Population codes of prior knowledge learned through environmental regularities

Silvan C. Quax[1*], Sander E. Bosch[1,2], Marius V. Peelen[1], Marcel A. J. van Gerven[1],

1 Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, The Netherlands

2 Vrije Universiteit, Amsterdam, The Netherlands

* s.quax@donders.ru.nl

Abstract

How the brain makes correct inferences about its environment based on noisy and ambiguous observations, is one of the fundamental questions in Neuroscience. Prior knowledge about the probability with which certain events occur in the environment plays an important role in this process. Humans are able to incorporate such prior knowledge in an efficient, Bayes optimal, way in many situations, but it remains an open question how the brain acquires and represents this prior knowledge. The long time spans over which prior knowledge is acquired make it a challenging question to investigate experimentally. In order to guide future experiments with clear empirical predictions, we used a neural network model to learn two commonly used tasks in the experimental literature (i.e. orientation classification and orientation estimation) where the prior probability of observing a certain stimulus is manipulated. We show that a model population of neurons learns to correctly represent and incorporate prior knowledge, by only receiving feedback about the accuracy of their inference from trial-to-trial and without any probabilistic feedback. We identify different factors that can influence the neural responses to unexpected or expected stimuli, and find a novel mechanism that changes the activation threshold of neurons, depending on the prior probability of the encoded stimulus. In a task where estimating the exact stimulus value is important, more likely stimuli also led to denser tuning curve distributions and narrower tuning curves, allocating computational resources such that information processing is enhanced for more likely stimuli. These results can explain several

different experimental findings and clarify why some contradicting observations concerning the neural responses to expected versus unexpected stimuli have been reported and pose some clear and testable predictions about the neural representation of prior knowledge that can guide future experiments.

Author summary

The probability with which certain events are observed in our environment plays an important role in how we perceive the world. In many situations humans use such knowledge about the environment efficiently to come to optimal inferences from the observations made. However, it remains unclear how the brain learns to incorporate such prior knowledge and what the neural representations are exactly. By simulating two tasks commonly used in experiments where we manipulate the probability of certain stimuli occurring we show that such prior knowledge can be acquired from feedback about the accuracy of individual observations , and does not need any explicit information about the probability of an event occurring. We identify different properties of the neural populations that can influence the neural responses in experiments and show how prior knowledge can influence these neural responses through different mechanisms. Interestingly, the networks learn to allocate neural resources efficiently to more likely stimuli in order to maximize task performance. These results give several interesting and testable predictions about the neural representation of prior knowledge in commonly used experimental paradigms that can guide future experiments.

Introduction

During perception, the brain is continually trying to infer the cause of the perceptual input that reaches the senses. Efficient neural machinery is required to successfully perform this inference given the noise and ambiguity in our observations. Prior knowledge about the occurrence of stimuli in our environment can help to make better perceptual decisions about these noisy and ambiguous observations. The influence of prior knowledge on the perceptual decisions made by humans and animals has been shown in many situations.

For example, the fact that light almost always comes from above, influences our perception of ambiguous shapes [1] and our prior knowledge of shadows causes equally bright surfaces to appear of different brightness [2]. Similarly, the perception of oriented

bars shows a bias towards the cardinal axis, the so called 'oblique' effect, which has been successfully explained by an optimal environmental prior which accounts for the frequency with which certain orientations occur in the natural environment [3]. Furthermore, owls systematically underestimate sound direction in their periphery, which has been explained as a bias induced by a prior that emphasizes directions near the center of gaze [4].

The prior knowledge that is used in such perceptual decision making should somehow match the environmental statistics in order to enable optimal inference. How the brain learns to represent such prior knowledge and incorporate it in its sensory processing, is still an open question. While one possibility is that prior knowledge is transferred through evolutionary means, this would lead to very inflexible priors, not optimally adapted to changing environments. Alternatively, the brain learns to represent prior probabilities about the environment during its lifetime through the frequency with which an object is observed over time [2].

Different ideas exist about the representation of prior knowledge in the brain. One popular theory is that uncertainty about a variable is represented together with the variable itself, by the variability in neural firing [5]. Thus prior knowledge should already be represented in the neural activity without sensory input [6]. Another theory is that specialized neurons encode prior uncertainty and sensory uncertainty [7]. The various predictions that these theories provide have not been conclusively confirmed though.

To gain more insight in the kind of representations of prior knowledge that could emerge in neural populations, we trained generic neural networks to perform two commonly-used tasks in the experimental literature. These networks can provide valuable insight into possible representations in neural populations. While it has been shown that generic neural networks can learn to perform probabilistic tasks that depend only on the current input, without a need to incorporate prior knowledge [8], it remains an open question to what extent these networks can learn to represent a statistically optimal prior and combine this with probabilistic input.

Here, we simulated a classification task and a stimulus estimation task, where the prior probability of observing a stimulus was manipulated. Both tasks use a fundamentally different prior; a discrete prior for the different classes in the classification task and a continuous prior for the stimulus estimation task. We show that a model population of neurons learns to correctly represent and incorporate prior knowledge in both, by only receiving feedback about the accuracy of their inferences and without probabilistic feedback. We find an opposing effect of prior and input uncertainty

on the sparseness of neural firing. In the classification task, prior knowledge lowers the activation thresholds for the more likely class, causing neurons to fire more easily. In the continuous task, a sharper prior leads to a shift and narrowing of tuning curves, as often observed in experimental data. These results lead to testable predictions that can help validate whether prior knowledge is represented similarly in humans and animals.

Methods

Neural population model

Our neural population model has a feedforward architecture consisting of an input layer, a hidden layer and an output layer (Fig 1). The input layer consists of a group of independent Poisson neurons that respond to a stimulus, s, with firing rates, $\mathbf{r} = [r_1, ..., r_N]$, each generated from its own independent Poisson distribution [8]. The probability that neuron i responds with firing rate r_i is given by

$$p(r_i = k) = \frac{\lambda_i^k}{k!} e^{-\lambda_i} \tag{1}$$

where the rate parameter λ_i is determined by Gaussian tuning curves $f_i(s)$, specific for each input neuron, scaled with a global contrast parameter c, i.e. $\lambda_i = cf_i(s)$. The Gaussian input tuning curves for each neuron are thus given by

$$\lambda_i = ce^{-\frac{(s-\phi_i)^2}{2\sigma^2}} \tag{2}$$

where the maxima of the tuning curves, $\phi = [\phi_1, \ldots, \phi_N]$, were uniformly distributed over a range of -20 to 20 and the width of every tuning curve was given by $\sigma^2 = 10$. The contrast parameter, c, determined the quality of the stimuli. The range of contrasts used depended on the task and is specified in the task descriptions below.

The hidden layer was densely connected to the input layer with weights, $\mathbf{W}$, and consisted of K rectified linear units (ReLUs) with bias parameters b^k and firing rates $\mathbf{h} = [h^1, \ldots, h^K]$

$$h^k = \max\left[0, \mathbf{W}^T\mathbf{r} + b^k\right] \tag{3}$$

for $1 \leq k \leq K$.

The output layer of the network was densely connected to the hidden layer via output weights, $\mathbf{U}$, and consisted either of two units, one for every class, in the classification task (Fig 1A), or a single unit, in the stimulus estimation task (Fig 1D). A

softmax function was applied to the output units of the classification task to convert them to probabilities.

In some experiments a cue was added to indicate which class was more likely to occur on a certain trial. Two important ways in which a cue can modulate a neural population to represent prior knowledge are through gain modulations or biases. To enable the network to potentially learn such representations, we equipped the ReLU units with both a gain parameter, g^k_{cue}, and an extra bias parameter, b^k_{cue}, that could be modulated by the cue. That is,

$$h^k = \max\left[0, g^k_{cue}\left(\mathbf{W}^T\mathbf{r} + b^k + b^k_{cue}\right)\right] \tag{4}$$

for $1 \leq k \leq K$.

Fig 1. Neural population model and task overview. Two different tasks were studied, a classification task (top row) and a stimulus estimation task (bottom row). The neural network for the classification task and estimation task had a similar input and hidden layer, only the output layer differed. **(A)** Two softmax outputs (one for each class) indicated the probability that the stimulus observed came from the respective class in the classification task. **(B)** The stimuli from each class came from a normal distribution with a mean and variance specific to that class. Different prior probabilities for the classes were tested. **(C)** The generative model used to generate a stimulus from a selected class and subsequently generate firing rates for the input Poisson neurons. **(D)** A single output indicated the estimated stimulus value in the estimation task. **(E)** The stimuli presented to the network were drawn from a normally distributed prior. Different priors were tested (indicated by gray scale). **(F)** The generative model for the estimation task, used to generate firing rates for the input Poisson neurons from a generated stimulus.

Network training

All simulations were performed in Python. To train our neural networks we used the PyTorch neural network library [9]. The networks were trained using stochastic gradient descent [10]. To optimize the network parameters we used Adam [11], with a learning rate of 2e−4 and a weight decay regularizer of 1e−5. We used a minibatch size of 10 per iteration. Every network was trained for 100 epochs with 1000 iterations per epoch. Stimuli were presented in a random order during training.

Classification task

To study the representations of prior knowledge emerging in a population of neurons, we designed a classification task. Stimuli could be drawn either from class 1 or class 2. The relative frequency with which a stimulus was drawn from a class determined the prior probability with which stimuli from this class occurred. The stimulus was converted to firing rates by a population of Poisson neurons, which in turn served as the input to a population of rate-based neurons that was trained using error-backpropagation between the predicted and the correct class label.

In this classification task the network had to decide whether observed stimuli, s, belonged either to the first class ($C = 1$) or the second class ($C = 2$). Both classes were determined by an one-dimensional Gaussian distribution with the same variance, $\sigma_1^2 = \sigma_2^2 = 25$, but different means, with $\mu_1 = -5$ and $\mu_2 = 5$ (Fig 1B,C). Stimuli with contrast values $c \in \{0.5, 1.2, 1.9, 2.6, 3.3, 4.0\}$ were presented. We used five different priors over the two classes chosen from $\pi \in \{0.25, 0.33, 0.5, 0.67, 0.75\}$ such that $p(C = 1) = 1 - p(C = 2) = \pi$. The prior probabilities for the two classes determined the frequency with which samples from these classes were shown during training.

To observe how optimal our network performed this task we compared it with the optimal Bayesian posterior probability. We determined these posterior probabilities for both classes. We wish to estimate the posterior probability of a class C given observations of neuronal responses $\mathbf{r}$:

$$P(C \mid \mathbf{r}) \propto p(\mathbf{r} \mid C)P(C) \ , \tag{5}$$

where the likelihood is given by:

$$p(\mathbf{r} \mid C) = \int p(\mathbf{r} \mid s)p(s \mid C)ds \ . \tag{6}$$

In our task we draw the stimulus from a normal distribution depending on the class:

$$p(s \mid C) = \mathcal{N}(s; \mu_C, \sigma_C). \tag{7}$$

The conditional probability of neuronal responses $\mathbf{r}$ of the Poisson neurons can be written in the following form:

$$\begin{aligned} p(\mathbf{r} \mid s) &= \prod_i \frac{e^{-\lambda_i(s)}\lambda_i(s)^{r_i}}{r_i!} \\ &= \frac{e^{\sum_i -\lambda_i(s)} e^{\sum_i \log \lambda_i(s) r_i}}{r_i!} \\ &\propto e^{\sum_i \log \lambda_i(s) r_i} \end{aligned} \tag{8}$$

where we dropped the parts that are independent of s and assumed that $\sum_i -\lambda_i(s)$ is constant and thus also independent of s.

If we let the tuning functions take a Gaussian form this results in

$$p(\mathbf{r} \mid s) = \mathcal{N}(s; \mu_r, \sigma_r^2) \tag{9}$$

with $\mu_r = \phi^T \mathbf{r}(\mathbf{1}^T \mathbf{r})^{-1}$ and $\sigma_r^2 = \sigma^2(\mathbf{1}^T \mathbf{r})^{-1}$.

Since both parts of the likelihood are normally distributed we can marginalize the likelihood to arrive at the following form

$$\begin{aligned} p(\mathbf{r} \mid C) &\propto \int \frac{e^{-\frac{(s-\mu_r)^2}{2\sigma_r^2}}}{\sqrt{2\pi\sigma_r^2}} \frac{e^{-\frac{(s-\mu_C)^2}{2\sigma_C^2}}}{\sqrt{2\pi\sigma_C^2}} ds \\ &= \frac{1}{\sqrt{2\pi\left(\sigma_r^2 + \sigma_C^2\right)}} e^{-\frac{\left(\mu_r^2 + \mu_C^2 + 2\mu_r\mu_C\right)}{2\left(\sigma_r^2 + \sigma_C^2\right)}} \end{aligned} \tag{10}$$

The log posterior ratio of the probability that an observation is of one of both classes is given by

$$\begin{aligned} d &= \log \frac{P\left(C = 1 \mid \mathbf{r}\right)}{P\left(C = 2 \mid \mathbf{r}\right)} \\ &= \log \frac{p\left(\mathbf{r} \mid C = 1\right)}{p\left(\mathbf{r} \mid C = 2\right)} + \log \frac{P\left(C = 1\right)}{P\left(C = 2\right)} \end{aligned} \tag{11}$$

For our task filling in Equation 10 in Equation 11 drops out the normalization term and gives

$$\begin{aligned} d &= \log \frac{e^{-\frac{(\mu_r - \mu_1)^2}{2\left(\sigma_r^2 + \sigma_C^2\right)}}}{e^{-\frac{(\mu_r - \mu_2)^2}{2\left(\sigma_r^2 + \sigma_C^2\right)}}} + \log \frac{P\left(C = 1\right)}{P\left(C = 2\right)} \\ &= \frac{2\mu_r\left(\mu_1 - \mu_2\right) + \mu_2^2 - \mu_1^2}{2\left(\sigma_r^2 + \sigma_C^2\right)} + \log \frac{N_1}{N_2} \end{aligned} \tag{12}$$

where μ_c is the mean response of class c and N_c is the number of training samples that have been shown for class c.

By plugging in the expressions for μ_r and σ_r^2 in Equation 12, we obtain the final expression for the log posterior ratio:

$$d = \frac{2\mu_r\left(\mu_1 - \mu_2\right) + \mu_2^2 - \mu_1^2}{2\left(\sigma_r^2 + \sigma_C^2\right)} + \log \frac{N_1}{N_2} \ . \tag{13}$$

The posterior probability for class 1 can then be determined by

$$P\left(C = 1 \mid \mathbf{r}\right) = \frac{1}{1 + e^{-d}} \ . \tag{14}$$

Estimation task

In the estimation task, the network had to estimate the value of the observed stimulus s. Stimuli were drawn from a Gaussian distribution with mean $\mu_s = 0$ and variance σ_s^2.

The variance of the Gaussian distributions determined the prior probability that stimuli were presented during training (Fig 1E,F). Priors with different variances were tested: $\sigma_s^2 \in \{100, 50, 25, 10, 5\}$. Stimuli with contrast values $c \in \{0.30, 0.72, 1.45, 2.26, 2.86, 3.2\}$ were presented.

To determine the optimal Bayesian estimate, we would like to estimate the most probable stimulus value s, given observations of neuronal responses $\mathbf{r}$

$$p(s \mid \mathbf{r}) = \frac{P(\mathbf{r} \mid s)\, p(s)}{p(\mathbf{r})} \tag{15}$$

For these continuous values we take a normally distributed prior as example, to keep our Bayesian approximation as simple as possible

$$p(s) = \mathcal{N}(s; \mu_s, \sigma_s) \tag{16}$$

The response to a certain stimulus value is again given by Poisson neurons (Eq 8-9). The posterior will then be proportional to the product of two normal distributions, which will again be normally distributed

$$\begin{aligned} p(s \mid \mathbf{r}) &\propto e^{-\frac{(s-\mu_r)^2}{2\sigma_r^2}} e^{-\frac{(s-\mu_s)^2}{2\sigma_s^2}} \\ &= e^{-\frac{\left(s - \frac{\mu_r \sigma_s^2 + \mu_s^2 \sigma_r}{\sigma_r^2 + \sigma_s^2}\right)^2}{2\left(\frac{\sigma_r^2 \sigma_s^2}{\sigma_r^2 + \sigma_s^2}\right)}} \end{aligned} \tag{17}$$

The posterior mean is given by the maximum-a-posteriori (MAP) estimate, which is used as optimal Bayesian estimate of the underlying stimulus value and compared against the stimulus estimate given by the network

$$s_{MAP} = \frac{\mu_r \sigma_s^2 + \mu_s \sigma_r^2}{\sigma_r^2 + \sigma_s^2} \tag{18}$$

We can also identify the variance of the posterior from eq. 17 as

$$\sigma_{pos} = \frac{\sigma_r^2 \sigma_s^2}{\sigma_r^2 + \sigma_s^2} \tag{19}$$

Results

To study the representations of prior knowledge emerging in a population of neurons, we designed two representative tasks with a distinctive type of prior covering many behavioral experiments. In the first task the network had to decide to which of two classes a stimulus belonged. A discrete prior determined the prior probability of observing a stimulus from one of these two classes. In the second task the network had to estimate the value of a stimulus. Here a continuous prior determined the prior

probability of observing a certain stimulus value (see Methods section for further details on the different tasks). Since the prior affects the optimal outputs in both tasks differently, there could be distinctive representations emerging, making both example tasks worth investigating. In the first part of the results we investigate the classification task, while the latter part of the results focuses on the estimation task.

