scientific reports

OPEN

Implicit and explicit learning of Bayesian priors differently impacts bias during perceptual decision-making

V. N. Thakur[1], M. A. Basso[2], J. Ditterich[3] & B. J. Knowlton[1]✉

Knowledge without awareness, or implicit knowledge, influences a variety of behaviors. It is unknown however, whether implicit knowledge of statistical structure informs visual perceptual decisions or whether explicit knowledge of statistical probabilities is required. Here, we measured visual decision-making performance using a novel task in which humans reported the orientation of two differently colored translational Glass patterns; each color associated with different orientation probabilities. The task design allowed us to assess participants' ability to learn and use a general orientation prior as well as a color specific feature prior. Classifying decision-makers based on a questionnaire revealed that both implicit and explicit learners implemented a general orientation bias by adjusting the starting point of evidence accumulation in the drift diffusion model framework. Explicit learners additionally adjusted the drift rate offset. When subjects implemented a stimulus specific bias, they did so by adjusting primarily the drift rate offset. We conclude that humans can learn priors implicitly for perceptual decision-making and depending on awareness implement the priors using different mechanisms.

Regularities in the environment can be learned implicitly and can bias judgments, preferences or fluency of movement[1–7]. For example, the statistical regularities of one's native language are readily acquired by infants[8]. Humans are also able to learn to use contextual information to search the location of a target in a display when the location is correlated with contextual features. Contextual cuing can occur without awareness of the correlation[9]. Implicit learning of regularities in finite-state artificial grammars also occurs through exposure to exemplars formed according to the grammar. After such exposure, participants classify new items as grammatical despite lacking awareness for the grammatical rules[5,10–12]. Perceptuo-motor behavior is also influenced by implicitly learned statistical regularities[13,14].

Here, we test the hypothesis that base rate priors that are implicitly acquired can influence orientation judgments. It is currently unclear whether implicit knowledge could become integrated with diagnostic information that is present to influence a judgement. In some circumstances, base rates, or Bayesian priors, influence judgements[15–23]. For example, a cloudy sky in Seattle might lead one to grab an umbrella when leaving home, while the same sky would not lead to this decision in Los Angeles. Much of our knowledge of regularities of the environment is learned without awareness, and thus it seems adaptive that such knowledge should contribute to judgments. On the other hand, such knowledge may not integrate readily with perceptual decisions that are based on sensory information in awareness such that only explicit knowledge of priors would be able to influence judgments.

If implicit knowledge of priors influences perceptual judgments, it may be that this knowledge affects judgments differently than explicit knowledge of priors. To compare decision-making under these different conditions, we applied the Drift Diffusion Model (DDM) to behavioral data and assessed how model parameters shift when subjects are aware of priors compared to when priors are implicitly learned. The DDM has been used extensively to effectively explain psychometric and neural data from perceptual choice tasks[22,24–26].

[1]Department of Psychology, University of California – Los Angeles, Los Angeles, CA, USA. [2]Fuster Laboratory of Cognitive Neuroscience, Departments of Neurobiology and Psychiatry and Biobehavioral Science, Jane and Terry Semel Institute for Neuroscience and Human Behavior, University of California – Los Angeles, Los Angeles, CA, USA. [3]Center for Neuroscience and Department of Neurobiology, Physiology, & Behavior, University of California – Davis, Davis, CA, USA. ✉email: knowlton@psych.ucla.edu

Results

Participants made judgments about the orientation (45° left or right) of a dynamic Glass pattern stimulus[27]. We parameterized the difficulty of the perceptual judgement on each trial by varying the dot pair correlations, referred to as coherence, ranging from 0 to 100% with 0% having no dot pair correlations and thus no orientation signal and with 100% having all dot pairs correlated and thus the strongest orientation signal. On each trial, after the appearance of a centrally-positioned fixation spot for a fixation time of 1000 ± 200 ms, two choice targets appeared, one on the left and the other on the right of the screen, to orient the participants. Then, a Glass pattern stimulus appeared and the fixation spot disappeared and the participants reported their perception with a key press; either 'o' for left or 'p' for right. If the participant chose correctly, an audible tone provided feedback. Incorrect choices received no tone and the 0% coherence trials received feedback consistent with the base-rate prior for that condition (e.g., 50% of the trials in the condition in which the two orientations were equally likely; Fig. 1a).

The critical aspect of this task was that each Glass pattern stimulus appeared in one of two different colors (red or green) and for each participant, the probability of the left or right orientation for each color differed. Thus, stimulus color provided information that could contribute to the perceptual decision about orientation, and in this sense knowledge of the color-orientation relationship acts as a Bayesian prior[25,28]. The stimulus color and prior probabilities were counterbalanced across participants, hence we converted the orientations to positive (right) and negative (left) orientations. In Experiment 1, referred to as the 75–50 condition, one colored Glass pattern had positive orientation priors and the other had equal orientation priors (Fig. 1b). Similar to Experiment 1, Experiment 2 also used two colored Glass patterns with different prior probabilities (Fig. 1c). In Experiment 2, each color was associated with a positive prior and negative prior with equal strength, so there was no base-rate difference for the two orientations across all stimuli. Previous studies on perceptual decision-making have not directly examined whether knowledge of priors was explicit or implicit. Here, we were able to classify participants' knowledge of the priors based on their responses to a questionnaire given after the session, and compare their performance to that of participants who were informed of the priors at the onset of the task.

To characterize differences based on the type of knowledge of the priors, we examined performance using the DDM. The DDM considers a process where the net evidence—the difference in evidence for or against a particular outcome—accumulates with time. When the amount of evidence reaches a level associated with one or the other bound or threshold, a decision is made (Fig. 1d). In this framework, priors can create biases in decision-making either by changing the starting point from which evidence begins to accumulate to be closer to the bound associated with the more frequent alternative[17,19,29,30] (Fig. 1e), or by increasing the rate of evidence accumulation for the more frequent alternative[1,23,24,31] (Fig. 1f), or both[21,25,32]. We modeled the data obtained from Experiments 1 and 2 using a simple DDM (see "Methods" section) to assess the mechanisms underlying how implicit and explicit knowledge of priors create perceptual decision biases.

Experiment 1—participants can learn priors implicitly.

Data from 67 participants were analyzed in Experiment 1, all with informed consent using procedures approved by the UCLA IRB. Based on responses to the questionnaire, participants were divided into three groups. Participants were classified as "Implicit Learners" if they reported no awareness that the different colors were associated with different orientation base-rates and that they perceived the orientations to occur with equal frequency (N = 23, 34%). Participants were classified as "Partial Explicit Learners" if they reported an overall difference in the orientation frequency but were unaware that this differed by color or could only identify the orientation prior for one color (N = 20, 30%). A group of participants (N = 24, 37%) classified as 'Explicit Learners' was informed of the priors for the different colored stimuli at the outset of the session and they all reported them correctly at the end of the session.

All participants performed the task well, using the orientation information to guide their choices (Fig. 2a–c). For non-zero coherence stimuli, percent correct was 83.14%, 81.16%, 79.54% for Implicit, Partial Explicit and Explicit groups respectively. This accuracy was highest in Implicit Learners and was significantly greater than Explicit learners: t(45) = 2.55, p = 0.014. All participants also showed a bias in making more positive choices. We calculated the bias by measuring the proportion of choices made at 0% coherence. For the Implicit Learners, the mean proportion of positive choices on the positive prior trials is 0.575, (Fig. 2a, difference from chance: t(22) = 31.88, p < 0.001). Furthermore, there was a 4.2% increase in positive choices compared to chance level even for the equal prior trials (Fig. 2a, 0.542; difference from chance: t(22) = 25.25, p < 0.001). The bias did not differ significantly between the two conditions (t(22) = 1.008, p = 0.289). The Partial Explicit Learners showed a similar trend in performance and the use of priors (Fig. 2b, mean proportion = 0.622; t(19) = 24.09, p < 0.001). For the equal prior trials, the Partially Explicit Learners also showed an increase in choosing the positive orientation (Fig. 2b, mean proportion = 0.551; t(19) = 20.25, p < 0.001), which was lower than for the positive prior trials (t(19) = 2.274, p < 0.05). Thus, implicit learning of both the general prior (Implicit group) and the color-specific priors (Partially Explicit group) occurred in the participants who were not informed of the base-rate differences. The mean bias for the positive prior trials in the Explicit group was 0.645, which was significantly above chance (Fig. 2c, t(23) = 28.08, p < 0.001). The Explicit Learners also showed an increase in bias for the equal prior condition that was significantly different from chance (Fig. 2c, 0.591; t(23) = 18.47, p < 0.001), but significantly lower than for the positive prior condition (t(23) = 2.316, p < 0.05). The Partial Explicit and Explicit groups performed similarly, with no significant difference between the ability to apply the priors associated with the two colors (p = 0.889), even though participants in the Partially Explicit group had minimal awareness of how stimulus color was associated with orientation probability.

