Published in final edited form as:

Nat Neurosci. 2010 January ; 13(1): 112–119. doi:10.1038/nn.2450.

Synaptic Computation Underlying Probabilistic Inference

Alireza Soltani^{1,2} and Xiao-Jing Wang²

¹Division of Biology, California Institute of Technology, Pasadena, California, 91125

²Department of Neurobiology and Kavli Institute of Neuroscience, Yale University School of Medicine, New Haven, Connecticut, 06510.

Abstract

In this paper we propose that synapses may be the workhorse of neuronal computations that underlie probabilistic reasoning. We built a neural circuit model for probabilistic inference when information provided by different sensory cues needs to be integrated, and the predictive powers of individual cues about an outcome are deduced through experience. We found that bounded synapses naturally compute, through reward-dependent plasticity, the posterior probability that a choice alternative is correct given that a cue is presented. Furthermore, a decision circuit endowed with such synapses makes choices based on the summated log posterior odds and performs near-optimal cue combination. The model is validated by reproducing salient observations of, and provide insights into, a monkey experiment using a categorization task. Our model thus suggests a biophysical instantiation of the Bayesian decision rule, while predicting important deviations from it similar to ‘base-rate neglect’ observed in human studies when alternatives have unequal priors.

Decision making often relies on our ability to combine information from different sources, and to make inferences even when the relationship between cues and outcomes is not deterministic. For instance, in the so called weather prediction task which is commonly used in cognitive neuro-science, a categorical choice (rain or sunshine) can be predicted only probabilistically based on a number of given cues ^{1–5}. Such a decision is challenging not only due to its probabilistic character but also because a single choice is preceded by a large number of cues, so it is not obvious how to deduce correctly cue-outcome associations (e.g. identify an allergenic or poisonous substance after consuming a few food items and getting sick). Little is known about the neural computations underlying this cognitive ability of probabilistic reasoning.

A recent study suggested that monkeys are capable of some forms of probabilistic inference and revealed neural correlates of this ability at the single-cell level in the lateral intra-parietal cortex (LIP) ⁶. In particular, this neural activity encodes the combination of information from different shapes in terms of the log likelihood ratio (log LR), a quantity which the monkeys appeared to use to combine information from experience with the cues in order to make a decision on each trial. This neurophysiological finding supports the theoretical proposal that log LR provides a quantity suitable for the accumulation of sensory evidence ^{7,8}. However it raises the question, how could such a quantity computed and leaned biophysically?

In this paper, we propose that quantities such as the likelihood or posterior probability can be learned and encoded by synapses which have bounded weights and undergo reward-dependent Hebbian plasticity ^{9–11}. The computational implications of bounded synapses

have only begun to be recognized in the theoretical community. In particular, how the maintenance of long term memory storage dramatically depends on whether synapses are bounded or not^{12,13}. In our model, trial-by-trial decision-making is determined by statistical sampling of stochastic neural dynamics^{14–19}; firing activity of single cells correlates with conditional reward probabilities because neurons are driven by plastic synapses that learn cue-outcome associations.

We show that in a simulated probabilistic inference task, these synapses can estimate the naive posterior probability, i.e. the posterior probability that a choice alternative is assigned with reward given that a cue is presented in any combination of cues. Furthermore, in a decision circuit, the choice behavior is determined by the difference in the inputs associated with each choice option, which is approximately proportional to the sum of the log posterior odds for the presented cues. The cue combination is thus near-optimal (i.e. according to the Bayes rule) when the prior probabilities of reward assignment on each choice alternative are equal. However, when priors are not equal, the model predicts specific deviations that can directly be tested experimentally. Such deviations from the Bayes rule can explain the ‘base-rate neglect’ effect observed in human behavioral studies²⁰. Overall, our model reproduces salient behavioral and single-unit neural data⁶ and provides insights into the neural mechanisms of three key computational processes: inference, cue combination, and probabilistic decision making.

Results

Learning posteriors by plastic synapses

In this section we show how plastic synapses are able to estimate probabilistic quantities such as posteriors. We assume that individual plastic synapses are binary (with a depressed and a potentiated state), hence the strength of a set of plastic synapses can be quantified by the fraction of synapses in the potentiated state^{9,11,21,22}. This quantity is called the ‘synaptic strength’ and denoted as c_{iA} or c_{iB} , for the set of synapses from sensory neurons selective for cue $\mathbb{S}_i$ onto action value-coding neurons selective for choice A or B , respectively (Fig. 1a).

Plastic synapses learn cue-outcome contingencies through stochastic reward-dependent Hebbian modifications^{9,11} (see Methods for details about the learning rule). That is, at the end of the trial, only the sets of plastic synapses from sensory neurons selective for the presented cues onto action value-coding neurons selective for the chosen alternative are updated, with the direction (potentiation or depression) that depends on the choice outcome: if the choice of the model is rewarded, synapses in the depressed state make a transition to the potentiated state with a probability q_+ ; otherwise the transition in the reverse direction occurs with a probability q_- .

Consider a simple situation where one cue ($\mathbb{S}_i$) alone is presented on each trial, and it determines the reward probability for each of the two alternative responses, $P(A|\mathbb{S}_i)$ and $P(B|\mathbb{S}_i) = 1 - P(A|\mathbb{S}_i)$. In this case, the reward assignment is independent of the choice selection and the probability that a set of synapses is potentiated (say for synapses selective for cue $\mathbb{S}_i$ and choice A) is equal to the product of three probabilities: the probability that cue $\mathbb{S}_i$ is presented, $P(\mathbb{S}_i)$; the probability that choice A is selected when cue $\mathbb{S}_i$ is presented, $P_A(\mathbb{S}_i)$; the probability that choice A is assigned with reward given cue $\mathbb{S}_i$ is presented, $P(A|\mathbb{S}_i)$. The probability of depression for the same set of synapses is $P(\mathbb{S}_i) \times P_A(\mathbb{S}_i) \times (1 - P(A|\mathbb{S}_i))$.

Through ongoing learning, the synaptic strength for each set of plastic synapses eventually reaches a steady-state value. If the learning rates are small, the steady-state of the synaptic strength can be computed by setting the overall change equal to zero,

$$\Delta c_{iA} = q_+ (1 - c_{iA}) \times P(\mathbb{S}_i) P_A(\mathbb{S}_i) P(A|\mathbb{S}_i) - q_- c_{iA} \times P(\mathbb{S}_i) P_A(\mathbb{S}_i) (1 - P(A|\mathbb{S}_i)) = 0$$

which gives an expression for the steady-state of the synaptic strength

$$c_{iA}^{ss} = \frac{rP(A|\mathbb{S}_i)}{1 + (r - 1)P(A|\mathbb{S}_i)} \quad (1)$$

where r is the learning rate ratio ($r \equiv q_+/q_-$). Therefore, when each cue is presented alone the steady-state is independent of the choice behavior (i.e. $P_A(\mathbb{S}_i)$).

