Supplementary Methods

Experimental Model and Subject Details: Genotype and Surgery

We used the following genotypes: Thy1-GCaMP6f (n = 37), Camk2-tTA TRE-GCaMP6s (n = 2), Emx1-tTA TRE-GCaMP6s (n = 1). 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.

During surgery, animals were anesthetized with gas anesthesia (Isoflurane 1.5-2.5%; Pfizer) and injected with an antibiotic (Baytril®, 0.5 ml, 2%; Bayer Yakuhin), a steroidal anti-inflammatory drug (Dexamethasone; Kyoritsu Seiyaku), an anti-edema agent (Glyceol®, 100 µl, 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 4 mm coverslip (120~170 µm 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 4 mm coverslip to stop any bleeding. The 6 mm coverslip was fixed to the bone with cyanoacrylic glue (Aron Alpha®, Toagosei). A round metal chamber (6.1 mm 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.

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 (**Fig. 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°.

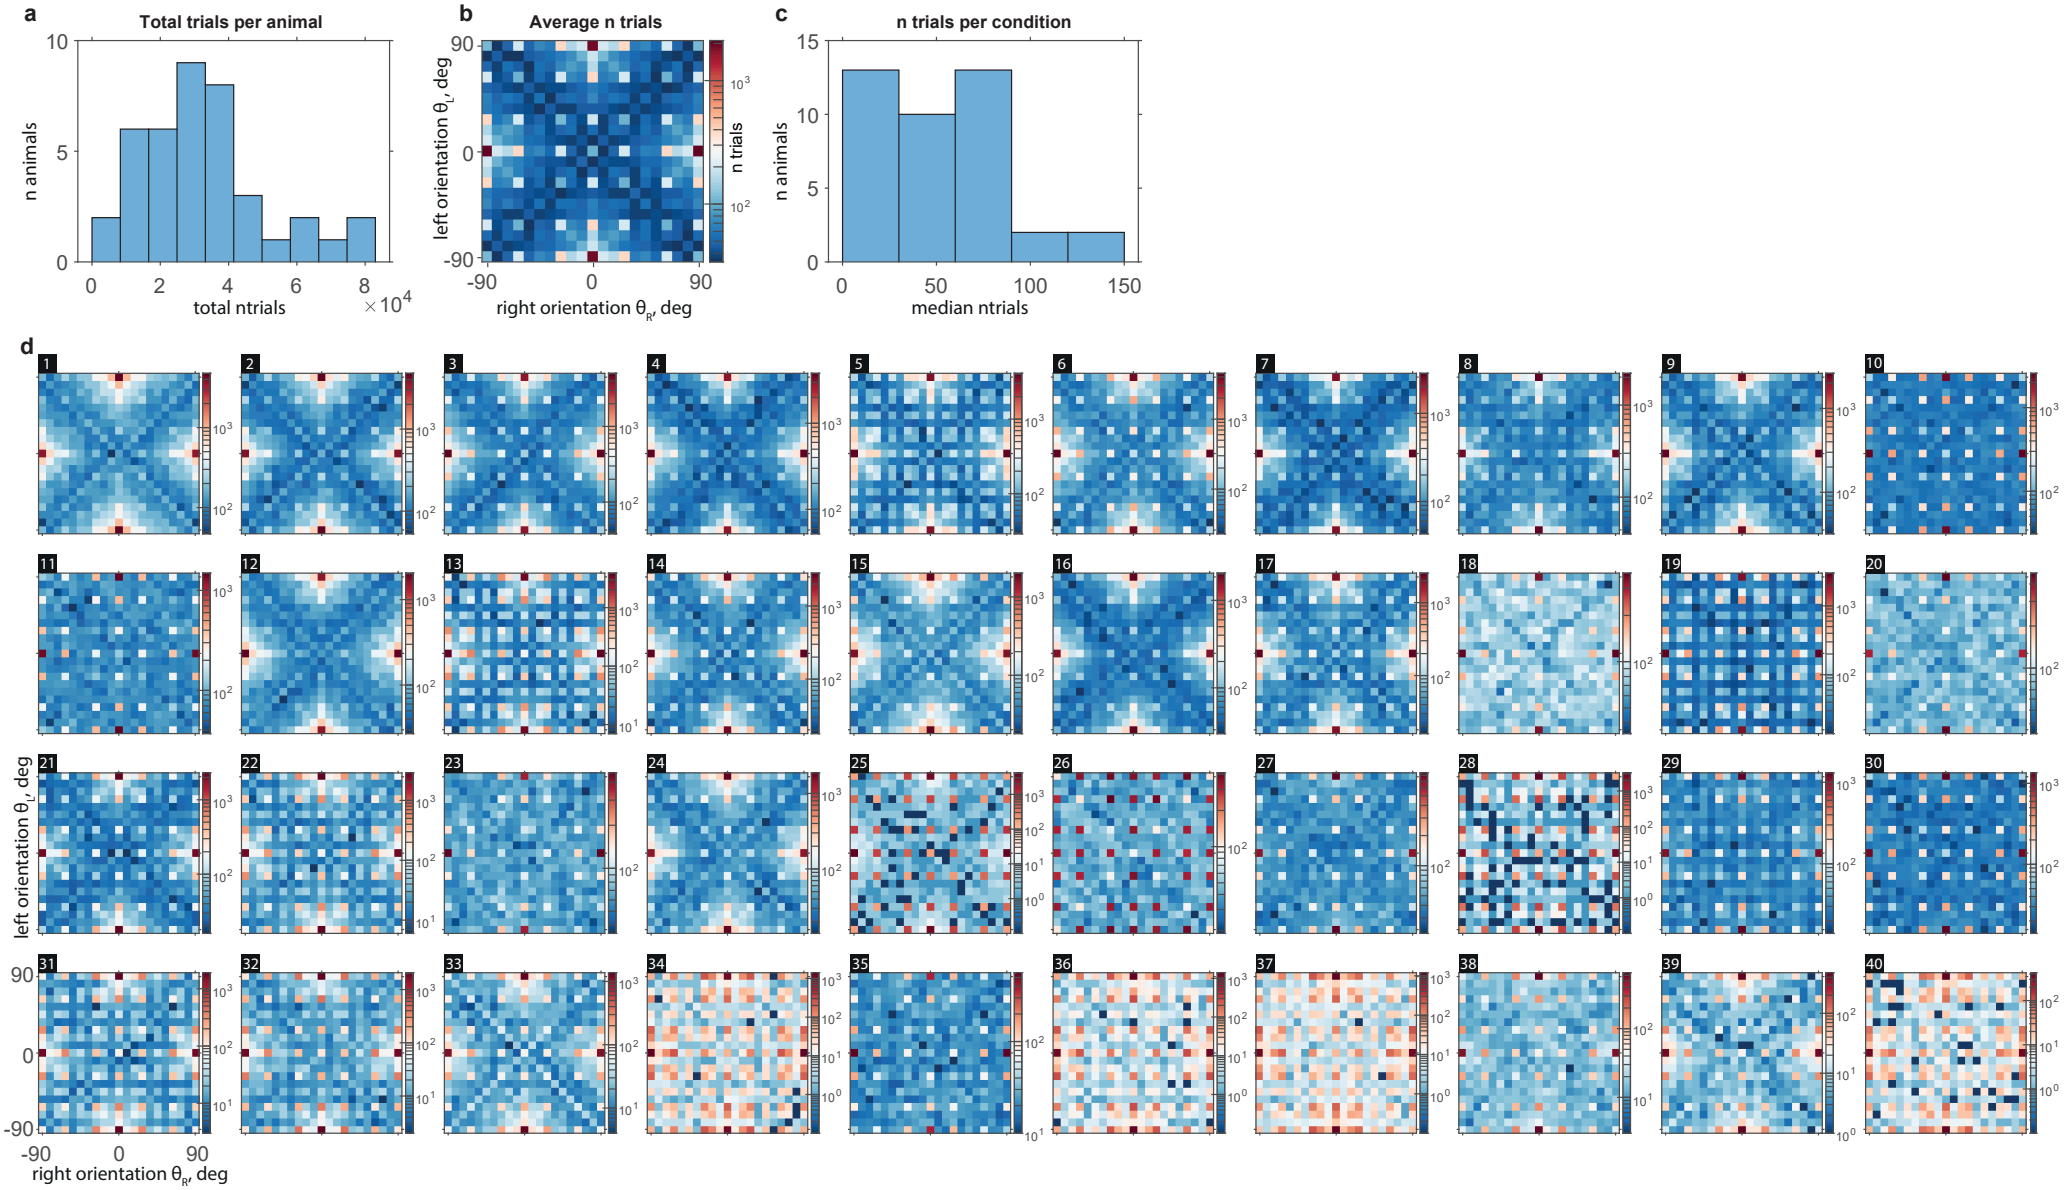

Supplementary Figure 1. A high number of trials was collected from n = 40 animals. **a.** Total number of trials collected for all animals. **b.** Population-average number of trials for every stimulus condition (pair of angles); color bar – number of trials, log scale. **c.** Median number of trials 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 Supplementary Table 1.