Given previous work demonstrating that a correct choice on a previous trial can bias the decision on a current trial (Urai et al.[23]), we examined whether such "history biases" were present in the current data and whether these could account for the effect of priors we observed. For each group, we examined choices on zero-coherence trials to examine whether or not they differed based on the direction of the previous choice when it was correct. We

2
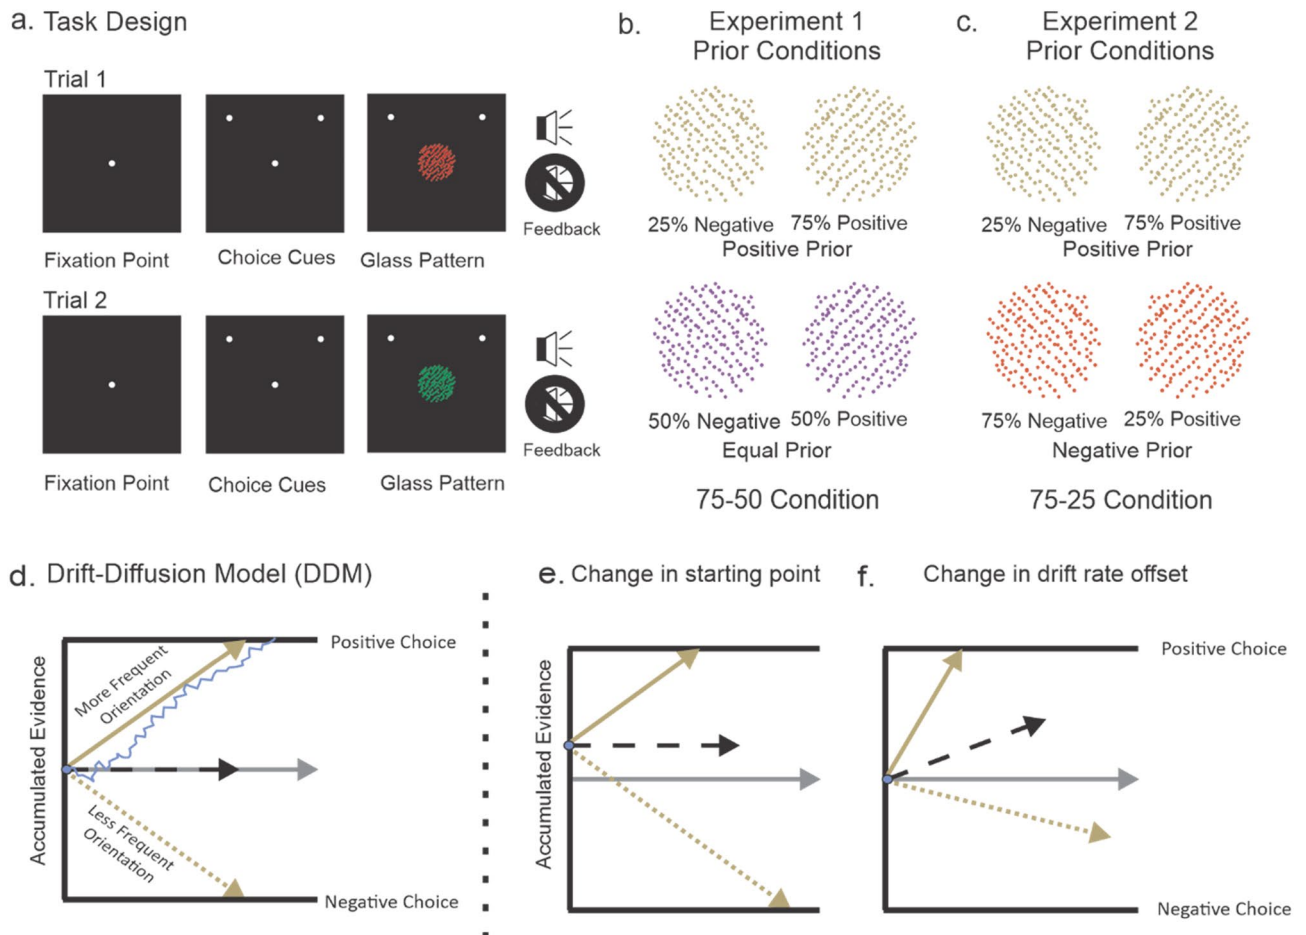
Figure 1. Schematics of the task design and the drift diffusion model (DDM). (**a**) The black squares show the screen that the participants viewed and illustrate the spatial and temporal arrangements of the task. The white circle in the center of the screen shows the fixation spot and the two additional white circles indicate the two possible choices; left or right. The fixation spot and the choice targets appeared sequentially with a delay of 1000 ms ± 200 ms between. Next the Glass pattern appeared and the fixation spot disappeared simultaneously, cueing the participant to report the perceived orientation with a key press ('o' for left and 'p' for right). A tone, indicated by the audio symbol, provided feedback only for correct choices. (**b**) Experiment 1: Khaki and lilac Glass patterns (100% coherence, 45 deg orientation negative and positive), illustrate the prior manipulation. The negative orientation occurred for the khaki Glass pattern on 25% of the trials and the positive orientation occurred on the remaining 75% of the khaki Glass pattern trials (Positive Prior). For the lilac Glass pattern stimuli, the negative and positive orientations occurred with equal probability (Equal Prior). All trials types were interleaved randomly and we counterbalanced the color and direction of the Glass pattern stimuli across participants and eventually converted to positive and negative orientation for analysis purpose. (**c**) Experiment 2: The negative orientation occurred on 25% of the khaki Glass pattern trials and the positive orientation occurred on the remaining 75% (Positive Prior). In the same experimental block, the negative orientation occurred on 75% of the orange Glass pattern trials and the positive orientation occurred on remaining 25% of the orange Glass pattern trials (Negative Prior). All trial types were randomly interleaved and we counterbalanced the color and direction of the Glass pattern stimuli across participants and eventually converted into positive and negative orientations for analysis purpose. See text and "Methods" section for further details. (**d**) The DDM of perceptual decision making proposes that sensory evidence is accumulated over time until a bound is reached, here labeled Positive and Negative. The solid khaki arrow shows the drift rate of evidence accumulation, labeled drift rate for more frequent orientations and the dashed khaki arrow shows the sane for less frequent orientations. The blue line shows the noisy evidence accumulation for an individual trial. In the absence of a bias, evidence for either the positive or the negative decision begins to accumulate at a point that is equidistant from the two bounds, referred to as the starting point of evidence accumulation (khaki arrow). Time is indicated along the horizontal axis. (**e**) Same as in (**d**) except the accumulation process begins at a point closer to Positive decision bound (shown by dashed black arrow). (**f**) Same as in (**d**) except with an increased drift rate offset for positive decisions (shown by dashed black arrow).
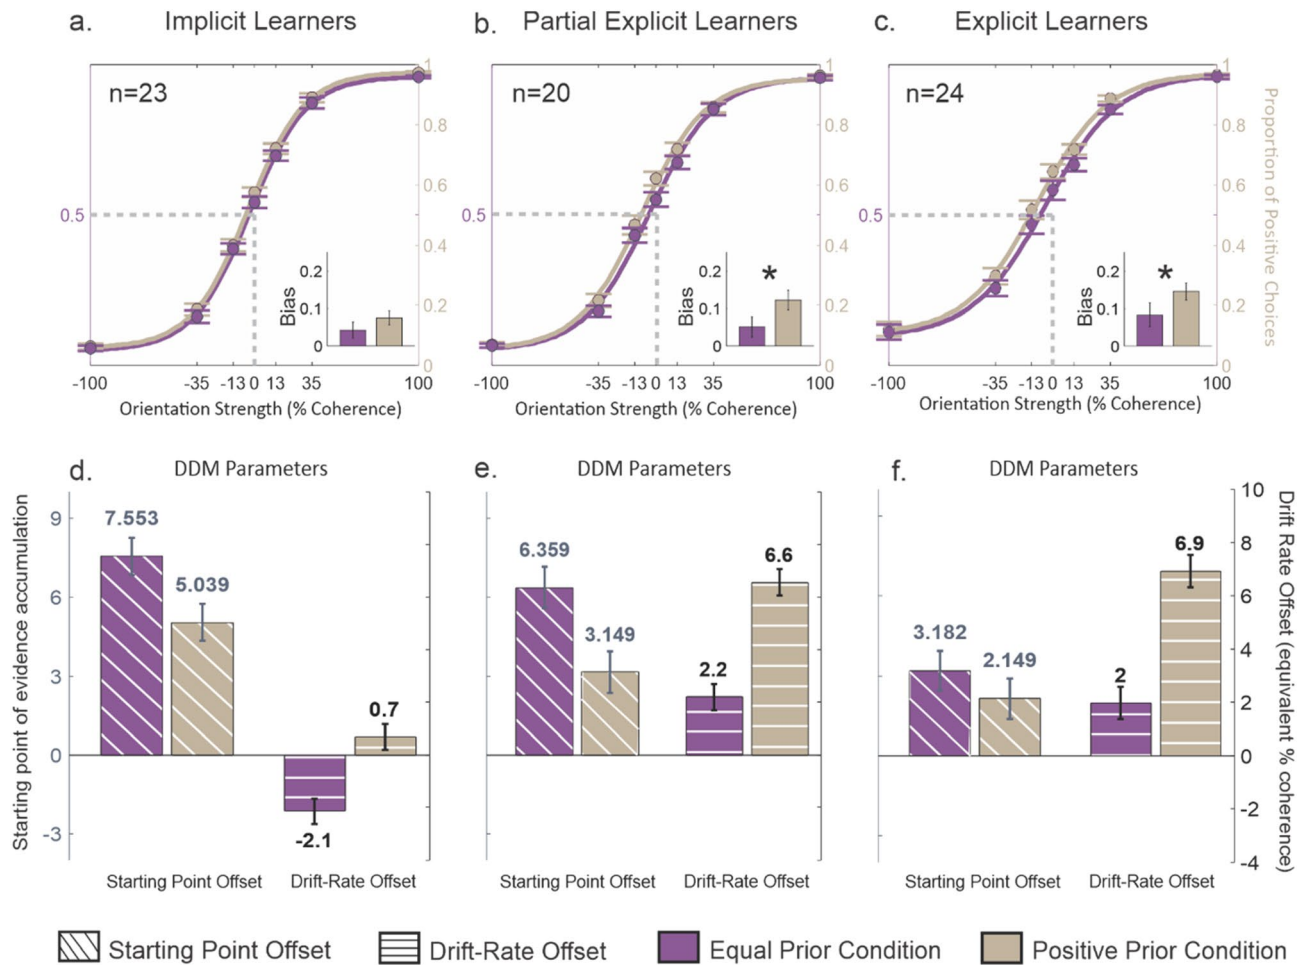
Figure 2. Participants learn multiple priors implicitly and explicitly. (**a**) Proportion of choices plotted against Glass pattern coherence for 23 participants. Filled lilac circles and error bars show the mean and SE for each coherence across all participants who reported no difference in orientation frequency (Implicit Learners) for the equal prior condition (50% positive and 50% negative). The lilac solid lines show the best fit logistic function (see "Methods" section). The khaki filled circles and lines show the same for the positive prior condition in which the more frequent orientation was positive, referred to as positive prior condition (75% positive and 25% negative). The inset shows the bias parameter for equal (lilac) and positive (khaki) priors represented as deviation of the proportion of positive choices at 0% coherence from unbiased condition (0.5). (**b**) Same as in (**a**) for the participants who reported a difference in orientation frequency regardless of color or for only one color (Partially Explicit Learners; n = 20). (**c**) Same as in (**a**) for the participants who were informed of the prior conditions (Explicit Learners; n = 24). (**d**) Parameter estimates from the DDM fits for the equal prior condition (lilac) and the positive prior condition (khaki) for all Implicit Learners. Diagonal hatched bars show the mean and 95% confidence intervals (CI) of the starting point offset and horizontal hatched bars show mean and 95% confidence interval for the drift-rate offset (equivalent % coherence). (**e**) Sarnia as in (**d**) for the Partially Explicit Learners. (**f**) Same as in (**d**) for the Explicit Learners.