In the special case of equal potentiation and depression rates ($r = 1$), the steady-state of the synaptic strength is equal to the posterior probability, $c_{iA}^{ss} = P(A|\mathbb{S}_i)$ (the lightest curve in Fig. 1b). In general, when the learning rates are not equal, the synaptic strength is a nonlinear monotonic function of the posterior probability (Fig. 1b).

Computation of log posterior odds

In our model, the decision circuit (Fig. 1a) generates a categorical choice (A or B) stochastically with a probability which is a sigmoid function of the difference in the overall inputs to its selective pools (the differential input)^{9–11,14}. Because cue-selective neurons fire at a similar rate, the differential input is solely determined by the difference in the synaptic strengths from the action value-coding neurons onto the decision neurons. Using equation(1), we can compute the difference in the synaptic strengths ($\Delta c_i^{ss} \equiv c_{iA}^{ss} - c_{iB}^{ss}$) for cue $\mathbb{S}_i$

$$\Delta c_i^{ss} = \frac{r(P(A|\mathbb{S}_i) - P(B|\mathbb{S}_i))}{r + (r - 1)^2 P(A|\mathbb{S}_i) P(B|\mathbb{S}_i)} \quad (2)$$

This formula can be simplified by observing that the second term in the denominator, $k \equiv (r - 1)^2 P(A|\mathbb{S}_i) P(B|\mathbb{S}_i)$, is zero when $r = 1$, and does not vary significantly (compared to the first term r) provided that r is not too large and that values of posterior probabilities are in an intermediate range (For $0.2 \leq P(A|\mathbb{S}_i) \leq 0.8$, $4(r-1)^2/25 \leq k \leq (r-1)^2/4$). Since k is roughly constant, the difference in the steady-state of the synaptic strengths is proportional to the difference in posterior probabilities for the two choice alternatives (Fig. 1c)

$$\Delta c_i^{ss} \approx \frac{r}{r+k} (P(A|\mathbb{S}_i) - P(B|\mathbb{S}_i)) \quad (3)$$

Furthermore, we note that $x - (1-x) \approx \log_{10}(x/(1-x))$, if $0.2 \leq x \leq 0.8$ (Fig. 1c). Therefore, for the intermediate range of posteriors where the model's choice behavior is stochastic, the difference in the synaptic strengths is linearly proportional to the log posterior odds (Fig. 1c)

$$\Delta c_i^{ss} \approx \frac{r}{r+k} \log_{10} \frac{P(A|\mathbb{S}_i)}{P(B|\mathbb{S}_i)} \quad (4)$$

For smaller or larger values of posteriors, the choice behavior is deterministic (the probability of choosing A is close to 0 or 100). Importantly, equation(4) holds in the general case of unequal learning rates, when c_{iA}^{ss} is a nonlinear function of $P(A|\mathbb{S}_i)$ (Fig. 1b).

In summary, when the choice outcome is based on a single cue, synapses endowed with realistic reward-dependent plasticity are capable of estimating quantities such as posteriors, and a decision network driven by such synapses can make decision according to log posterior odds.

Summation of log posterior odds

What happens if a choice outcome is preceded by several cues? To address this question, we consider a probabilistic categorization task known as the weather prediction task in which four shapes precede a selection between two (A =red, B =green) response targets on each trial⁶. These shapes were selected randomly from a set of 10 distinguishable shapes ($\mathbb{S}_i$, $i = 1, 2, \dots, 10$), each of which was allocated a unique weight of evidence about the probability of

reward assignment on one of the two choice targets ($\text{WOE} = \log_{10} \frac{P(A|\mathbb{S}_i)}{P(B|\mathbb{S}_i)}$). The computer assigned a reward to one of the two alternative choices with a probability that depended on the sum of the WOE's from all the shapes presented on a given trial (see Methods for more details).

In model simulations of the weather prediction task, on each trial, pools of sensory neurons selected for the presented cues are activated and converge onto value-coding neurons (Fig. 1a). At the end of the trial, all sets of plastic synapses from neurons selective for the presented shapes onto value-coding neurons selective for the chosen alternative are updated independently of their role in decision making. As a result, synaptic changes for different shapes become correlated. Interestingly, even though there are 10^4 stimulus patterns, we found that it took only a few hundred trials for the synaptic strengths to reach their average (steady-state) values (Fig. 2a), on a timescale largely set by the learning rates. This means that, after a few hundred trials, the model is able to correctly perform the task while plastic synapses continue to fluctuate (around their steady states) due to ongoing learning.

Similar to equation(1), we can find an expression for the steady-state of the synaptic strength (see section 1 of the Supplementary Materials). Simulation results revealed that it is approximately a linear function of the 'naive' posterior probability (Fig. 2b). Naive posterior probability, $\tilde{P}(A|\mathbb{S}_i)$, is the conditional probability that alternative A is assigned with reward given that shape $\mathbb{S}_i$ is presented in any pattern. It is the generalization of posterior when more than one cue precede an outcome, assuming independence between the evidence provided by each cue. It follows from the above mathematical reasons that the difference in the synaptic strengths is approximately a linear function of the log naive posterior odds (Fig. 2c)

$$\Delta c_i^{ss} = \alpha \log_{10} \frac{\tilde{P}(A|\mathbb{S}_i)}{\tilde{P}(B|\mathbb{S}_i)} \quad (5)$$

where the linear fit yielded $\alpha = 0.48$.

Because the convergence of sensory neurons onto action value-coding neurons naturally summates the currents through sets of plastic synapses related to presented cues, the overall differential input onto decision neurons is given by the sum of log naive posterior odds.

Thus, this model provides a natural mechanism for integrating evidence in terms of log posterior odds.

Simulation of weather prediction task: behavioral results

The results reported above are general, suggesting that our model endowed with the proposed reward-dependent learning rule is broadly applicable to probabilistic inference tasks. To test whether this model can account for behavioral performance as well as neural activity data, we simulated the monkey experiment of Ref.⁶.

