OPEN

Correspondence and
requests for materials
should be addressed to
M.T. (mariomtv@
hotmail.com)

Stimulus similarity determines the prevalence of behavioral laterality in a visual discrimination task for mice

Mario Treviño

Laboratorio de Plasticidad Cortical y Aprendizaje Perceptual, Instituto de Neurociencias, Universidad de Guadalajara, México.

Animal choices depend on direct sensory information, but also on the dynamic changes in the magnitude of reward. In visual discrimination tasks, the emergence of lateral biases in the choice record from animals is often described as a behavioral artifact, because these are highly correlated with error rates affecting psychophysical measurements. Here, we hypothesized that biased choices could constitute a robust behavioral strategy to solve discrimination tasks of graded difficulty. We trained mice to swim in a two-alterative visual discrimination task with escape from water as the reward. Their prevalence of making lateral choices increased with stimulus similarity and was present in conditions of high discriminability. While lateralization occurred at the individual level, it was absent, on average, at the population level. Biased choice sequences obeyed the generalized matching law and increased task efficiency when stimulus similarity was high. A mathematical analysis revealed that strongly-biased mice used information from past rewards but not past choices to make their current choices. We also found that the amount of lateralized choices made during the first day of training predicted individual differences in the average learning behavior. This framework provides useful analysis tools to study individualized visual-learning trajectories in mice.

The mouse constitutes a practical model system for studying the neuronal circuits underlying visual discrimination, decision-making and perceptual learning[1]. Neurons in the mouse primary visual cortex have highly tuned receptive fields[2,3], and mice can discriminate simple[4,5] and complex[6] shapes. Mainly for these reasons, we are now aiming to study the interaction between the visual-learning[5] and visual-discrimination[7] capabilities of rodents.

Mouse vision can be indirectly assessed by using a variety of behavioral methods. Some of them involve measuring reflexive eye movements to moving bars[8], whereas others require the animals to perform a visual discrimination task where they are trained to select a discriminative visual stimulus[4,5,7,8]. In such tasks, individual lateral biases (*i.e.* a stereotypy of response location) constitute a challenge for psychophysical measures (*v.gr.* visual acuity) because they are highly-correlated with error rates. One common practice among experimenters is to remove such lateralized biases by introducing a spatial bias towards the opposite arm of the maze[4,5]. Maybe not surprisingly, some authors consider exploratory lateralization an artifact of laboratory conditions or even a neuropathological reflection. Some other researchers suggest, however, that behavioral asymmetries could be ubiquitously distributed among animals. Moreover, it has been proposed that exploratory biases could be advantageous for learning and could also increase success in foraging and escaping from predators (for review see ref. 9).

Lateralization can be described both at the 'individual level' or at the 'population level'. In the first case, the population is made up of two equally sized subgroups of lateralized individuals. Although consisting of lateralized individuals, the population as a whole is not lateralized (*i.e.* on average). Conversely, when lateralization occurs at the population level, this means that the population is formed by coherently lateralized individuals[10].

Here, we hypothesized that lateralization in mice could constitute an efficient and stable strategy within a broad spectrum of behavioral responses to solve a discrimination problem of graded difficulty. To test this idea, we trained mice in a visual discrimination swimming task that we implemented recently[7]. We found that that the mice solved the task not only by using the discriminative stimulus as a relevant source of information for behavioral control but also displayed idiosyncratic lateral biases which increased in frequency with higher stimulus similarity. Individual biases varied from mouse to mouse: some individuals showed an intrinsic preference for the left option, while others did so for the right option, but there was no net bias when averaging the

choices from all the mice. We propose that the mice could benefit from having different degrees and sides of individual lateralization without suffering the disadvantages conveyed by directional asymmetries at the population level, such as predictability of behavior[9]. The analytical tools we provide here will allow to compare exploratory lateralization between different mouse strains and under a variety of experimental conditions.

Methods

We used behaviorally naïve, unselected, wild-type C57BL/6 male mice (P40–50, 21 ± 3 g at the beginning of the experiment). All mice were reared in mixed-sex family groups until weaning at 21 days of age. The mice were housed individually at 22°C under a 12/12 h light/dark cycle (lights on at 8:00), 35–40% relative humidity, *ad libitum* access to food and water, and a fast-track device for physical exercise (PLexx, Netherlands). All mice were held under identical conditions. Groups of 4-to-6 animals were carefully handled by a single experimenter and habituated to the training room 3 days before starting with the experiments. Training was performed in a single daily session held during the light phase (between 10:00 and 16:30), 5 days a week. All animal experiments were carried out at the Max Planck Institute for Medical Research in accordance with the animal welfare guidelines of the Max Planck Society and were approved by the regional commission in Karlsruhe (G-171/10).

To train the mice, we used a well-established, two-alterative, forced-choice water discrimination task under a 'free response' paradigm that allowed the animals to control the decision time autonomously[4,5] (Fig. 1A). Water temperature (21 ± 1°C) and room illumination were kept constant throughout the experiments and the pool was wiped down daily with 70% ethanol. We confirmed that the mice did not see the hidden transparent platform (see ref. 5). For each trial, the animals were considered to have made a choice once they had crossed a line that delineated a decision area which offered visual access to both images. To encourage faster discrimination learning, we increased the cost of errors by immediately repeating the swimming trials that produced incorrect choices up to a maximum of 5 times until the animal made the correct choice (Fig. 1B). We defined these sets of swims, ranging from 1-to-5, as a 'training unit', and it was considered as being correct only when the mouse made a correct choice during the first trial. The mice remained on the platform for 30 s before being carefully removed from the pool by the experimenter. The inter-trial interval was 10 s, and the period between training units was 1–2 min (*i.e.* distributed practice). We used daily training sessions consisting of 3 blocks of 10 'training units' with 10-min breaks. During rest periods, the mice were transferred to individual chambers with a warm plate.

Each experiment consisted of two phases. During phase 1 ('pre-training', duration: 1 week; training units 1–150, see below), the mice were familiarized with the swimming task, and learned to associate a discriminative image (S_0^D) with a predictive value. In this pre-training phase, the animals learned that swimming towards the S_0^D and reaching a transparent, submerged platform was rewarded with removal from the water, whereas swimming towards the non-reinforced image (S_0^Δ; 50% gray) was not. During phase 2 ('training', duration: 2 weeks; training units 151–450), the reinforced S^D image was changed (yet, it was identical for all groups) and 3 sub-groups of 10 mice each were trained with three different S^Δ images. The structural similarity index (SSIM) between pairs of images was measured using parametric descriptions derived from image quality metrics, as described before[7]. Briefly, the images consisted of white shapes on a black background, or *vice versa* (*i.e.* shape was the only relevant 'feature'); they were of similar size and were further standardized using a symmetric Gaussian low-pass filter (60 pixel size, 30-pixel standard deviation; 0.30 cycles per degree [c/d]), to remove all frequency components that exceeded the average mouse's visual acuity of 0.48 c/d[5]. The data in Figure 1 were published previously[7] and serve as the basis for the current analysis.

For each training unit, we calculated the mean probability (±S.E.M) of making a correct choice on the first presentation of the S^D image (%correct), and of making 5 consecutive errors (%error; this means that %correct choices and %errors are not complementary in our working conditions). The average changes in these scattered patterns over time were visualized by using a Savitzky-Golay filter (span = 30 trials, degree = 1) that served as a low-parameter estimate to compare group data from different training regimens. We used a digital video camera mounted above the pool to record each swimming path throughout the experiments. For each trajectory, we analyzed continuous measures of path length (*i.e.* cumulative Euclidean distance) and escape latency (*i.e.* time from release from the chute to task completion). Learning was inferred both by correct choice and continuous measures from conditioned responses (see below). To avoid positional learning, the sides of the discriminative stimulus (S^D) and the platform (left or right) were changed continuously according to a Gellerman-like schedule: *i.e.*, a pseudo-random pattern in which no more than 3 trials are repeated on one side, and that produces a score of 50% correct choices if a subject shows simple or double alternation[4,5].

The strength of the side-bias was quantified as the number of consecutive swims towards the same arm of the pool (right or left). For the quantitative assessment of behavioral laterality, we implemented a laterality index (L_i):

$$L_i = \frac{[R-L]}{[R+L]} \quad (1)$$

where R and L denote right and left choices, respectively. This index was used either as the average of choices for every 10 trials for each mouse, or as the group average per training unit. To explore the possibility that biased, alternating or more complex choice sequences emerged during the acquisition, we implemented a pair-wise alignment method that consisted in sliding each query sequence along the choice records from each mouse. A sequence was considered to be present in the choice record when the alignment matched perfectly for the entire length of the query sequence (*i.e.* sequence similarity = 1). The probability of occurrence of a choice sequence was then calculated by dividing the number of times that it was found in the choice records by the maximum number of times that it could fit within the total length of the choice record without interfering with any identical sequence that generated a count in previous trials (very important for alternating sequences). Comparing the probabilities instead of the number of cases was crucial because: i) the choice records varied in length for each mouse and ii) the alignment method implicitly increases the counts of shorter sequences contained in longer ones.

In operant conditioning, the matching law is a quantitative relationship that holds true between relative response and reinforcement rates in concurrent reinforcement schedules[11]. We tested whether biased and alternating choice sequences complied with matching behavior using a generalized matching law with the following form:

$$\frac{C_L}{C_R} = C \left(\frac{R_L}{R_R}\right)^a \quad (2)$$

where C and R denote the number of steady-state responses and the number of steady-state reinforcers for the left, or right option, respectively. The coefficient a denotes the sensitivity to the reinforcement ratio, while c is a bias term unrelated to reinforcer frequency or magnitude[12,13]. We fit the data for each aligned sequence with the logarithmic form of the equation using least-squares regression:

$$\log\left(\frac{C_L}{C_R}\right) = a\log\left(\frac{R_L}{R_R}\right) + \log c \quad (3)$$

and extracted the slope (a), the intercept ($\log c$), and the coefficient of determination for each regression (R^2; *i.e.* the square of the sample correlation coefficient between outcomes and predicted values).