found no evidence of history biases in any of the groups (t's < 1.6, p's > 0.12). Thus, it appears unlikely that biases were driven by the choice on the previous trial rather than across trials.

Implicit and explicit learners use different mechanisms to apply decision biases.   Using the DDM, we previously observed that healthy participants adjust both the starting point of evidence accumulation and the drift rate offset to express decision-making biases, and that it was the change in the drift rate offset that provided feature specificity to the bias[25]. However, in that study we did not know whether the healthy participants were aware of the feature-prior association or whether awareness would be associated with the use of different mechanisms to implement the bias. Here, we assessed whether the mechanism underlying the expression of the bias differed for the different types of learners, implicit, partially explicit and explicit.

We fitted two feature-specific DDMs to the second-half of the experimental session data independently (600 trials), allowing only the starting point and the drift rate offset to vary between the two features (Table S1). Figure 2d–f shows the mean parameter estimates and confidence intervals of the parameter estimates for the starting point and drift rate offsets for both prior conditions. For the Implicit Learners, the starting point shifted toward the more frequent orientation for both the equal and the positive prior trials. The shift in starting point

was similar for the two conditions (Equal Prior- estimate: 7.6%, S.E.: 0.7%, Positive Priors- estimate: 5%, S.E.: 0.7%). The change in starting point was numerically larger for the equal prior condition compared to the positive prior condition, but this was counteracted by an opposite change in the drift rate offset for equal prior trials (estimate: − 2.1%, S.E.: 0.5%). The drift rate offset for the positive prior condition was in the same direction as the more frequent orientation but with a small magnitude (estimate: 0.7%, S.E.: 0.5%). In summary, for the Implicit Learners, there was a significant increase in starting point for both prior conditions accompanied by a slight decrease in drift offset when priors were equal, leading to a bias in decision-making that was undifferentiated by stimulus color.

For the Partially Explicit Learners, we observed similar changes in the starting point of evidence accumulation and quantitatively larger changes in the drift rate offset compared to the Implicit Learners (Fig. 2e). The starting point of evidence accumulation for equal prior trials was 6.4% (S.E.: 0.8%) and 3.1% for the positive prior trials (S.E.: 0.8%). Similarly, the drift rate offset was positive for both prior conditions but in contrast with the starting point, the drift rate was much higher for the positive prior condition (6.6%, S.E.: 0.5%) compared to the equal prior condition (2.2%, S.E.: 0.5%). Thus, both Implicit and Partially Explicit Learners expressed a decision bias by changing the starting point, however, only the Partially Explicit Learners developed a sensitivity to priors associated with stimulus colors that appears to be mediated primarily by changes in the drift rate offset.

