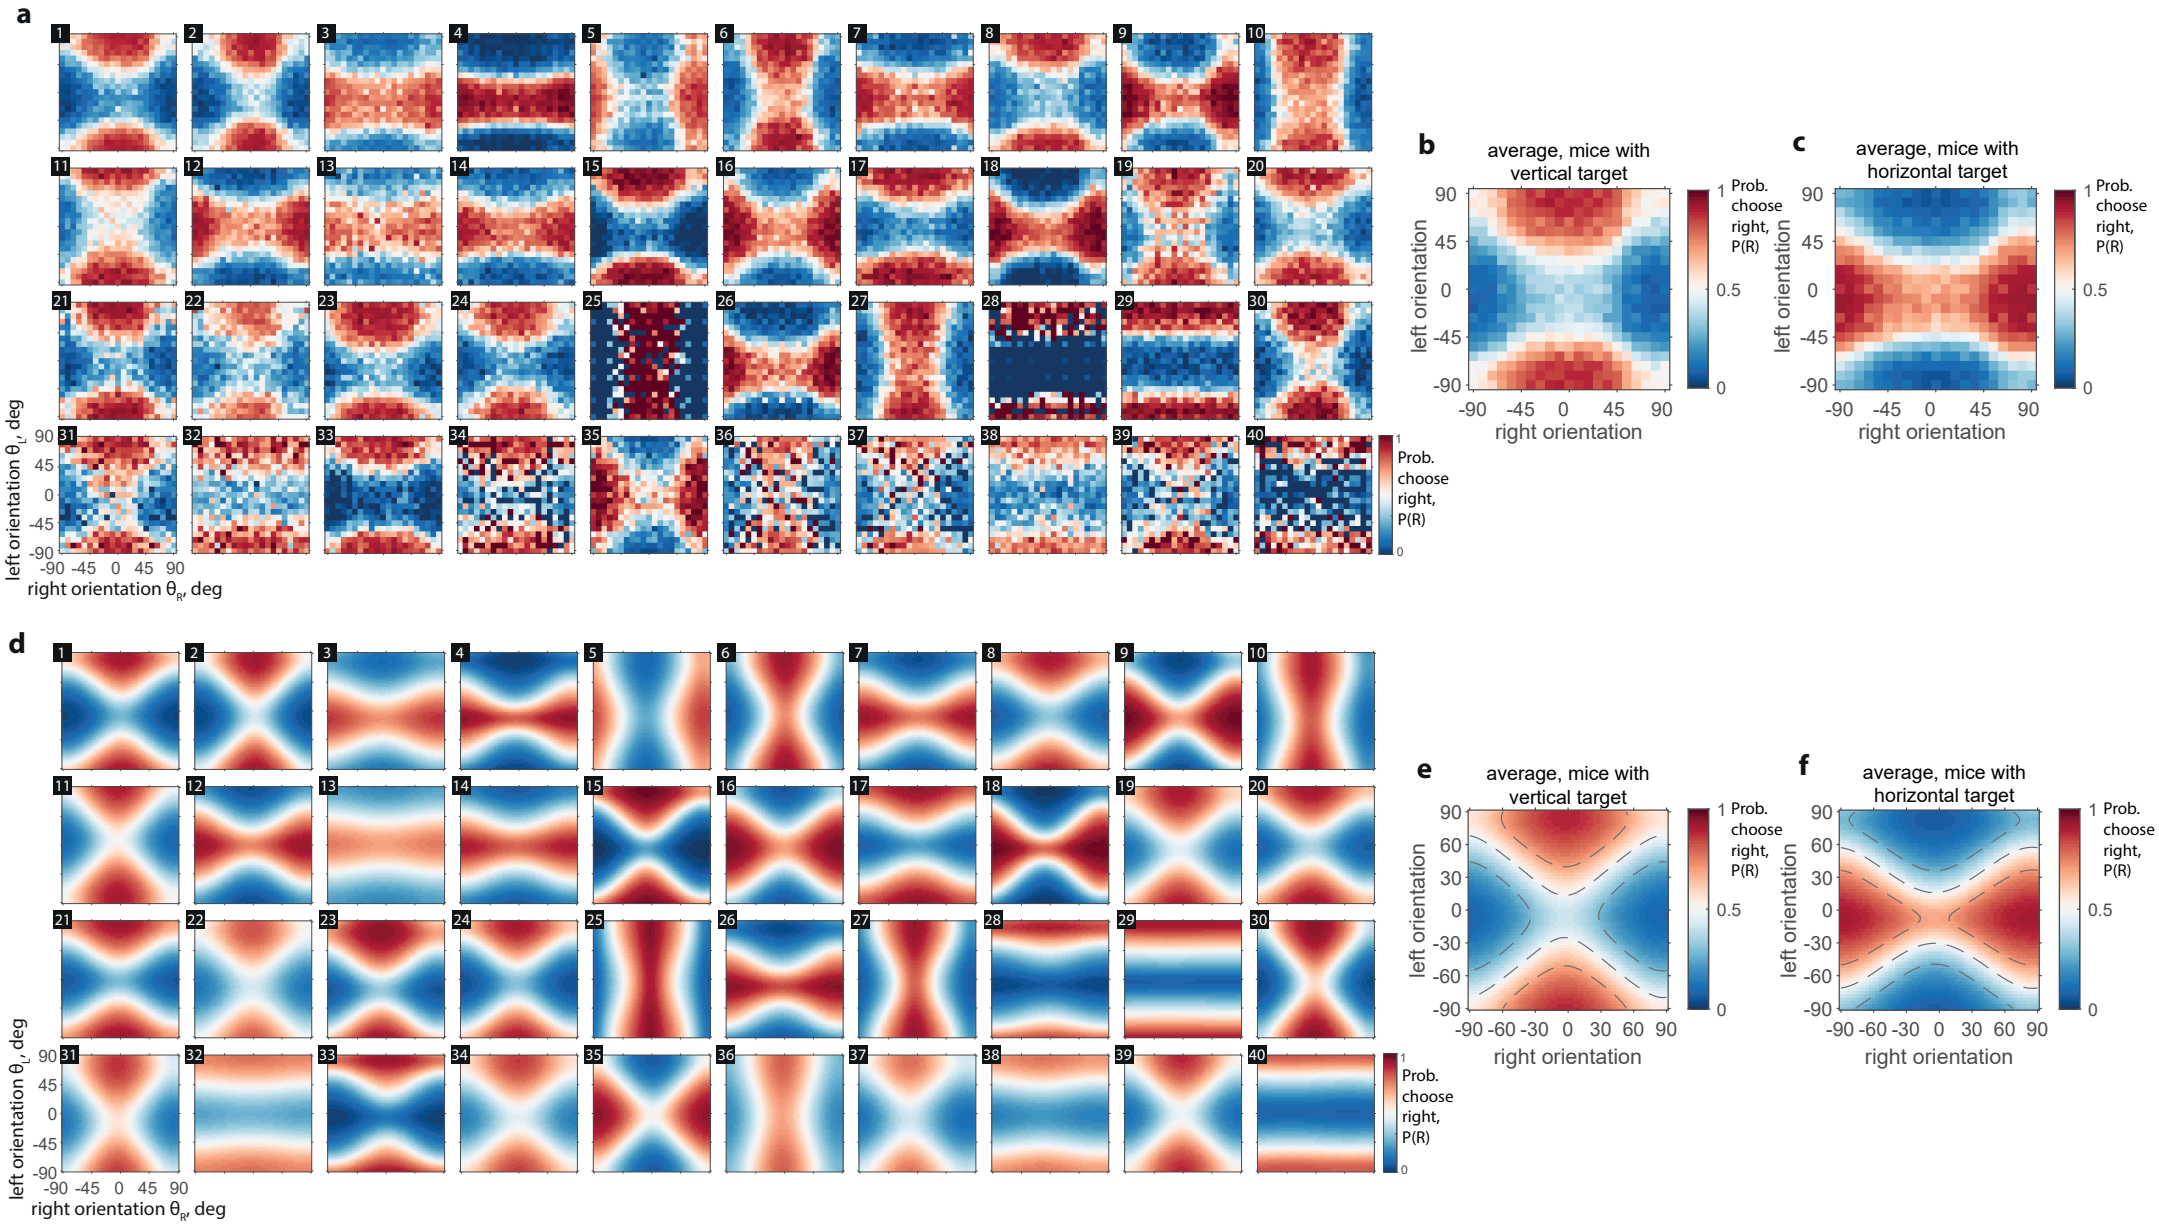

Figure S3. Choices of mice are largely determined by the rewarded side in the two-dimensional stimulus space, and choice model recapitulates choice probability. **a.** Probability of right choice, P(R), for all mice. Stimulus conditions are binned to 9°. Color limits are the same in all panels and in b-c. Animal IDs (number in a black square) are as in Supplementary Figure 2 and Table 1. **b.** Average P(R) across animals trained to find a more vertical orientation. **c.