We also analyzed the response-by-response behavior with a multiple linear regression model that used past reinforcers as well as past choices to predict the mice's choices in each trial[14,15]. Assuming symmetric effects of past reinforcers (*i.e.* for right or left choices), the model was then reduced to solving the following generalized logistic regression:

$$\log\left(\frac{p_{R,i}}{p_{L,i}}\right) = \sum_{j=1} \alpha_j \left(r_{R,i-j} - r_{L,i-j}\right) + \sum_{j=1} \beta_j \left(c_{R,i-j} - c_{L,i-j}\right) + \gamma \quad (4)$$

where $p_{R,i}$ is the probability of choosing the right alternative (*i.e.* probability when $c_{R,i} = 1$), and $p_{L,i} = (1 - p_{R,i})$ is the probability of choosing the left option (binomial distributions). $c_{R,i}$ and $c_{L,i}$ are binary variables that represent the choice of the right and left alternatives on the ith trial, and r is the magnitude of reinforcement received for choosing a particular alternative on the jth past trial, which otherwise is zero ($r = 1$ is fixed for each correct choice). The α and β coefficients measure the influence of past reinforcers and choices, and the intercept term γ captures preference that is not accounted for by past reinforcers or choices (similar to the bias term in the generalized matching law). Here, a unit reinforcer obtained in j trials in the past increases the log odds of choosing an alternative by α_j if the reinforcer was received for choosing that alternative; otherwise, it decreases the log odds by α_j. This applies similarly to the effects of past choices, where a significant β_j means that the current choice depends on a choice made j trials ago[14]. In other words, this model provides a convenient way to test the null hypothesis (H_0) which states that the factor in question does not affect the measured response[16]. Excluding the terms associated with the β parameters (forcing all $\beta_j = 0$) yields a model that depends on the history of obtained reinforcers only, whereas excluding the terms associated with the α parameters (forcing all $\alpha_j = 0$) yields a regression that depends only on the history of choices. Including both α and β allowed us to assess the effects that reinforcer and choice history have on current choice. The intercept term γ shifts preference towards one of the alternatives, irrespective of reinforcement (*i.e.* it captures a bias that is not due to either reinforcer frequency or reinforcer magnitude).

We derived a local estimation of task efficiency as follows:

$$\varepsilon = 1 - \left[\frac{X-1}{X_{random}-1}\right] \quad (5)$$

where X is the group average for the number of swimming trials per training unit, and X_{random} corresponds to the average number of swimming trials required to complete each training unit by making random choices (R_{1000}: binomial distribution, $n = 1000$ 'subjects'). Next, we implemented two different methods to estimate the relative change in swimming efficiency. In the first one, we divided the average efficiency of the query sequence by the average efficiency of a sequence of trials of equal length taken just before executing the query sequence (*i.e.* to estimate the baseline). In the second method, we divided the average efficiency of the query sequence by the
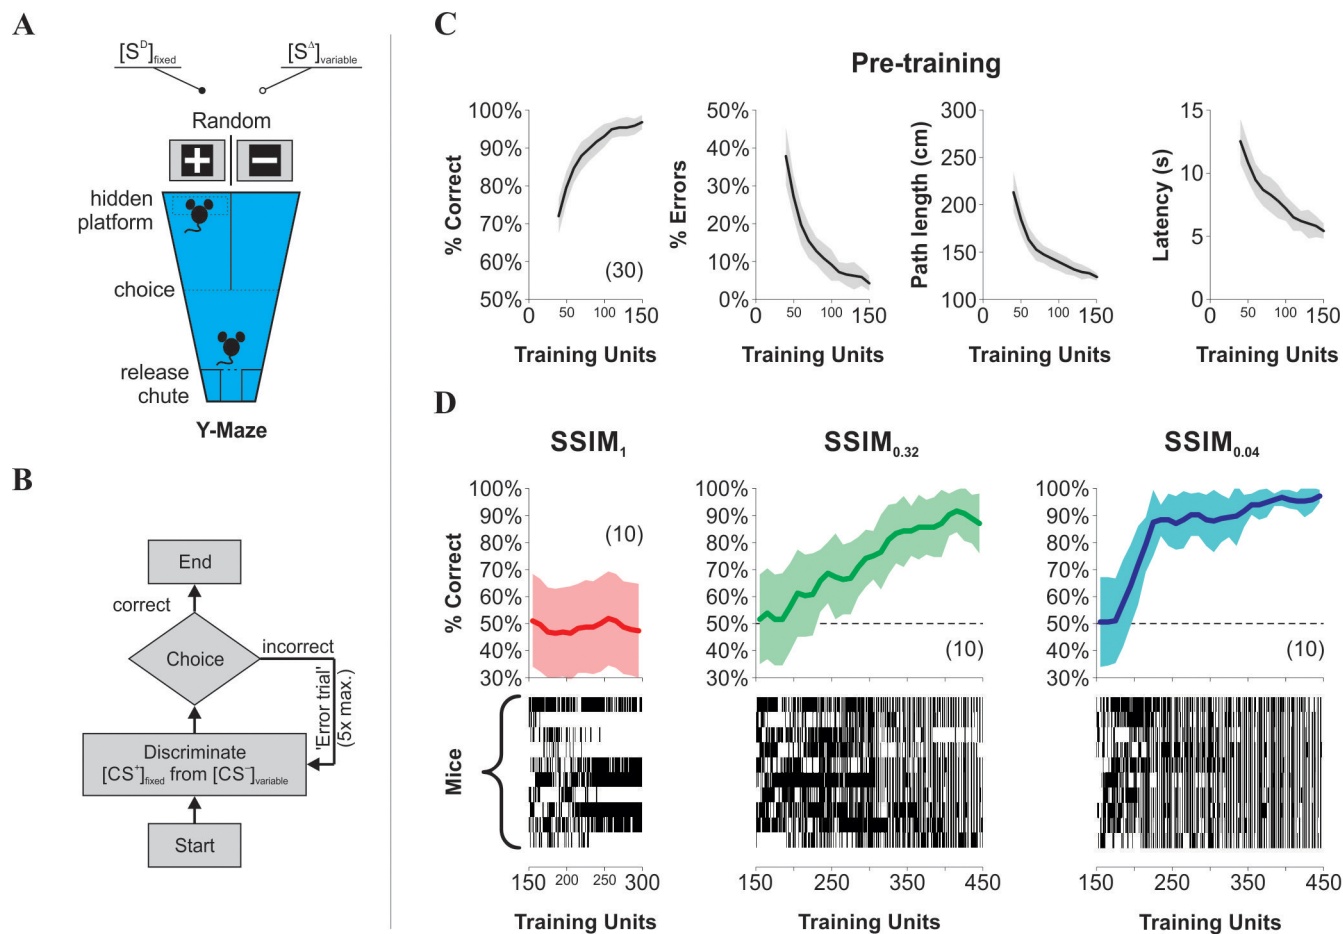
Figure 1 | Emergence of side-biased choice behavior during a visual discrimination task. (A) A drawing that we've made depicting the visual discrimination task where two monitors facing the ends of the arms of a Y-watermaze simultaneously display the discriminative (S^D, reinforced) and non-reinforced (S^A) images (100% contrast). A submerged transparent platform below the S^D serves as the unconditioned stimulus (US). The position of both the platform and S^D in either arm varies pseudo-randomly over consecutive trials. During training, mice are released into the pool from a release chute and they learn to swim towards the S^D (correct choice) in order to reach the platform and escape from the water. (B) Flowchart of a 'training unit' where the mice are presented with a given pair of S^D/S^A, a trial that can be repeated up to 5 times if the mouse makes incorrect choices. (C) %correct choice increased as learning progressed during pre-training whereas the %errors, length of the swimming path (*i.e.* to reach the platform; 'Path length'), and the time from the beginning to the end of the trial ('Latency') decreased asymptotically. (D) After pre-training was completed, three sub-groups of mice were trained with different degrees of similarity between S^D and S^A: SSIM = 1 ($SSIM_1$; red), SSIM = 0.32 ($SSIM_{0.32}$; green) and SSIM = 0.04 ($SSIM_{0.04}$; blue). The average learning curves (continuous lines) were approximated by a Savitzky-Golay filter (see Methods). At the bottom of each plot: the choices of individual mice (y-axis) are displayed as panels with either black (right choices) or white (left choices) rectangles as a function of the first trial of each training unit (x-axis). Individual side-biases are reflected as horizontal white or black blocks of different lengths. Number of mice per group in parentheses.

average efficiency within a block of 10 trials occurring 10 trials before executing the query sequence.

Analysis algorithms were written in MATLAB 7.8 (MathWorks, Inc.; Natick, USA). Learning was assessed with repeated measures ANOVA tests (consecutive blocks of 30 training units), all followed by Bonferroni *post hoc* tests. We switched to non-parametric tests (Wilcoxon Signed Rank test; Kruskal-Wallis test followed by Dunn post hoc tests) whenever the assumptions required to apply the parametric versions were not met. All results are shown as averages ± S.E.M; significance was set at $P < 0.05$.