Computer simulations followed the experimental protocol (see Methods for details). Due to intrinsic noise in the neural circuit, the model's choice can vary from trial to trial even if the synaptic strengths are identical. On the other hand, synapses are updated at the end of each trial, hence the leverage that each presented shape has on decision making changes from trial to trial, leading to dynamic adjustment of choice behavior over time. The adaptive choice behavior of the model is described by the psychometric function (Fig. 3a), where the probability of selecting A is plotted against the sum of the WOE's assigned to individual shapes in a pattern. Therefore, the model reproduces the main behavioral observation of Ref.⁶ namely, the monkeys select each alternative stochastically based on the combined evidence provided by all presented shapes in a pattern, with a probability that is approximately a sigmoid function of the summated WOE's (see Fig. 1b in Ref.⁶).

The psychometric function quantifies the influence of combinations of shapes (and not the individual shapes) on the choice behavior. Following Ref.⁶, we used a logistic regression model (see equation(13) in Methods) to estimate the influence of individual shapes on the choice behavior. These regression coefficients, called the subjective weights of evidence (SWOE), are shown in Fig. 3b. Similar to experimental findings, the SWOE's are smaller than the assigned WOE's (see Fig. 1c in Ref.⁶). In particular, the SWOE's for the trump shapes are finite while the assigned WOE's for these shapes are infinite (compare Fig. S7 in the Supplementary Materials with Fig. S2 in Ref.⁶). Moreover, we found that the SWOE's do not depend on the epoch as observed experimentally (compare Fig. S8 in the Supplementary Materials with Fig. S3 in Ref.⁶).

It is instructive to compare our model with an ideal Bayesian observer. The latter would combine the likelihood ratio associated with a given pattern and the prior odds (that each alternative is assigned with reward), in order to obtain the posterior odds for that pattern (see section 3 of the Supplementary Materials for another alternative). The posterior odds then can be used to make decision according to different decision rules; e.g. strict Bayesian (selecting the alternative with the larger posterior) or probabilistic Bayesian (matching the probability of choice with the posterior). The choice behavior of these two types of Bayesian observers is shown in Fig. 3a. We found that the average reward rate for our model, probabilistic and strict Bayesian observers to be equal to 0.79, 0.84, and 0.89, respectively (see also Figs. S3 and S4 in the Supplementary Materials). Therefore, the reward rate of our model is lower than a Bayesian observer that has direct access to the posteriors associated with each pattern without underestimating the WOE's. However, the model performance is only slightly different from a probabilistic Bayesian observer.

Our results explain the following two main observations regarding the choice behavior. Firstly, why the model underestimates the WOE's? This happens because the SWOE of the model is proportional to the log naive posterior odds, which is less than the WOE because of the concurrence of different shapes on each trial (see Fig. S9 in the Supplementary Materials for more details). Secondly, how the model combines information from multiple cues, and why the choice behavior is approximately a sigmoid function of the summated evidence provided by shapes in a pattern (Fig. 3a)? For a given pattern, the choice behavior of the

model is approximately a sigmoid function of the overall differential synaptic input,

$\sum_i c_{iA}^{ss} - c_{iB}^{ss}$ (equation(S5) and Fig. S1 in the Supplementary Materials). Consequently, the choice behavior is a sigmoid function of the sum of the log naive posterior odds (using equation(5))

$$P_A(\mathbb{C}^t) = \frac{100}{1 + \exp\left(-\frac{1}{\sigma} \sum_i \log_{10} \frac{\tilde{P}(A|S_i)}{\tilde{P}(B|S_i)}\right)} \quad (6)$$

where $P_A(\mathbb{C}^t)$ is the probability of selecting A given pattern $\mathbb{C}^t$ (i.e. all patterns with the sum WOE equal to t) is presented, and $1/\sigma$ quantifies the sensitivity of the choice behavior on the differential synaptic input. We found that the log naive posterior odds is linearly proportional to the WOE for the non-trump shapes (Fig. S9C in the Supplementary Materials). Therefore, the psychometric function shown in Fig. 3a is a sigmoid function of the summated WOE of all shapes (for patterns which do not contain trump shapes).

Neural activity correlates of probabilistic inference

To link decision making with its underlying neural activity, we examined how shape presentation influences the firing rates of decision neurons, hence the model's choice behavior. As shown in sample neural traces (Fig. 4), the activity of decision neurons is influenced by the WOE of the presented shape on each trial.

We analyzed how the neural activity and the cumulative evidence co-vary in time. Among different ways to measure evidence provided by presented shapes, we chose to perform the same analysis as in Ref.⁶, using the logLR that selection of red or green target is accompanied by reward after the presentation of n shapes (see Methods for details). We observed a graded dependence on the logLR (Fig. 5a). Moreover, the average activity in each epoch is a linear function of the average logLR in that epoch (compare with Figs.4 and 5a in Ref.⁶) We also computed the incremental change in the population firing rate across successive epochs of shape presentation, and we found that this change was proportional to the average change in the logLR ($\Delta\logLR$) caused by the presentation of a new shape (see Fig. S10 in the Supplementary Materials). All of these simulation results are similar to those observed experimentally in LIP neurons of behaving monkeys⁶.

Because the neuronal activity was also strongly modulated by the choice on each trial, we performed the same analysis on the data divided into two groups depending on the choice of the model on each trial (Figs. 5c and S11 in the Supplementary Materials). We found that the baseline of the neural activity was higher when the choice was the preferred target. Conversely, modulation by evidence of neural activity was weaker when the choice was the preferred target. This is qualitatively similar to the observed modulations of LIP neurons when the choice was towards the response field (RF) versus away from it (see Figs. S7 and S8 in Ref.⁶).

To conclude, neural activity in our model reproduces the main physiological observations from LIP in the monkey experiment⁶. These results demonstrate that empirically observed neural correlates of probabilistic quantities such as likelihoods may be interpreted in terms of synaptic rather than neuronal computations.

Model prediction: effect of prior probability

The behavioral and neural data of the weather prediction experiment were reported in terms of likelihood ratios ⁶. Because in this experiment the prior reward probability was the same for the two choice alternatives, these data can equivalently be expressed in terms of posterior odds. In our model, evidence from multiple cues is combined by effectively summing the log naive posterior odds, which are different from log likelihood ratios when priors are not equal. Therefore, we next explored model's behavior in simulations when the priors were not equal (see Methods for details).

The model makes decisions based on the differential input which computes the sum of the log naive posterior odds of shapes in a given pattern. The latter is proportional to the log posterior odds for that pattern (Fig. S2A in the Supplementary Materials), a general quantity that we used to express the psychometric function (Fig. 6a).