Classification task

Optimal prior learned on trial-by-trial basis

To investigate whether the network was able to learn a prior about the occurrence of the different classes and incorporate this optimally, we compared a linear softmax readout of the network with an optimal Bayesian solution of the task (Fig 2).

Fig 2. Network posterior matches optimal Bayesian estimates. The neural network learned to choose each class with a certain probability by getting feedback about correct or incorrect classification. Based on Bayesian statistics the optimal posterior estimate depends on the prior probability of a class occurring. This prior probability was influenced by the relative frequency that a stimulus from a class was shown during training. Five prior probability ratios were tested: [0.25, 0.75], [0.33, 0.67], [0.5, 0.5], [0.67, 0.33] and [0.75, 0.25] for class 1 and class 2 respectively. The posterior that the network learned is plotted against the optimal Bayesian estimates that accounts for the prior (solid line), or that ignores the prior probability, i.e. a uniform prior (dashed line). The network posterior matches the optimal Bayesian estimate that accounts for the prior and thus learns to correctly incorporate prior knowledge in its connections. Error bars represent standard deviation over the validation trials.

While the network did not receive any probabilistic feedback about the prior probabilities of either class, it formed a correct hidden representation of these prior probabilities and adjusted its choices in a Bayes-optimal way. If the network would only perform a maximum likelihood estimation (i.e. use a uniform prior), the network posterior would have matched the dashed line (see Fig 1). This result suggests that no higher order cognitive functions that keep track of the relative frequency by which an object class is observed are needed to form a Bayes optimal prior in the brain. The use of this prior knowledge led to more accurate responses for the class with higher prior probability (Fig 3), a behaviour that humans have also shown in many studies [3, 12].

Fig 3. Prior enhances accuracy for more likely class. The accuracy with which stimuli were correctly classified was influenced by the prior probability of a certain class. A uniform prior (panel 3) led to equal accuracy on stimuli from both classes. A network trained on more stimuli from class 2 than class 1 (panel 1 and 2) showed higher accuracy on class 2 stimuli, but lower accuracy on class 1 stimuli, compared to the uniform prior. On the other hand a network trained on more stimuli from class 1 than class 2 (panel 4 and 5) showed higher accuracy on class 1 stimuli, but lower accuracy on class 2 stimuli, compared to the uniform prior.

Representations of prior probability

To investigate how the neural population encoded the prior probability of the different classes, we first looked at the neural activation in response to different stimuli. Both uncertainty in the input neurons and in the prior probability of the different classes were manipulated. Less uncertainty in the input led to sparser population responses, as previously reported by [8]) (Fig 4A). However, when the posterior uncertainty is considered, we observed that more certainty about the presented class actually led to stronger responses, where more neurons fire. Uncertainty in the input strongly changes responses to stimuli that lead to uncertain posterior responses (around 0.5). These responses could be explained by a template matching code, where the neural population develops templates for both classes against which the input is matched. Very certain inputs led to clear matches with either of the classes, or a clear mismatch when a stimulus came from the decision boundary. On the other hand, more noisy inputs led to partial matches with both templates, even for stimuli around the decision boundary.

Fig 4. Neural activity changes due to input and prior uncertainty. (A) Number of neurons with non-zero activity in response to both stimuli from class 1 and class 2. Fewer neurons in the hidden layer responded when the input contrast was higher. On the other hand, more neurons responded when the posterior was certain which class was presented. **(B)** Number of neurons with non-zero activity in response to stimuli from class 1. While fewer neurons responded to higher contrast stimuli, more neurons were active when the prior probability of the presented stimulus class was higher. **(C)** The average activity of the population increased both with stimulus contrast and prior probability. **(D)** The weights in the network require less adaptation in response to likely stimuli, while the contrast has no influence on the strength of the weight updates.

Next, we considered the effect of uncertainty in the prior probability of the classes. A higher prior probability for a certain class led to more neurons firing in response to stimuli of this class. Uncertainty in the input and uncertainty in the prior thus have opposing effects on the sparseness of the neural responses. While a higher contrast in the input (less uncertainty) led to sparser neural responses, a higher prior probability for a class (also less uncertainty) led to less sparse neural responses for this class (Fig 4B).

The average firing rate for the entire neural population increased both with more certainty in the input, as well as with more certainty in the prior (Fig 4C). These enhanced responses for stimuli that have higher prior probability have experimentally been observed in several studies [13–15]. However, other studies have shown decreased responses for more likely stimuli [12, 16–18]. A possible explanation for decreased neural responses to more likely stimuli, could stem from the fact that fewer errors are made on these trials (Fig 3), requiring less updating of synaptic weights. The weights between neurons in our model require indeed less adaptation in response to more likely stimuli (Fig 4D). While the weights in our model are updated artificially by gradients, the brain would need to convey such learning signals through actual neural activation, requiring lower learning signals for more likely stimuli. Thus if such learning signals make up an important part of the activity of the brain, an overall lower activation could be expected for more likely stimuli.

Prior probability encoded by changing activation thresholds

The neural responses of the population offer measurable predictions that can be experimentally validated. A more difficult, but nonetheless interesting, question to answer experimentally is which mechanism gives rise to these neural responses. Changes in the properties of individual neurons and wiring between neurons in the population should somehow represent the prior probability of diverse stimuli.

To see which how the parameters in our network encoded this prior information we determined for every neuron in the population whether it contributed more to class 1 or class 2. We analyzed how the weights of input to hidden units, $\mathbf{W}$, the bias of the hidden units, $\mathbf{b}$, and the weights from hidden to output units, $\mathbf{U}$, changed with different prior probabilities for a class. The neurons in the population responsive to the class with higher prior probability developed a higher bias, effectively lowering their activation threshold (Fig 5A). Lower thresholds explain the fact that more neurons fire when a class has higher prior probability. We observed no effects of the prior probability on the gain of the neurons (Fig 5B). As a control we also checked for the readout

weights of the neural population, but found no effect of the prior probability on the readout weights to either class (Fig 5C).

Fig 5. Prior probability encoded by changing activation thresholds. The hidden units could change the decisions of the output units using the bias parameters ($\mathbf{b}$), hidden weights ($\mathbf{W}$) or output weights ($\mathbf{U}$). For every hidden neuron, it was determined whether it contributed stronger to class 1 or class 2 by comparing the output weights to class 1, $\mathbf{U}(C=1)$, against the output weights to class 2, $\mathbf{U}(C=2)$. That resulted in two subpopulations of hidden units: $\mathbf{H}_1 = \mathbf{H}[\mathbf{U}(C=1) > \mathbf{U}(C=2)]$ and $\mathbf{H}_2 = \mathbf{H}[\mathbf{U}(C=1) < \mathbf{U}(C=2)]$. The biases and weights of the hidden units were averaged accordingly and are plotted for different prior probability ratios ($P(C=1) \in [0.25, 0.33, 0.5, 0.67, 0.75]$). **(A)** Only the biases of the hidden neurons changed as a result of different prior probabilities, where a higher probability of a class led to higher bias values in neurons coding for this class, thus lowering the activation threshold in these neurons. **(B)** The hidden weights did not change with prior probability. **(C)** Neither did the output weights. Error bars represent standard error over the units. (* $p < 0.01$, ** $p < 0.001$).

Prior probabilities indicated by contextual cues

In our everyday environment the prior probability of a stimulus occurring often depends strongly on the context. Since in these situations the context has to be able to manipulate the prior on a very short time scale, it is not possible that long term synaptic changes underlie such representations. Different hypotheses exist about the neural mechanisms underlying the encoding of prior knowledge dynamically in such cases. A popular hypothesis is that prior knowledge is encoded by modulating the gain of neurons. Neurons encoding for a more probable stimulus fire more strongly in response to this stimulus. Another possibility is that prior knowledge modulates the activation threshold of neurons. A neuron encoding a more probable stimulus is then more easily activated in response to this stimulus.

To see if our network is able to learn such a prior that has to change flexibly depending on the context, we added an extra cue input feeding into the hidden layer, with connections to both the gain and bias of the neurons. A single neural network was trained that had to learn to change its prior probabilities depending on the cue. Given the cue the probability of observing a stimulus from class 1 was either 25% or 75% (vice versa for class 2). The network was able to learn to correctly account for the prior

probabilities for the different context cues (Fig 6A,B). Since all network parameters are

Fig 6. Contextual cue modifies activation thresholds. The frequency with which each class of stimuli was shown depended on a contextual cue. Depending on the cue the prior probability of class 1 and class 2 could either be [0.25, 0.75] or [0.75, 0.25]. **(A,B)** The network learned to adapt its output probability to reflect the optimal Bayesian estimate given the cue. Error bars represent standard deviation over the validation trials. **(C,D)** For every hidden neuron, it was determined whether it contributed stronger to class 1 or class 2. The input from the cue to the bias and gain of the hidden neurons was averaged accordingly for both both cues. The cue for the context where class 2 is more likely caused both a higher input to the bias and gain of hidden neurons encoding class 2 and a lower input to the hidden neurons encoding class 1. The opposite effect occurred for the cue for the context where class 1 is more likely. Error bars represent standard error over the units. (* $p < 10^{-7}$, ** $p < 10^{-10}$).

the same for both cue conditions, the network had to learn a flexible mechanism to account for the prior probability. Closer inspection of the effect of the cue on the hidden neuron input shows that the hidden neurons encoding the more likely class received both a stronger baseline input from the cue (Fig 6C) and had a stronger gain (Fig 6D).

Since both the bias and the gain are influenced by the cue, we investigated whether one of these played a more important role in representing the correct prior knowledge. We analyzed the accuracy of the network on classifying both classes under the different cues. When the cue influenced both the gain and the bias, the prior caused a large difference in accuracy between both classes (Fig 7, blue bars). Next we turned off either the influence of the cue on the gain (Fig 7, orange bars) or the influence of the cue on the bias 7, green bars). Turning off the influence of the cue on the gain diminished the effect of the prior on the accuracy only slightly. Turning of the influence of the cue on the bias on the other hand, diminished the effect of the prior on the accuracy almost completely. The bias thus seems to play a greater role in representing the prior knowledge.

Estimation task

Prior probability encoded by efficient allocation of neural resources

In the previous section we studied how prior knowledge about a discrete class can be learned and encoded in a neural population. However, many tasks in the brain require a

Fig 7. Cue influence mainly through bias. The accuracy for both classes was determined for the cued priors (left, [0.25, 0.75], right, [0.75,0.25]). **(A,B)** The class with the highest prior probability according to the cue, was decoded much more accurately when the cue influenced both the gain and the bias (blue bars). Turning off the cue effect on the gain diminished this effect only slightly (orange bars). Turning off the cue effect on the bias diminished this effect much more, almost completely removing any effect of the cued prior (green bars).

prior representation of a continuous environmental variable. E.g., an often used example in the experimental literature is the estimation of the orientation of gratings [3], but also sound localization has been explained this way [4]. Since the effect of such a continuous prior on the Bayesian posterior is mathematically more complex, this might also require a different neural representation. To study whether such a continuous prior can also be successfully learned on a trial-by-trial basis and whether the representations differ from the discrete case, we trained our neural population to estimate the value of a noisy stimulus. Different Gaussian priors, $\mathcal{N}(\mu = 0, \sigma^2 \in \{100, 50, 25, 10, 5\})$, were tested, where the probability that a certain stimulus value was presented to the network was determined by this prior distribution. The same neural population model with Poisson input neurons and ReLU hidden units as for the discrete case was used. Instead of a softmax output layer, we used a linear readout for the continuous case. The mean squared error loss was used to optimize the parameters of the network through backpropagation.

The neural population learned to estimate the stimulus in a Bayesian way, correctly incorporating the prior knowledge about the relative frequency with which a stimulus was presented (Fig 8). A uniform prior or maximum likelihood estimate did not match the network output for priors differentiating strongly from uniform (Fig 8E).

The network showed a 'perceptual' bias similarly to those that are observed in humans and animals (Fig 9) [4]. The size of this bias depended on the sharpness of the prior distribution, with a sharper prior having a more pronounced bias. The bias in the network estimates matched those predicted by the optimal Bayesian estimate.

One of the effects that is often observed in humans and animals is that the limited neural resources that are available to the organism, are used efficiently, such that more resources are allocated to ranges of stimuli that occur more often [3, 4]. This improves performance for stimuli in these ranges, while consequently decreasing performance for stimuli that do not occur often. Our network showed a similar performance pattern,

Fig 8. Network learns to incorporate prior knowledge optimally. Stimuli were drawn from Gaussian priors with different variances (i.e. $\sigma_p^2 \in 100, 50, 25, 10, 5$). During testing, stimuli in the range $[-10, 10]$ were presented, making the $\sigma_p^2 = 100$ prior effectively uniform. The posterior that the network learned is plotted against the optimal Bayesian estimates that accounts for the prior (solid line), or that ignores the prior probability, i.e. a uniform prior (dashed line). While for $\sigma_p^2 = 100$ both Bayesian estimates are similar in the range $[-10, 10]$, there is a big difference between both estimates for $\sigma_p^2 = 5$. The network estimates match the Bayesian estimates accounting for the prior distribution, but not Bayesian estimates for the prior distribution. The network thus learns to incorporate the prior knowledge correctly. Error bars represent standard deviation over the validation trials.

Fig 9. Learned prior causes perceptual bias that is statistically optimal. The network biased its estimates of the stimulus towards the mean of the prior (solid line). This effect was strongest for the sharpest prior ($\sigma_p^2 = 5$). This perceptual bias matched the bias predicted by an optimal Bayesian observer (dashed line). Error bars represent standard deviation over the validation trials.

where stimuli around the maximum prior probability were most accurately estimated, while those further away from the maximum had a larger estimation error (Fig 10).

Fig 10. Prior probability lowers estimation error. When a stimulus value has a higher prior probability the error in the estimation produced by the network is lower, while stimulus values with low probability result in higher estimation errors. The sharpest prior, $\sigma_p^2 = 5$, has the lowest error for estimations around the mean of the prior.

The neural activation patterns show a different influence of the prior for the estimation task. While a higher prior probability caused an increase in the number of neurons that fire and the overall activity in the neural population (Fig 4, we see the opposite effect in the estimation task (Fig 11A,B). Responses become sparser when the prior distribution is more narrow (Fig 11A). The effect of the stimulus contrast remains the same as for the classification task, with sparser responses on stimuli with higher contrast. The average activity in the neural population decreases when the prior distribution is more narrow (Fig 11B). The effect of the stimulus contrast on the average activity also remains the same, with higher activity in the neural population for higher stimulus contrast. For the lowest contrast value the average activity increases

slightly again.

Fig 11. Neural activity changes due to prior uncertainty. (A) Number of neurons with non-zero activity in response to stimuli from different prior distributions. Fewer neurons fired when the prior distribution was more narrow. **(B)** The average activity of the population decreased with a more narrow prior probability distribution. While higher contrast led to increased neural activity. For the lowest contrast value the activity in the network also rises again.

To understand what the underlying representations are that encode this prior, we investigated the tuning curve properties of the hidden neurons. These tuning curves could change their properties such that more likely stimuli get allocated more of the limited resources. One popular hypothesis based on Gaussian tuning curves is that the density of the peaks of the tuning curves correlates with the prior probability of a stimulus, such that the highest density of tuning curves is at the stimulus value with the highest prior probability [3]. We plotted a histogram of the tuning curve peaks to see whether the highest density is found around the maximum of the prior distribution. We found a bimodal distribution where the highest density of tuning curve peaks is found slightly away from the maximum prior probability (Fig 12).

Fig 12. Density of tuning curves does not follow prior probability. A histogram of the peaks of all hidden neuron tuning curves is shown for the different prior conditions. While the distribution of tuning curve peaks becomes denser for the sharper prior (right panel), there is not a direct match with the Gaussian prior probability density. The tuning curve peaks rather concentrate around the flanks of the prior probability distribution.

One possible explanation is that the neural population learns to allocate the most sensitive parts of the tuning curves to the most probable stimuli range. This is part of the tuning curve where the slope is steepest, leading to the highest Fisher information. To test this alternative, we determined the slopes of the tuning curves by taking the second order derivative of the tuning curves and calculating the maximum of this slope. The histogram of these maximal slope distributions closely matches the underlying prior probability, indicating that the neural population indeed optimizes its tuning curves according to this criterion (Fig 13).

Besides shifting tuning curves to position their maximal slope around more likely stimuli values, the tuning curves themselves can be made sharper in order to get more

Fig 13. The slope of the tuning curves follow the prior distribution. A histogram of the maximal slope of every tuning cure is plotted for the different prior distributions. The density of maximal slopes concentrates more for the sharper prior probability distribution (right panel), forming a Gaussian bell shape.

accurate responses around more likely stimuli values. We measured the width of the tuning curves as the full-width-at-half-maximum. The distribution of tuning curve widths shows a clear narrowing of the tuning curves for the prior distribution with the lowest variance (Fig 14). Narrower tuning curves for more likely stimuli can result in the sparser activations observed in Fig 11A.