We found that the Explicit Learners also showed an overall bias but only a numerically greater feature specific sensitivity to the priors unlike the Partially Explicit Learners (Fig. 2b,c). The starting point offsets in the Explicit Learners were similar for both priors (Equal Prior: 3.2%, S.E.: 0.8%; Positive Prior: 2.1%, S.E.: 0.8; Fig. 2f). Similar to Partially Explicit Learners, the Explicit Learners also adjusted their drift rate offset asymmetrically to express the bias (Equal Prior: 2.0%, S.E.: 0.6%; Positive Prior: 6.9%, S.E.: 0.6%; Fig. 2f). Hence, similar to the Partially Explicit group, the feature-specific biases were primarily driven by drift rate offset in the Explicit group.

To see the effect of starting point offset or drift rate offset individually, we simulated the DDM model using optimized parameters with minor modifications. To identify the contribution of the starting point offset, we manually kept the drift offset at zero and measured biases at 0% coherence for both prior conditions. Similarly, to identify the contribution of the drift rate offset to the bias, we simulated the model with optimized parameters and manually set the start point offset to zero. We observed that with the full model the biases in the Implicit group, Partially Explicit group and Explicit group were: [Equal Prior, Positive Prior]: [0.531, 0.555]; [0.567, 0.602]; [0.585, 0.622] respectively. When we looked at only the starting point's contribution, we found simulated biases as follows: [Equal Prior, Positive Prior]: [0.555, 0.547]; [0.543, 0.531]; [0.567, 0.558] for the Implicit group, Partially Explicit group and Explicit group respectively. When only the drift offset was kept in the model, we found the biases for the Implicit, Partially Explicit and Explicit groups to be: [Equal Prior, Positive Prior]: [0.484, 0.524]; [0.534, 0.586]; [0.567, 0.610]. As we can see from the above analysis, the start point contributes toward an increase in the overall bias of the choice behavior and the drift offsets contribute toward the color specificity in all three groups including the Implicit group. However, the drift rate offsets in the Implicit group were smaller than start point offset leading to a lack of a significant color-specific bias.

In Experiment 1, participants were able to learn a Bayesian prior implicitly based on different base-rates of orientation and could apply this prior in a perceptual decision-making task. Participants were also able to apply priors differentially depending on stimulus color even if they were unaware that the colors were associated with different base rates. Implicit knowledge of the general prior, that one of the orientations occurred more frequently, was captured by a start point offset in the DDM. Explicit knowledge of the color-specific priors was captured mainly by changes in drift rate offset in the DDM. Participants who gained some knowledge of the frequency of the orientations across the task but did not become aware of the color-specific priors also showed differences in drift rate offset for the two conditions.

Experiment 2—learning feature-specific priors in the absence of a general prior. The design of Experiment 1, with an equal base rate prior for one stimulus color (50% for each orientation) and a positive prior for the other (75% for one orientation) led to a general prior of 67% towards the positive orientation across all stimuli. Participants in all groups generalized across the stimulus colors to some extent with biased responses toward the positive prior. Thus, it was difficult to tease apart learning feature-specific and general prior knowledge in the results. To examine learning of feature-specific priors in the absence of a task general prior, we conducted a second experiment in which the base rate orientation of the two stimulus colors were in opposite directions. So, The Glass patterns of one color occurred with 75% of trials with a positive orientation and 25% of the trials with a negative orientation, whereas the Glass patterns of the other color occurred with a positive orientation on 25% of the trials and a negative orientation on 75% of the trials (Fig. 1c). This arrangement ensured that the overall positive and negative orientation probabilities remained constant.

We analyzed the data for a different group of 41 participants in Experiment 2, all with informed consent and using procedures approved by the UCLA IRB. 21 participants were assigned to the Implicit Learner group as they were not told about the differences in base rates. None of these participants reported noticing base rate differences on the questionnaire given after the session. 20 participants were assigned to the Explicit group, and were informed of the different base rates for the two colors. Participants accurately classified 81.03% and 83.40% of non-zero coherence stimuli in Implicit and Explicit groups respectively and there were no differences between the groups (t(39) = − 1.48, p = 0.146). Unlike in Experiment 1, the Explicit group performed numerically better than the Implicit group, suggesting that the lower performance observed in the Explicit group in Experiment 1 was not reliable. Figure 3a shows that even though Implicit Learners were unaware of the priors, they showed a feature-specific bias. The mean bias for the positive prior trials was 0.540 and for the negative prior trials was 0.423 and this difference in biases were statistically significant (Fig. 3a: t = 2.68, p < 0.001). As in Experiment 1, we found no evidence that correct choices on the previous trial influenced choices at zero coherence for either
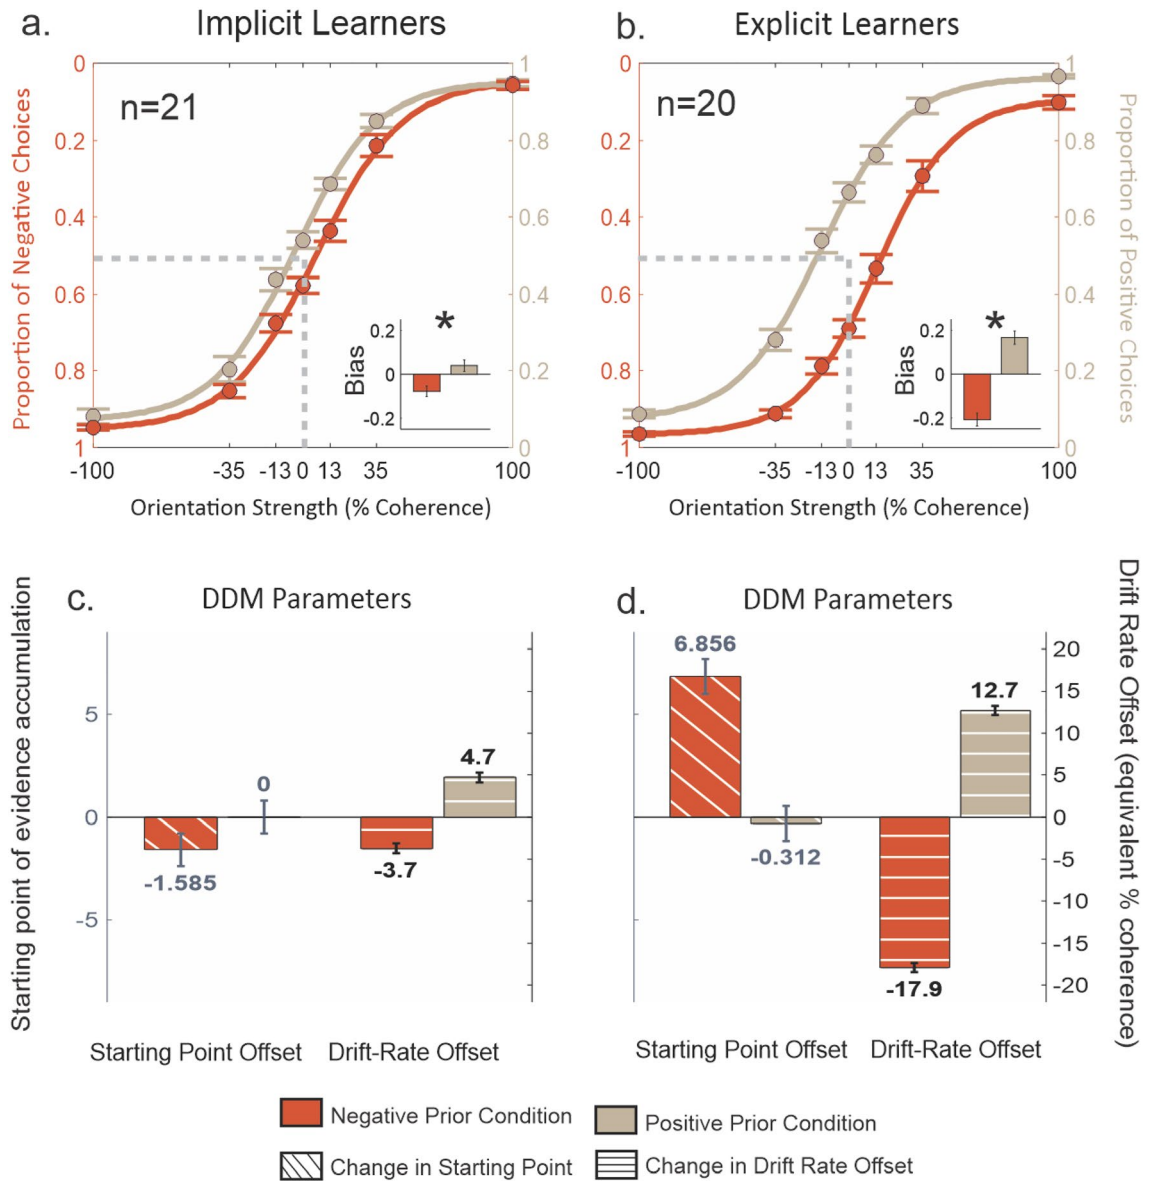
Figure 3. Participants learn and use stimulus specific orientation priors that bias perceptual decisions. (**a**) Proportion of choices plotted against Glass pattern coherence from 21 participants. Filled orange circles and error bars show the mean and one SE for each coherence across all participants who reported no difference in orientation frequency (Implicit Learners) unequal prior condition in which the more frequent orientation was negative referred to as Negative Prior condition (75% negative and 25% positive). The orange lines show the best fit logistic function (see "Methods" section). The khaki filled circles and lines show the same for the unequal prior condition in which the more frequent orientation was positive referred to as Positive Prior condition (25% negative and 75% positive). The inset shows the bias parameter for negative (orange) and positive (khaki) priors. It represents the deviation of proportion of positive choices for 0% coherence from the unbiased condition (0.5). (**b**) Same as in (**a**) for the participants who reported a difference in orientation frequency (Explicit Learners; n = 20). (**c**) Parameter estimates from the DDM fits for the negative (orange) and positive (khaki) prior conditions for all implicit learners. Diagonal hatched bars show the mean and 95% CI of the parameter estimates for the starting point offset and horizontal hatched bars shows the mean and 95% CI of the parameter estimate for the drift-rate offset (in equivalent % coherence). (**d**) Same as in (**c**) for the Explicit Learners.