Results

Learning with different degrees of visual discriminability. We exposed 30 mice to an initial one-week period of 'pre-training' (150 training units), which allowed them to become familiar with the swimming pool and the task[5]. During this phase, the animals learned that a highly discriminative image (S_0^D) predicted the location of the transparent platform inside the pool. The average correct choice increased towards an asymptotic level of 95% ± 1% (average of last 30 units; $n = 30$ mice), whereas the number of errors, the total swimming distance and the escape latency for correct

choices decreased as a function of training (Fig. 1C). All these values were highly correlated ($r > 0.9$ for all groups) and the within-group variability decreased with learning.

Next, we formed three random sub-groups of 10 mice each and trained them to discriminate images with maximum (SSIM = 1, red), intermediate (SSIM = 0.32, green), and low (SSIM = 0.04, blue) similarity levels during the second and third weeks of the experiments (see Methods, Fig. 1D). As expected, the mice failed to discriminate between identical S^D and S^A images with SSIM = 1 ($SSIM_1$; Wilcoxon test, $P = 0.17$, $n = 10$), whereas training with a SSIM of 0.32 and 0.04 yielded above-random choice levels (Wilcoxon tests, $P < 0.01$; Fig. 1D). These two groups had different learning rates ($SSIM_{0.32}$: 0.63%/training unit, $n = 10$; $SSIM_{0.04}$: 2.02%/training unit, $n = 10$), but reached similar correct choice levels at the end of the training phase (one-way ANOVA, $F_{2,29} = 123.5$, $P < 0.001$, Bonferroni's post hoc test, $P > 0.05$; Fig. 1D). These results indicate that the learning rate increased when stimulus similarity was lowered.

By inspecting the individual responses, we noticed that the mice displayed different choice sequences during training. Some of these sequences consisted of swimming repeatedly to the same arm of the pool during the training phase[4,5,17]. To estimate whether side-biased choices could be influenced by stimulus similarity, we first labeled and plotted the trials in which the mice swam to the right (black squares) or left (white squares) arms (lower panels in Fig. 1D). The lower diagrams in Figure 1D show that alternation in the side of the spatial bias, from right-to-left or *vice versa*, could occur quickly, within just a few trials. Subsequently, we counted the number of side-biased sequences of different length derived from these three experimental groups.

Stimulus similarity determines the probability of side-biased swimming. The probability of displaying side-biased swimming decayed mono-exponentially with the length of the biased sequence (Fig. 2A–B). Yet, the decay constant (λ) for the probability of side-biased behavior decreased, non-linearly, with stimulus similarity (pre-training: $\lambda = 0.2$, $R^2 = 0.98$; $SSIM_1$: $\lambda = 0.04$, $R^2 = 0.86$; $SSIM_{0.32}$: $\lambda = 0.05$, $R^2 = 0.89$; $SSIM_{0.04}$: $\lambda = 0.08$, $R^2 = 0.94$; inset in Fig. 2B). Thus, the prevalence of side-biased swimming was influenced, in a graded manner, by stimulus similarity.

We then wondered whether the bias from each individual showed any predominance towards a 'preferred side'. We averaged the laterality of choices per subject (left $= -1$, right $= +1$) in blocks of 10 trials and computed their frequency distribution referenced to the preferred side (*i.e.* referenced to the mode). We found that the mice tended to alternate and balance the laterality of their choices when they were trained with high discriminability (symmetric distribution for $SSIM_{0.04}$; Fig. 3A). However, a preferred side for side-biased swimming gradually emerged as stimulus similarity was increased (skewed distributions for $SSIM_{0.32}$ and $SSIM_1$; Fig. 3A). The skewness in the $SSIM_1$ distribution is a direct consequence of frequent individual lateralization because, although alternations could occur very fast, the probability of swimming towards the preferred arm was allays bigger than that of swimming towards the non-preferred arm (t-test, $P < 0.001$). Notably, the preferred arm was specific for each mouse (see below).

The concern that side-biased swimming could be influenced by additional sources of sensory information (*v.gr.* asymmetric illumination settings along the longitudinal axis of the pool) or by innate population biases, was ruled out by plotting the group average laterality for all mice as a function of training (Fig. 3B). The average laterality of group choices was balanced between both arms of the pool throughout training (blocks of 30 training units; Wilcoxon test, $P = 0.95$ for all groups), leading to no net side preference when considering all swimming trials from all the mice ($SSIM_1$: laterality index $= 0.02\% \pm 0.08\%$; $SSIM_{0.32}$: laterality index $= -0.14\% \pm 0.05\%$; $SSIM_{0.04}$: laterality index $= -0.07\% \pm 0.04\%$; one-way ANOVA, $F_{2,26} = 2.24$, $P = 0.32$). Thus, the population, taken as a whole, showed no net bias. These results illustrate how a zero bias at the population level does not imply that individual laterality must be zero.

The probability of choosing specific swimming sequences is graded by stimulus similarity. We next used a sequence alignment method to explore whether the emergence of different swimming strategies depended on stimulus similarity. We searched for the occurrence of specific sequences with different amounts of choices to the left (LL…L), the right (RR…R), or alternating choices ('LR…L' or 'RL…R'). Figure 4A shows an example of the matching of such sequences (red squares) to the choice records of individual mice from the $SSIM_{0.04}$ group. To refine the analysis, we added the counts of complementary sequences (*v.gr.* ['LL…L' + 'RR…R'], ['LR…L' + 'RL…R']) and normalized them to an equivalent number of training trials, thus allowing group comparisons. We quantified their
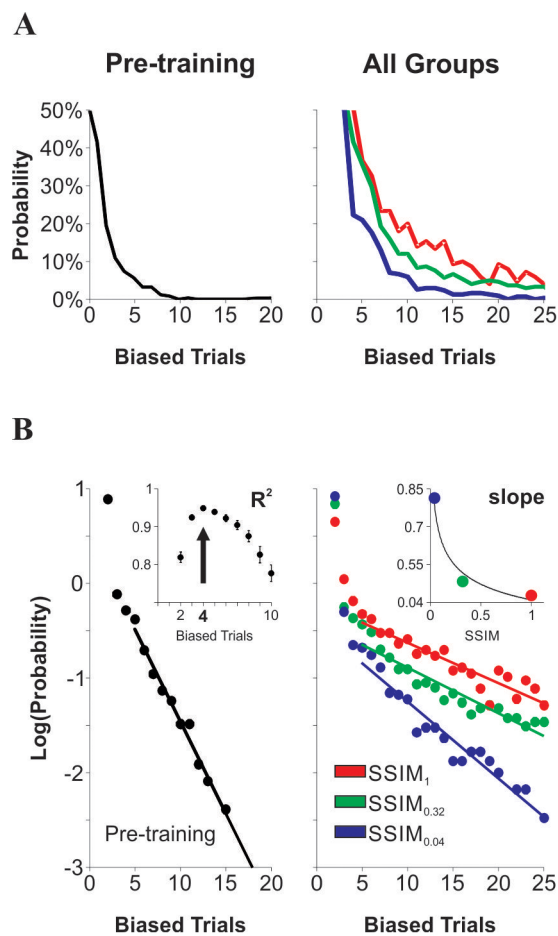
Figure 2 | Stimulus discriminability determines the amount of side-biased choices during training. (A) The probability of finding side-biased sequences of different lengths (x-axis) increases with stimulus similarity: $SSIM_1$ (red), $SSIM_{0.32}$ (green) and $SSIM_{0.04}$ (blue). (B) The log(probability) of side-biased swimming behaves linearly with respect to the length of the biased sequences. We defined biased sequences as having ≥ 4 trials because removing the first 3 trials maximized the coefficient of determination for all linear regressions (derived from all groups), depicted by the black arrow in the left inset. The three decay constants for the probability of side-biased swimming decrease, in a non-linear fashion, with stimulus similarity (right inset).