It is evident that the model's choice behavior is strongly biased toward the more probable alternative (i.e. the alternative which is assigned with reward more often, *A* in these simulations). We define the bias in the choice behavior to be the probability of choosing *A* when the log posterior odds is zero. We found the bias in the choice behavior of our model to monotonically increase with the log prior odds (open circles in Fig. 6b). In contrast, an ideal Bayesian observer would display no bias (Fig. 6b, but see the Supplementary Materials for alternative Bayesian observers that show bias in the choice behavior).

To elucidate the bias in the choice behavior, we examined how the synapses learn about evidence related to each shape in this modified task. As shown in Fig. 6c, the difference in the synaptic strengths is increased with the prior, while it is still a linear function of the log naive posterior odds, with a nearly identical slope (hence independent of the prior). Interestingly, we found that the difference in the synaptic strengths can be fit as a linear function of the log naive posterior and log prior odds

$$\Delta c_i^{ss} = \alpha \log_{10} \left(\frac{\tilde{P}(A|\mathbb{S}_i)}{\tilde{P}(B|\mathbb{S}_i)} \right) + \beta \log_{10} \left(\frac{P(A)}{P(B)} \right) \quad (7)$$

where linear fitting yielded $\alpha = 0.48$ and $\beta = -0.31$.

Hence, there is also a bias in the information stored in plastic synapses about the evidence provided by each shape stimulus, when priors are not equal. Let us define the bias in the learned evidence as the probability of choosing *A* when one shape is presented alone (after learning) and the log naive posterior odds is zero (note that there is no shape with zero log naive posterior odds, so this measure of bias is based on extrapolation). As shown in Fig. 6b (filled circles), this probability (i.e. choice bias) is smaller than 50% and is proportional to the log prior odds (see equation(S8) in the Supplementary Materials). Therefore, if after learning the model is asked to make a judgment based on a single cue, the choice behavior is biased toward the less probable alternative (*B* in this case). At first sight, this result seems to contradict the observed bias in the psychometric function toward the more probable alternative, but it can be explained as follows (for complete explanation see section 2 of the Supplementary Materials).

When a pattern consisting of four shapes is presented, the sum of the log naive posterior odds of shapes in that pattern is proportional to the log posterior odds for that pattern, but is also positively biased by the log prior odds (Fig. S2a in the Supplementary Materials)

$$\sum_i \log_{10} \left(\frac{\tilde{P}(A|\mathbb{S}_i)}{\tilde{P}(B|\mathbb{S}_i)} \right) = \gamma \log_{10} \left(\frac{P(A|\mathbb{C}^t)}{P(B|\mathbb{C}^t)} \right) + \lambda \log_{10} \left(\frac{P(A)}{P(B)} \right) \quad (8)$$

where $P(A|\mathbb{C}^t)$ is the posterior probability that A is assigned with reward given that a set of patterns $\mathbb{C}^t$ is presented, $\gamma = 0.36$, and $\lambda = 3.64$. The sum of the log naive posterior odds provides an estimate of the log posterior odds because the reward assignment is based on the sum of the WOE of shapes in each pattern. Combined with equation(7), we see that the total differential synaptic inputs, $\sum_i \Delta c_i^{ss}$, which determines the choice behavior of the model, is given by

$$\sum_i \Delta c_i^{ss} = \alpha \gamma \log_{10} \left(\frac{P(A|\mathbb{C}^t)}{P(B|\mathbb{C}^t)} \right) + (\alpha \lambda + 4\beta) \log_{10} \left(\frac{P(A)}{P(B)} \right). \quad (9)$$

At zero posterior odds, because $(\alpha + 4\beta) > 0$, the overall effect of prior is positive and the choice behavior is biased toward the more probable alternative (Fig. 6a-b). These results are robust and are not sensitive to the model parameters (see section 4 and 5 of the Supplementary Materials).

Intuitively, with increasing prior, plastic synapses onto decision neurons selective for the more probable alternative (A) are potentiated more often, and the difference in the synaptic strengths for all shapes becomes more positive. However, according to our learning rule synapses are updated collectively on each trial, independently of their exact roles in the ultimate decision, and the prior information is mixed with, and attenuated by, the evidence provided by different sets of synapses. Therefore, the influence of prior on each set is smaller than it should be and results in a bias toward the less probable alternative in the estimate of predictive power of each shape. On the other hand, when four shapes are presented together, the influence of prior on each of four shapes in a pattern adds up and amounts to a bias toward the more probable alternative.

We have shown that the differential synaptic input is linearly proportional to the log naive posterior odds, with a slope approximately independent of the prior probability (Fig. 6c). Consequently, the average neural activity in each epoch (dictated by the differential synaptic input) depends linearly on the logLR in that epoch, and the slope is only weakly influenced by the prior probability (Fig. 6d). The effect of the prior, however, is manifested in the range of firing rates, as the prior induces a shift in the differential synaptic input. As can be seen by comparing Fig. 5c with $P(A) = 0.5$ and Fig. 6d with $P(A) = 0.8$, the firing rates are more markedly different with a larger prior, when the choice is the preferred or nonpreferred target of the decision neurons. This difference in the neural activity gives rise to biased choice behavior toward the alternative with a larger prior.

Discussion

The main findings of this paper are threefold. First, summing log posterior odds, a seemingly complicated calculation, can be readily realized, through approximations, by a plausible plasticity mechanism with bounded synapses in a decision circuit. Bounded synapses have previously been shown to impose limits on memory storage^{12,13}, here we found that they enable our model to perform probabilistic computations. Second, in order to test our model by empirical data, we investigated a biophysically based neural circuit model

implementation for the monkey weather forecasting task. Our model was shown to quantitatively account for many behavioral and single-unit neurophysiological observations of ref. ⁶, with a small number (3) of free parameters. Third, we considered situations when the choice alternatives have unequal priors, which led us to nontrivial predictions about deviations from the Bayes decision rule.

Inference and combination of information

The weather prediction task exemplifies complex decision making in which one has to acquire information about the predictive power of each sensory cue as well as to combine evidence from multiple cues to make a choice. In this work, we showed that a decision neural circuit model endowed with a simple form of reward-dependent synaptic plasticity is capable of such probabilistic reasoning. Hence, such a high-level cognitive function may be instantiated by reward-dependent learning, rather than by sophisticated strategies assumed in some human studies ²³.

We showed that plastic synapses dynamically learn and store the association between each shape and outcome in just a few hundred trials, in spite of the large number of patterns in this task (10^4 patterns). This happens because a synaptic plasticity rule that assumes independence between sources of information enables the system to learn regularities in the external world quickly and robustly. This also makes plastic synapses to encode the predictive power related to each shape in the form of the naive posterior probability. As a result, the predictive weight assigned by the model to each shape is smaller than the assigned weight of evidence, as observed experimentally. Therefore, we propose that synaptic computation of the naive posteriors provides a cellular mechanism for the brain to perform inference.