Fig 14. Tuning curves become narrower for sharper prior. A histogram of the full-width-at-half-maximum for all tuning curves is shown. For a sharper prior (right panel) tuning curves became narrower, effectively increasing the slopes of the tuning curves and thus the discriminability of different stimuli.

To investigate whether the biases in the neural population also play a role in encoding prior information we divided the neurons into a group encoding more likely, or 'expected', stimuli and more unlikely, or 'unexpected' stimuli. Looking at the results from Fig 13, we selected the neurons with their maximal slope within one standard deviation of the prior probability distribution as those sensitive to 'expected' stimulus values and neurons with their maximal slope outside one standard deviation of the prior probability distribution as those sensitive to 'unexpected' stimulus values. The neurons sensitive to 'expected' stimuli had a stronger bias, or baseline activation, then neurons sensitive to 'unexpected' stimuli for all different prior distribution (Fig 15).

Fig 15. Bias influence by prior probability of stimuli. For every trained network we determined which neurons were sensitive to 'expected' stimuli and which were sensitive to 'unexpected' stimuli. The bias was much stronger for neurons sensitive to expected stimuli for all prior probability distributions. (* $p < 0.02$, ** $p < 0.005$).

Discussion

Summary of results

We have shown that a neural population is able to learn to incorporate prior knowledge optimally by error based learning without any probabilistic information, both for discrete priors (classification task) and continuous priors (stimulus estimation task). For both discrete and continuous priors, errors for likely stimuli are lower than errors for unlikely stimuli, as often observed in empirical literature. However, the mechanism through which this is achieved differs for both types of priors. In the classification task, a higher prior probability for a class led to higher activity in neurons encoding this class. On the other hand, stimuli with higher input contrast led to lower activity in the network. The increase in neural activity by the prior was caused by a lowered activation threshold of the neurons encoding the more likely class. When a contextual cue was added, a similar effect was achieved by a strong excitatory input to neurons encoding the more likely class and an inhibitory input to the neurons encoding the less likely class. In the stimulus estimation task the opposite effect of prior probability on the neural activity was observed. A narrower prior distribution led to lower activity in the neural population. The prior changed the distribution of tuning curves, increasing the density and narrowing the tuning curves. Not the peak of the tuning curves, but the maximal slope followed the prior distribution. Neurons with their maximal slopes around 'expected' stimuli values had a lower activation threshold then neurons with their maximal slope around 'unexpected' stimuli values. Thus the effects of prior probability on neural activations are affected by a range of different factors, like lower activation thresholds, higher tuning curve density or narrower tuning curve width, which can play a different role depending on the exact task that is being performed.

Relation to previous work

These results give a strong indication that it is possible to learn an efficient incorporation of prior knowledge based on errors in individual observations, and without the need of higher cognitive regions keeping track of explicit probabilities. This learning through observations is supported by experimental studies showing the formation of an environmental prior over the lifetime [6, 19] and could potentially be how humans and animals learn to from optimal Bayesian priors [4, 20, 21].

The representations of prior knowledge have been empirically studied in various settings. Conflicting observations reporting either increasing or decreasing neural

responses to expected or unexpected stimuli have been found [12–18, 22]. Here we showed that different factors can influence such neural responses, often indistinguishable from purely the average neural activity in a population. In the classification task, our model predicts stronger activity in response to expected stimuli, which has been observed in multiple studies [13–15]. However, in the estimation task, our model predicts lower activity in response to expected stimuli in line with experiments showing lower blood-oxygen-level-dependent (BOLD) responses in humans [12, 16] and electrophysiological responses in non-human primates [17, 18] to expected stimuli. Besides the particular task that is being performed, differences in neural activations in response to expected or unexpected stimuli can also result from the particular neurons that are measured during experiments. More detailed experimental measurements revealed that responses are higher in neurons selective to expected stimuli versus those not selective to expected stimuli [12, 14], similar to our results in Fig 4.

The properties of individual neurons, such as the activation threshold, gain, sharpness of the tuning curves or density of the tuning curves are all factors that determine the average neural activity of a population. We showed that all these factors can be influenced by the prior probability of a certain stimulus. Which factors are contributing strongest depends on the actual task that is being performed, leading to different effects on the neural activity.

Our model shows a clear role for the bias and to lesser extent the gain of individual neurons in the encoding of expected stimuli. Interestingly, lowered activation thresholds as suggested by our model have also been implied by observations that neurons encoding a stimulus fire already more strongly in response to a cue indicating that the stimulus is expected [23–25], where the latter study also observed a subgroup of neurons showing gain effects in response to expected stimuli.

Besides the effects on the bias and gain of individual neurons we also found that the tuning curves of individual neurons adapted to efficiently encode more likely stimuli in the estimation task. The efficient allocation of limited neural resources to more likely stimuli has been suggested as a driving force underlying the formation of neural representations of prior knowledge [3, 26, 27]. Our results agree with such an efficient allocation of resources, with the important notion that the slopes of the tuning curves are actually concentrated around the most likely stimulus values. While some studies have suggested that the peak of a tuning curve is most important for encoding a stimulus [28], others claim that the steepest slope of a tuning curve is most important [29, 30]. Interestingly, a theoretical analysis of tuning curve information

properties has indicated that maximal-slope encoding is most efficient for two-alternative-forced-choice tasks (such as our classification task) as well as an estimation task under reasonable levels of neuronal noise [31], whereas peak-tuning-curve encoding has been shown to be more efficient only in an estimation task with high neuronal noise and a small population of neurons. The representation developed by our model agrees well with these findings and makes interesting suggestions for empirical validation.

Predictions for empirical validation

We have shown that the particular effect that the prior has on the neural activity depends on the particular task that is being performed. Different mechanisms can give rise to the observed effects on neural activity. To verify which mechanisms are at play under which circumstances, more detailed experiments are necessary, that not only investigate the effects of prior knowledge on neural activations, but also properties of individual neurons such as their receptive fields and activation thresholds.

We find a large effect on the activation thresholds of neurons, with only a small effect due to gain modulations when there is a cue to indicate which stimuli are expected and which are unexpected. Changes in the activation threshold as a result of neural activity have been observed in real neurons [32], and could potentially play a role in the encoding of prior knowledge. This mechanism we observed also allows for flexibly adjusting this activation threshold by increasing the baseline input to a neuron. This enables the neural population to adjust its prior expectations based on the context. While such lowered activation thresholds have also been implicitly observed by stronger firing rates for neurons encoding a stimulus in response to a cue indicating that the stimulus is expected [23–25], future studies could further confirm these changes in activation threshold. Validating this prediction requires intracellular recordings of sub-threshold activity, to see whether a contextual cue increases sub-threshold membrane potentials in neurons encoding stimuli that have a higher prior probability.

In the estimation task our model suggests that neural tuning curves are organized such that more neurons have their maximal slope around more likely stimuli. This could be validated in the by mapping tuning curves in response to oriented gratings in the visual cortex. An overrepresention of the cardinal axis due to prior exposure, as empirically observed [3], suggests that more neurons should have the maximal slope of their tuning curves around the cardinal axis, while less neurons have their maximal slope around the oblique orientations.

Future work

One important aspect of cognitive processing that is not included in our model, but can also have a major effect on neural activations, is the role that feedback processing plays in the brain. A popular theory that describes how prior knowledge influences our perceptual processing is predictive coding. According to this theory, the brain forms predictions about incoming sensory information based on its internal prior knowledge. Since prior knowledge 'explains away' some of the sensory input through feedback, this also influences neural responses to expected stimuli class [12]. Since we used only feed-forward neural networks to model neural responses, this prevents feedback mechanisms from playing a role in encoding prior knowledge. An important step for future simulations would be to incorporate such feedback connections.

One way of doing this is through recurrent computations by adding a time component to our networks. Such a time component also allows for distinguishing neural activity in response to a cue, from neural activity in response to actual stimuli. This would allow to answer the question how a neural population could dynamically adjust to expected stimuli in order to improve sensitivity around expected stimuli values.

Conclusions

We have shown that a model population of neurons can learn to correctly represent and incorporate prior knowledge, by only receiving feedback about the accuracy of their inference from trial-to-trial and without any probabilistic feedback. The prior probability of a stimulus affects the neural responses of the neural populations to this stimulus, where activations can either increase or decrease depending on the task. We show different mechanisms that influence the neural activations. We find that the activation threshold of neurons encoding a certain stimulus play an important role in representing its prior probability, with more likely stimuli having a lower activation threshold. In a task where estimating the exact stimulus value is important, more likely stimuli also lead to denser receptive field distributions and more narrow receptive fields, allocating computational resources such that information processing is enhanced for more likely stimuli. These results pose some clear and testable predictions about the neural representations of prior knowledge in commonly used experimental paradigms that can guide new experiments.

Acknowledgements

This research was supported by VIDI grant number 639.072.513 of The Netherlands Organization for Scientific Research (NWO).

References

1. Mamassian P, Goutcher R. Prior knowledge on the illumination position. Cognition. 2001;81(1):B1–B9.
2. Purves D, Beau Lotto R, Mark Williams S, Nundy S, Yang Z. Why we see things the way we do: evidence for a wholly empirical strategy of vision. Philosophical Transactions of the Royal Society of London Series B: Biological Sciences. 2001;356(1407):285–297.
3. Girshick AR, Landy MS, Simoncelli EP. Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. Nature Neuroscience. 2011;14(7):926.
4. Fischer BJ, Peña JL. Owl's behavior and neural representation predicted by Bayesian inference. Nature Neuroscience. 2011;14(8):1061.
5. Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. Nature Neuroscience. 2006;9(11):1432.
6. Berkes P, Orbán G, Lengyel M, Fiser J. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. Science. 2011;331(6013):83–87.
7. Vilares I, Howard JD, Fernandes HL, Gottfried JA, Kording KP. Differential representations of prior and likelihood uncertainty in the human brain. Current Biology. 2012;22(18):1641–1648.
8. Orhan AE, Ma WJ. Efficient probabilistic inference in generic neural networks trained with non-probabilistic feedback. Nature Communications. 2017;8(1):138.
9. Paszke A, Gross S, Chintala S, Chanan G, Yang E, DeVito Z, et al. Automatic differentiation in Pytorch. 2017;.
10. Rumelhart DE, Hinton GE, Williams RJ, et al. Learning representations by back-propagating errors. Cognitive Modeling. 1988;5(3):1.

11. Kingma DP, Ba J. Adam: A method for stochastic optimization. arXiv preprint arXiv:14126980. 2014;.

12. Kok P, Jehee JF, De Lange FP. Less is more: expectation sharpens representations in the primary visual cortex. Neuron. 2012;75(2):265–270.

13. Furmanski CS, Engel SA. An oblique effect in human primary visual cortex. Nature Neuroscience. 2000;3(6):535.

14. Jaramillo S, Zador AM. The auditory cortex mediates the perceptual effects of acoustic temporal expectation. Nature Neuroscience. 2011;14(2):246.

15. Doherty JR, Rao A, Mesulam MM, Nobre AC. Synergistic effect of combined temporal and spatial expectations on visual attention. Journal of Neuroscience. 2005;25(36):8259–8266.

16. Alink A, Schwiedrzik CM, Kohler A, Singer W, Muckli L. Stimulus predictability reduces responses in primary visual cortex. Journal of Neuroscience. 2010;30(8):2960–2966.

17. Kumar S, Kaposvari P, Vogels R. Encoding of predictable and unpredictable stimuli by inferior temporal cortical neurons. Journal of Cognitive Neuroscience. 2017;29(8):1445–1454.

18. Meyer T, Olson CR. Statistical learning of visual transitions in monkey inferotemporal cortex. Proceedings of the National Academy of Sciences. 2011;108(48):19401–19406.

19. Kaiser D, Quek GL, Cichy RM, Peelen MV. Object Vision in a Structured World. Trends in Cognitive Sciences. 2019;.

20. Mamassian P, Landy MS. Interaction of visual prior constraints. Vision Research. 2001;41(20):2653–2668.

21. Weiss Y, Simoncelli EP, Adelson EH. Motion illusions as optimal percepts. Nature Neuroscience. 2002;5(6):598.

22. Hansen K, Hillenbrand S, Ungerleider L. Effects of prior knowledge on decisions made under perceptual vs. categorical uncertainty. Frontiers in Neuroscience. 2012;6:163.

23. Peelen MV, Kastner S. A neural basis for real-world visual search in human occipitotemporal cortex. Proceedings of the National Academy of Sciences. 2011;108(29):12125–12130.

24. Battistoni E, Stein T, Peelen MV. Preparatory attention in visual cortex. Annals of the New York Academy of Sciences. 2017;1396(1):92–107.

25. Ikeda T, Hikosaka O. Reward-dependent gain and bias of visual responses in primate superior colliculus. Neuron. 2003;39(4):693–700.

26. Wei XX, Stocker AA. A Bayesian observer model constrained by efficient coding can explain'anti-Bayesian'percepts. Nature Neuroscience. 2015;18(10):1509.

27. Wei XX, Stocker AA. Efficient coding provides a direct link between prior and likelihood in perceptual Bayesian inference. In: Advances in Neural Information Processing Systems; 2012. p. 1304–1312.

28. Bosking W, Maunsell J. The correlation between the firing of individual MT neurons and behavioral response across different directions of motion. Society for Neuroscience Abstracts. 2004;31:935–937.

29. Seung HS, Sompolinsky H. Simple models for reading neuronal population codes. Proceedings of the National Academy of Sciences. 1993;90(22):10749–10753.

30. Purushothaman G, Bradley DC. Neural population code for fine perceptual decisions in area MT. Nature Neuroscience. 2005;8(1):99.

31. Butts DA, Goldman MS. Tuning curves, neuronal variability, and sensory coding. PLOS Biology. 2006;4(4):e92.

32. Desai NS, Rutherford LC, Turrigiano GG. Plasticity in the intrinsic excitability of cortical pyramidal neurons. Nature Neuroscience. 1999;2(6):515.