probability of occurrence dividing the number of cases by the maximum number of repetitions per sequence that could fit within the choice record for each mouse, this without interfering with any sequence that generated a count in previous trials (crucial for alternating sequences). We applied this analysis to: i) the original choice record (*i.e.* what the mice actually did; black circles); ii) a randomized choice pattern equal in length to the original choice record (binomial distribution; yellow circles); and, iii) the platform location record (*i.e.* how a choice pattern with perfect discrimination would have looked like; gray circles; Fig. 4B–C). In accordance with our previous results, all probabilities decayed with sequence length whereas the probability of choosing side-biased sequences increased with stimulus similarity (Kruskal-Wallis test, $F_{2,29} = 29.2.9$, $P < 0.0001$, Dunn's multiple comparison test, right plot, Fig. 4B). By using the cumulative probabilities for side-biased swimming (length of sequences from 3 to 9) we confirmed that biased swimming could not be accounted for by using a random-choice maker. In contrast, the probability of finding alternating sequences decreased with stimulus similarity (Kruskal-Wallis test, $F_{2,29} = 32.2$, $P < 0.0001$, Dunn's multiple comparison test, right plot, Fig. 4C).
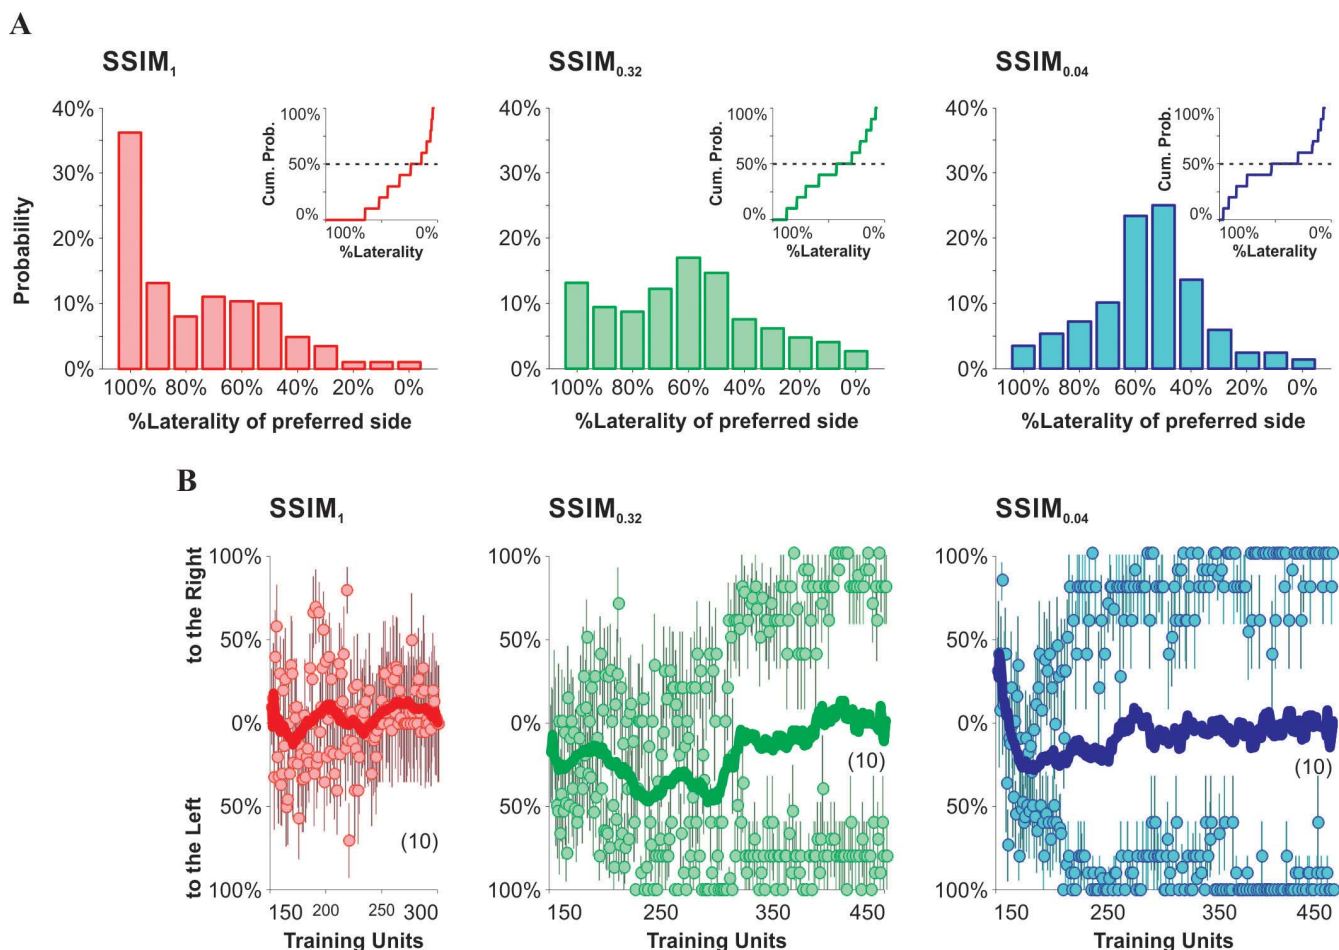
Figure 3 | Lateralization at the individual level, but not at the population level. (A) Probability distributions for the laterality of choices (blocks of 10 trials) extracted from the choice records of individual mice for the $SSIM_1$ (red), $SSIM_{0.32}$ (green) and $SSIM_{0.04}$ (blue) groups. The arrangement of these distributions was referenced with respect to the preferred side for lateral bias from each individual, identified as the side of the pool towards which the statistical mode was closest. The skewness in the two first distributions indicate that mice in $SSIM_1$ and $SSIM_{0.32}$ showed a tendency to have a preferred arm during side-biased behavior. (B) Balanced group average laterality for all experimental groups during acquisition. Number of mice per group in parentheses.

This analysis demonstrates that the mice adapted their task-solving strategies depending upon stimulus similarity.

**Side-biased behavior complies with the generalized matching law.** We asked whether the mice matched the proportion of biased and alternating swims to the proportion of reinforcers received from choosing them (see Methods). For this analysis, we selected blocks of 30 training units with choice-records in steady-state with respect to the last 30 trials of each group (Kruskal-Wallis test, $P < 0.05$)[14]. Next, we counted the number of times that each sequence occurred within the choice record for each mouse (i.e. frequency of sequences), and calculated the corresponding reward value as the sum of the correct choices for those sequences (i.e. frequency of reward). Figure 5 shows the log frequency for biased (Fig. 5A) and alternating (Fig. 5B) choice sequences as a function of the log net amount of reinforcer received from choosing them. Independently of stimulus similarity, biased sequences had better linear regressions to the generalized matching law than alternating sequences of equal length (biased sequences: $R^2 = 0.89 \pm 0.02$; alternating sequences: $R^2 = 0.53 \pm 0.09$; t-test, $P < 0.0001$). The sensitivity to reinforcement (the slope of the linear regression models) was positive for all biased sequences (t-test, $P < 0.001$ in all cases) and independent of stimulus similarity (one-way ANOVA, $F_{3,23} = 0.5595$, $P = 0.91$; the coefficients of the linear fits to the data shown in Fig. 5 are listed

in Table 1). The slopes (i.e. the sensitivity) smaller than one and the negative intercept of the regression models (i.e. the bias term) indicate that mice under-matched for reinforcement (i.e. less behavior was allocated to the alternative that provided greater reinforcement). These results demonstrate that biased choice sequences obeyed the generalized matching law during steady-state behavior.

The side-biased sequences from the $SSIM_{0.04}$ group that was trained in conditions of high discriminability also complied with the matching law (Fig. 5A). The average %correct choice of those sequences decreased towards chance level (i.e. 50%) as their frequency increased; the opposite occurred for alternating sequences (Fig. 6). Thus, biased choice sequences occur without using discriminative information even in conditions of high discriminability.

Side-biased behavior depends on reinforcer history but not on past choices. The generalized matching law predicts average choice behavior for arbitrary combinations of reinforcer frequency, but it does not specify how the animals produce matching behavior at a 'response-by-response' level. To estimate the dependence of biased behavior on the history of past reinforcers and choices, we applied a multiple linear logistic regression model to approximate the average choice behavior for each trial (see Methods). The coefficients of the model were computed using a history of 10 trials and their statistical
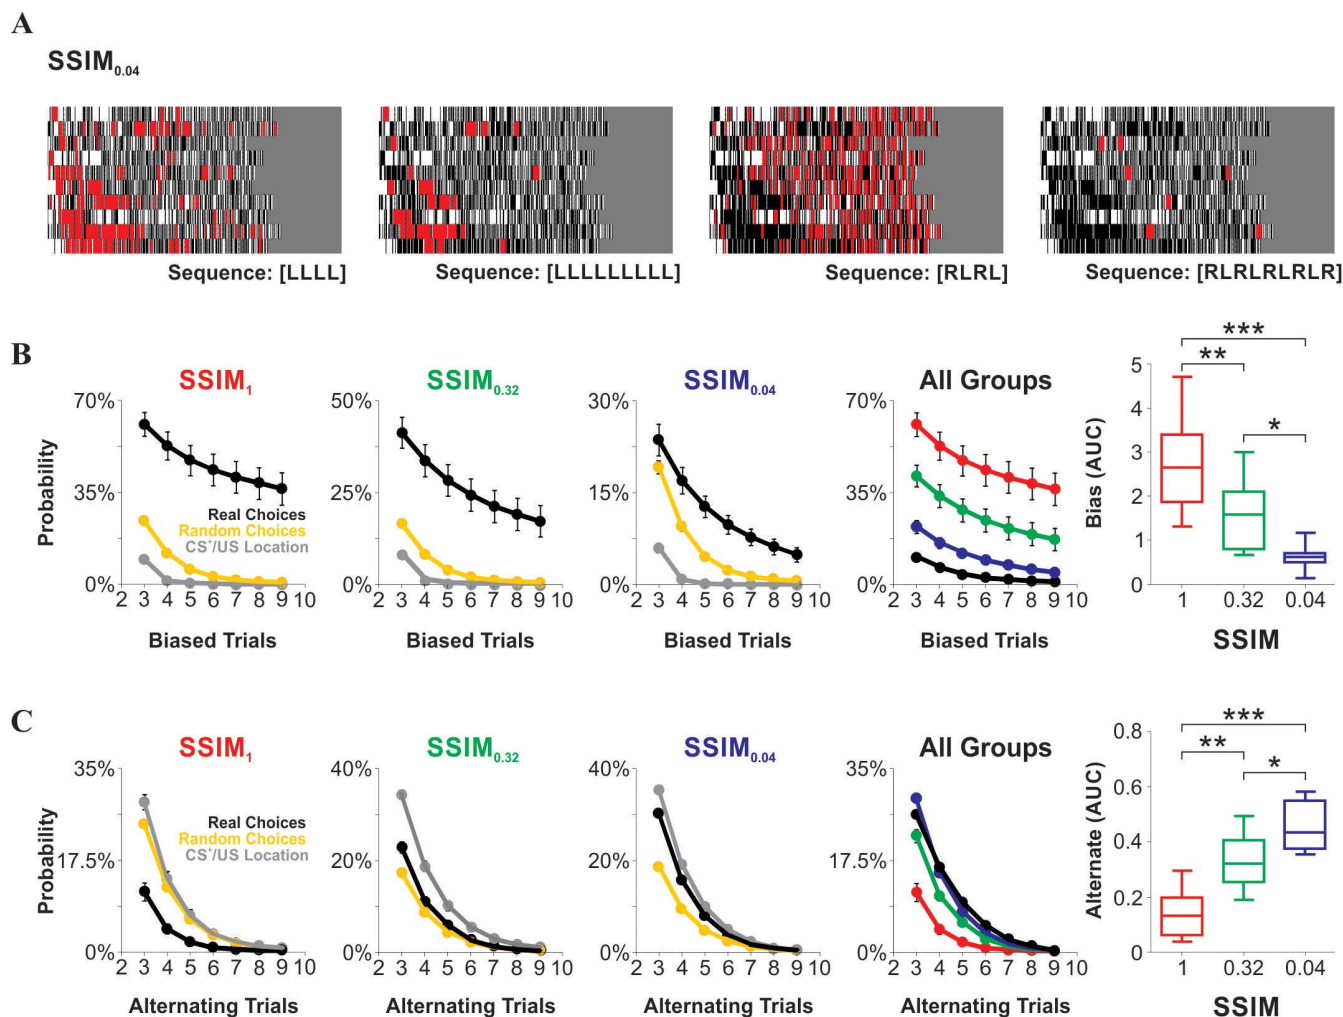
Figure 4 | Graded probability for biased and alternating sequences by stimulus similarity. (A) Example of the sequence analysis we performed on the choice records from the $SSIM_{0.04}$ mice. The occurrence of sequences is depicted by red squares superimposed on the choice diagrams (black: right choices; white: left choices) from individual mice (y-axis) as a function of the training unit (x-axis). The choice records of different total length are shown over a gray background, with the query sequences displayed below the choice diagrams. (B–C) Probability of occurrence of different sequences based on real choices (black circles), random choices (binomial distribution; yellow circles), and the real platform location (gray circles) for biased- (B) and alternating- (C) sequences. The probability axes have different scales depending on peak probabilities. For comparative purposes, colored panels on the fourth column display the probability of actual choices for each group: pre-training (black), $SSIM_1$ (red), $SSIM_{0.32}$ (green) and $SSIM_{0.04}$ (blue). On the right column: statistical tests using cumulative probabilities (i.e. the area under the probability curves, AUC) for all the mice from each group using Dunn's Multiple Comparison test.