Decision neurons integrate evidence from different cues, simply through convergence of synaptic inputs from value-coding neurons. The strong recurrent dynamics of the decision circuit is critical for generating choices stochastically on single trials, with the aggregate of choice probability as a sigmoid function of the difference in the inputs associated for each choice option ^{9–11}. The latter, as we have shown, is approximately proportional to the sum of the log naive posterior odds. Therefore, by summing the log naive posterior odds of shapes in a pattern, our model uses a different strategy than ideal Bayesian observers. Nevertheless, its choice behavior is close to the probabilistic Bayesian observer which follows ‘probability matching’ ^{24–26}, and so it provides a biophysical instantiation of the probabilistic Bayesian decision rule for two-alternative choice tasks.

Influence of prior information and ‘base-rate neglect’

In order to differentiate the effect of log likelihood ratios and log posterior odds on decision making, we simulated the weather prediction task with unequal prior probabilities. We found that the model’s choice behavior is biased toward the more probable alternative, while the information encoded by plastic synapses about each shape is biased toward the less probable alternative. These predictions are supported by evidence from the weather prediction task in a human study, subjects predict one of the two outcomes (e.g. rain or sunshine) after observing one, two, three, or four tarot cards, and receive feedback at the end of each trial ^{1,20}. At the end of the experiment, the subjects are asked to estimate the strength of association between a card and an outcome. When the prior probabilities are not equal, it was found that a card which was equally predictive of each outcome was perceived to be more predictive of the less probable outcome, a phenomena known as the base-rate neglect and has been described as a judgement fallacy ^{27,28}. At the same time, the choice behavior was biased in favor of the more probable alternative (see **Table 1** in ref.²⁰).

Using our model, we can explain these counterintuitive results in terms of a plausible biophysical mechanism. Although a few models have been proposed to explain the base-rate neglect^{20, 29}, all these mechanistic models assume learning mechanisms which require access to all connection weights in the network and moreover, do not pertain to any biophysical mechanisms and constraints. Our model prediction on combination of information from different sources through addition of the log naive posterior odds can be tested more directly experimentally. For instance, if after learning the subject has to predict the outcome of a number of cues together (e.g. one, two, or three) we expect that the bias in

these predictions is proportional to $\beta \log_{10} \left(\frac{P(A)}{P(B)} \right)$ times the number of cues (equation(9)).

LIP neural activity and probability representation

Various oculomotor experiments in monkeys have shown that activity of LIP neurons encodes decision variables⁸ and is correlated with reward values of choices^{19,30–33}. It has been proposed theoretically that a neural population, such as in LIP, can represent probability distributions about sensory information on each trial, which in turn can be used to perform optimal decision making^{34–36}. Adding to this body of literature, Yang and Shadlen⁶ demonstrated that activity of LIP neurons reflects probability integration, namely the summation of the logLR. This quantity, however, can be only computed by tabulating the frequency of occurrence of each shape combination and the outcome in different epochs of the task (see Methods). Our model suggests that LIP neurons may reflect reward probability, such as log posterior odds (rather than logLR), but posteriors are encoded at synapses onto action value-coding neurons that project to LIP. Consequently, on every trial, LIP neurons integrate reward information and contribute to decision making. Our model prediction can be tested experimentally using unequal priors, which would make it possible to differentiate the log naive posterior odds and the logLR.

The plastic synapses proposed in our model should be found in neural circuits involved with representation of stimulus-reward or action-reward associations, such as parts of the prefrontal cortex^{37,38} and basal ganglia^{39–41}. Moreover, corticostriatal synapses exhibit long-term potentiation and depression that depend on the presence or absence of dopamine modulation^{42,43}, and dopamine neurons are involved in reward signaling^{44,45}. All together, these findings provide strong neurophysiological support for the synaptic plasticity rule used in this paper. Consistently, using functional brain imaging in humans, it has been shown that the striatum gradually became active as learning progressed in the weather prediction task⁴⁶. It would be worthwhile, in future experiments, to test whether this and other brain areas encode reward probabilities in the form of log posterior odds.

Our model is general and can be applied to different probabilistic decision making tasks. Indeed, we have used a similar learning rule and decision making mechanism to capture a foraging behavior known as matching law⁹, as well as choice behavior in a competitive game¹⁰. This work shows how complicated inference and cue combination can be performed by a recurrent decision circuit endowed with a plausible synaptic plasticity rule. Perhaps, other high level judgments can be instantiated by simple neural mechanisms as well.

Methods

Description of the weather prediction task

In the simulated experiment, monkeys were trained to choose between two color targets (green and red) after observing 4 shapes which were presented on a screen sequentially with 500 ms intervals⁶. These shapes were selected randomly from a set of 10 distinguishable

shapes (with replacement), each of which was assigned a unique weight of evidence (WOE). The WOE for each shape was defined as the log LR that the red or green target was assigned with reward or equivalently, the selection of red or green target was accompanied by reward. The WOE for these 10 shapes, $\{w_1, w_2, \dots, w_{10}\}$, were chosen to be $[-\infty, -0.9, -0.7, -0.5, -0.3, +0.3, +0.5, +0.7, +0.9, +\infty]$ in favor of the red target. For example, the presentation of a shape with the WOE equal to 0.9 by itself predicted that the red target was assigned with

reward 89% $\left(= \frac{10^{0.9}}{1+10^{0.9}} \right)$ of the times.

On each trial, four shapes (which we call a pattern) were presented on the screen and either of the two targets was assigned with reward with a probability depending on the sum of the WOE of all shapes in that pattern. More specifically, the probability that the red target was assigned with reward given that four shapes were presented was equal to

$$P(R|s_1, s_2, s_3, s_4) = \frac{10^{\sum_{i=1}^4 w_i}}{1 + 10^{\sum_{i=1}^4 w_i}} \quad (10)$$

and the probability that the green target was assigned with reward was $1 - P(R|s_1, s_2, s_3, s_4)$.

In order to introduce unequal prior, we first generated a set of patterns according to the described paradigm and then randomly removed a portion of trials in which the reward was assigned to the less probable alternative. This alteration changes the prior probability that each alternative is rewarded without changing the structure of the task.

Description of the model and learning rule

The model is an extended version of our previous biophysically-based model of probabilistic decision-making network^{9,10,14}. The decision making circuit of the model is a firing-rate model which has been shown to reproduce the choice and neural activity of the detailed spiking network model¹⁵. All details about the decision circuit of the model and its parameters were reported elsewhere¹⁵.