group (t's < 0.42, p's > 0.68). Using the DDM, we found that both the starting point and the drift rate offsets were asymmetric with the task feature i.e., both were negative for negative prior and non-negative for the positive prior condition. The change in starting point for the negative priors was − 1.6% (S.E.: 0.8%; Fig. 3c) and for positive priors was 0% (S.E.: 0.8%; Fig. 3c). The drift rate offset parameter, like the starting point, changed such that the negative prior had a negative drift rate offset and the positive prior had positive drift rate offset: Negative Prior: 3.7%, S.E.: 0.6%; Positive Prior: 4.7, S.E.: 0.6%; Fig. 3c). These results are consistent with Experiment 1 in that

an implicitly learned feature-selective prior biases decision-making through both starting point and drift rate offset with a dominant effect from the drift rate offset.

Participants in the Explicit Group showed robust expression of feature-specific biases. The mean bias measured for the 0% coherence, positive prior trials was 0.662 and for the negative trials was 0.292 and these differences were significantly different (Fig. 3b; t(19) = 7.13, p < 0.001) and the difference in bias for the two colors was significantly different from those observed in the Implicit Learners (t(39) = 3.782, p < 0.001). The modeling results showed that the Explicit Group used drift rate offsets to express stimulus feature-prior associations, and these changes were much larger than those obtained in the Implicit Learners.

In the Explicit group, the starting point offset was actually in the opposite direction to the priors. The starting point estimates for the Explicit group were, Negative Prior: 6.9%, S.E.: 0.8%; Positive Prior: 0.3%, S.E.: 0.8%; (Fig. 3d—diagonal stripes). Similar to the Implicit Learners, the Explicit group also showed feature-specific changes in the drift rate offset and these changes were much more prominent: Negative Prior: – 17.9%, S.E.: 0.6%; Positive Prior: 12.7%, S.E.: 0.5%; (Fig. 3d—horizontal stripes). This result suggests that the change in starting point for the negative prior condition could be a compensatory mechanism for a high change in the drift rate offset as reported previously[25]. These results indicate that according to the DDM, people implement explicit priors through the enhancement of drift rate toward the decision boundaries.

In Experiment 2, we wanted to measure the individual contribution of starting point and drift rate offset when there were priors that were stimulus specific with no overall prior. We observed that with the full model the biases in Implicit Learners and the Explicit group were: [Negative Prior, Positive Prior]: [0.448, 0.550]; [0.318, 0.646] respectively. When we looked at only the starting point contribution, we found simulated biases as follows: [Negative Prior, Positive Prior]: [0.482, 0.507]; [0.500 0.516] for the Implicit and Explicit groups respectively. When the drift offset was kept in the model, we found the biases for the Implicit and Explicit groups to be: [Negative Prior, Positive Prior]: [0.457, 0.550]; [0.287, 0.648]. Similar to Experiment 1, the drift rate offset contributes to the stimulus specific biases. The starting point had a smaller contribution that appeared compensatory for the large drift rate offsets in the Explicit group.

Based on the results of the two experiments, implicit learning of base rate information can occur in both a general and a feature-specific manner, and this knowledge appears to influence decision-making through the more efficient accumulation of sensory evidence consistent with the priors. When participants are aware of the base rate differences, there are greater increases in the drift rate offset. Starting point offsets can also contribute to the stimulus specificity either by increasing the bias of smaller changes due to drift rate offset, as in the Implicit Learners, or by counteracting the large changes in drift rate offset seen in the Explicit group. Thus, starting point may be adjusted to both bias decision-making implicitly and possibly to compensate for large drift rate increases that implement explicit knowledge of priors.

Additionally, we also explored whether a single mechanism of bias implementation could explain the data better. For this we fitted two different models, one with only starting point offset and another with only drift rate offset and calculated their BIC scores which are shown in Table S3 of Supplementary Data. The BIC score for the two free-parameter model was lower for four out of five groups. This analysis suggests that the model with both start point offset and drift-rate offset as a free parameters explains the data better after accounting for the increase in the number of parameters. The fifth group, the Explicit group in Experiment 2, shows a lower BIC score for the model with only drift-rate offset. As discussed above, the Explicit group in Experiment 2 shows the highest separation of stimulus-specific biases that appears almost exclusively driven by the manipulating the drift-rate offset. Importantly, for implicit learners, behavior is captured most effectively by the model including both drift-rate and start point parameters.

Discussion

We assessed the ability of participants to learn base-rate priors implicitly and apply them in a perceptual decision-making task. In both experiments, participants acquired a bias in an orientation judgment task that was apparent when participants judged items with zero coherence. Participants who did not report awareness of the priors nevertheless exhibited evidence of bias in their decisions about the orientations of the stimuli. In Experiment 1, two different kinds of priors were present; an overall prior based on the frequency of the different orientations across all stimuli and feature-specific priors that were specific to stimulus color. We found that both types of priors could be learned implicitly. In the group that reported no explicit knowledge of the different priors of the two orientations (Implicit Learners), there was still a significant bias toward the more frequent orientation across all stimuli, which was also seen in participants who were informed of this prior or became aware of it during testing. Implicit learning of feature-specific priors was also seen in those participants who became aware of only the general difference in frequency between the orientations (Partially Explicit Learners). In this group, explicit knowledge of the structure of the task lagged behind implicit knowledge in that they did not become aware of differences in priors for the two colors despite applying the differential priors in their choice behavior. In Experiment 2, the two orientations were equally likely overall and the priors for the two colors were more distinct than in Experiment 1. Here, we also found implicit learning of feature-specific priors.