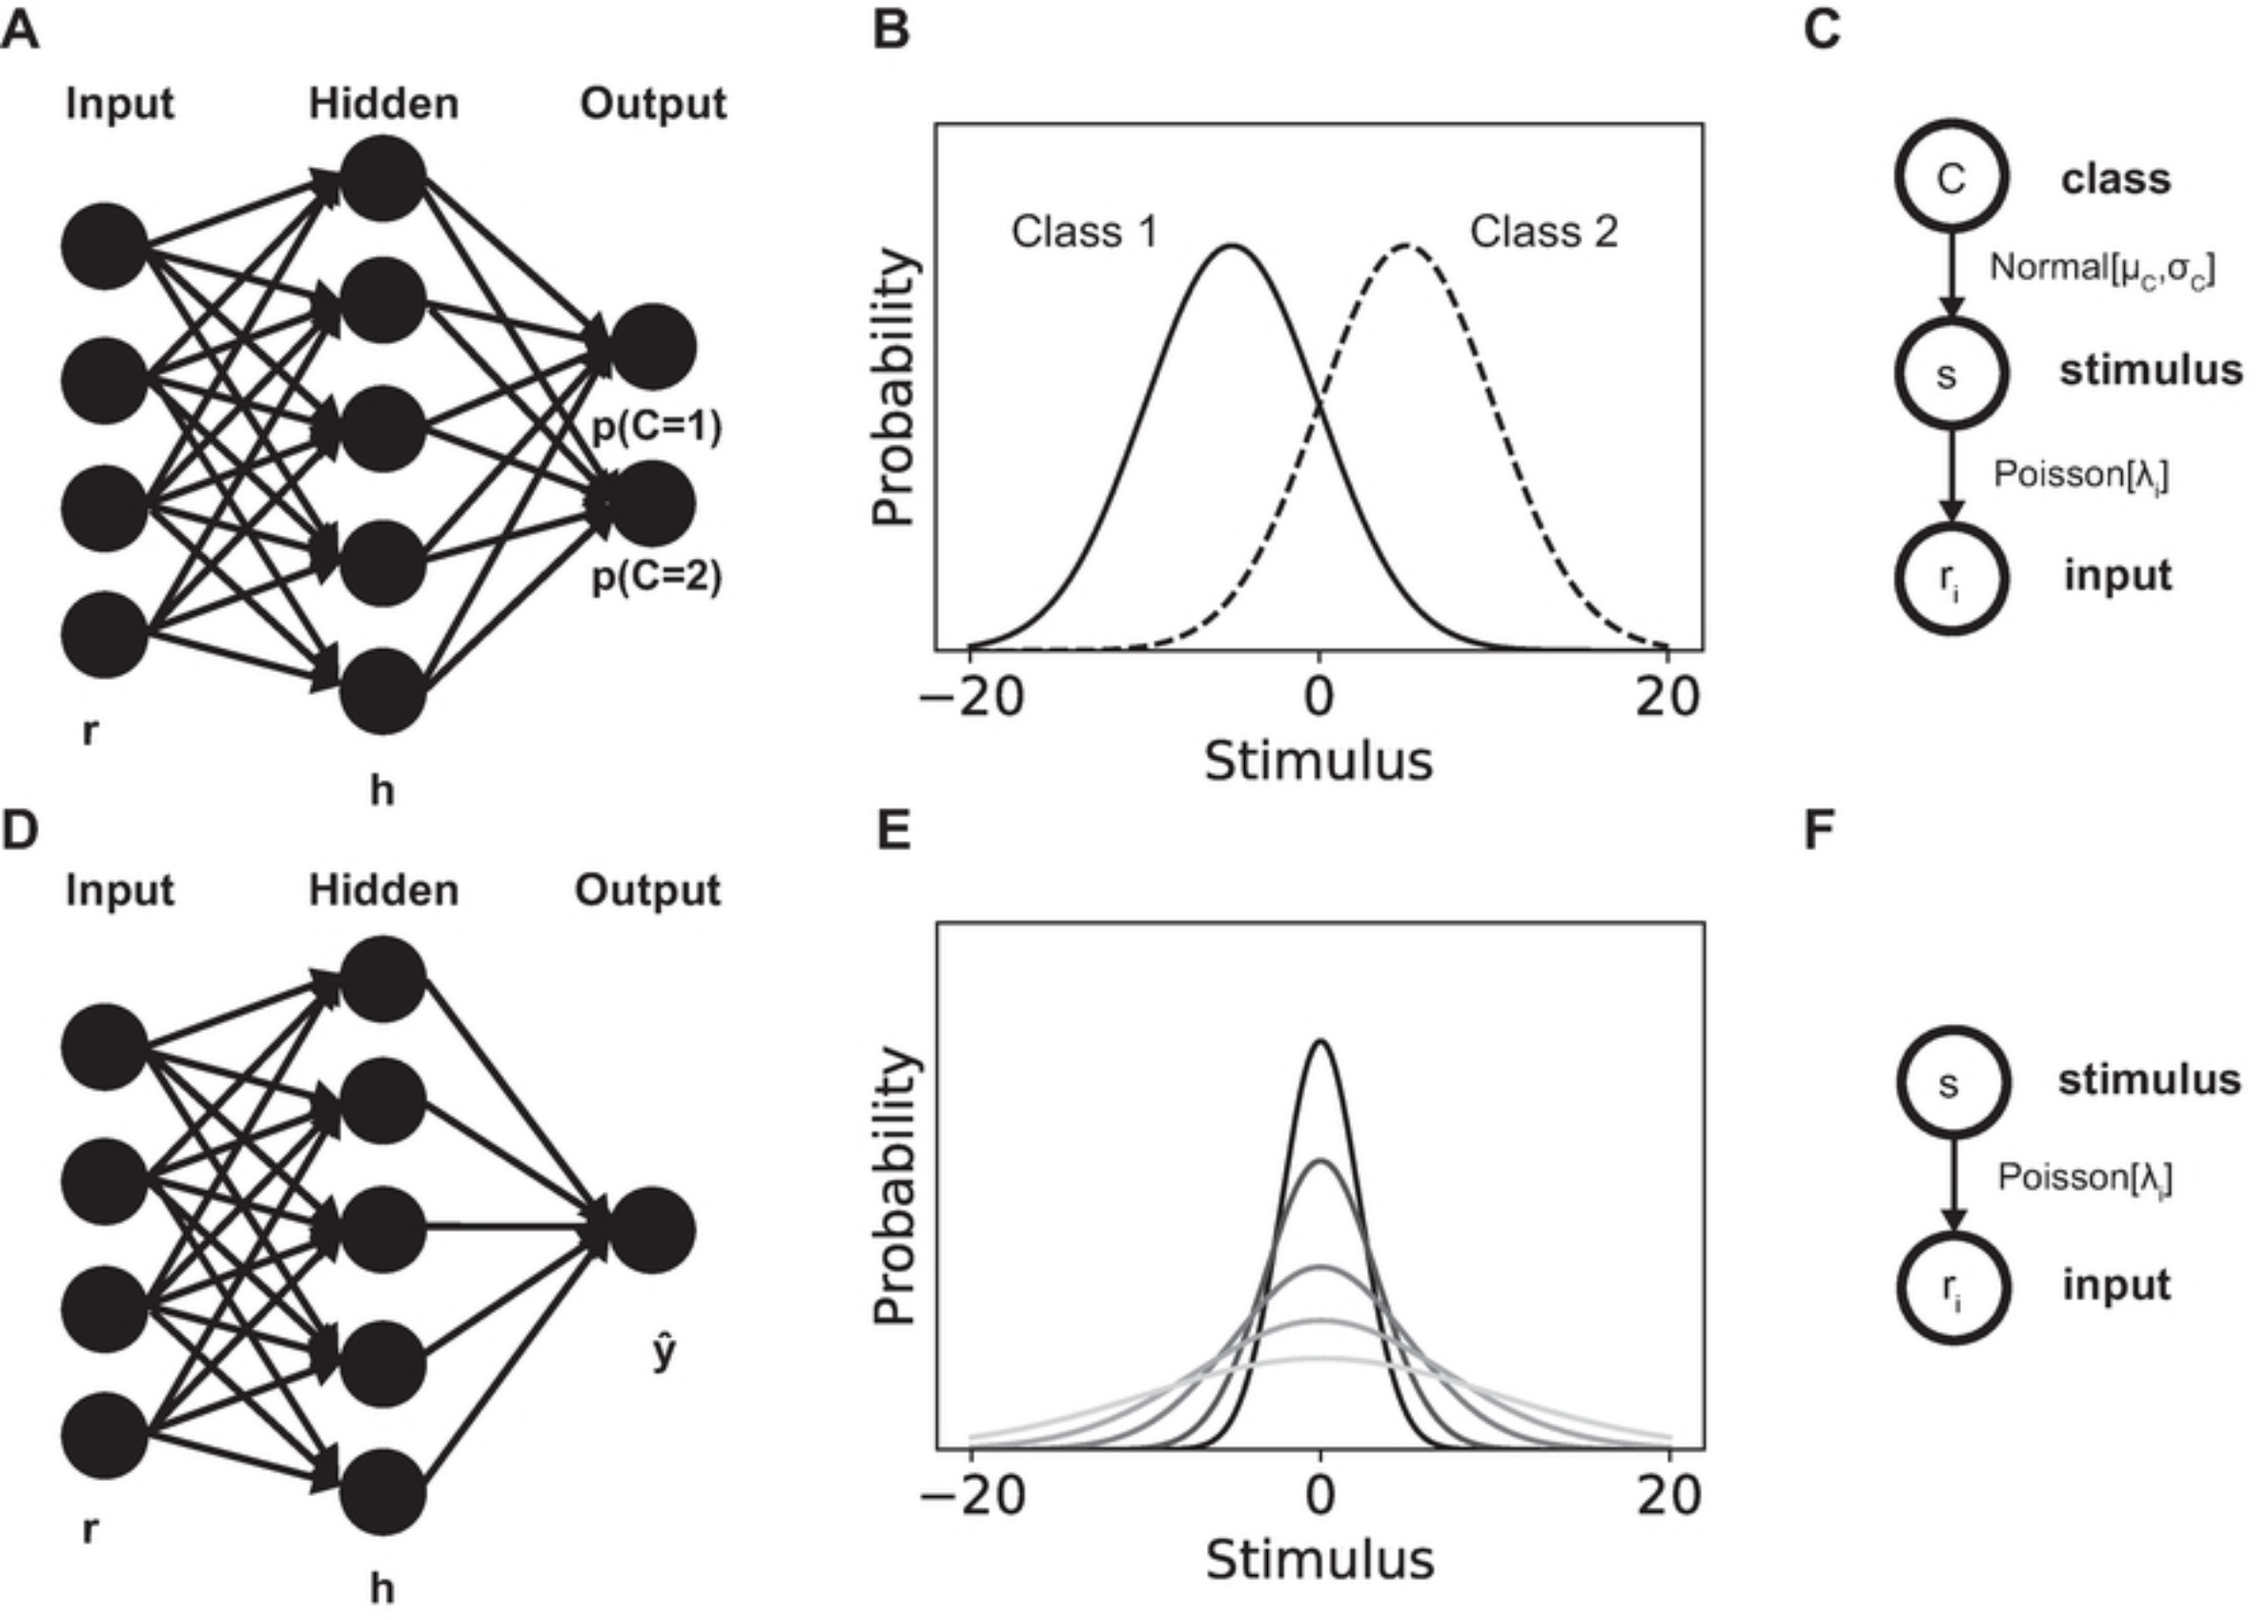
A
Input
Hidden
Output
p(C=1)
p(C=2)
r
h
B
Class 1
Class 2
Probability
−20
0
20
Stimulus
C
C
class
Normal[μc,σc]
s
stimulus
Poisson[λi]
ri
input
D
Input
Hidden
Output
ŷ
r
h
E
Probability
−20
0
20
Stimulus
F
s
stimulus
Poisson[λi]
ri
input

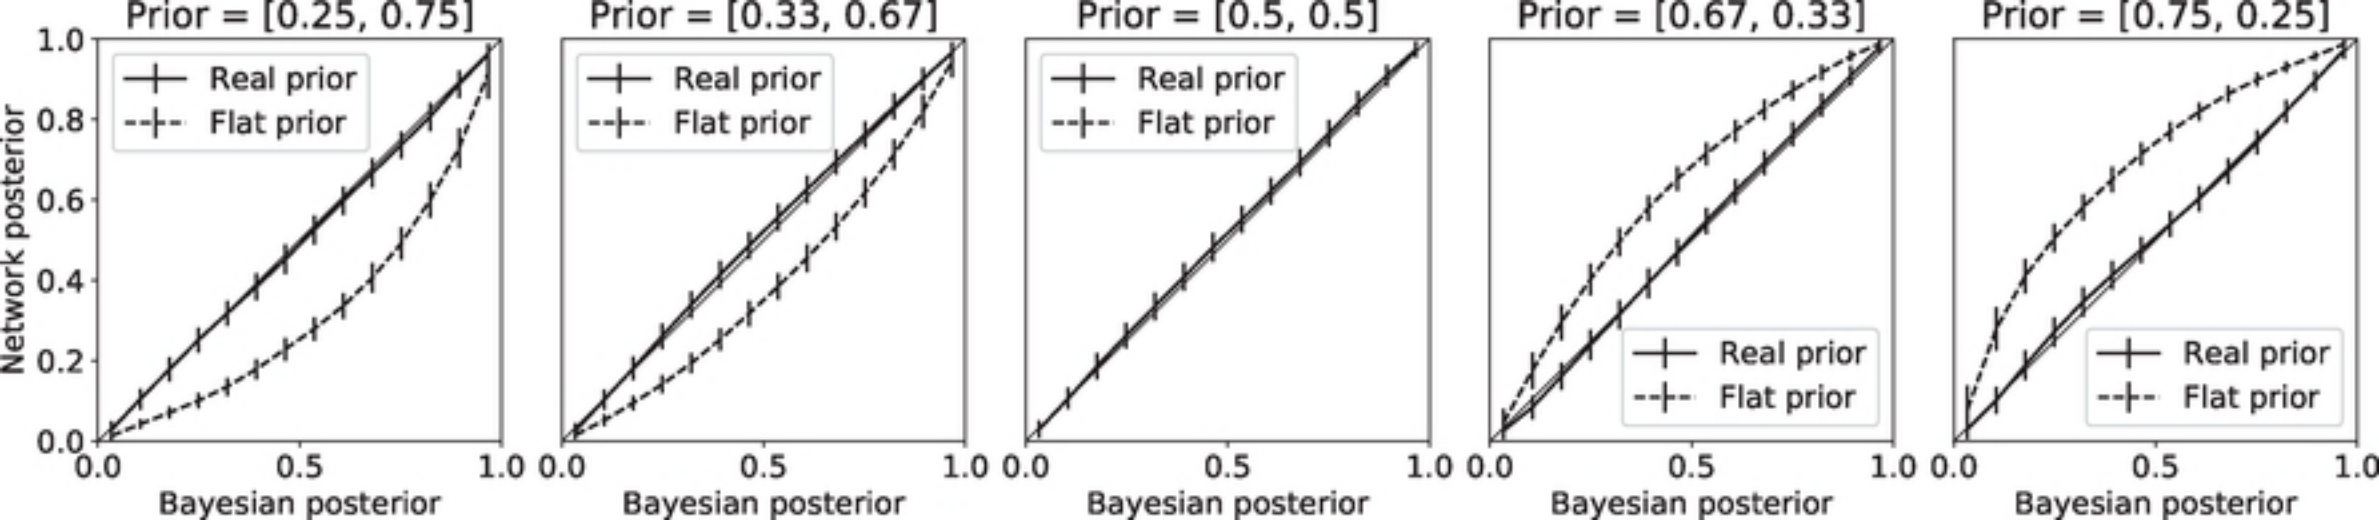
Prior = [0.25, 0.75]
Prior = [0.33, 0.67]
Prior = [0.5, 0.5]
Prior = [0.67, 0.33]
Prior = [0.75, 0.25]
Real prior
Flat prior
Network posterior
Bayesian posterior
0.0
0.2
0.4
0.6
0.8
1.0
0.5

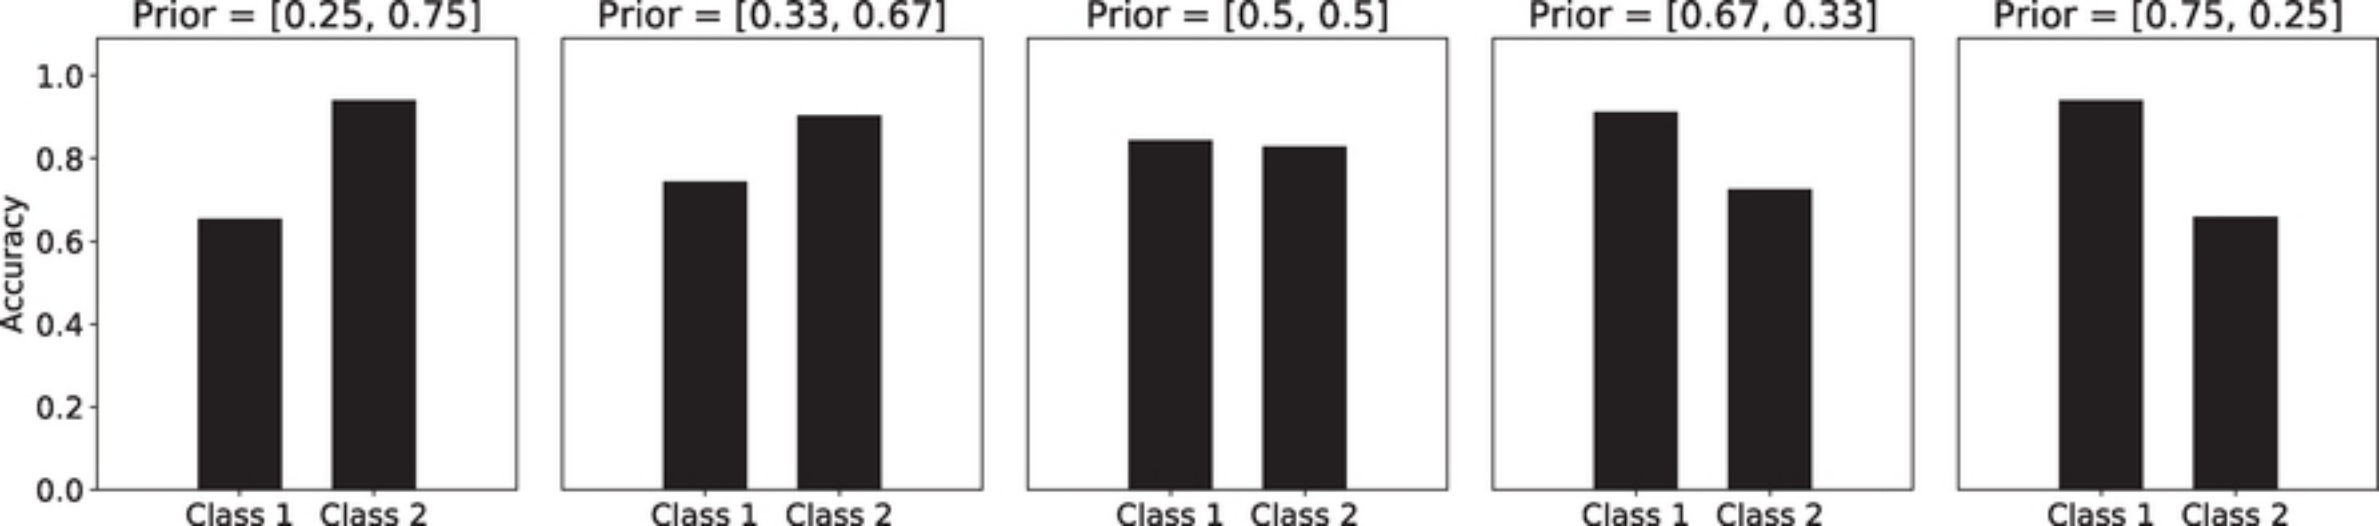
Prior = [0.25, 0.75]
Prior = [0.33, 0.67]
Prior = [0.5, 0.5]
Prior = [0.67, 0.33]
Prior = [0.75, 0.25]
Accuracy
1.0
0.8
0.6
0.4
0.2
0.0
Class 1
Class 2