As schematically shown in Fig. 1a, the model consists of three layers. The first layer contains sensory neurons that are selective for visual cues (shapes). These cue-selective neurons can be located in the inferotemporal cortex, which has been shown to contain neurons that encode different shapes by a combination of active and nonactive columns selective for individual features^{47, 48}. The cue-selective neurons project to the second layer, where neurons learn to encode reward values of the two alternative responses (action values) through plastic synapses that undergo reward-dependent Hebbian modifications^{9,11} (see below). The second layer therefore presumably corresponds to certain frontal areas, such as the anterior cingulate cortex, that are known to be involved with representation and learning of action values^{37,38}. Note that the convergence from cue-selective neurons enables value-coding neurons to combine information from different cues. The third layer is a decision circuit with two competing neural pools that are selective for choice ('preferred target') A and B, respectively. This decision circuit is modeled the same way as in our previous work^{9,14,15}. The firing activity of these decision neurons was compared with neural data recorded from LIP⁶.

Computer simulations followed the experimental protocol. On a simulated trial, at the onset of the first shape stimulus, the visual inputs (representing fixation point as well as the visual

cue) trigger a brief transient response of neurons in the decision-making network that decreases to a moderate level, similar to LIP neurons⁶. Upon the presentation of each shape, the activity of sensory neurons selective for that shape was increased from zero to a constant value and this activity was sustained throughout the trial (if a shape is repeated on a trial, the activity of the corresponding population is multiplied by the number of repetitions of the shape on that trial). At the end of the trial, when the fixation point goes off, the activity of these two populations drops due to a decrease in the overall inputs (see below). This is followed by a divergence of activity between the two populations which, as a distinctive feature of competition in the decision-making network, signals the choice of the model on each trial¹⁵. Specifically, the model's decision is determined by the neural population that is the first to reach a fixed firing rate threshold of 30 Hz. In general, a population of neurons receiving larger inputs reaches a higher level of activity and consequently has a higher chance to win the competition and determines the model's choice on a trial.

Since the sensory responses of cue-selective neurons are similar, the only factor that differentiates the inputs to decision neurons is the strength of plastic synapses from sensory neurons onto value-coding neurons. In addition to these inputs, decision neurons also receive a large background input as well as purely visual inputs which mimics the visual response of neurons in the visual cortex and keeps the decision circuit from entering the competition regime during the presentation of four shapes and before the extinction of fixation point (see Supplementary Materials for more details).

The inputs to decision neurons are determined by the firing activity of sensory neural populations coding the presented shapes, and by the strength of plastic synapses from these populations onto value-coding populations. We assumed these plastic synapses are binary (i.e. they only have two stable states)^{21,22}. There is increasing evidence that plastic synapses have discrete (binary) states^{49,50}. Here, we used binary synapses, but our results still hold with multiple discrete states. For binary synapses, the average strength of these synapses can be defined as the fraction of synapses in the potentiated state, denoted by c_{iA} and c_{iB} (for synapses from neurons selective for shape i onto value-coding neurons selective for alternative A or B , respectively).

At the end of each trial, plastic synapse were modified according to a stochastic reward-dependent Hebbian learning rule^{9,11,21,22}. First, Hebbian plasticity required a high level of activity in both pre- and post-synaptic neurons so only synapses from cue-selective neurons selective for presented shapes onto the value-coding neurons selective for the chosen alternative were modified. Secondly, depending on the outcome (reward or no reward) on a given trial, plastic synapses were potentiated or depressed. Third, these modifications took place stochastically.

Note that at the end of each trial, only neural populations selective for all presented shapes (through working memory) and for the selected alternative were active. If the choice of the model was rewarded, all sets of synapses selective for presented shapes onto the value-coding neurons selective for the chosen alternative (say A) were potentiated with probability q_+ . That is, all synapses in the depressed state make a transition to the potentiated state with probability q_+ . As a result, the synaptic strength for each set of plastic synapses (say i) was updated as follows

$$c_{iA} \rightarrow c_{iA} + q_+ (1 - c_{iA}) \quad (11)$$

Alternatively if the choice of the model was not rewarded, plastic synapses were depressed with probability q_- so the synaptic strength was updated as

$$c_{iA} \rightarrow c_{iA} - q_- c_{iA} \quad (12)$$

For all simulation presented in the paper we set $q_+ = 0.02$ and $q_- = 0.02$, but in section 4 of the Supplementary Materials we showed how model's choice behavior depends on the learning rates.

Data analysis

All of the data analysis and average values reported here were computed over 100000 trials of the simulated experiment, except values related to the neural activity of decision circuit which are computed using 20000 simulated trials. In order to estimate the influence of each shape on the model's choice behavior we used a logistic regression fit similar to one used in Ref.⁶. We assumed that probability of selecting a choice is influenced by the presence of each shape as

$$P_A = \frac{10^Q}{1 + 10^Q} \quad \text{where} \quad Q = \sum_{i=1}^{10} q_i N_i \quad (13)$$

where the N_i is the number of appearances of shape i in each pattern. The regression coefficients, q_i 's, are called the subjective weight of evidence (SWOE) and measure the influence of each shape on decision making.

We used the logLR through the paper as a measure of evidence provided by the presented shapes until a certain epoch or the end of a trial. The logLR that a target is assigned with reward is equal to

$$\log LR_n = \log_{10} \frac{P(s_1, \dots, s_n | \text{reward at A})}{P(s_1, \dots, s_n | \text{reward at B})} \quad n=1, 2, 3, 4 \quad (14)$$

This quantity was computed by tabulating the frequency of reward being assigned to each shape combination for each epoch. In order to compute the change in the logLR due to presentation of each shape, we calculated the average change in the logLR due to the presentation of another shape from one epoch to the next, while excluding the trump shapes if they predicted one of the outcomes deterministically (similar to Ref.⁶).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by NIH grants 2-R01-MH062349 and MH073246. We are thankful to David Andrieux, Alberto Bernacchia, and Robert Wilson for useful comments on the manuscript.