significance was determined by using a permutation test that shuffled 1,000 times the trial order for the variable of interest to yield a *P* value for the permutation test (dotted lines in Fig. 7).

We first applied the model considering past reinforcers only (Fig. 7A). The positive coefficients for the $SSIM_1$ group (Fig. 7A) indicate that biased choices increased the log odds of picking the same alternative on the next trial due to the effects of reinforcer history. In contrast, the negative coefficients for the groups that were trained with lower stimulus similarity (*i.e.* pre-training, $SSIM_{0.32}$ and $SSIM_{0.04}$) indicate that choosing a particular alternative decreased the odds of picking the same alternative again on the subsequent trial (first past reinforcer, pre-training: -0.79 ± 0.01; $SSIM_1$: 0.76 ± 0.03; $SSIM_{0.32}$: -0.09 ± 0.03; $SSIM_{0.04}$: -0.47 ± 0.03; Fig. 7A). The decaying effect of past reinforcers was much slower and more persistent in $SSIM_1$ than in any other group (τ_{decay}, pre-training: 0.16 ± 0.10; $SSIM_1$: 0.39 ± 0.05; $SSIM_{0.32}$: 0.09 ± 0.06; $SSIM_{0.04}$: 0.16 ± 0.08).

To better capture the response patterns of choice behavior, we incorporated both past reinforcers and past choices into the model[14], and found that mostly past reinforcers (first past reinforcer, pre-

training: -1.66 ± 0.06; $SSIM_1$: 0.61 ± 0.07; $SSIM_{0.32}$: -0.67 ± 0.08; $SSIM_{0.04}$: -0.99 ± 0.11), but not past choices (first past choice, pre-training: 0.59 ± 0.02; $SSIM_1$: 0.14 ± 0.02; $SSIM_{0.32}$: 0.38 ± 0.04; $SSIM_{0.04}$: 0.37 ± 0.04), influenced choice behavior for the $SSIM_1$ group (Fig. 7B).

Side-biased sequences produce a local increase in task-solving efficiency with zero discriminability. Using sliding epochs of 10 trials, we determined the group probability of solving the task with side-biased sequences as a function of training (Fig. 8A). All groups began the training phase with a non-zero average probability for side-biased swimming of ~70% (biased sequences ≥ 4 trials; one-way ANOVA, $F_{2,26} = 3.51$, $P = 0.17$), but the prevalence of making side-biased choices either increased to ~82% for the $SSIM_1$ group, or decreased to ~32% and to ~7% for the $SSIM_{0.32}$ and $SSIM_{0.04}$ groups, respectively (paired t-tests, $P < 0.005$; Fig. 8A).

We implemented a measure of efficiency based on the number of swimming trials that were required to solve the task (see Methods). Solving the task in a side-biased manner throughout the entire phase
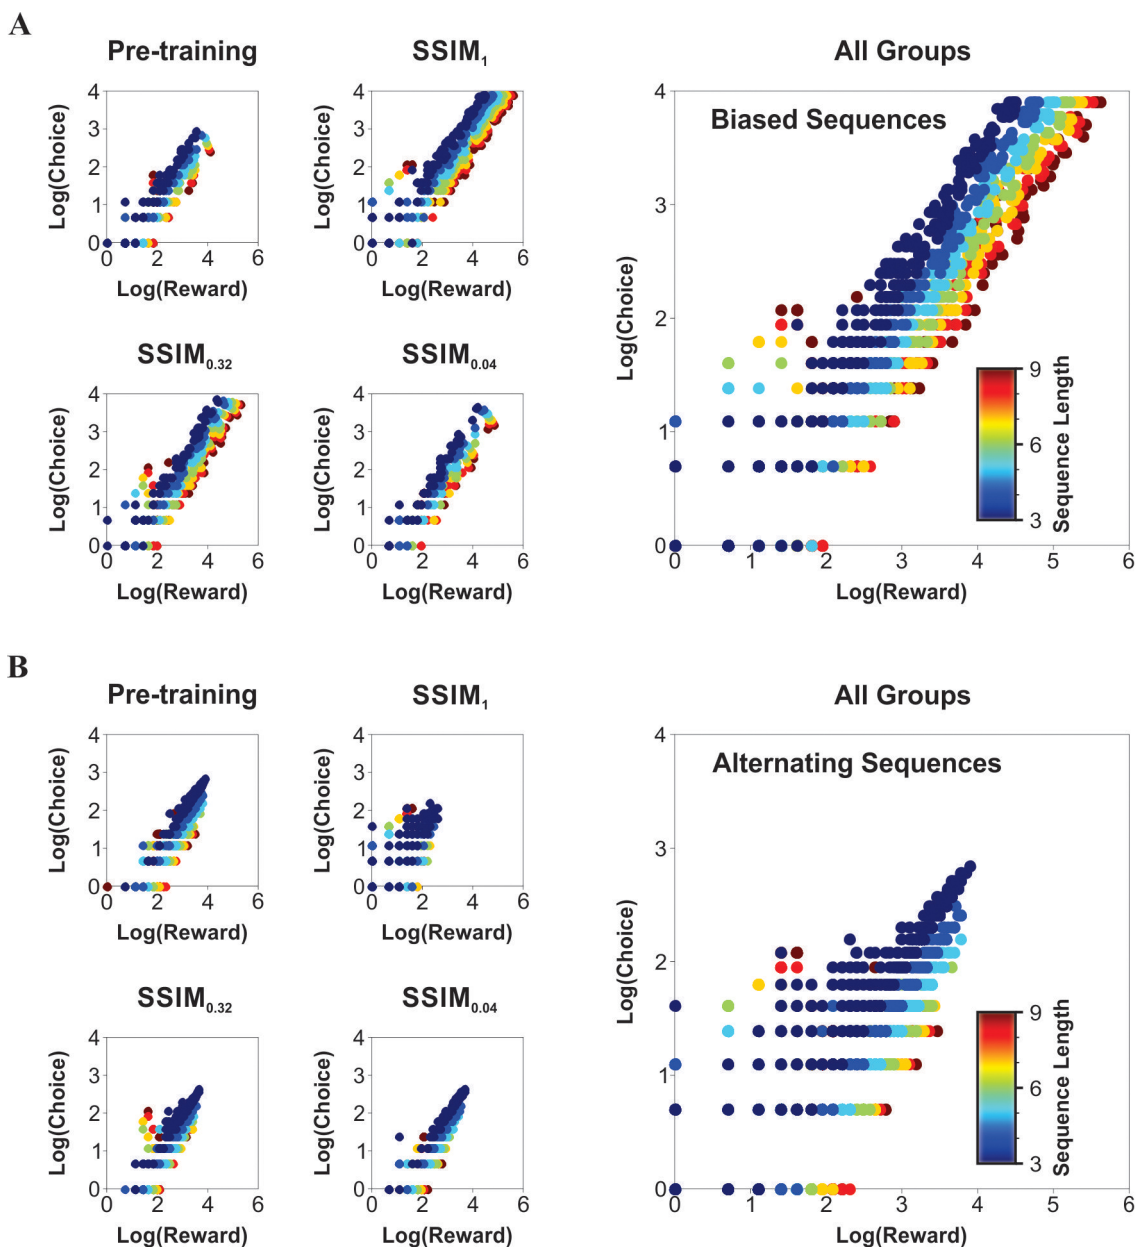
Figure 5 | Steady-state side-biased behavior complies with the generalized matching law. Each point represents the log frequency of the occurrence of a particular choice sequence *vs.* the log sum of reward values for that specific sequence, for the different experimental groups. (A) The log frequencies of choices are linearly-related to their log reinforcer frequency for all biased sequences from all groups. Note how the logarithmic transformations portrayed the data and their variability into linear functions. (B) Alternating sequences display worse linear regressions than biased ones (see Table 1 for further details). Color represents the length of the choice sequence.

was ∼10% more efficient than making random or other choices (completely biased: 537.5 ± 0.2 trials/300 training units, $n = 1000$ 'subjects'; random choosing: 590.5 ± 0.7 trials/300 training units, $n = 1000$ 'subjects'; 'choosing the arm where the platform was in the previous trial': 765.5 ± 0.1 trials/300 training units, $n = 1000$ 'subjects', one-way ANOVA, $F_{2,2996} = 2684.2$, $P < 0.001$). Next, we measured average swimming efficiency when mice adopted side-biased or alternating sequences of different lengths (see Methods). This analysis revealed that adopting side-biased sequences of 6 or more trials increased task efficiency by ∼5%, whereas adopting alternating sequences decreased it by ∼5–20% for the SSIM$_1$ group (Kruskal-Wallis test, $P < 0.05$; colored asterisks in Fig. 8B–C). In contrast, adopting side-biased sequences reduced task efficiency by ∼6–9% during pre-training and for the SSIM$_{0.04}$ group. Therefore, adopting side-biased sequences yielded a local increase in task effi-

ciency for the SSIM$_1$ group, but the opposite occurred for the SSIM$_{0.32}$ and SSIM$_{0.04}$ groups.