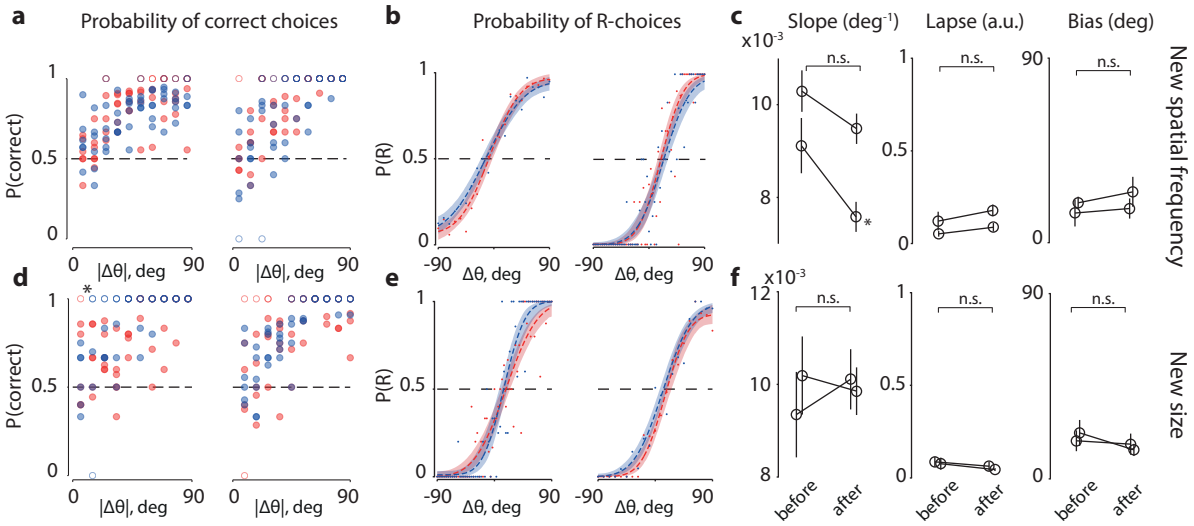

Supplementary Figure 2. Performance is invariant relative to size and spatial frequency transformations of the stimulus. **a.** Correct rate as a function of trial easiness |Δθ|, comparing 3 sessions (divided into 6 groups of trials—one dot per group, dots may overlap—with ~75 trials/group) before a 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 A, left panel, and SF = 0.0016 -> 0.032 cpd for mouse B, right panel). Open circles for correct rates {0, 1}. Data for the right panel (mouse B, minimal angular difference 3°) was 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 from 3 sessions before (red) and after (blue) changing the spatial frequency of stimuli. Same data as in a, used as right/left choices; dots for average P(R) as in the data; dotted lines for the fits; colored bands for bootstrap confidence intervals. **c.** Comparison of fitting parameters: slope, lapse rate, and bias, before and after changing spatial frequency of stimuli (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 B, C). 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.