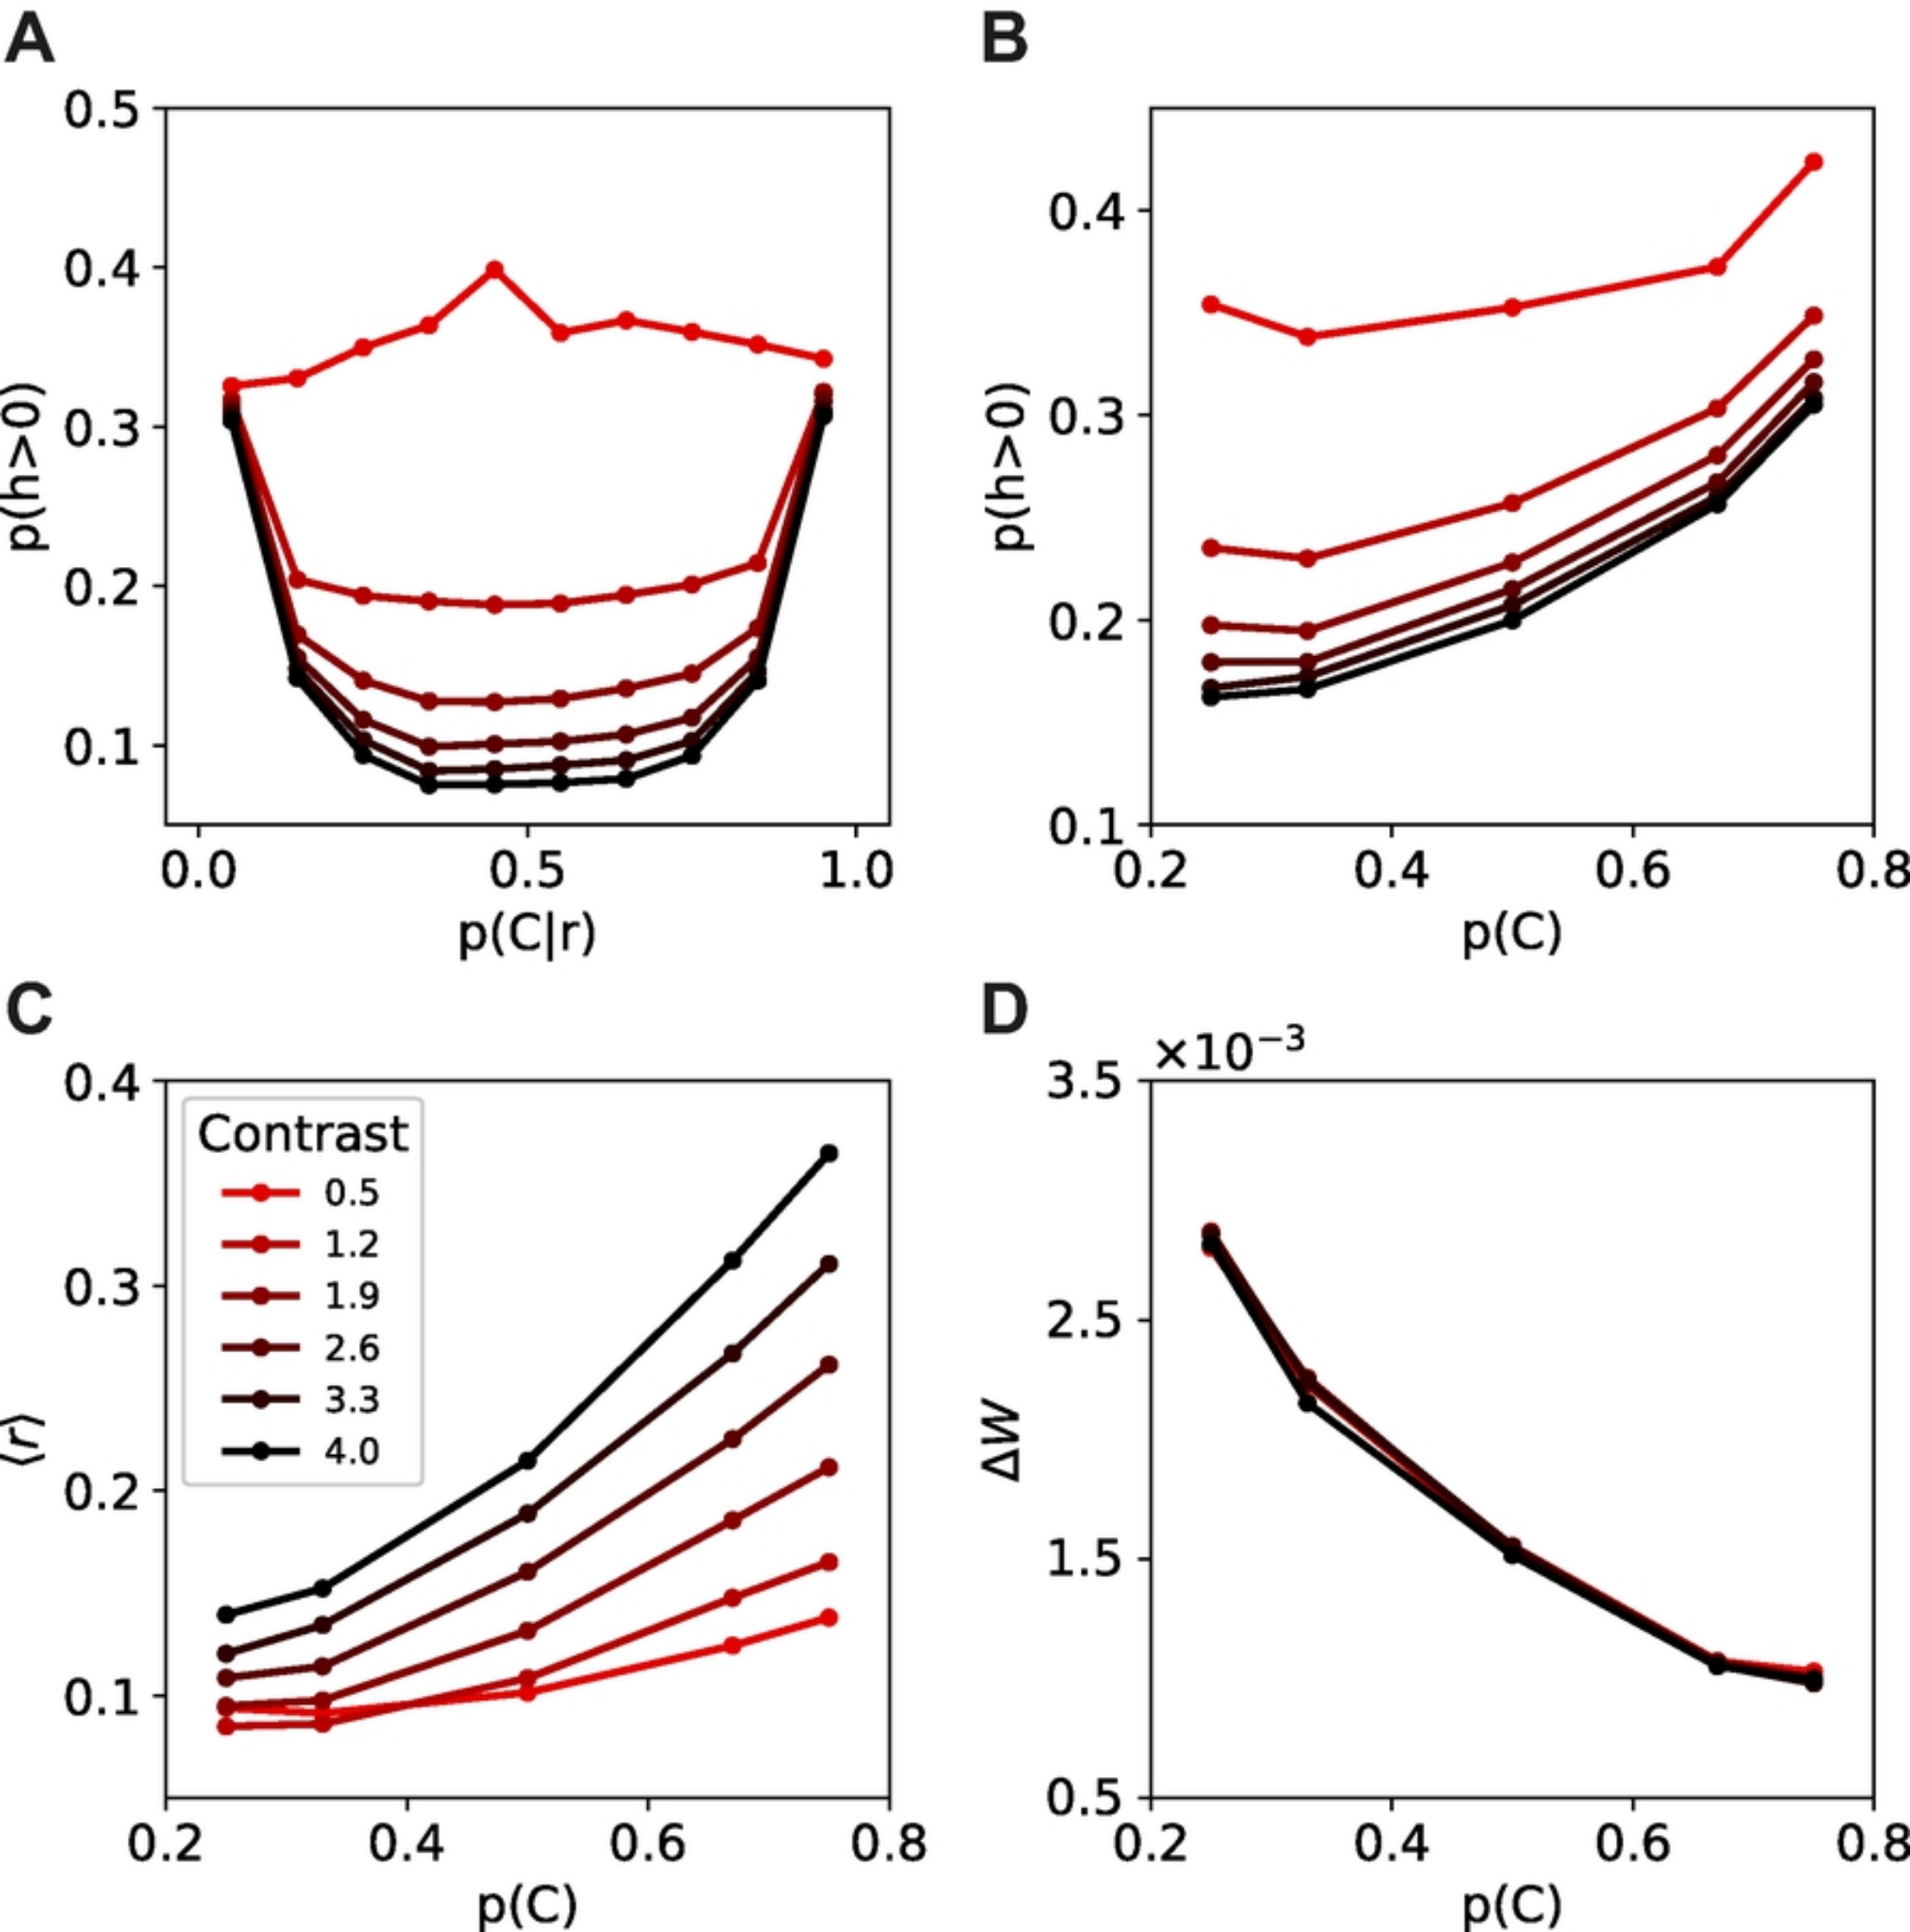
A
B
C
D
p(h>0)
p(C|r)
p(C)
⟨r⟩
ΔW
×10⁻³
Contrast
0.5
1.2
1.9
2.6
3.3
4.0

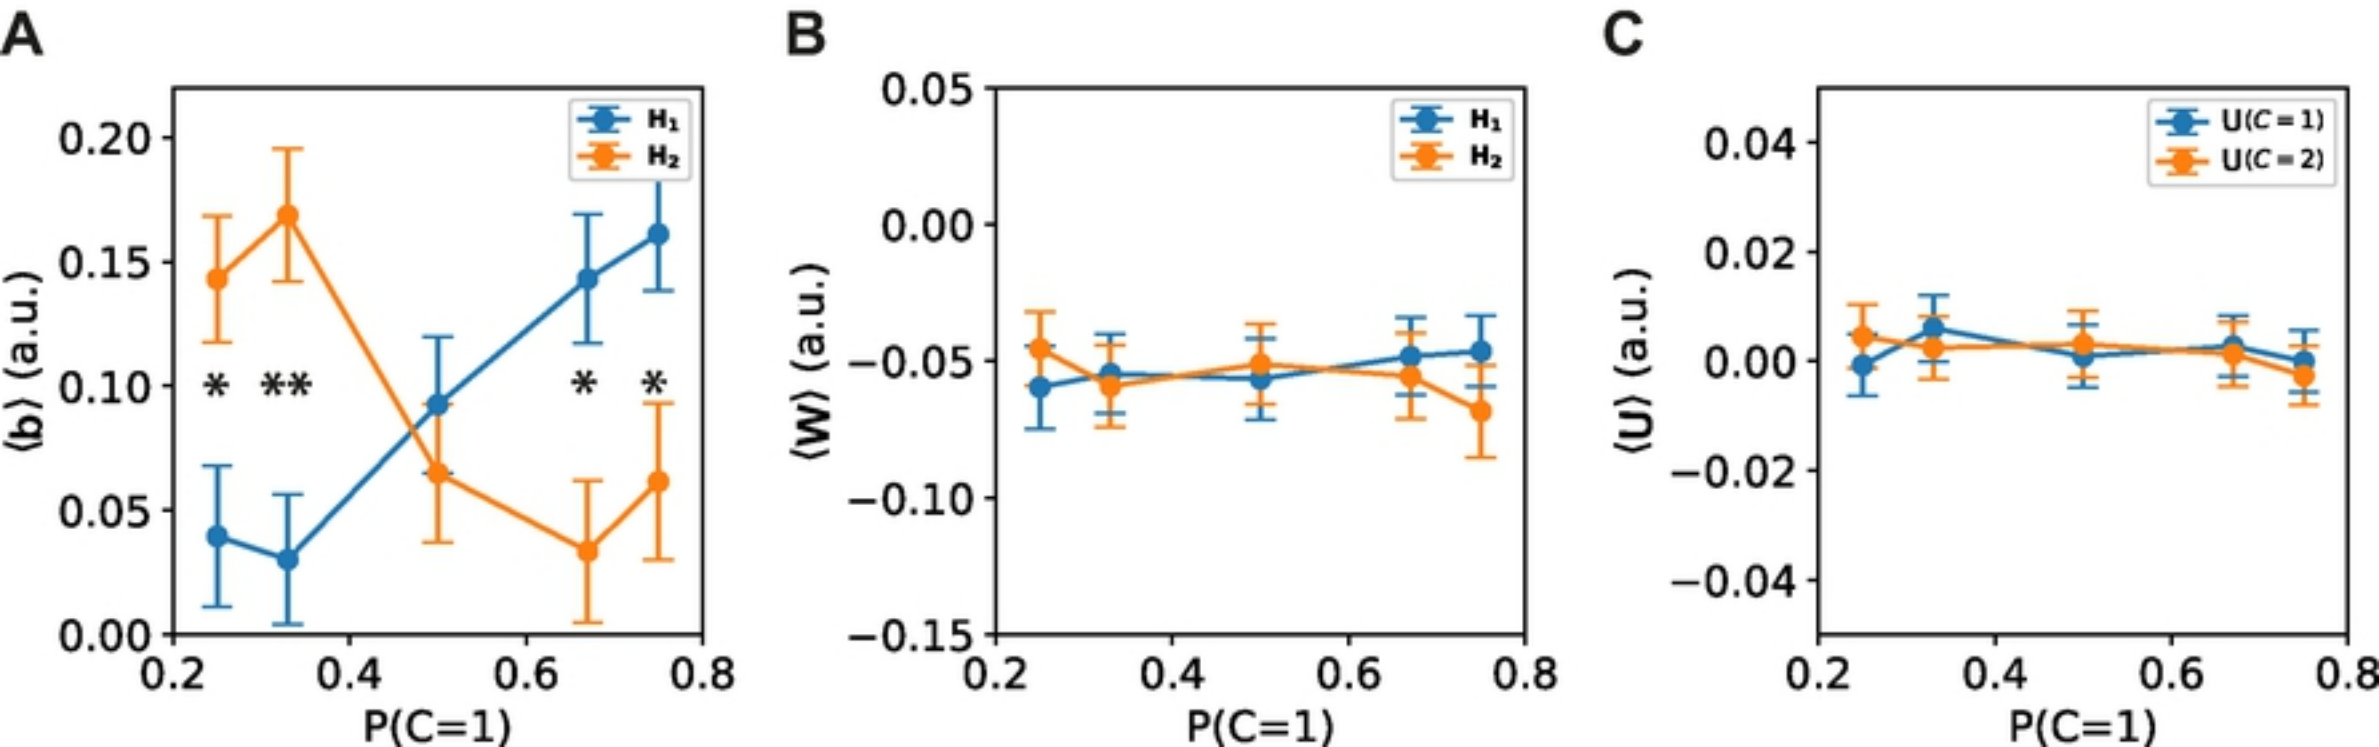
A
B
C
H₁
H₂
U(C = 1)
U(C = 2)
⟨b⟩ (a.u.)
⟨W⟩ (a.u.)
⟨U⟩ (a.u.)
P(C=1)
0.00
0.05
0.10
0.15
0.20
−0.15
−0.10
−0.05
0.00
0.05
−0.04
−0.02
0.00
0.02
0.04
0.2
0.4
0.6
0.8
*
**
*
*