References

1. Knowlton BJ, Squire LR, Gluck MA. Probabilistic classification learning in amnesia. *Learn Mem* 1994;1:106–120. [PubMed: 10467589]
2. Knowlton BJ, Mangels JA, Squire LR. A neostriatal habit learning system in humans. *Science* 1996;273:1399–402. [PubMed: 8703077]

3. Moody TD, Bookheimer SY, Vanek Z, Knowlton BJ. An implicit learning task activates medial temporal lobe in patients with parkinson's disease. *Beh Neurosci* 2004;118:438–42.
4. Fera F, et al. Neural mechanisms underlying probabilistic category learning in normal aging. *J of Neurosci* 2005;25:11340–8. [PubMed: 16339029]
5. Ashby FG, Maddox WT. Human category learning. *Annu Rev Psychol* 2005;56:149–178. [PubMed: 15709932]
6. Yang T, Shadlen MN. Probabilistic reasoning by neurons. *Nature* 2007;447:1075–80. [PubMed: 17546027]
7. Gold J, Shadlen M. Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 2001;5:10–16. [PubMed: 11164731]
8. Gold JJ, Shadlen MN. The neural basis of decision making. *Ann Rev Neurosci* 2007;30:535–74. [PubMed: 17600525]
9. Soltani A, Wang X-J. A biophysically-based neural model of matching law behavior: melioration by stochastic synapses. *J Neurosci* 2006;26:3731–3744. [PubMed: 16597727]
10. Soltani A, Lee D, Wang X-J. Neural mechanism for stochastic behavior during a competitive game. *Neural Networks* 2006;19:1075–90. [PubMed: 17015181]
11. Fusi S, Asaad WF, Miller EK, Wang X-J. A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple timescales. *Neuron* 2007;54:319–33. [PubMed: 17442251]
12. Fusi S, Drew PJ, Abbott LF. Cascade models of synaptically stored memories. *Neuron* 2005;45:599–611. [PubMed: 15721245]
13. Fusi S, Abbott LF. Limits on the memory storage capacity of bounded synapses. *Nat Neurosci* 2007;10:485–493. [PubMed: 17351638]
14. Wang X-J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 2002;36:955–968. [PubMed: 12467598]
15. Wong K-F, Wang X-J. A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci* 2006;26:1314–1328. [PubMed: 16436619]
16. Wong K-F, Huk AC, Shadlen MN, Wang X-J. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front. Comput Neurosci* 2007;1:6. [PubMed: 18946528]
17. Furman M, Wang X-J. Similarity effect and optimal control of multiple-choice decision making. *Neuron* 2008;60:1153–68. [PubMed: 19109918]
18. Liu F, Wang X-J. A common cortical circuit mechanism for perceptual categorical discrimination and veridical judgment. *PLoS Comput Bio* 2008;4:e1000253. [PubMed: 19112487]
19. Wang X-J. Decision making in recurrent neuronal circuits. *Neuron* 2008;60:215–34. [PubMed: 18957215]
20. Gluck MA, Bower GH. From conditioning to category learning: an adaptive network model. *J Exp Psychol Gen* 1988;117:227–247. [PubMed: 2971760]
21. Amit DJ, Fusi S. Dynamic learning in neural networks with material synapses. *Neural Comput* 1994;6:957–982.
22. Fusi S. Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biological Cybernetics* 2002;87:459–70. [PubMed: 12461635]
23. Meeter M, Myers CE, Shohamy D, Hopkins RO, Gluck MA. Strategies in probabilistic categorization: results from a new way of analyzing performance. *Learn Mem* 2006;13:230–9. [PubMed: 16547162]
24. Myers, JL. Probability learning and sequence learning.. In: Estes, WK., editor. *Handbook of Learning and Cognitive Processes*. Erlbaum; Hillsdale, NJ: 1976. p. 171-205.
25. Vulkan N. An economist's perspective on probability matching. *J Econ Surv* 2000;14:101–118.
26. Shanks DR, Tunney RJ, McCarthy JD. A re-examination of probability matching and rational choice. *J. Behav. Dec. Making* 2002;15:233–250.
27. Kahneman D, Tversky A. On the psychology of prediction. *Psych Rev* 1973;80:237–51.

28. Tversky, A.; Kahneman, D. Evidential impact of base rates.. In: Kahneman, D.; Slovic, P.; Tversky, A., editors. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press; 1982. p. 153-60.
29. Kruschke JK. ALCOVE: an exemplar-based connectionist model of category learning. *Psych Rev* 1992;99:22–44.
30. Glimcher PW. The neurobiology of visual-saccadic decision making. *Ann Rev Neurosci* 2003;26:133–79. [PubMed: 14527268]
31. Sugrue LP, Corrado GC, Newsome WT. Matching behavior and representation of value in parietal cortex. *Science* 2004;304:1782–1787. [PubMed: 15205529]
32. Sugrue LP, Corrado GS, Newsome WT. Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat Rev Neurosci* 2005;6:363–375. [PubMed: 15832198]
33. Soltani A, Wang X-J. From biophysics to cognition: reward-dependent adaptive choice behavior. *Curr Opin Neurobiol* 2008;18:209–16.
34. Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. *Nat Neurosci* 2006;9:1432–8. [PubMed: 17057707]
35. Beck JM, et al. Probabilistic population codes for bayesian decision making. *Neuron* 2008;60:1142–52. [PubMed: 19109917]
36. Ma WJ, Beck JM, Pouget A. Spiking networks for bayesian inference and choice. *Curr Opin Neurobiol* 2008;18:217–22.
37. Rushworth MFS, Behrens TEJ. Choice, uncertainty and value in prefrontal and cingulate cortex. *Nat Neurosci* 2008;11:389–97. [PubMed: 18368045]
38. Lee K-M, Keller EL. Neural activity in the frontal eye fields modulated by the number of alternatives in target choice. *J Neurosci* 2008;28:2242–51. [PubMed: 18305257]
39. Lauwereyns J, Watanabe K, Coe B, Hikosaka O. A neural correlate of response bias in monkey caudate nucleus. *Nature* 2002;418:413–7. [PubMed: 12140557]
40. Samejima K, Ueda Y, Doya K, Kimura M. Representation of action-specific reward values in the striatum. *Science* 2005;310:1337–1340. [PubMed: 16311337]
41. Lau B, Glimcher PW. Value representations in the primate striatum during matching behavior. *Neuron* 2008;58:451–63. [PubMed: 18466754]
42. Reynolds JN, Hyland BI, Wickens JR. A cellular mechanism of reward-related learning. *Nature* 2001;413:67–70. [PubMed: 11544526]
43. Shen W, Flajolet M, Greengard P, Surmeier DJ. Dichotomous dopaminergic control of striatal synaptic plasticity. *Science* 2008;321:848–51. [PubMed: 18687967]
44. Schultz W. Predictive reward signal of dopamine neurons. *J Neurophys* 1998;80:1–27.
45. Schultz W. Multiple dopamine functions at different time courses. *Ann Rev Neurosci* 2007;30:259–88. [PubMed: 17600522]
46. Poldrack RA, et al. Interactive memory systems in the human brain. *Nature* 2001;414:546–50. [PubMed: 11734855]
47. Tanaka K. Inferotemporal cortex and object vision. *Ann Rev Neurosci* 1996;19:109–139. [PubMed: 8833438]
48. Tsunoda K, Yamane Y, Nishizaki M, Tanifuji M. Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. *Nat Neurosci* 2001;4:832–838. [PubMed: 11477430]
49. Petersen CC, Malenka RC, Nicoll RA, Hopfield JJ. All-or-none potentiation at CA3-CA1 synapses. *Proc Natl Acad Sci U S A* 1998;95:4732–37. [PubMed: 9539807]
50. O'Connor DH, Wittenberg GM, Wang SS-H. Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *Proc Natl Acad Sci U S A* 2005;102:9679–84. [PubMed: 15983385]

Figure 1.