Our results are consistent with the recent study by Rungratsameetaweemana et al.[33] in which the ability of patients with amnesia following hippocampal damage to acquire a prior in a motion direction judgment task was measured. The sensitivity to prior information was similar in amnesic patients and control participants, supporting the idea that explicit learning of the prior does not benefit performance. We assessed awareness of priors in participants after the session to determine the level of explicit knowledge gained during the task. We also demonstrated that different priors specific to different sets of stimuli based on color could be learned implicitly. These data suggest that changes in perceptual processing can occur based on implicit knowledge of priors rather than only in overall response bias.

Using the DDM, we found that explicit and implicit knowledge of the priors appeared to influence decisions differently. Substantial drift rate modulation was associated with knowledge of the color-specific priors and was much more robust in groups with explicit knowledge than in groups with only implicit knowledge. The influence of starting point on decision-making was confined to implicitly learned knowledge of the priors. The Implicit group in Experiment 1, who learned to apply a bias towards the more frequent orientation, showed a significant starting point offset in the positive direction for both colors. In Experiment 2, where the priors were in equal and opposite directions, the bias in decision-making occurred as a result of significant changes primarily in drift rate offset for stimuli of the different colors. The difference in the drift rate offsets for the two stimulus colors was relatively modest for the Implicit group. In contrast, in the group that was aware of the different priors, there were substantial increases in drift rate offset towards the positive and negative direction for the two colors of stimuli. In Implicit Learners, the starting point supported further increases in bias by offsetting in the direction of the prior for one of the colors. In contrast, in the Explicit group the high change in drift offset rate was counteracted in part by a start point offset in the opposite direction for one of the priors. These results suggest that explicit knowledge of the priors leads to faster accumulation of evidence in the direction of the prior across coherences. Implicit knowledge of stimulus-specific priors appears to also affect drift rate offset. Start-point offsets appeared to reflect implicit knowledge of priors, and may even compensate for large drift rate offset changes implemented by explicit knowledge of priors.

The present work shows that people can implicitly learn base-rate differences and apply this knowledge in a perceptual-decision making task. While classic work in decision-making shows that participants often ignore explicit base-rates[34], base-rates may have a greater impact when they are learned through experience. An implicit learning mechanism may be engaged with experiential learning that may lead to knowledge that is more effective in influencing judgments than explicit knowledge[35]. We argue that implicit learning of general base-rate information shifts the starting point of evidence accumulation toward the more frequent outcome. Consistent with this idea, local fluctuations in base rates influence decision-making through starting point modulation, even when these local fluctuations do not reflect the explicit priors for the entire task[36]. Thus, starting point, perhaps reflecting biasing activity in neural representations of alternatives, may be modulated fairly automatically by reinforcement history.

Additionally, we also observed changes in the drift rate offset associated with implementing feature-specific biases in decision-making. When participants had knowledge that there were different priors that depended on stimulus color, there was a significant modulation of drift rate, indicating that awareness of feature-specific priors led to greater efficiency in accumulating perceptual evidence in support of the priors. Drift rate changes were also observed despite lack of awareness that the different stimulus colors were associated with different priors in both experiments. It thus appears that implicit knowledge of stimulus specific priors enhances the perceptual processes contributing to decision-making.

Modulation of starting point and diffusion rate are likely supported by different brain regions in the service of decision-making. Starting point offsets may reflect increased activity in neural representations corresponding to more frequent alternatives. In the seminal work of Basso and Wurt[37] and Platt and Glimcher[38], monkeys performing a two-alternative forced-choice task showed increased pre-trial activity in SC and area LIP representation for the alternative that was cued to be more likely. In Experiment 1, a similar increase in activity in representations corresponding to the more frequent alternative may have occurred. This increase may have occurred very early in the trial, before the integration of evidence about orientation. In the present study, subjects were not cued with base-rate priors but were able to acquire them implicitly through experience. Thus, it appears that modulation of start point does not require awareness of priors. Perugini et al.[25] found that patients with Parkinson's disease are impaired at adjusting starting point based on priors, suggesting that basal ganglia might also be involved in changes in starting point and adjusting decision threshold generally[39]. Future work may clarify whether modulation of starting point in decision-making relies on similar neural mechanisms for explicitly cued and implicitly learned priors.

The adjustment of drift rate offset in the present study may reflect enhanced processing of elements of the display consistent with the prior. Adjustment of drift rate has been shown to occur in decision-making paradigms in which there is variability of signal to noise across trials[40]; as variability of difficulty across trials increased, fitting the DDM to data relies more heavily on adjustments to the drift rate[41]. In the present study, where trials varied substantially in terms of difficulty, it is therefore not surprising that drift rates were offset resulting in decision biases. Although present in the Implicit group in Experiment 2, this finding was much more striking in subjects who gained partial knowledge of the base rates or were informed of the different priors at the onset of the task. One possibility is that drift rate offsets reflect top-down attentional effects that enhance gain in perceptual or decision-making regions.

Conclusion

In sum, our results demonstrate that knowledge about base-rate priors can be acquired implicitly and that this knowledge can be specific to stimuli with different features within a stimulus set. Within the DDM framework, the implicit acquisition of an overall bias was implemented through a change in starting point offset. However, feature-specific biases were primarily implemented through changing the rate of evidence accumulation. Future studies can help to identify the brain regions involved in implementation of general and feature-specific biases, and whether implicitly acquired biases are supported by different neural mechanisms than explicit biases.

Methods

Participants.   We recruited 140 people for this study. Experiment 1 included 89 participants (23 male and 66 female) and Experiment 2 included 51 participants (13 male and 38 female). All participants were right-handed UCLA students with age ranging from 18 to 30 years. None of the participants had active medical, neurological, or psychiatric diagnoses, and were not taking any chronic medication that could affect sensory processing, movement, or cognition. All participants had normal/corrected to normal vision. All the participants were tested for colorblindness using an online version of the Ishihara color blindness test. All the participants gave written informed consent as per the guidelines of the Institutional Review Board of the University of California and received course credit for their participation. In each experiment, subjects were assigned to either an Explicit group, in which they were informed of the different orientation probabilities for red and green stimuli or an Implicit group, in which they were not informed. All subjects were given instructions by the experimenter and completed at least 50 practice trials of stimuli that were 100% coherent with 85–90% accuracy to familiarize themselves with the task before beginning.

Visual stimuli.   Dynamic translation Glass patterns were used for this study, which was developed previously[25]. This task uses two identical dot patterns, one of which is translated in position with respect to the other. When these patterns are superimposed, multiple correlated dot pairs are formed, which results in a strong perception of the orientation of dots[27]. The strength of this orientation can be varied using the number of dot-pairs formed. For both experiments, we used 0%, 13%, 35% and 100% coherence, referring to the percentage of correlated dot pairs. In the 100% coherence condition, all dots pairs are correlated, and the orientation of the stimulus (right or left) is readily apparent. In contrast, the 0% coherence condition does not have any dot pairs which are correlated and thus subjects have no perceptual information upon which to base their orientation decision. The stimuli were made dynamic by generating 30 frames of such images and showing them with a frequency of 85 frames/s.

Experimental paradigm.   Each trial began with the appearance of a centrally-located white fixation box. Subjects were asked to keep looking at the fixation box. After fixation ranging from 1000 ms ± 200 ms, two target boxes appeared on the screen at left-upper and right-upper positions. Immediately after that, a Glass pattern with one of the coherences mentioned above was shown to the subjects. In each trial, the Glass pattern was shown either in red (luminance 0.6 cd/m^2) or green (luminance 0.54 cd/m^2), which was randomized with equal probability. Subjects had to gather the information about the orientation of the stimulus and record their response as quickly as possible by pressing a key on the keyboard. If they judged the stimulus to be oriented to the left, they pressed the 'O' key, and if they judged it to be oriented to the right, they pressed the 'P' key. The Subjects were given 3 s after the onset of the Glass pattern to decide the orientation. If they did not respond within 3 s, the trial was aborted. If they chose the correct orientation, the feedback was given as a high tone by the computer. For both experiments, each subject performed at least 1200 trials.