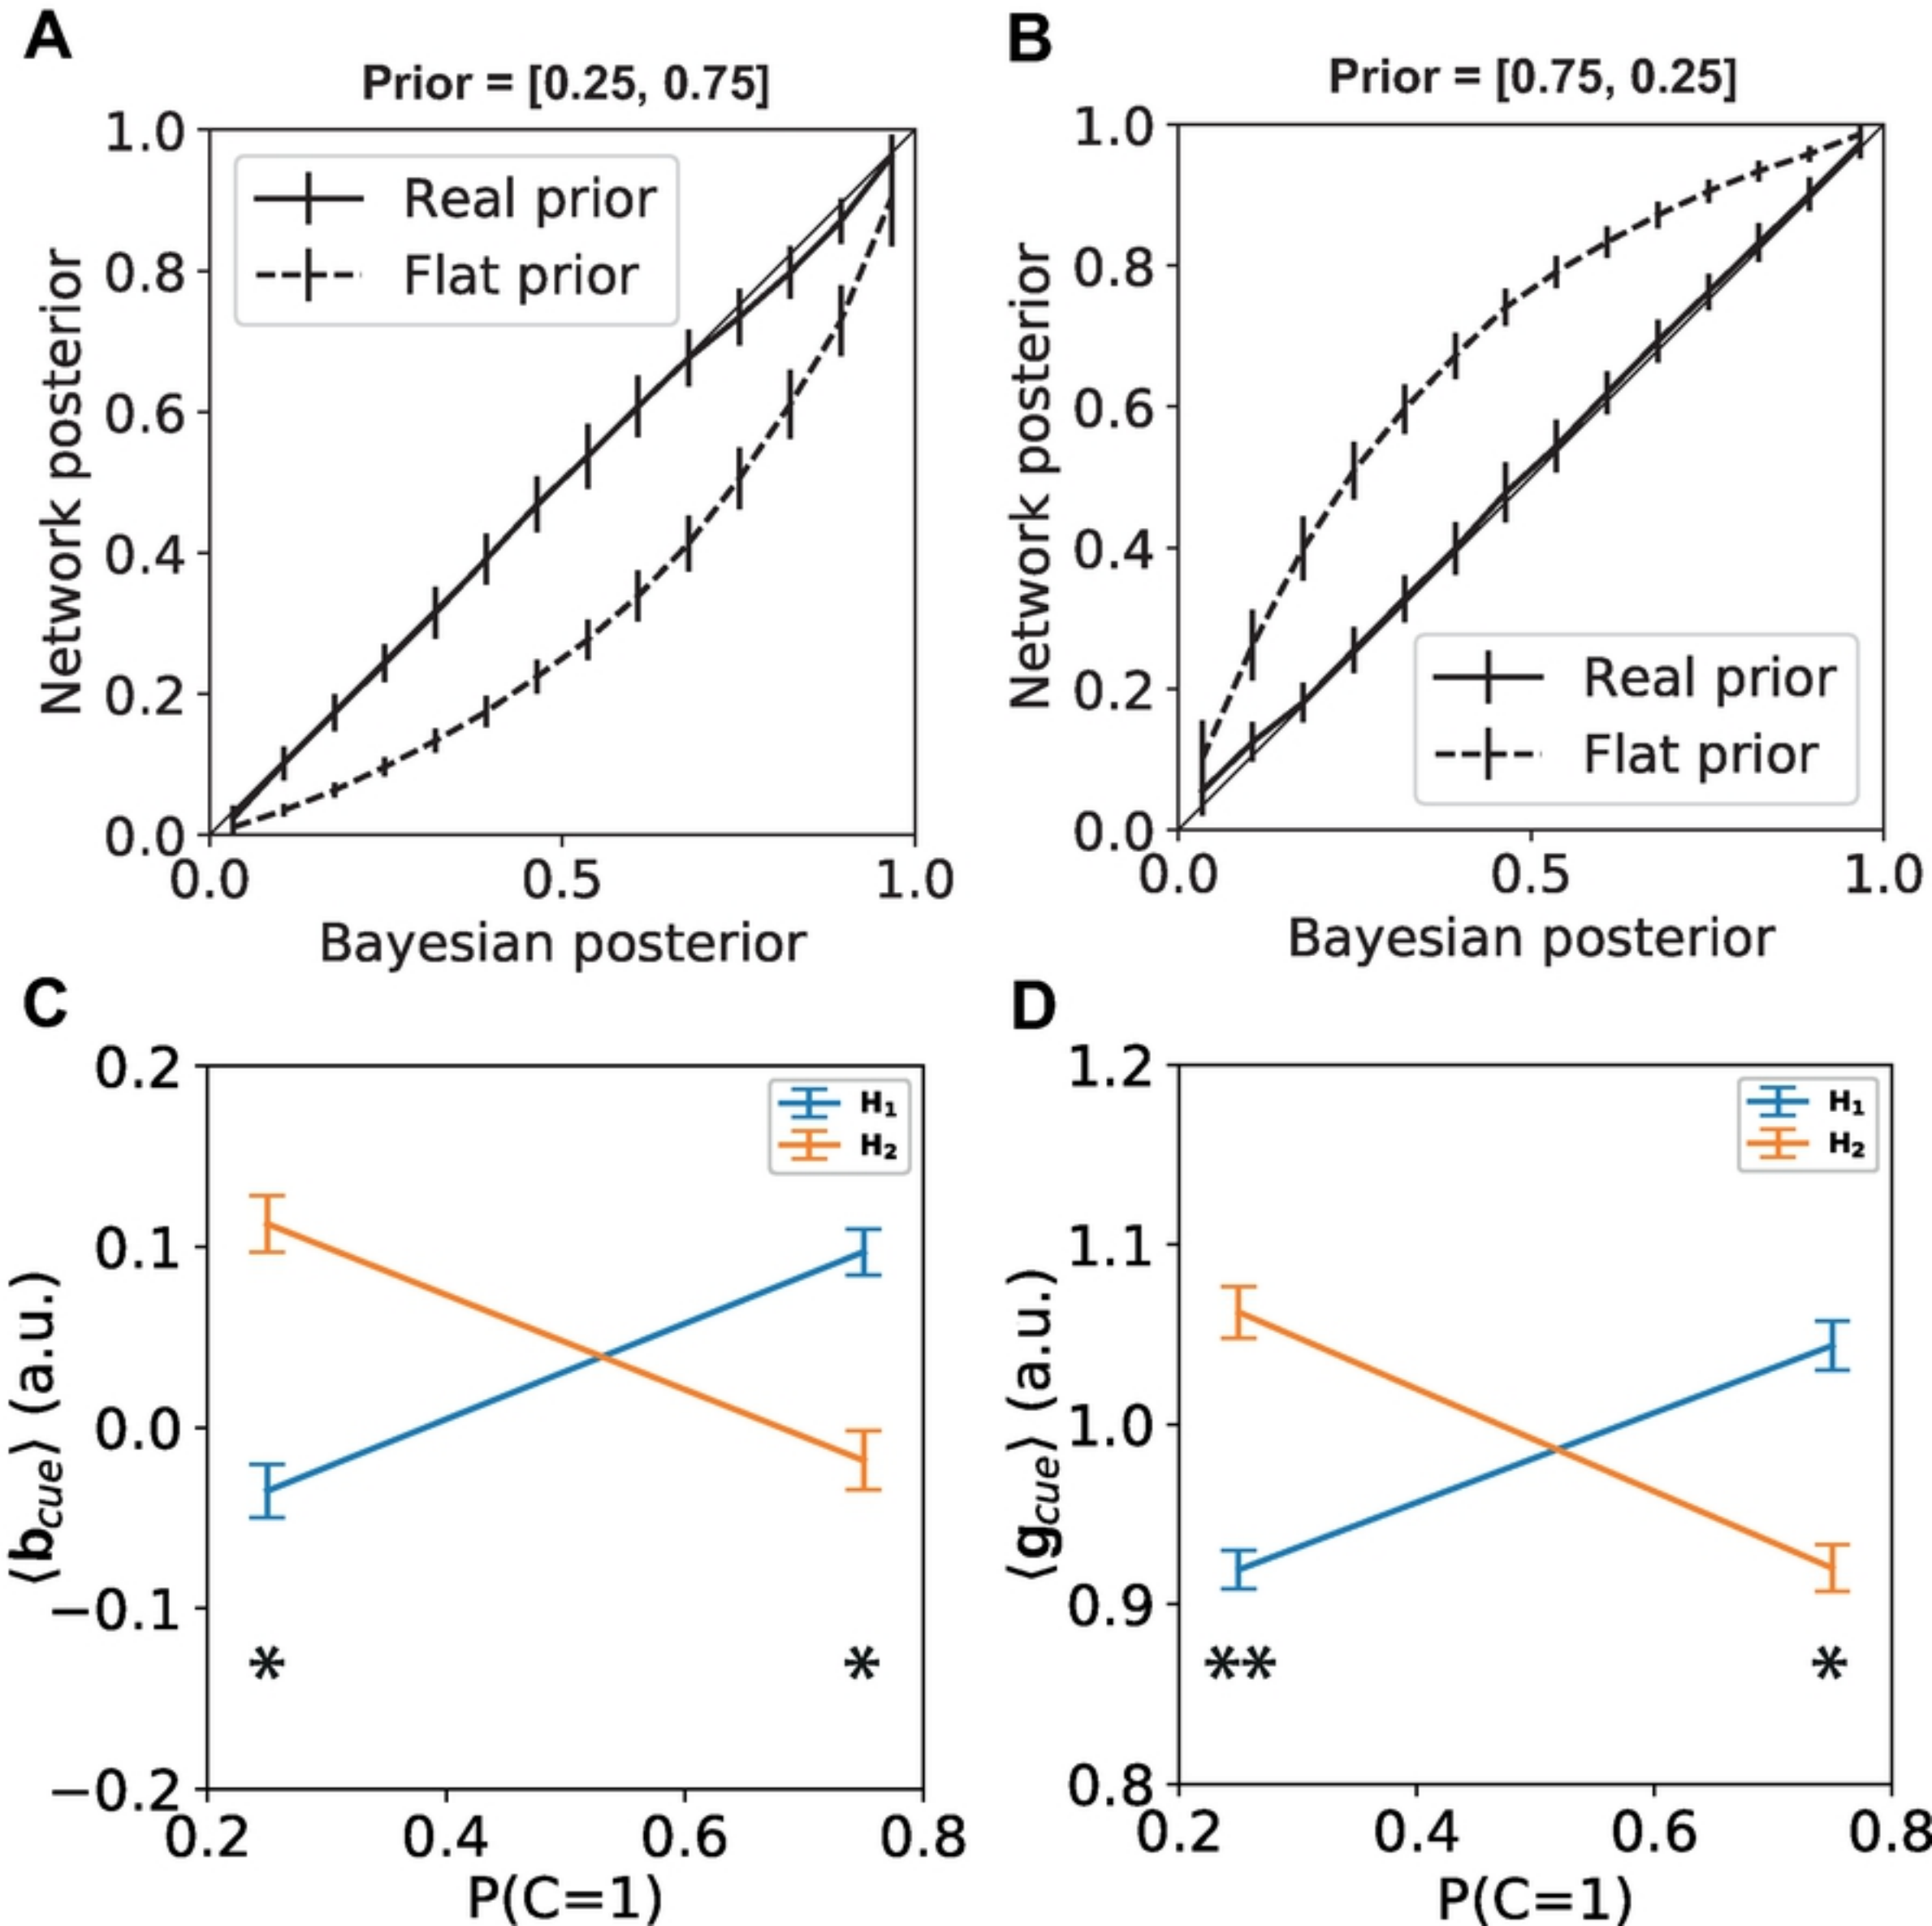
A
Prior = [0.25, 0.75]
Real prior
Flat prior
Network posterior
Bayesian posterior
B
Prior = [0.75, 0.25]
Real prior
Flat prior
Network posterior
Bayesian posterior
C
H1
H2
⟨b_cue⟩ (a.u.)
P(C=1)
D
H1
H2
⟨g_cue⟩ (a.u.)
P(C=1)

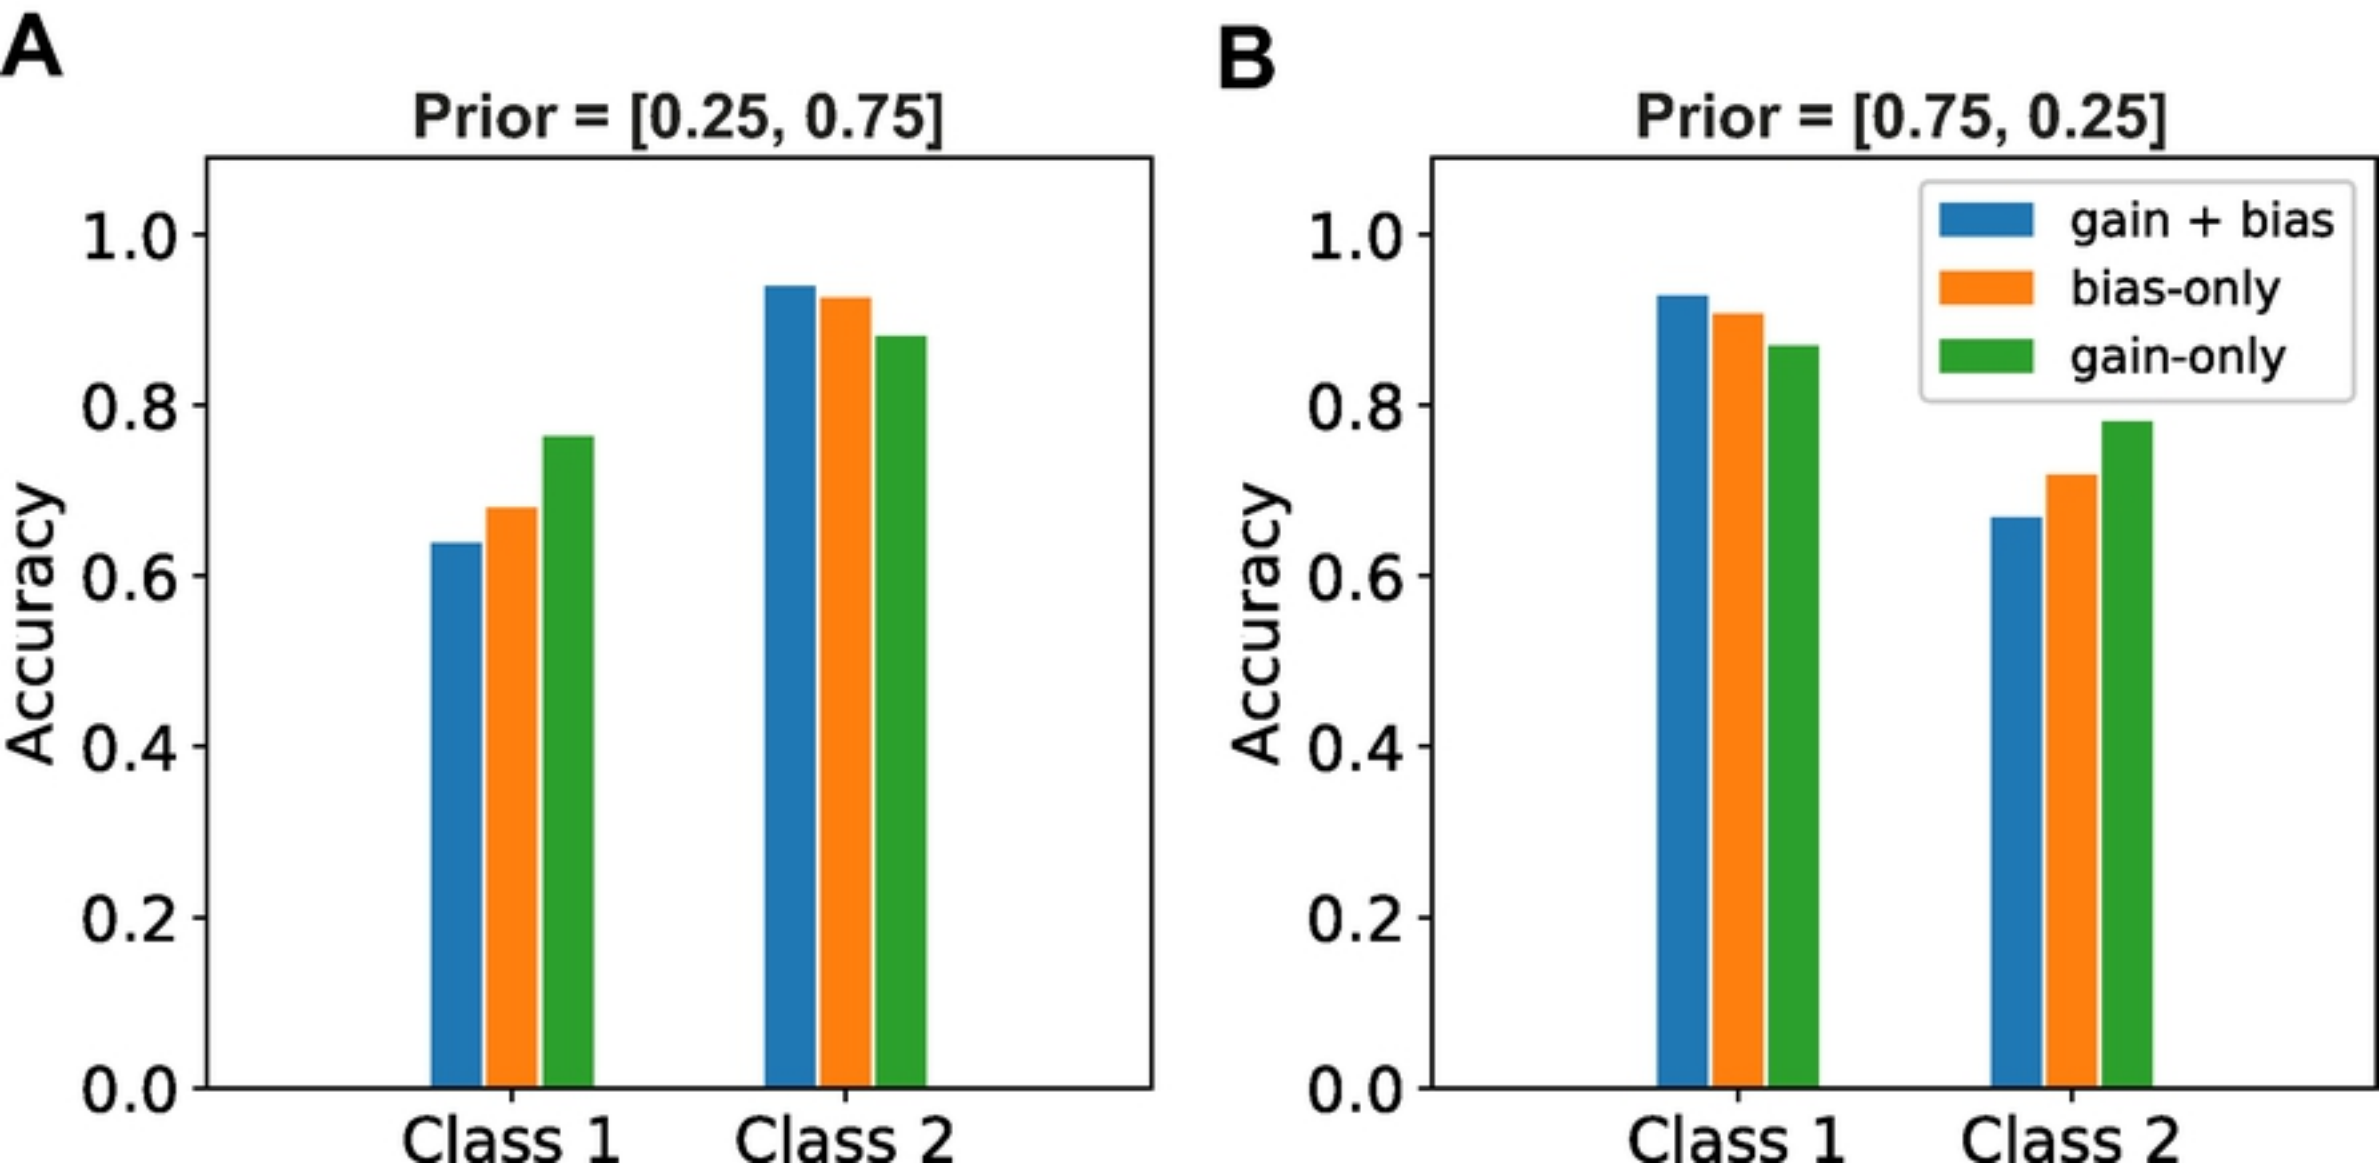
A
Prior = [0.25, 0.75]
Accuracy
1.0
0.8
0.6
0.4
0.2
0.0
Class 1
Class 2
B
Prior = [0.75, 0.25]
gain + bias
bias-only
gain-only
Accuracy
1.0
0.8
0.6
0.4
0.2
0.0
Class 1
Class 2