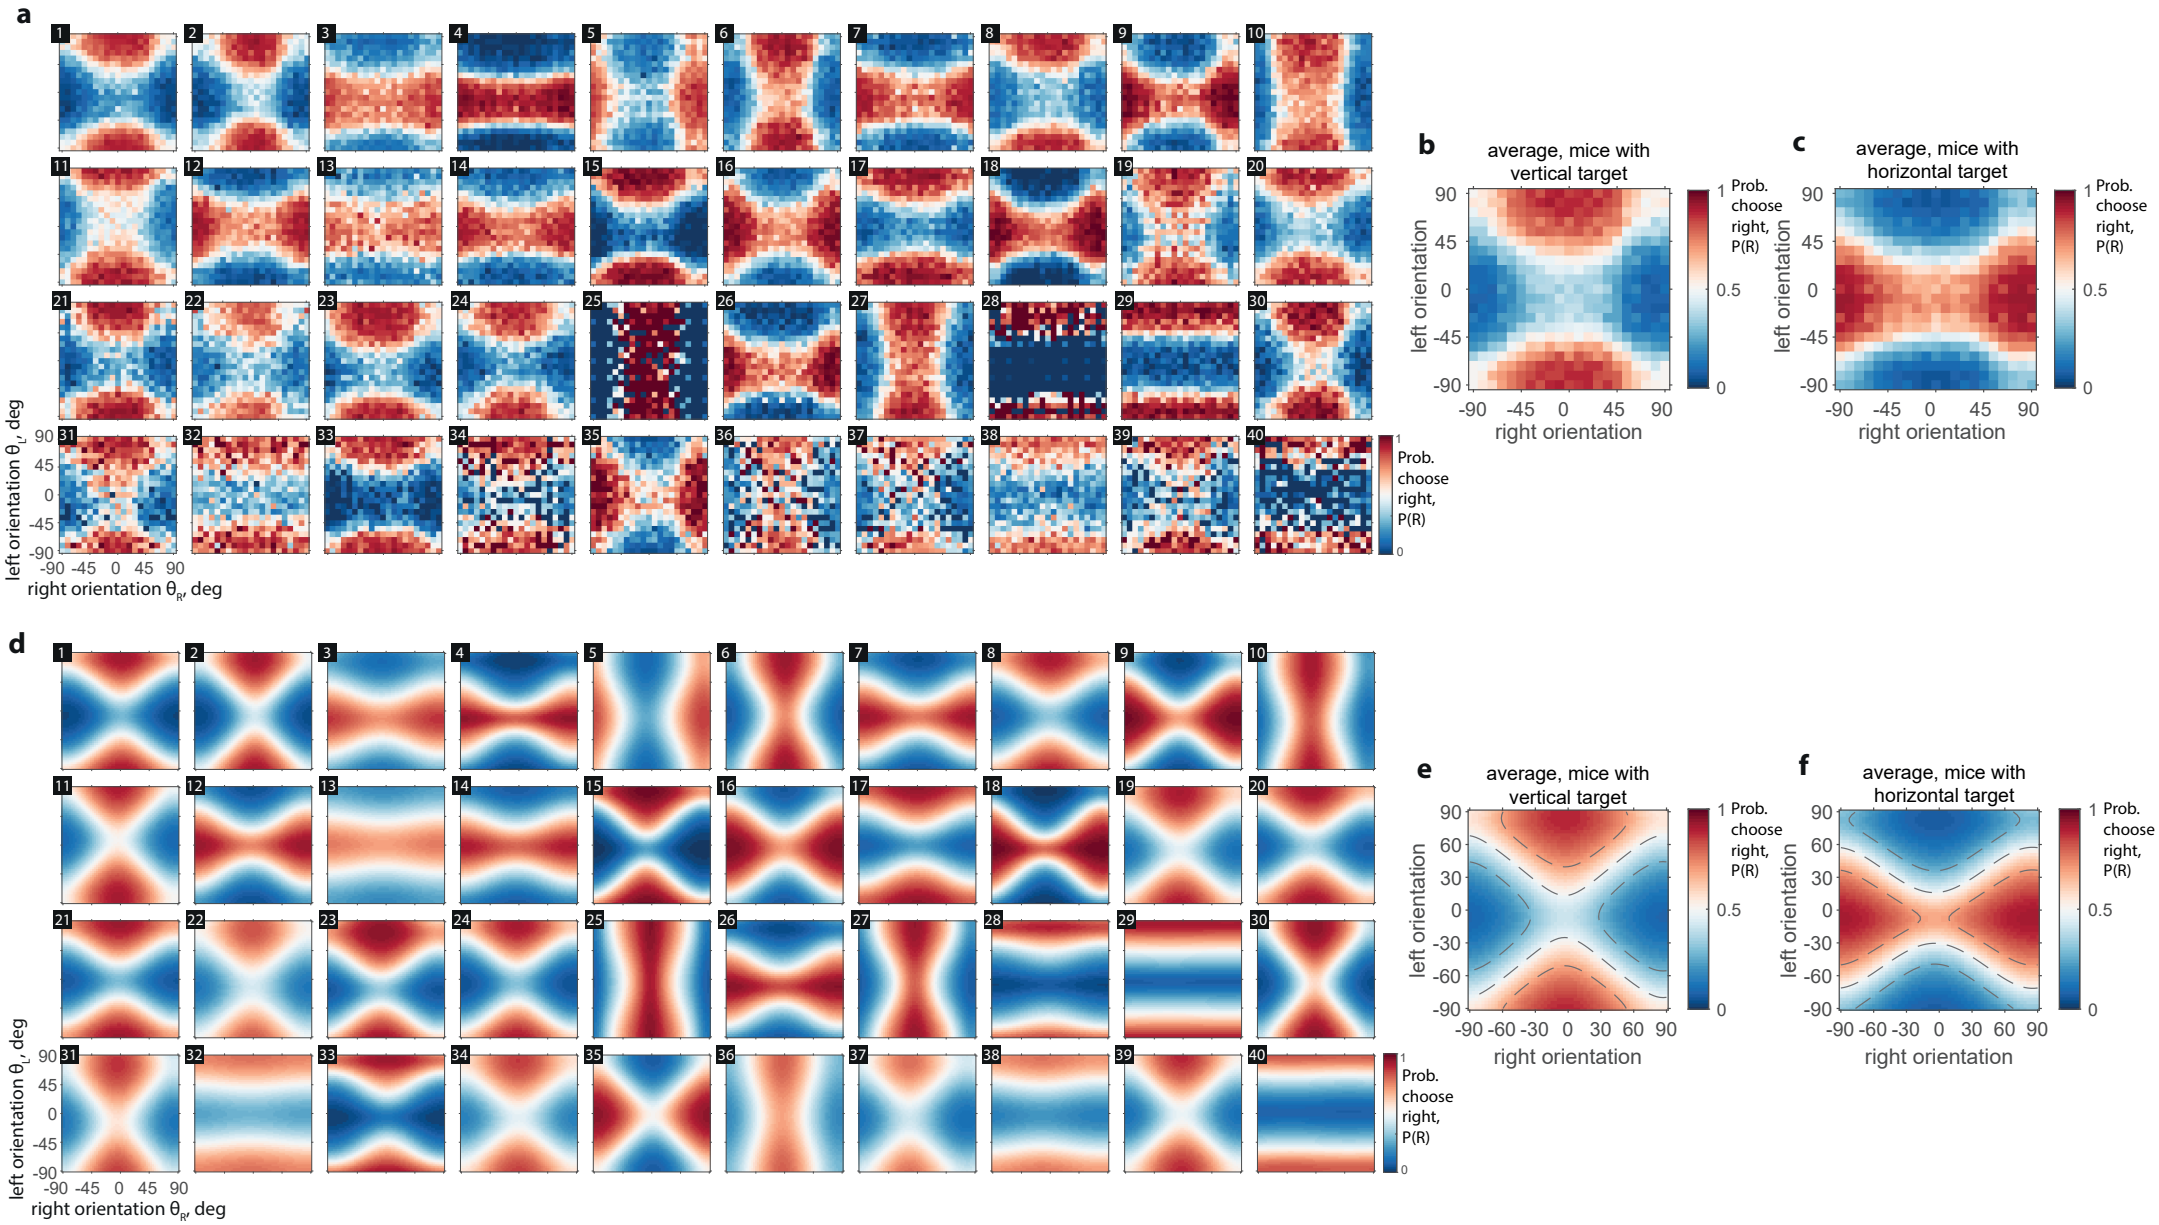

Supplementary Figure 3. Choices of mice are largely determined by the rewarded side in the two-dimensional stimulus space, and choice model recapitulates choice probability. **a.** Probability of right choice, P(R), for all mice. Stimulus conditions are binned to 9°. Color limits are the same in all panels and in b-c. Animal IDs (number in a black square) are as in Supplementary Figure 1 and Supplementary 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.