For Experiment 1, one color of Glass pattern was unbiased regarding orientation (both left and right were equally likely) and for the other color, the orientation occurred 75% in one orientation and 25% in the other. For Experiment 2, both the red and green patterns were biased, with one orientation occurring 75% of the time and the other 25% of the time, and the opposite probabilities occurred for the other color. Thus, the two orientations (left and right) occurred equally often across the stimulus set. The assignment of colors to conditions was counterbalanced across subjects for both experiments and later converted to positive orientation and negative orientation. In this way, we were able to examine whether both general biases (one outcome is more likely to occur in the task) and feature-specific biases (red and green stimuli have different base rates) can be learned in the absence of awareness.

Following the completion of the task, participants' awareness of the priors was assessed using a follow-up questionnaire. The questionnaire asked participants if they thought that each colored stimulus (red and green) was equally distributed between two orientations. Based on their response, participants were divided into two categories: (1) implicit learners and (2) partial explicit learners. Participants who responded that both orientations occurred equally for each color, i.e., there was no bias, were categorized as Implicit Learners. Participants who identified the prior orientation of only one stimulus category correctly were categorized as Partial Explicit Learners. In the Explicit group, participants were informed about the feature-prior associations at the start of the experiment.

Data analysis.   Behavior.   The reaction time (RT) was defined as the duration between onset of Glass pattern stimulus and response time (keypress) onset. All trials with RT less than 200 ms were excluded from the final analysis due to the possibility of anticipatory trials. Additionally, all participants whose performance at 100% coherence was below 90% accuracy were removed from the final analysis (n = 32, across all five groups). Hence, data was analyzed for 67 participants in Experiment 1 and 41 participants in Experiment 2. For both experiments, reaction times for all remaining participants were within two standard deviations at all coherences, and thus data were pooled for subsequent analyses.

Next, a psychometric function of accuracy with respect to coherency was calculated. This psychometric function was fitted using the following logistic function of the form:

$$p(P) = \lambda_1 + \frac{1 - \lambda_1 - \lambda_2}{1 + \exp(-\beta \cdot (C - \alpha))},$$

where $p(P)$ is the proportion of choices towards more frequent orientation, i.e., for 75–50 condition with positive prior, $p(P)$ is the proportion of positive choices; C is dot pair coherence; α and β are free to model parameters

and are determined by the maximum likelihood method. They also represent response bias and sensitivity to stimulus, respectively. λ_1 and λ_2 represent the lapse rates for the two sides, which equal the difference between perfect perfomance and actual performance. It indicates the transient lapse in attention during the experiment[42]. As 0% coherence is an utterly ambiguous stimulus, subjects' biases can be assessed at that coherence condition. The use of priors for each stimulus color was calculated as the proportion of choices towards the more frequent orientation at 0% coherence level per subject. The statistical analysis was carried out using pairwise t-tests.

Modeling. *Drift diffusion model* One well-known example of the class of evidence accumulation models is the Drift Diffusion Model (DDM). This model explains both, the reaction time and the choice of every trial. We used one of the variants of this model with a collapsing boundary, i.e., as time passes, participants require less accumulation of evidence to make a decision to account for urgency. We fit the DDM model with 11 parameters on pooled data across all participants in a given condition of the experiment. These 11 parameters were: the non-decision time, starting point of accumulation, proportionality factor between coherence and drift rate, the diffusion coefficient, a scaling parameter for collapsing bound, delay in collapsing bound, drift rate offset, proportion of uninformed positive choices, proportion of uninformed negative choices, mean and standard deviation of reaction time for uninformed choices. The biases can be implemented in the DDM-like model by either changing the starting point of evidence accumulation (Fig. 1e) or by changing the rate of evidence accumulation across all coherences i.e., drift rate offset. This drift rate offset can also be interpreted as the apparent coherence during 0% coherence stimulus (Fig. 1f). Each of these methods predict slightly different behavior. The change in starting point predicts an asymmetric change in RT between positive and negative choices. On the other hand, a drift rate offset predicts a similar change in RT for positive and negative choices, such that zero coherence trials are no longer the slowest ones. To assess these two mechanisms, we fitted two DDM models, one for each stimulus color, such that five of the model parameters, namely, non-decision time, proportionality factor, the diffusion coefficient, scaling and delay of collapsing bound were shared by both models. The other two parameters, starting point and drift rate offset, were allowed to differ between both models, thus allowing us to assess which parameter change or combination of changes could better explain the behavioral data. The DDM model was fitted to only the second-half of the session for each subject, the rationale being that during the first half of the session subjects are still learning these priors and decision behavior will be more stable during the second half of the session.

Under drift–diffusion models, a decision-making process is initiated at the starting point and continues to accumulate noisy evidence, thus continues to drift until it reaches one of the decision bounds. The decision bounds are defined as:

$$B_{upper}(t) = \frac{1}{1 + \exp(s \cdot (t - d))} + \frac{\exp(-s \cdot d)}{1 + \exp(-s \cdot d)},$$

where, s is a decay parameter of the boundary, d is a delay parameter to start the decay and t is time. B_{lower} is defined as $- B_{upper}$. Hence at $t = 0$, the bounds are at ± 1 and decay towards zero as time progresses. The drift rate (rate of accumulation) r is defined as

$$r = k \cdot C + o,$$

where C is the signed coherence of the stimulus such that positive C means coherence towards positive orientation, and o is the offset in the drift rate. The reaction time can be obtained by

$$RT = t_{decision} + t_0,$$

where $t_{decision}$ is the time required to cross a decision bound, and t_0 is the non-decision time to account for the sensory and motor processing delays.

The traditional DDM model does not account for lapses in the data. To account for this, we added four parameters to our model such that on random proportions of trials subjects provide an uninformed response with a RT derived from an unknown Gaussian distribution. The mean and the standard deviation of the Gaussian distribution are estimated as model parameters from the data. Hence, the additional four parameters are: proportion of uninformed positive choices, proportion of uninformed negative choices, mean and standard deviation of reaction time for uninformed choices. The proportion of uninformed positive and negative choices were fitted separately for each color.

All the behavioral data analyses and modeling was performed using MATLAB. The DDM model was implemented from the Stochastic Integration Modeling Toolbox for MATLAB which was developed by J. Ditterich (https://www.peractionlab.org/software).

The parameter optimization was carried out using maximum likelihood estimation accounting for both reaction time and choice on each trial. We divided the data into 16 subsets based on the coherence (4×), direction (2×), and choice (2×). The likelihood of choice and RT (using QMLE method on normalized histogram of RT distribution) for each subset was calculated independently and log summed at the end for all subsets.

The reaction time likelihood was calculated using modified Quantile Maximum Likelihood Estimation (QMLE)[43]. To get better resolution of estimated RT distribution, we interpolated RT distribution from model and data to 1 ms resolution. Next, we computed the log likelihood from model for each millisecond and took its weighted sum based on RT distribution from data.

$$Log(Likelihood_{RT}) = \sum_{i=\min(RT)}^{\max(RT)} RT_{data,i} \cdot \log(RT_{model,i}).$$

Once RT likelihood was calculated for all 16 subset of conditions, we calculated a weighted sum based on the number of trials in each subset. The choice likelihood was calculated using the binomial probability equation[44], as this was a 2AFC task:

$$Likelihood_{choice} = p(Positive_{model})^{p(Positive_{data})} \cdot p(Negative_{model})^{p(Negative_{data})}.$$

Similar to RT likelihood, choice likelihoods were log summed across all 16 subsets of conditions, and, finally, both RT log likelihood and choice log likelihood were summed. This total likelihood was later optimized using MATLAB functions.