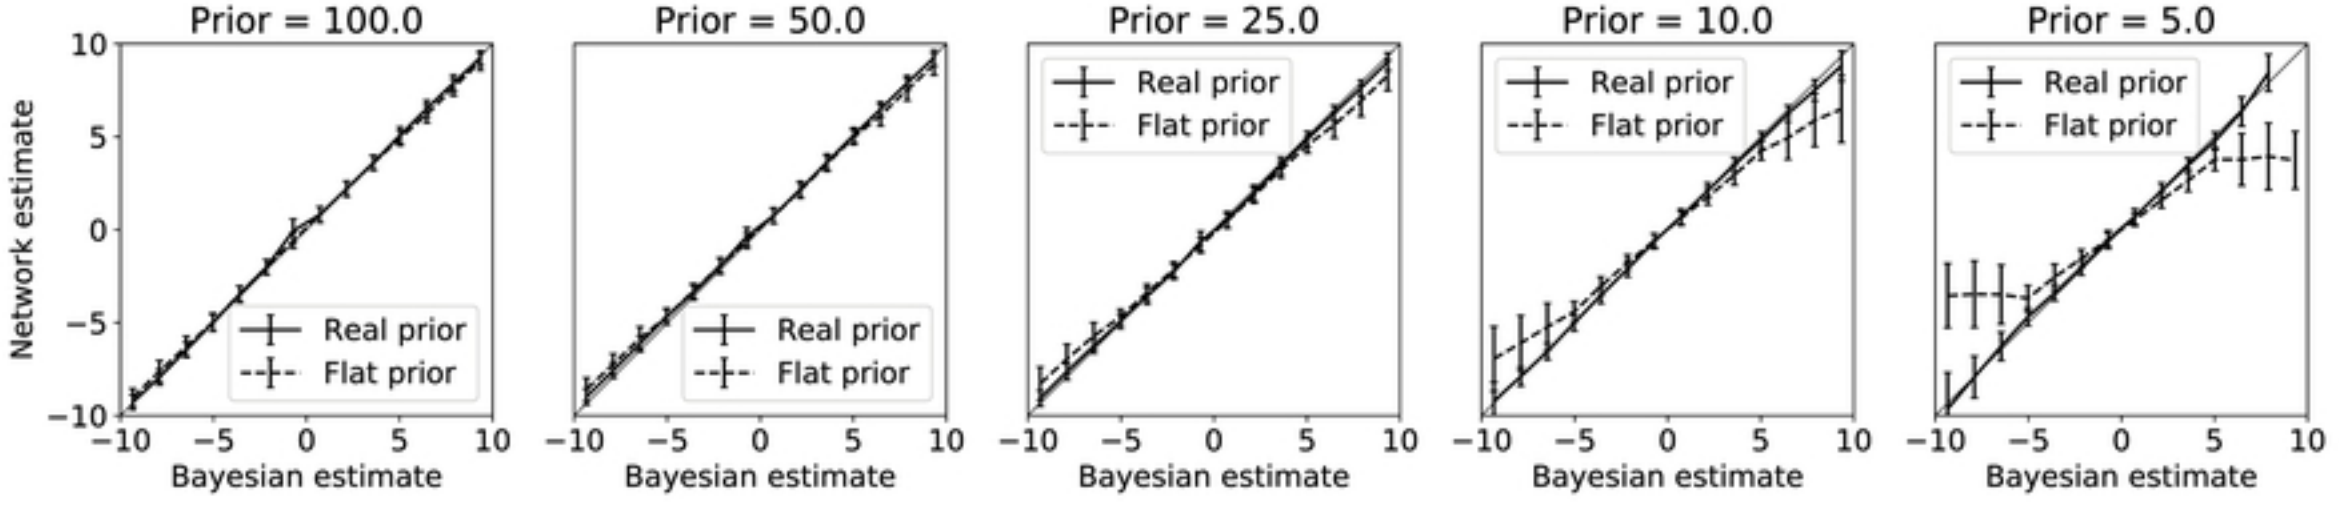
Prior = 100.0
Prior = 50.0
Prior = 25.0
Prior = 10.0
Prior = 5.0
Network estimate
Bayesian estimate
Real prior
Flat prior
10
5
0
−5
−10

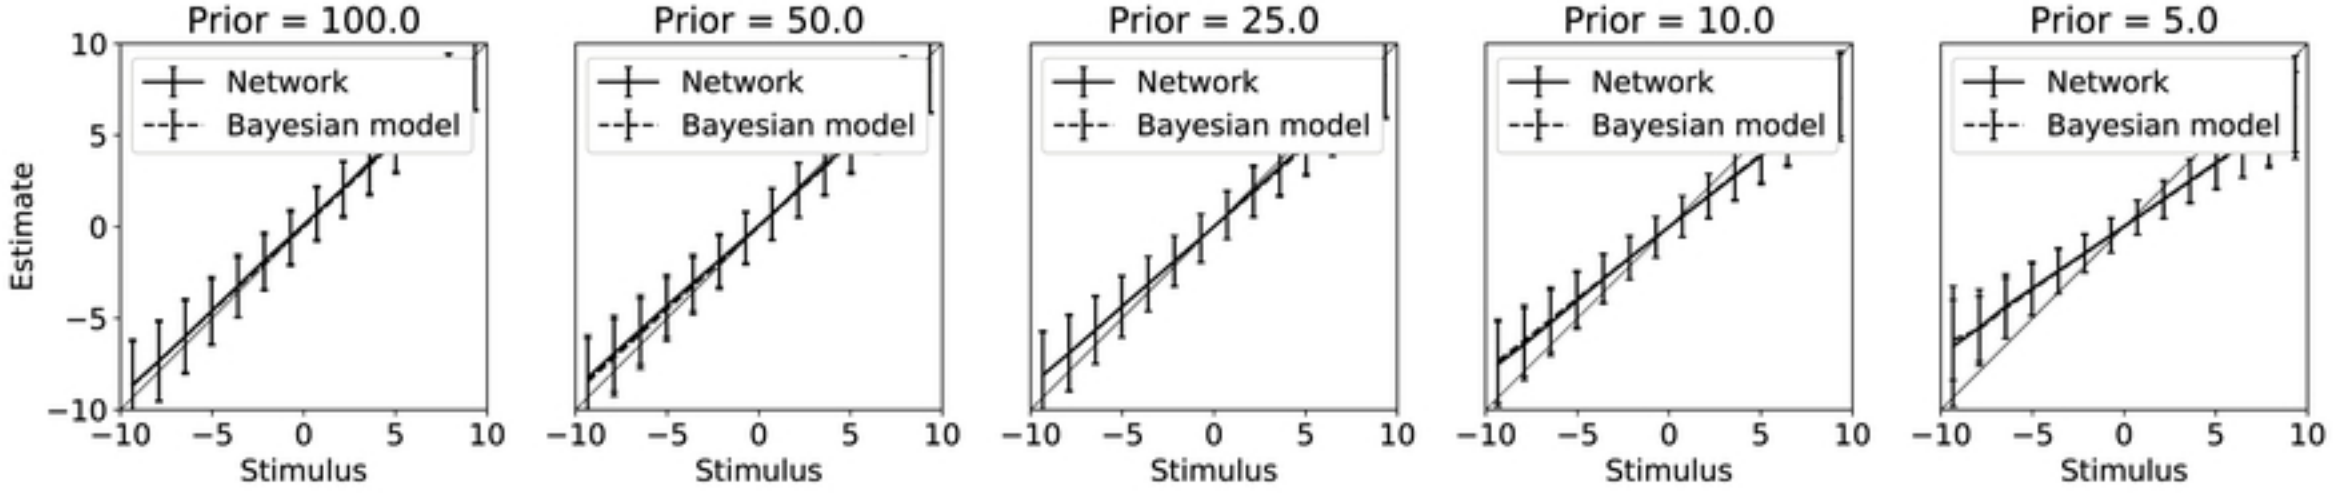
Prior = 100.0
Prior = 50.0
Prior = 25.0
Prior = 10.0
Prior = 5.0
Network
Bayesian model
Estimate
Stimulus
10
5
0
−5
−10

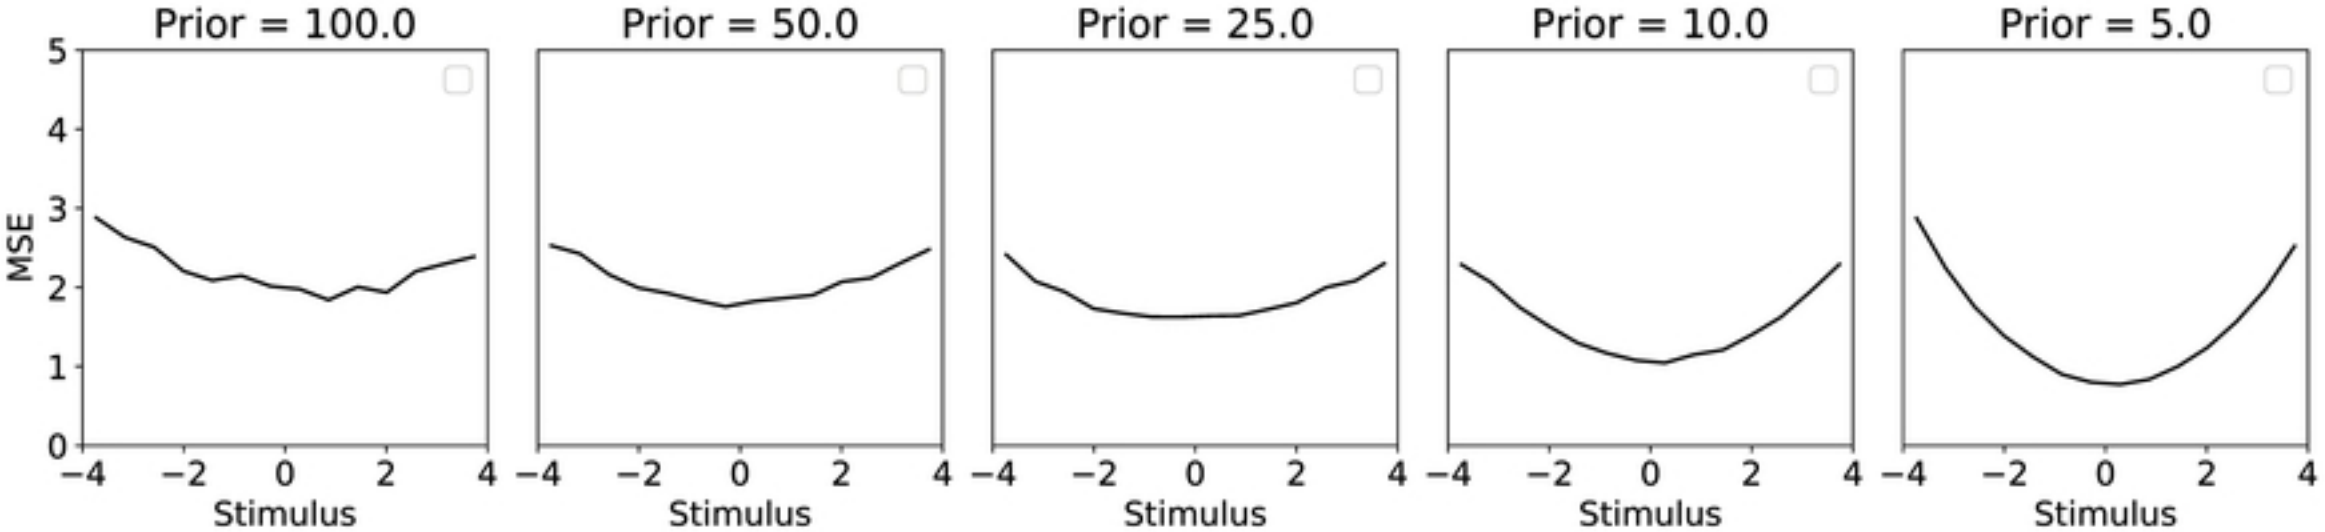
Prior = 100.0
Prior = 50.0
Prior = 25.0
Prior = 10.0
Prior = 5.0
MSE
0
1
2
3
4
5
−4
−2
0
2
4
Stimulus