** Average P(R) across animals trained to find a more horizontal orientation. **d.** Model P(R) surfaces for every animal, same color bar on all panels, and as in e and f. **e.** Average P(R) surface of all animals trained to find the more vertical target. Dashed lines at P(R) values of 0.25, 0.5, 0.75. **f.** Average P(R) surface of all animals trained to find the more horizontal target.

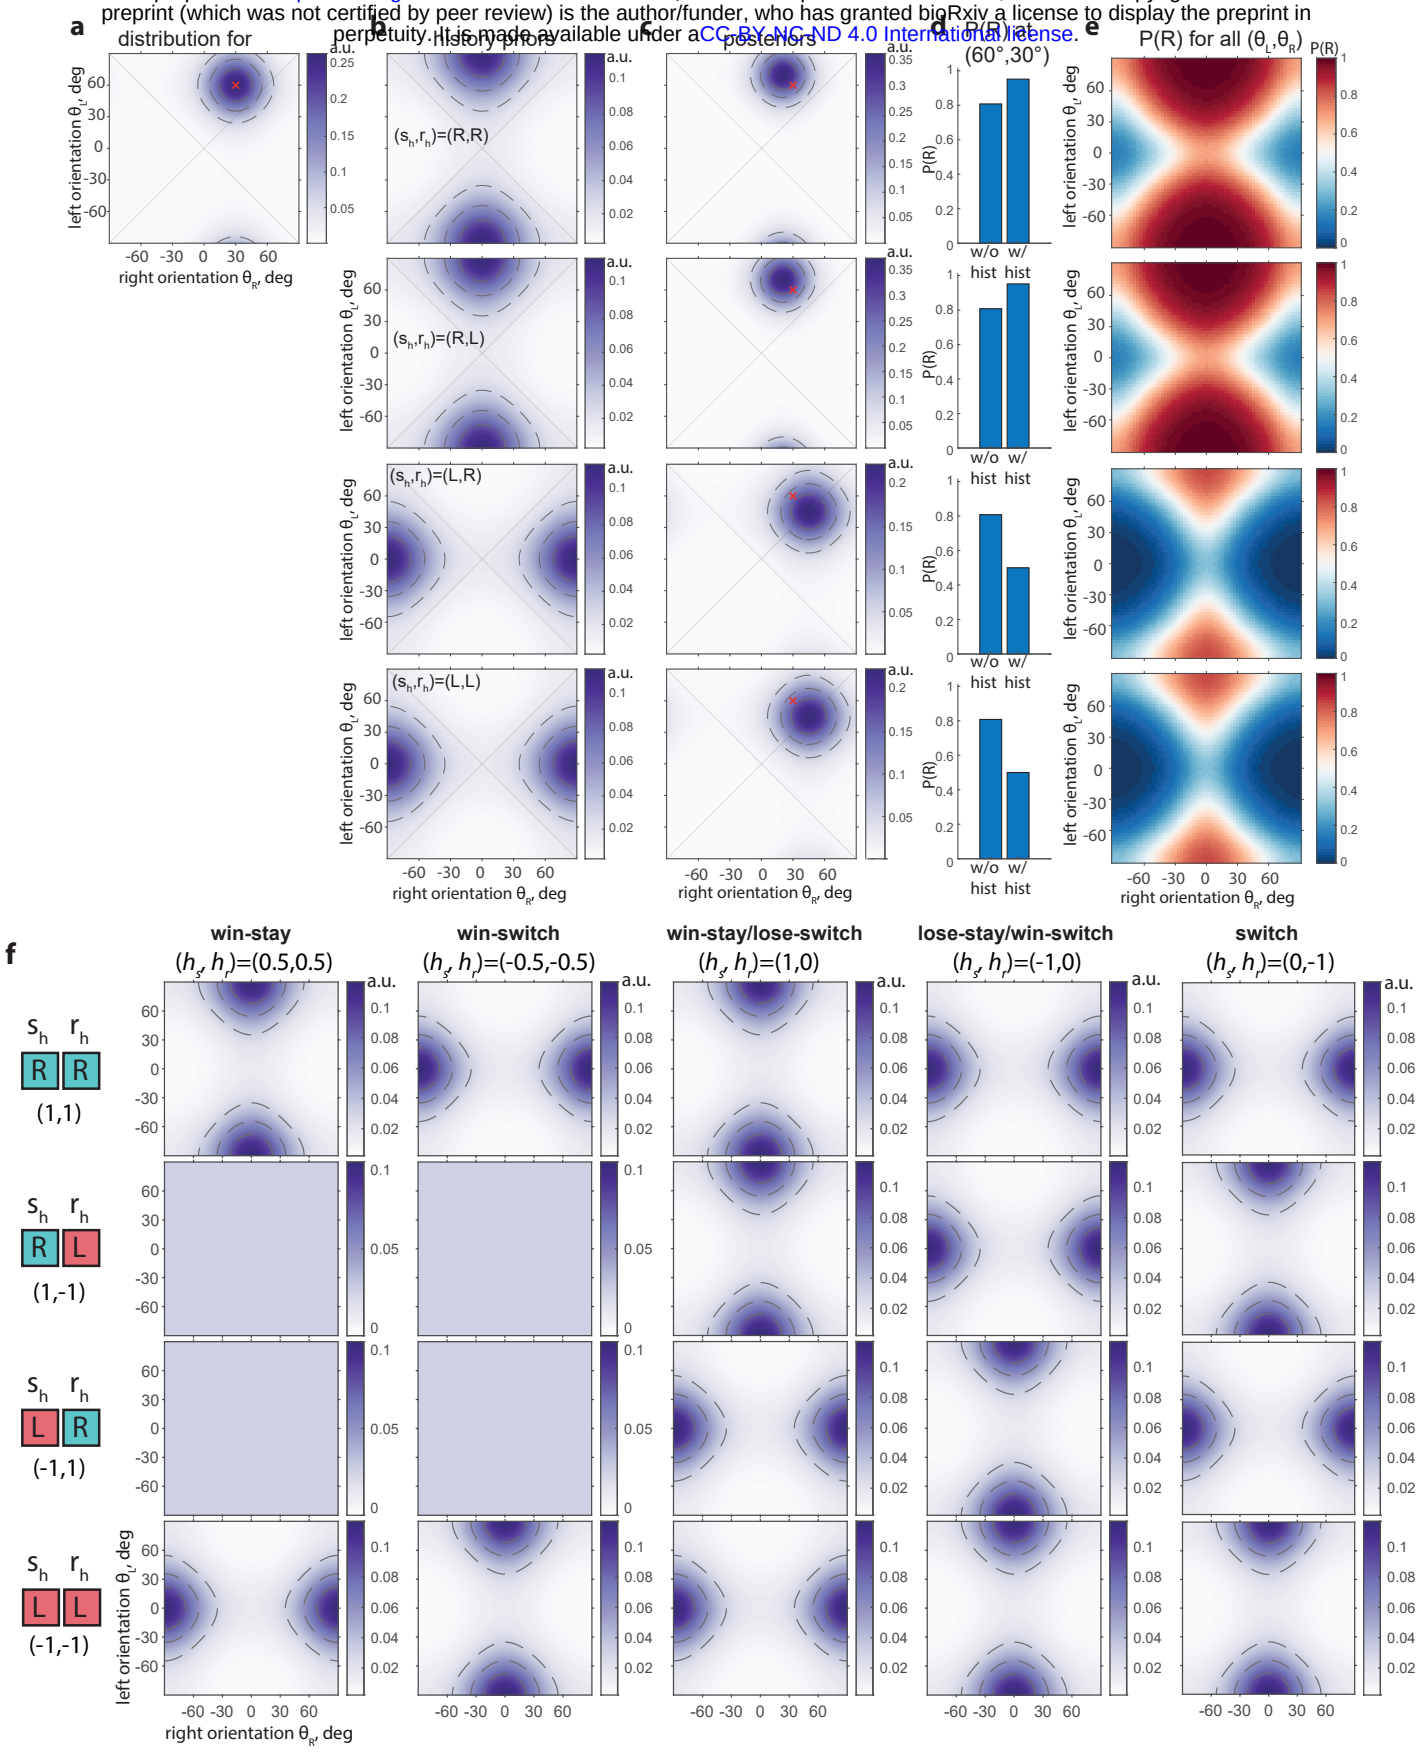

Figure S4. History priors can represent many possible strategies, and affect choice probability , P(R), by shifting the probability density p(x,y) inside or outside the |x|<|y| region **a.