We used a combination of a multidimensional simplex approach ("fminseach" in MATLAB Optimization Toolbox) and a pattern search algorithm ("patternsearch" in MATLAB Global Optimization Toolbox) to find the global optima. Standard errors of the estimated parameters were obtained using 1-dimensional local Gaussian approximation of the likelihood function (L) around the optimal value p_{opt} of each parameter p:

$$L(p) \approx L_{max} \cdot \exp\left(-\frac{(p - p_{opt})^2}{2 \cdot \sigma_p^2}\right).$$

In addition to the above model, we also explored whether a single mechanism of bias implementation could explain the data better after controlling for the number of parameters. We fitted two additional models to the data, one allowing only change in starting point and other where change is allowed only in the drift-rate offset. To quantify which model performs better we calculated the BIC scores of each of these three models in all five groups. The scores for all experimental conditions are shown in Table S3 in the Supplementary Data. The BIC score indicates which model performs better after accounting for the number of parameters based on the likelihood estimation, with lower scores indicating better performing models.

References

1. Ratcliff, R., Hockley, W. & McKoon, G. Components of activation: Repetition and priming effects in lexical decision and recognition. *J. Exp. Psychol. Gen.* **114**, 435–450 (1985).
2. Berry, D. C. & Broadbent, D. E. The combination of explicit and implicit learning processes in task control. *Psychol. Res.* **49**, 7–15 (1987).
3. Berry, D. C. & Broadbent, D. E. Interactive tasks and the implicit-explicit distinction. *Br. J. Psychol.* **79**, 251–272 (1988).
4. Schacter, D. L. Implicit memory: History and current status. *J. Exp. Psychol. Learn. Mem. Cogn.* **13**, 501–518 (1987).
5. Reber, A. S. Implicit learning and tacit knowledge. *J. Exp. Psychol. Gen.* **118**, 219–235 (1989).
6. Lee, Y. S. Effects of learning contexts on implicit and explicit learning. *Mem. Cognit.* **23**, 723–734 (1995).
7. Stefan, K. *et al.* Formation of a motor memory by action observation. *J. Neurosci.* **25**, 9339–9346 (2005).
8. Saffran, J. R., Aslin, R. N. & Newport, E. L. Statistical learning by 8-month-old infants. *Science* **274**, 1926–1928 (1996).
9. Goujon, A., Didierjean, A. & Thorpe, S. Investigating implicit statistical learning mechanisms through contextual cueing. *Trends Cogn. Sci.* **19**, 524–533 (2015).
10. Reber, A. S. & Lewis, S. Implicit learning: An analysis of the form and structure of a body of tacit knowledge. *Cognition* **5**, 333–361 (1977).
11. Mathews, R. C. *et al.* Role of implicit and explicit processes in learning from examples: A synergistic effect. *J. Exp. Psychol. Learn. Mem. Cogn.* **15**, 1083–1100 (1989).
12. Mathews, R. C. Abstractness of implicit grammar knowledge: Comments on Perruchet and Pacteau's analysis of synthetic grammar learning. *J. Exp. Psychol. Gen.* **119**, 412–416 (1990).
13. Nissen, M. J. & Bullemer, P. Attention requirements of learning evidence from performance measures. *Cogn. Psychol.* **19**, 1–32 (1987).
14. Reber, P. J. & Squire, L. R. Parallel brain systems for learning with and without awareness. *Learn. Mem.* **1**, 217–229 (1994).
15. Koehler, J. J. The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behav. Brain Sci.* **19**, 1–53 (1996).
16. White, C. N. & Poldrack, R. A. Decomposing bias in different types of simple decisions. *J. Exp. Psychol. Learn. Mem. Cogn.* **40**, 385–398 (2014).
17. Leite, F. P. & Ratcliff, R. What cognitive processes drive response biases? A diffusion model analysis. *Judgm. Decis. Mak.* **6**, 651–687 (2011).
18. Ziori, E. & Dienes, Z. How does prior knowledge affect implicit and explicit concept learning?. *Q. J. Exp. Psychol.* **61**, 601–624 (2008).
19. Mulder, M. J., Wagenmakers, E. J., Ratcliff, R., Boekel, W. & Forstmann, B. U. Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *J. Neurosci.* **32**, 2335–2343 (2012).
20. Huang, Y., Friesen, A. L., Hanks, T. D., Shadlen, M. N. & Rao, R. P. N. How prior probability influences decision making: A unifying probabilistic model. *Adv. Neural Inf. Process. Syst.* **2**, 1268–1276 (2012).
21. Diederich, A. Bound-change, drift-rate-change, or two-stage-processing hypothesis. *Percep. Psychophys.* **68**(2), 194–207 (2006).
22. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
23 Urai, A. E., De Gee, J. W., Tsetsos, K. & Donner, T. H. Choice history biases subsequent evidence accumulation. *Elife* **8**, e46331 (2019).
24. Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E. & Shadlen, M. N. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.* **31**, 6339–6352 (2011).
25. Perugini, A., Ditterich, J. & Basso, M. A. Patients with Parkinson's disease show impaired use of priors in conditions of sensory uncertainty. *Curr. Biol.* **26**, 1902–1910 (2016).

26. Kiani, R., Hanks, T. D. & Shadlen, M. N. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.* **28**, 3017–3029 (2008).
27. Glass, L. Moire effect from random dots. *Nature* **223**, 578–580 (1969).
28. Nagai, Y., Suzuki, M., Miyazaki, M. & Kitazawa, S. Acquisition of multiple prior distributions in tactile temporal order judgment. *Front. Psychol.* **3**, 1–7 (2012).
29. Carpenter, R. H. S. & Williams, M. L. L. Neural computation of log likelihood in control of saccadic eye movements. *Nature* **377**, 59–62 (1995).
30. Rao, V., Deangelis, G. C. & Snyder, L. H. Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas. *J. Neurosci.* **32**(29), 10063–10074 (2012).
31. Ashby, F. G. A biased random walk model for two choice reaction times. *J. Math. Psychol.* **27**, 277–297 (1983).
32. Diederich, A. Dynamic stochastic models for decision making under time constraints. *J. Math. Psychol.* **41**, 260–274 (1997).
33. Rungratsameetaweemana, N., Squire, L. R. & Serences, J. T. Preserved capacity for learning statistical regularities and directing selective attention after hippocampal lesions. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 19705–19710 (2019).
34. Tversky, A. & Kahneman, D. Availability: A heuristic for judging frequency and probability. *Cogn. Psychol.* **5**, 207–232 (1973).
35. Spellman, B. A. The implicit use of base rates in experiential and ecologically valid tasks. *Behav. Brain Sci.* **19**, 38–38 (1996).
36. Cho, R. Y. *et al.* Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task. *Cogn. Affect. Behav. Neurosci.* **2**, 283–299 (2002).
37. Basso, M. A. & Wurtz, R. H. Modulation of neuronal activity in superior colliculus by changes in target probability. *J. Neurosci.* **18**, 7519–7534 (1998).
38. Platt, M. L. & Glimcher, P. W. Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238 (1999).
39. Herz, D. M., Bogacz, R. & Brown, P. Neuroscience: Impaired decision-making in Parkinson's disease. *Curr. Biol.* **26**, R671–R673 (2016).
40. Dunovan, K. E., Tremel, J. J. & Wheeler, M. E. Prior probability and feature predictability interactively bias perceptual decisions. *Neuropsychologia* **61**, 210–221 (2014).
41. Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. D. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **113**, 700–765 (2006).
42. Wichmann, F. A. & Hill, N. J. The psychometric function: I. Fitting, sampling, and goodness of fit. *Percept. Psychophys.* **63**, 1293–1313 (2001).
43. Heathcote, A., Brown, S. & Mewhort, D. J. Quantile maximum likelihood estimation of response time distributions. *Psychon. Bull. Rev.* **9**, 394–401 (2002).
44. Selvin, S., Clayton, D. & Hills, M. Statistical models in epidemiology. *J. Am. Stat. Assoc.* https://doi.org/10.2307/2291094 (1995).

Author contributions

V.T., M.B. and B.K. designed experiments and wrote the manuscript text. V.T. prepared figures and conducted analyses of behavioral data. J.D. and V.T. conducted modeling analyses. All authors reviewed and edited the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1038/s41598-021-95833-7.

Correspondence and requests for materials should be addressed to B.J.K.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.