Schematic of the model and posterior computation by plastic synapses when a single cue is presented on each trial **(a)** Schematic of the three-layer model. The first layer consists of cue-selective neural populations, each is activated upon the presentation of a cue (a shape in the weather prediction task). The sensory cue-selective neurons provide, through some synapses endowed with reward-dependent Hebbian plasticity, inputs to two neural populations in an intermediate layer that encode reward values of two choice alternatives (action values). Combination of cues is accomplished through convergence of cue-selective neurons onto reward value-coding neurons. The latter project to excitatory and inhibitory neural populations in a decision making circuit (gray box). The choice (A or B) is determined by which of the two decision neural populations wins competition on a trial. Depending on the reward schedule, a chosen action may be rewarded or not. The presence (respectively absence) of a modulatory reward signal leads to potentiation (respectively depression) of plastic synapses. **(b)** Steady-state of the synaptic strength as a function of the posterior probability for different values of the learning rate ratio $r = q_+/q_-$ (equation(1)). In panels **b** and **c** light to dark curves correspond to the learning rate ratio from 1 to 5, respectively. When the learning rate ratio is equal to one (lightest gray) the synaptic strength is equal to the posterior. **(c)** Difference in the steady-state of the synaptic strengths as a function of the difference in the posteriors and of the log posterior odds, for different learning rate ratios. Blue and red dashed lines show linear fits for the values of posterior between 0.2 and 0.8. Note that $P(B|S_i) = 1 - P(A|S_i)$.

Figure 2.

Posterior computation by plastic synapses when multiple cues are presented on each trial.

(a) Time course of learning in plastic synapses. Shown is the difference in the strengths of synapses from a cue-selective neural population onto the two action value-coding neural populations (*A* and *B*). In all figures different color shades from blue to red correspond to shape $\mathbb{S}_1$ to $\mathbb{S}_{10}$. All the synaptic strengths are initially set to 0.5; they reach their steady-states after a few hundred trials. (b) Average synaptic strength for each set of synapses is plotted as a function of the naive posterior probability that alternative *A* is assigned with reward, given that shape $\mathbb{S}_i$ is presented in any pattern (note that $\tilde{P}(B|\mathbb{S}_i) = 1 - \tilde{P}(A|\mathbb{S}_i)$). (c) Difference in the average synaptic strengths is linearly proportional to the log naive posterior odds.

Figure 3.

Choice behavior of the model and the subjective weight of evidence in the weather prediction task. **(a)** Probability of choosing alternative A as a function of the evidence favoring this alternative for all patterns with finite WOE. The evidence is equal to the sum of the WOE of all shapes in a pattern. For clarity, the sum of the WOE values for all patterns were binned into 0.1 intervals. Gray curve: fitted logistic function. Black curves: performance of strict and probabilistic Bayesian observers. **(b)** Subjective weight of evidence for each shape as a function of the assigned WOE to that shape.

Figure 4.

Model neural population activity during the weather prediction task. (a) firing activity of two (black: A, grey: B) choice-selective populations in the decision-making network on a few sample trials. On these trials, the trump shape favoring A (with infinite WOE) is presented in the first three epochs; in the last epoch any one of the five shapes providing evidence in favor of A is presented. Time zero: onset of choice targets, fixation point, and the first visual shape stimulus. Four epochs: sequential presentation of four shapes. Last dashed line: offset of fixation point. If presented shapes in consecutive epochs are strongly predictive of an alternative, the activity of the population selective for that alternative increases, whereas the other competing population is suppressed. (b) Sample trials where the

same shape with the WOE equal to 0.7 is presented in the first three epochs; in the last epoch any of the three shapes with the WOE larger or equal to 0.7 is presented. (c) Sample trials where the same shape with the WOE equal to 0.3 is presented in the first three epochs; in the last epoch any shape with the WOE larger or equal to 0.3 is presented. In this case, the population *B* wins the competition and determines the choice of the network on some trials. If presented shapes are not strongly predictive of either alternative, the difference between the activity of two populations remains small during the trial.

Figure 5.

Neural population activity is parametrically correlated with the logLR. **(a)** Effect of the logLR on the firing rate of decision neurons. The population activity in each epoch is aligned on the onset of each shape stimulus presentation, and the average over many trials is computed for five quintiles of the logLR in that epoch (plotted with different color shading, more red means larger logLR favoring alternative A). The logLR in each epoch is equal to the sum of the logLR of shapes which are presented before and during that epoch. **(b)** The average population firing rate as a function of the logLR for four epochs. Average firing rate during the last 250 ms is computed by grouping the logLR into 10 equal bins in each epoch. **(c)** Average activity as a function of the logLR in each epoch, plotted separately for trials where the choice is the preferred (black) or nonpreferred (gray) of the neural population. Legends: the slope (with the estimated s.e.m.) of the linear fit for each set of data points.

Figure 6.

Effect of prior probability on the choice behavior and neural activity. **(a)** Psychometric function for patterns with finite log posterior odds is plotted for four values of prior probability. The gray lines are logistic function fits. **(b)** Bias of the psychometric function (open circles), and of the differential synaptic input (learned evidence) about each shape (filled circles), as a function of the log prior odds. These biases are defined as the probability of selecting choice A at zero log posterior odds. An ideal Bayesian observer shows no bias. The dashed lines are only to guide the eyes. **(c)** The difference in the synaptic strengths for each shape as a function of the log naive posterior odds for different values of prior probability. Same symbol convention as in **a** and the dashed lines show linear fits. Different shapes are coded with different colors. **(d)** Average population activity in each epoch as a function of the logLR in that epoch in the case of $P(A) = 0.8$. The black (gray, respectively) points show data from trials where the choice is the preferred (nonpreferred, respectively) of the neural population. The slope with the estimated s.e.m. is shown for each linear fit separately.