The strength of the side-bias predicts individual learning trajectories. We asked how side-biased choices interact with discriminative learning by using the same visual task described in Figure 1A. We pre-trained a group of 88 naïve mice (Fig. 9A) and, using the first 30 training trials from day one of exposure to the task (average %correct choice record of around chance level; Wilcoxon test, $P > 0.5$, $n = 88$; not shown), extracted a side-bias index for each mouse as the average length of its biased sequences. Subsequently, we used the frequency distribution of these side-bias indexes to median-split the group of mice into faster and slower learners of this discrimination task (mean: 2.07 ± 0.11 trials; median: 1.8 trials; Fig. 9B). The values for correct-choice (Fig. 9C) and error records

Table 1 | Coefficients of the linear fits to the response rates of biased sequences against their reinforcement rates. We extracted the sensitivity (*i.e.* the slope), the intercept (*i.e.* the bias term) and the coefficient of determination for each regression (R^2; *i.e.* the square of the sample correlation coefficient between outcomes and predicted values)

Biased Sequences			
Sequence length	Sensitivity (m)	Bias (b)	R-squared
3	0.57 ± 0.02	0.26 ± 0.45	0.98 ± 0.00
4	0.42 ± 0.01	0.07 ± 0.38	0.98 ± 0.00
5	0.34 ± 0.01	0.05 ± 0.32	0.98 ± 0.01
6	0.28 ± 0.02	0.31 ± 0.29	0.96 ± 0.02
7	0.24 ± 0.01	0.53 ± 0.08	0.93 ± 0.04
8	0.21 ± 0.01	0.80 ± 0.19	0.91 ± 0.06
9	0.19 ± 0.01	1.25 ± 0.19	0.87 ± 0.07
Alternating Sequences			
Sequence length	Sensitivity (m)	Bias (b)	R-squared
3	0.37 ± 0.03	0.81 ± 0.36	0.90 ± 0.09
4	0.28 ± 0.03	0.50 ± 0.15	0.86 ± 0.12
5	0.19 ± 0.02	0.75 ± 0.26	0.72 ± 0.23
6	0.12 ± 0.04	1.33 ± 0.38	0.47 ± 0.21
7	0.06 ± 0.07	1.97 ± 0.70	0.38 ± 0.15
8	0.01 ± 0.12	2.66 ± 1.17	0.22 ± 0.13
9	0.04 ± 0.15	3.23 ± 1.24	0.22 ± 0.10

(Fig. 9D) remained well separated between the two groups (repeated measures ANOVA test, $P = 0.0001$). Faster learners had a higher discrimination performance than slower ones at the end of training (faster learners: 96% ± 1% correct; slower learners: 90% ± 2% correct; paired t-test, $P = 0.02$; two-sample Kolmogorov-Smirnov test, $P < 0.01$), providing performance extremes suitable for additional comparisons.

We assessed the robustness of the classification method by making random combinations of the mice's membership in each sub-group (without altering the characteristics of the original data set) and testing whether the sub-groups could result from a random classification of the mice; *i.e.* the null hypothesis, H_0, without taking the classification index into consideration. A repeated measures ANOVA test rejected the H_0 in favor of the alternative hypothesis because only 4 out of 9999 tests presented equal or higher F values than the one obtained by using the side-bias index as a classifier ($P = 0.0004$; $F_{obs} = 7.11$; Fig. 9E). Next, we calculated the probability of detecting group differences between randomly-selected faster and
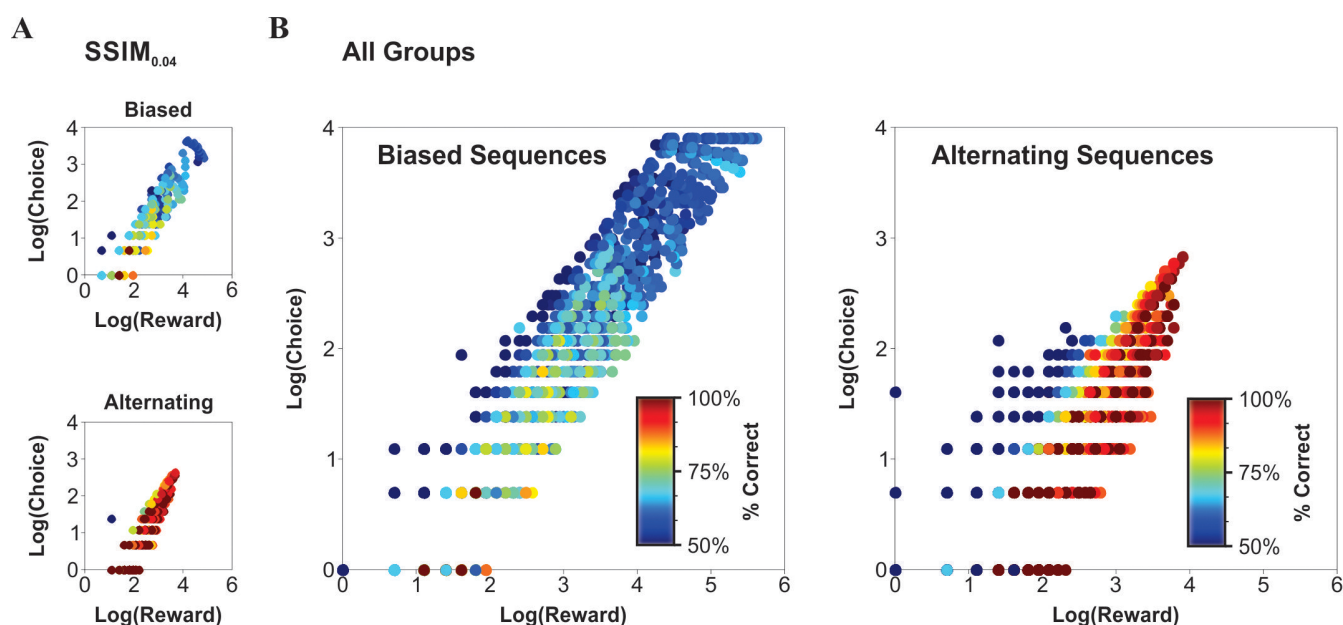
Figure 6 | Side-biased sequences have a correct choice probability around chance level. Log frequency of choices *vs.* log reinforcer frequency with color representing the group average %correct choice probability. (A) Although sporadic, the mice displayed some biased sequences with %correct choice values around chance level in conditions of high discriminability (SSIM$_{0.04}$). In contrast, they showed higher performance levels with alternating sequences (lower panel). (B) Pooling the sequences from all groups reveals that the mice were not discriminating when choosing biased sequences, but they did discriminate when choosing some of the alternating ones.
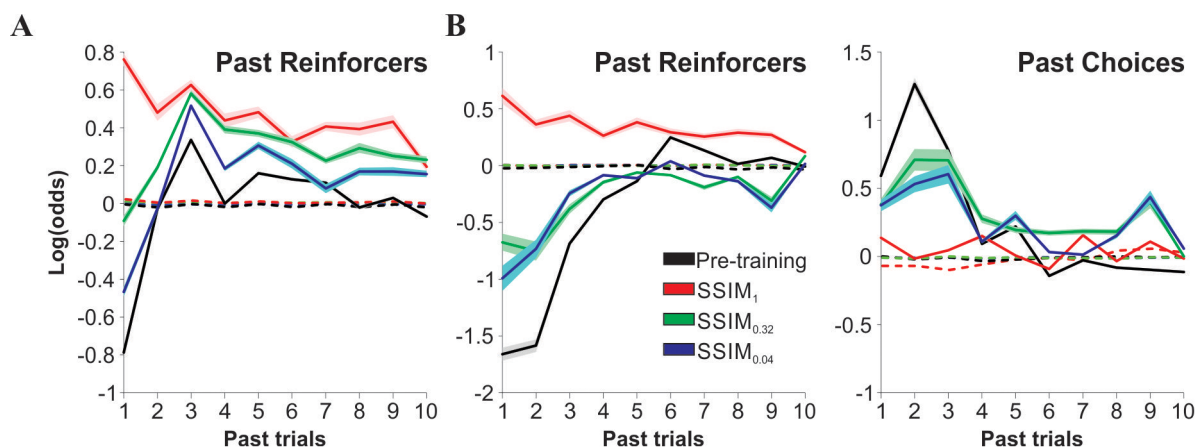
Figure 7 | Side-biased behavior mainly depends on past reinforcers. Weighted coefficients for two multi-regression linear models as a function of the number of past trials relative to the current trial. These coefficients estimate how past reinforcers (A) and past reinforcers plus past choices (B) influence current choice behavior (see Methods). The coefficients of the model were computed using a history of 10 trials and their statistical significance was determined by using a permutation test that shuffled 1,000 times the trial order for the variable of interest to yield a *P* value for the permutation test (dotted lines).

slower learners as a function of sample size (Fig. 9F). By generating 1000 random combinations per group, we found that these differences could be detected for groups of ≥20 mice.