** Probability density (p.d., shown by color saturation) p(x,y) [Eq.1] induced by stimuli (θ_R,θ_L)=(30°,60°) (red cross) in a model with κ_R=κ_L=2, b_R=b_L=0, and κ_b=0; dashed lines show distribution quartiles. **b.** History prior p_h(x,y) (Methods, Eq. 3) corresponding to the win-stay/lose-switch strategy, (h_s,h_r)=(0,1), with κ_h=5, and four possible target-response combinations (s_h,r_h) on the previous trial. Top to bottom: (s_h,r_h) = (R,R); (R,L); (L,R); (L,L). **c.** Posterior p.d.: normalized product of p(x,y) and p_h(x,y) before integration over |x|<|y|, with (s_h,r_h) same as in b in the same row. **d.** probability of right choice P(R) for (θ_R,θ_L)=(30°,60°) with and without history bias. **e.** P(R) with strategy for all (θ_R,θ_L) corresponding to (s_h,r_h) in b in the same row. **f.** History priors p_h(x,y) for five example history-based strategies (columns) shown for all four possible combinations of target and choice on the previous trial (rows); κ_h=1 in all cases.

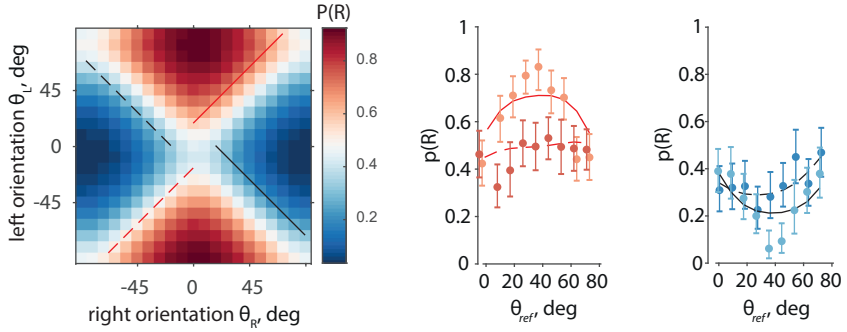

Figure S5. Variation of P(R) with reference orientation θ_{ref} is larger in the data than in the model. Left. Example mouse, selected here for its low translational bias and approximately equal concentrations for right and left stimuli, which results in a regularly shaped P(R) dependency on θ_{ref} (cf. Figure 2c). Center and right. P(R) along θ_{ref} for the $\Delta\theta$=const conditions marked on the left panel, as predicted by the model (lines) and as in the data (dots with whiskers). Dots of a lighter shade (orange, light blue) correspond to the solid lines.