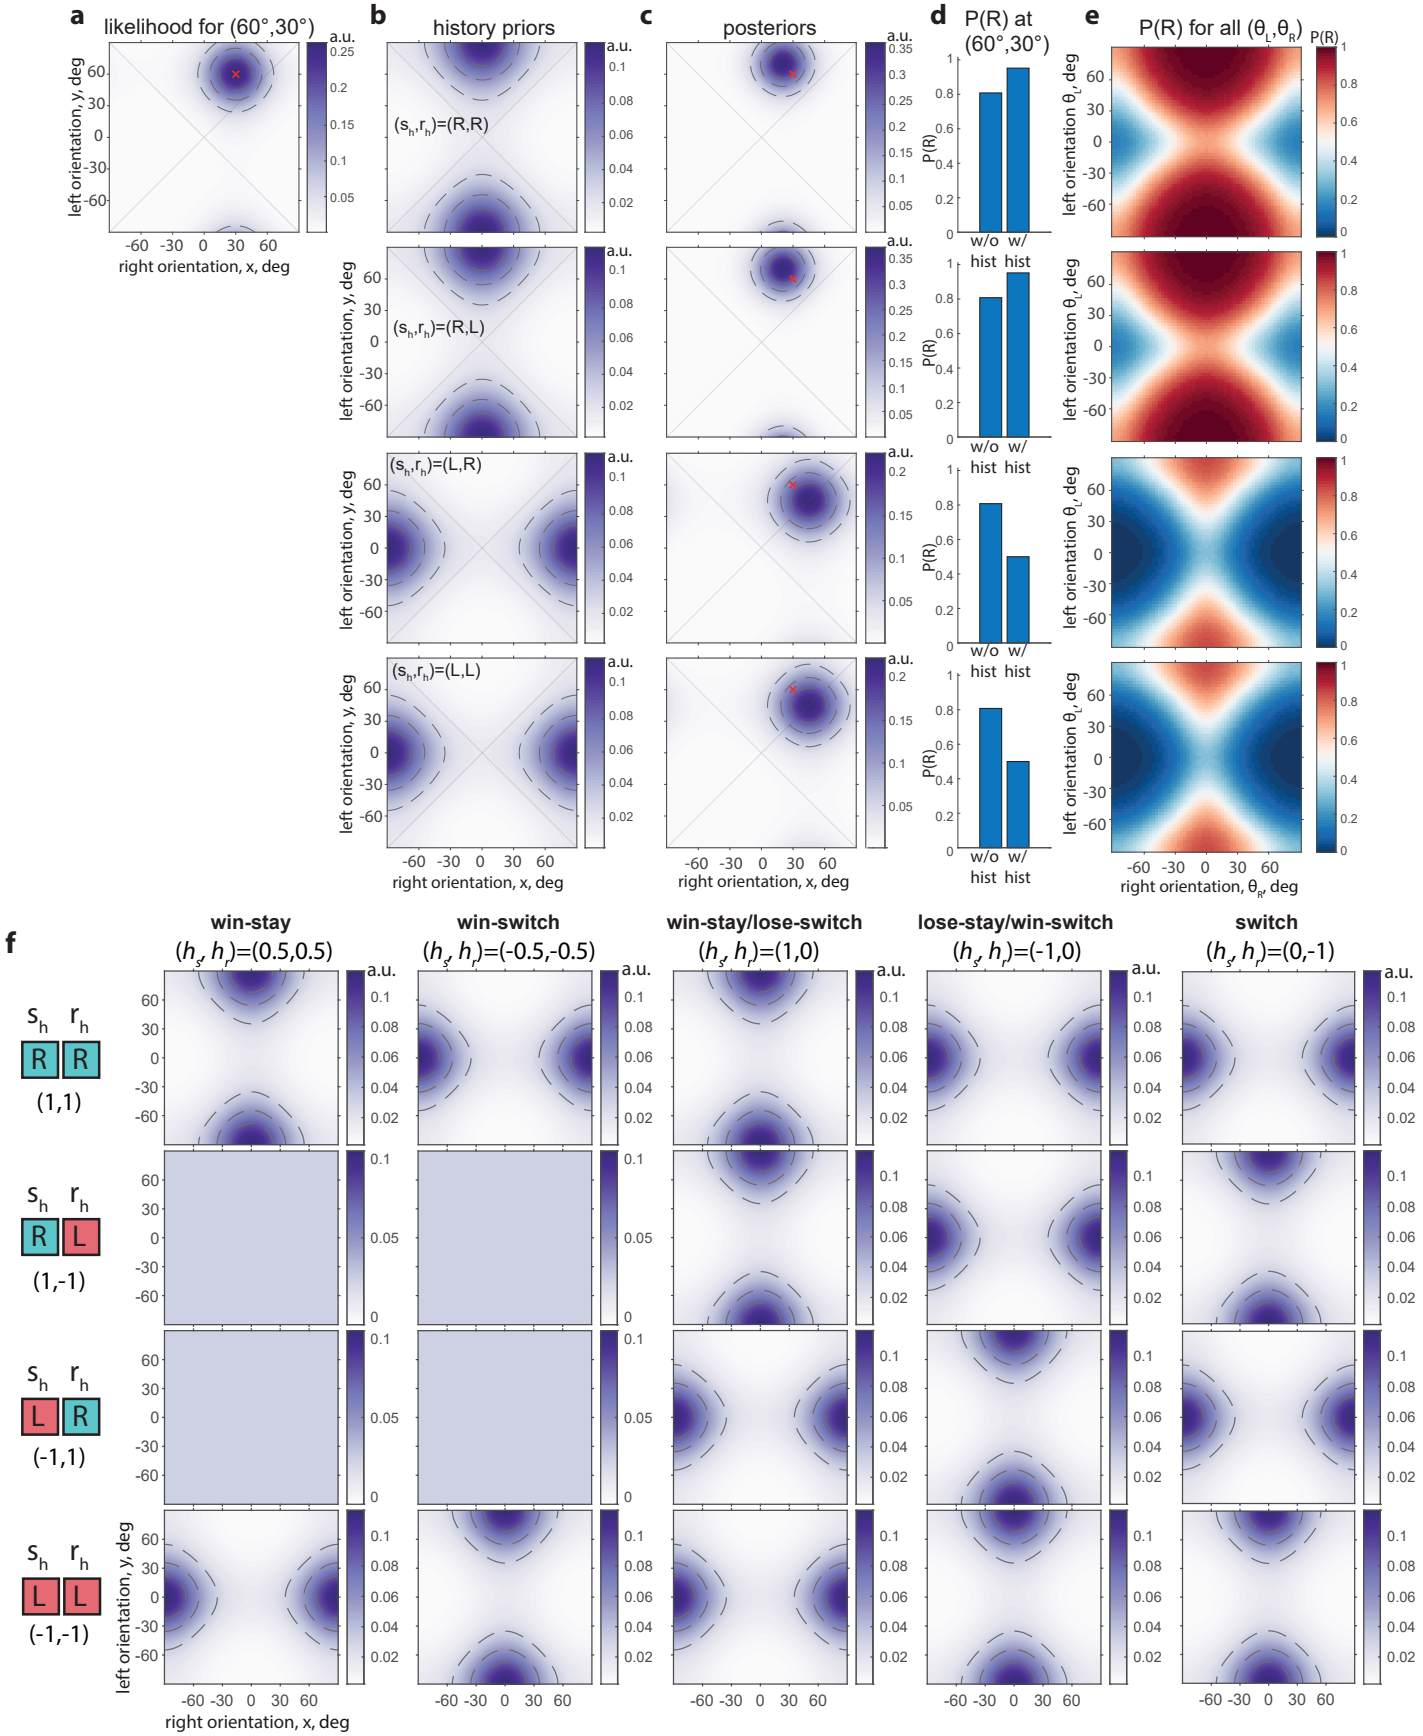

Supplementary Figure 4. History priors can represent many possible strategies, and affect choice probability , P(R), by shifting the probability density p(x, y) inside or outside the |x| < |y| region **a.** Probability density (p.d., shown by color saturation) p(x, y) (Methods, Eq. 1) induced by stimuli (θ_R, θ_L) = (30°, 60°) (red cross) in a model with $\kappa_R = \kappa_L = 2$, $b_R = b_L = 0$, and $\kappa_b = 0$; 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 $\kappa_h = 5$, 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. $p(\theta_R^*, \theta_L^*)$: 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); $\kappa_h = 1$ in all cases.