Discussion

The mice used multiple strategies with different efficiencies to solve our two-alterative, forced-choice, visual discrimination task[5,17]. One task-solving strategy consisted in adopting a side-biased behavior in

which they repeatedly swam to the same side of the pool. Indeed, we found that the mice could display idiosyncratic individual biases when choosing between alternative routes; the prevalence of such side-biased behavior increased with stimulus similarity. Our metrics for side-biased choices indicate that such behavior was consistent across animals and prevalent in all groups. Clearly, in conditions of high visual discriminability (SSIM$_{0.04}$), the mice decreased the usage of both biased and random search strategies in favor of visually-
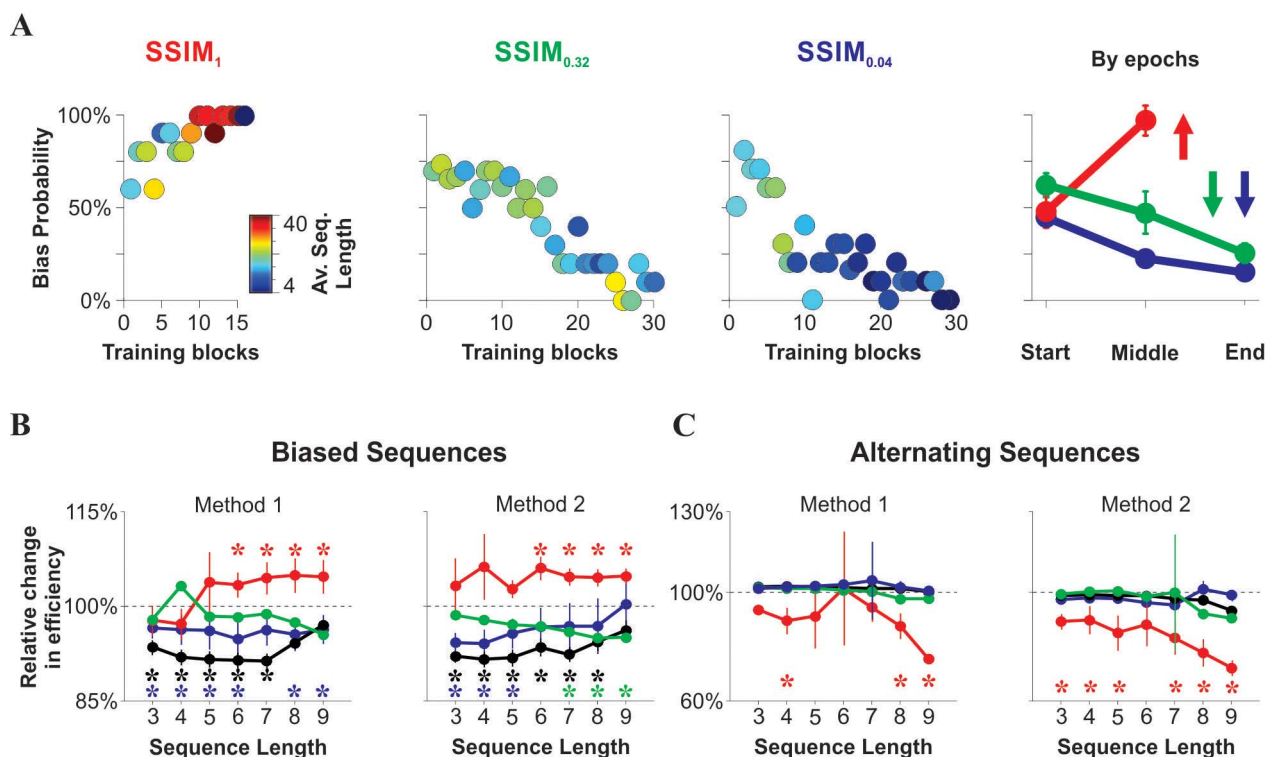
Figure 8 | Dynamic changes in the probability of bias formation and its local efficiency. (A) The average probability of finding biased sequences of ≥4 trials per mouse as a function of training (blocks of 10 training units). The panel on the right depicts the group probability of finding biased sequences of ≥4 trials within epochs of 30 training units placed either at the beginning, middle, or end of training. The relative changes in task efficiency depending on the choice of biased (B) and alternating sequences (C) were assessed with two different methods. In the first one, we divided the average efficiency of the query sequence by the average efficiency of a sequence of trials of equal length taken just before executing the query sequence (*i.e.* to estimate the baseline). In the second method, we divided the average efficiency of the query sequence by the average efficiency of 10 trials observed 10 trials before executing the query sequence. Asterisks depict significant changes (Kruskal-Wallis test, *P* < 0.05).
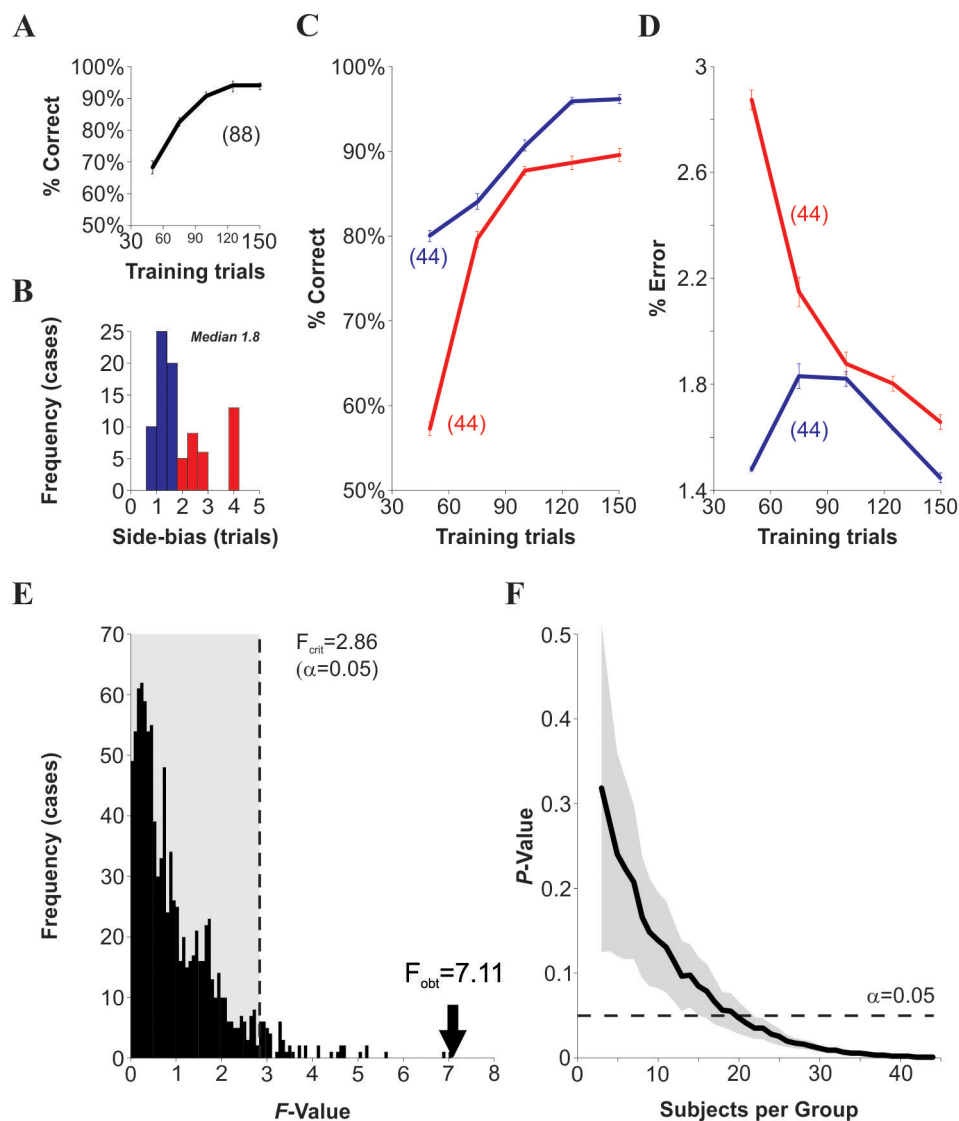
Figure 9 | Side-bias predicts individual differences throughout learning. (A) Performance of 88 behaviorally-naïve mice during pre-training. (B) Frequency histogram for the average side-bias index during the first day. Two sub-groups were formed by a median split (black and gray bars). Learning curves (C) and % errors (D) for these groups yielded inferior and superior learners that remained well-separated throughout training. (E) We tested the null hypothesis (H_0) that random permutations of the mice might produce sub-groups with similar behavioral differences against the alternative hypothesis (H_a) which claimed that side-bias was an effective index to classify the population. The frequency distribution for F is positively skewed. The one-sided P-value of the test was calculated as the proportion of sampled randomizations where F was greater than or equal to the observed F_{obs} values. The results indicate that reject H_0 in favor of H_a ($\alpha = 0.05$). (F) P-value for the one-way ANOVA comparison between randomly-selected sub-groups of different amount of classified mice.

guided ones. Yet, the prevalence of side-biased behavior increased with lower discriminability. We have shown recently that abrupt drops in the amount of side-biased choices precede the onset of a successful discrimination in individual mice[7]. One attractive possibility is that changes in the amount of lateral choices could indirectly reflect the involvement of attentional mechanisms engaged in the task.

A recent review by G. Vallortigara and L.J. Rogers provides a variety of interesting examples of perceptual and behavioral asymmetries among vertebrates. Such examples include the preferential use of a visual hemifield during activities such as foraging or escape from predators[9]. Interestingly, these authors suggest that lateralization on a population level can provide animals with some clear disadvantages. At the individual level, any bias favoring left or right position could leave the animal less able to attend, or to respond to stimuli that appear on the non-preferred side. On the other hand, at

the group level, if more than 50% of the individuals in a population were to show a similar direction of bias, then their behavior would become predictable to others[9]. In contrast, individually lateralized subjects could benefit from having individually different biases, because their escape responses would be less predictable to a predator. Furthermore, lateralized individuals, particularly those with lateralized eyes, could perform two tasks controlled by opposite brain hemispheres at the same time[10]. Our analysis revealed that lateralization varied in strength and polarity from mouse to mouse, but had a value close to zero when averaged across the population. Although our results indicate that average population laterality was zero, this does not imply that individuals were not lateralized or that individual lateralization does not have a detrimental influence in the estimation of psychophysical measures.

Our experimental paradigm was based on a training schedule that prohibited placing the discriminative stimulus and platform in the

same arm of the pool for more than 3 consecutive trials. This implied that choosing the same arm systematically would have produced a correct choice within a training unit after a maximum of three error repetitions, whereas other less efficient strategies, such as making random choices, could have reached up to five error repetitions. By implementing a measure based on counting the number of swimming trials that were required to solve each training unit, we found that the mice displayed multiple response variants with different swimming efficiencies that arose from the selective usage of discriminative information (in our specific working conditions: efficiency and %correct choices were not linearly related). In this respect, we found that choosing side-biased sequences with zero discriminability optimized the task efficiency locally, whereas choosing alternating sequences decreased it. In contrast, adopting a biased behavior in conditions of high discriminability reduced the task efficiency. Moreover, the prevalence of side-biased strategies was dynamic and could either increase or decrease, depending on the amount of discriminative information present. These observations strictly depend on the implicit rules, specific to each behavioral task. To our knowledge, this is the first report that shows an increase in task efficiency using lateralization at the individual level in a visual discrimination task for mice.

To optimize their choices, animals must continually update their behavioral strategies according to changes in their surroundings. Reinforcement learning theories provide a powerful theoretical framework for the understanding of choice behavior in dynamic environments, for they hold that future actions are chosen so as to maximize a long-term sum of positive outcomes, which can be accomplished through a set of value functions that represent the amount of expected reward associated with particular actions[15]. To match behavior to income, animals must integrate the rewards earned earlier from specific behaviors, and maintain an appropriate representation of the value of competing alternatives (i.e. reward frequency). Quantitatively, this is captured by the matching law[11], which states that the long-term average ratio of choices matches the long-term average ratio of reinforcers. Our results showed that biased-choices fully complied with the matching law, suggesting that the mice were able to discover the implicit rules of the task (v.gr. the statistical properties of the spatial distribution of the S^D) and update their behavior to improve their income.

Recent studies with monkeys[14] and rats[18] show that the history of past rewards exerts a strong influence on current choice. We implemented a response model to predict individual choices based on weighted combinations of recently-obtained reinforcers as well as previous choices[14]. Our findings showed that the location of past reinforcers, but not past choices, strongly influenced subsequent choices made by the mice from the $SSIM_1$ group (i.e. with stimulus similarity of 1) which had a high prevalence of side-biased behavior. For these mice, the positive coefficients in the logistic model implied a tendency to persist on biased behavior, whereas the negative coefficients for the $SSIM_{0.32}$ and $SSIM_{0.04}$ groups indicated a tendency to alternate. Also, the decaying effect of past reinforcers in the $SSIM_1$ group was much slower than in $SSIM_{0.32}$ and $SSIM_{0.04}$[15]. Therefore, these results demonstrate that mice can adjust the degree to which memories of past reinforcers influence their behavior. We propose that there should be an adaptive balance between the usage of discriminative information and side-biased strategies to solve discrimination problems.

Using a similar task, Carandini and co-workers provide a clear example that illustrates how the presence of spatial-biases can be detrimental when psychometric curves are being characterized for individual mice[17]. Their results indicate that the mice's discriminative choices are influenced not only by sensory information, but also by estimates of reward value, recent failures and past rewards. Their mice followed sub-optimal strategies influenced by non-visual factors and showed large spatial biases which varied slowly over the

daily sessions. Our data are in agreement with their observations, but also revealed that biased-strategies can emerge in the presence of a well-balanced, constant reward regime, and that alternation in the laterality of biases can occur in just a few trials. A question that remains open: what exactly triggers and influences the alternation in the preferred side when mice are strongly lateralized?

Our analysis further evidenced that the choices were not independent during steady-state behavior and that large side-biases were not random because they were consistent across animals (see also ref. 17). How do discriminative (sensory) and biased (non-sensory) strategies interact during discrimination learning in this behavioral task? We developed a robust method to sort mice into fast and slow learners using the side-bias index as a classifier. We found that the strength of the side-bias, collected during the first day of training, predicted individual differences in the average learning of the mice performing this task.

Variability in behavior provides the means by which new behaviors can be developed and individual factors are gaining recognition in behavioral neuroscience because they tend to correlate with the presence and severity of many neurobiological alterations[19]. Altogether, this framework constitutes a powerful tool to dissect the learning trajectories of individual mice performing this and other discrimination tasks.

1. Huberman, A. D. & Niell, C. M. What can mice tell us about how vision works? *Trends Neurosci.* **34**, 464–473 (2011).
2. Bonin, V., Histed, M. H., Yurgenson, S. & Reid, R. C. Local diversity and fine-scale organization of receptive fields in mouse visual cortex. *J. Neurosci.* **31**, 18506–18521 (2011).
3. Niell, C. M. & Stryker, M. P. Highly selective receptive fields in mouse visual cortex. *J. Neurosci.* **28**, 7520–7536 (2008).
4. Prusky, G. T., West, P. W. & Douglas, R. M. Behavioral assessment of visual acuity in mice and rats. *Vision Res.* **40**, 2201–2209 (2000).
5. Treviño, M., Frey, S. & Kohr, G. Alpha-1 adrenergic receptors gate rapid orientation-specific reduction in visual discrimination. *Cereb. Cortex* **22**, 2529–2541 (2012).
6. Bussey, T. J., Saksida, L. M. & Rothblat, L. A. Discrimination of computer-graphic stimuli by mice: a method for the behavioral characterization of transgenic and gene-knockout models. *Behav. Neurosci.* **115**, 957–960 (2001).
7. Treviño, M. *et al.* Controlled variations in stimulus similarity during learning determine visual discrimination capacity in freely moving mice. *Sci. Rep.* **3**, 1048 (2013).
8. Prusky, G. T. & Douglas, R. M. Characterization of mouse cortical spatial vision. *Vision Res.* **44**, 3411–3418 (2004).
9. Vallortigara, G. & Rogers, L. J. Survival with an asymmetrical brain: advantages and disadvantages of cerebral lateralization. *Behav. Brain Sci.* **28**, 575–589 (2005).
10. Csermely, D. & Regolin, L. *Behavioral Lateralization in Vertebrates: Two Sides of the Same Coin* (Springer, 2012).
11. Herrnstein, R. J. Relative and absolute strength of response as a function of frequency of reinforcement. *J. Exp. Anal. Behav.* **4**, 267–272 (1961).
12. Corrado, G. S., Sugrue, L. P., Seung, H. S. & Newsome, W. T. Linear-Nonlinear-Poisson models of primate choice dynamics. *J. Exp. Anal. Behav.* **84**, 581–617 (2005).
13. Stuttgen, M. C., Yildiz, A. & Gunturkun, O. Adaptive criterion setting in perceptual decision making. *J. Exp. Anal. Behav.* **96**, 155–176 (2011).
14. Lau, B. & Glimcher, P. W. Dynamic response-by-response models of matching behavior in rhesus monkeys. *J. Exp. Anal. Behav.* **84**, 555–579 (2005).
15. Sugrue, L. P., Corrado, G. S. & Newsome, W. T. Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat. Rev. Neurosci.* **6**, 363–375 (2005).
16. Roitman, J. D. & Shadlen, M. N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475–9489 (2002).
17. Busse, L. *et al.* The detection of visual contrast in the behaving mouse. *J. Neurosci.* **31**, 11351–11361 (2011).
18. Huh, N., Jo, S., Kim, H., Sul, J. H. & Jung, M. W. Model-based reinforcement learning under concurrent schedules of reinforcement in rodents. *Learn. Mem.* **16**, 315–323 (2009).
19. Pawlak, C. R., Ho, Y. J. & Schwarting, R. K. Animal models of human psychopathology based on individual differences in novelty-seeking and anxiety. *Neurosci. Biobehav. Rev.* **32**, 1544–1568 (2008).

Acknowledgments

Author contributions

M.T. conceived research, designed analysis algorithms, analyzed the data, wrote the manuscript and made the figures; all authors revised and approved the final version of the manuscript.

Additional information

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Treviño, M. Stimulus similarity determines the prevalence of behavioral laterality in a visual discrimination task for mice. *Sci. Rep.* **4**, 7569; DOI:10.1038/srep07569 (2014).