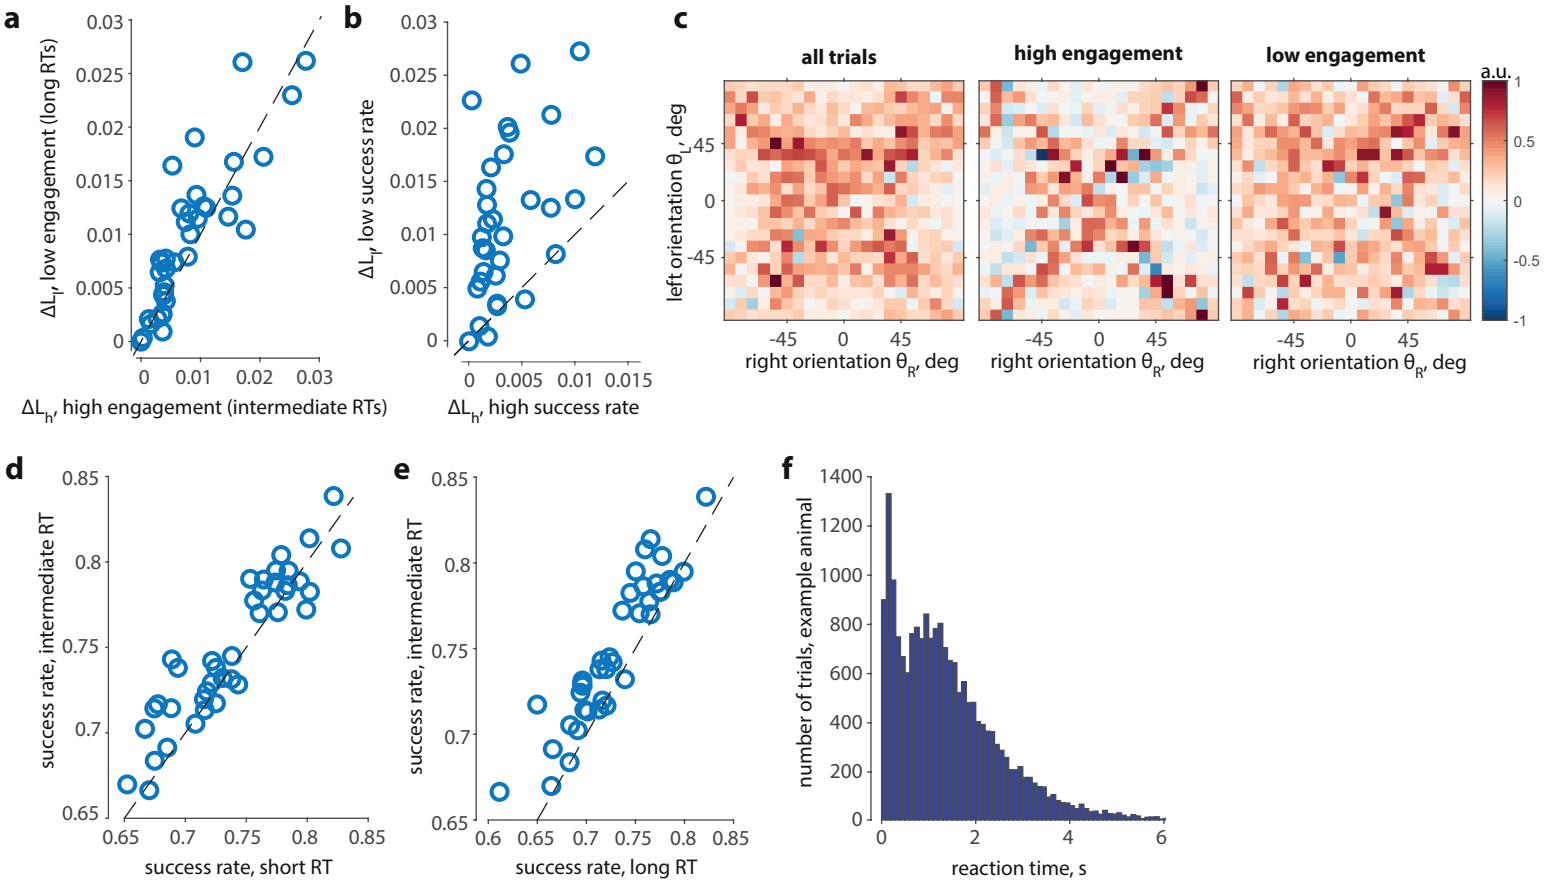

Supplementary Figure 5. Inclusion of history biases increases explanatory power of the model more in the low engagement trials, with criteria of engagement alternative to Figure 4e. **a.** Increase in the average likelihood with inclusion of history terms is larger in low engagement (long RT) trials than in high engagement (intermediate RT) trials, each circle is one of n=40 animals. **b.** Periods of high and low engagement are determined based on the running estimate of success rate, (we previously found co-variation of success rate with pupil size and reaction times for a significantly smaller subset of sessions (Abdolrahmani et al., 2021)). Abscissa – difference (ΔL_h) between the average log-likelihood of the model with and without history on high-engagement trials, ordinate – log-likelihood difference (ΔL_l) on low-engagement trials. **c.** With success rate-based engagement periods (as in b), increase in the log-likelihood due to inclusion of history terms is largely restricted to difficult stimuli in high engagement, and affect all stimuli in low engagement. Left to right: [1] ΔL_θ – average log-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. **d.** Trials with short RTs have significantly worse performance than trials with intermediate RTs ($p=9.9 \cdot 10^{-4}$, Wilcoxon signed rank test). **e.** Trials with long RTs have significantly worse performance than trials with intermediate RTs ($p=1.9 \cdot 10^{-7}$, Wilcoxon signed rank test). **f.** RT distribution for an example animal (clipped at 6s for presentation purposes).

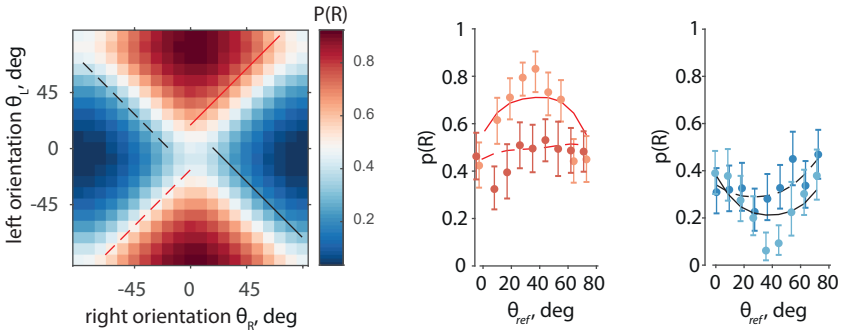

Supplementary Figure 6. Variation of P(R) with reference orientation θ_{ref} is larger in the data than in the model. Left. Example mouse, selected here for its low translational bias and approximately equal concentrations for right and left stimuli, which results in a regularly shaped P(R) dependency 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.

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

Supplementary Table 1. Total number of trials per animal. "ID" columns show animal identification numbers as in Supplementary Figures 1 and 3, "N Trials" columns show total number of trials of the corresponding mouse used in the analysis throughout the paper.