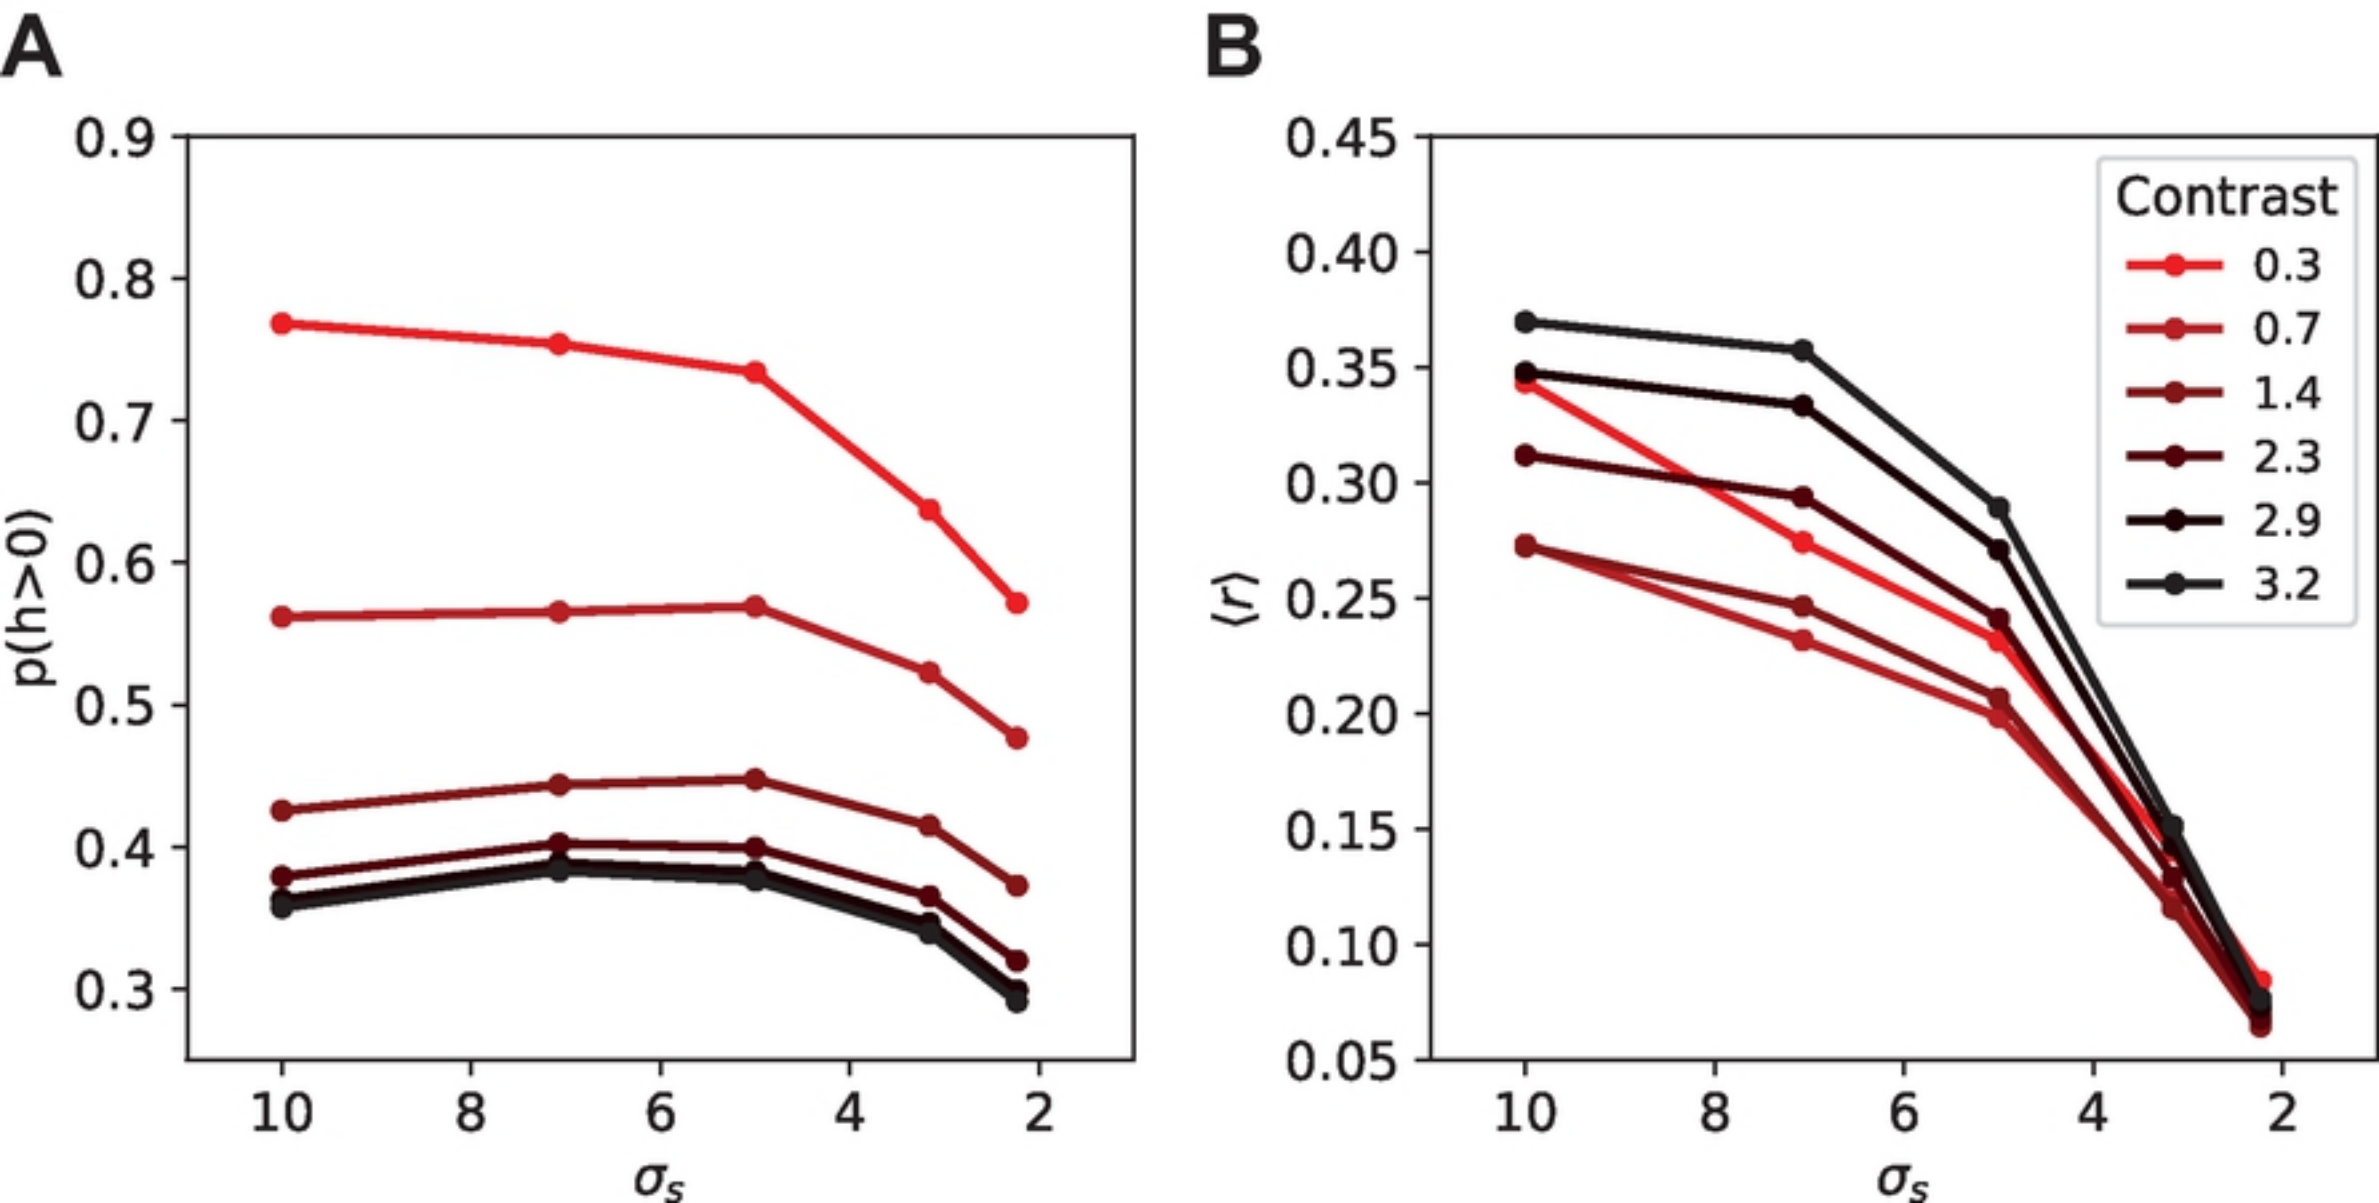
A
B
0.9
0.8
0.7
0.6
0.5
0.4
0.3
p(h>0)
10
8
6
4
2
σ_s
0.45
0.40
0.35
0.30
0.25
0.20
0.15
0.10
0.05
⟨r⟩
Contrast
0.3
0.7
1.4
2.3
2.9
3.2

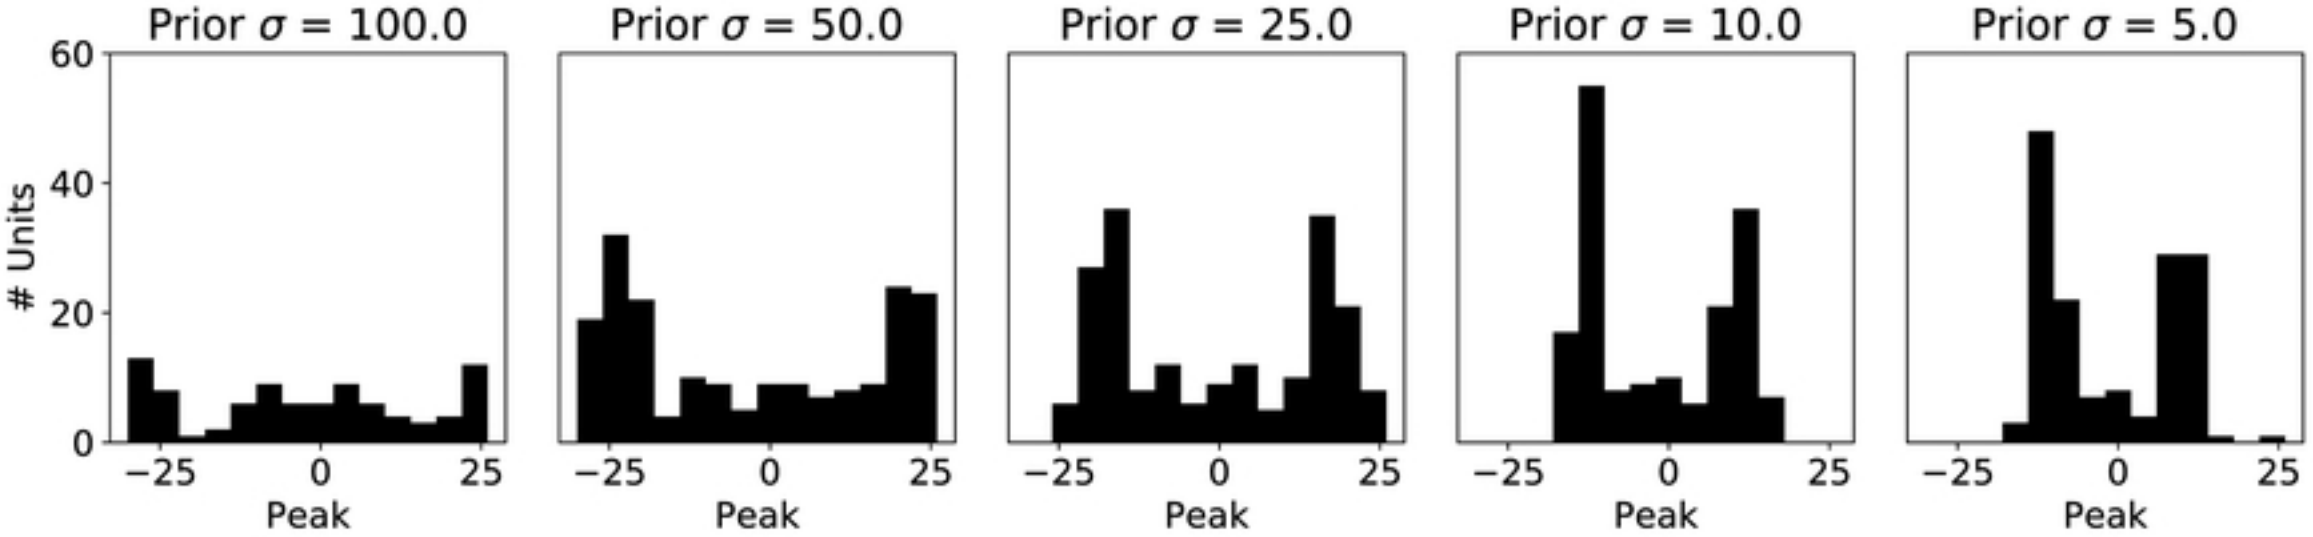
Prior σ = 100.0
Prior σ = 50.0
Prior σ = 25.0
Prior σ = 10.0
Prior σ = 5.0
Units
60
40
20
0
−25
0
25
Peak

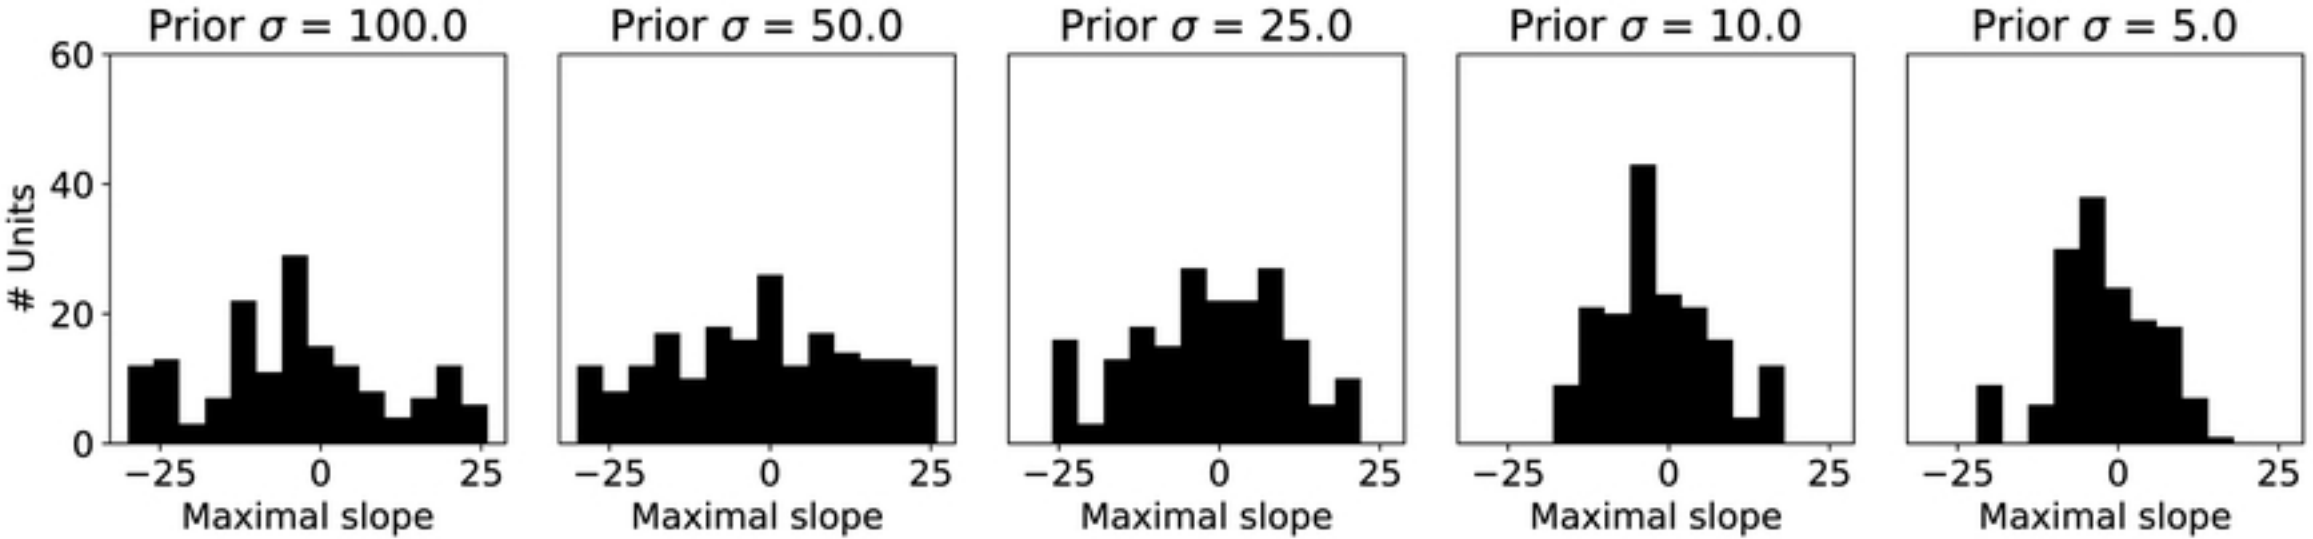
Prior σ = 100.0
Prior σ = 50.0
Prior σ = 25.0
Prior σ = 10.0
Prior σ = 5.0
Units
60
40
20
0
−25
0
25
Maximal slope

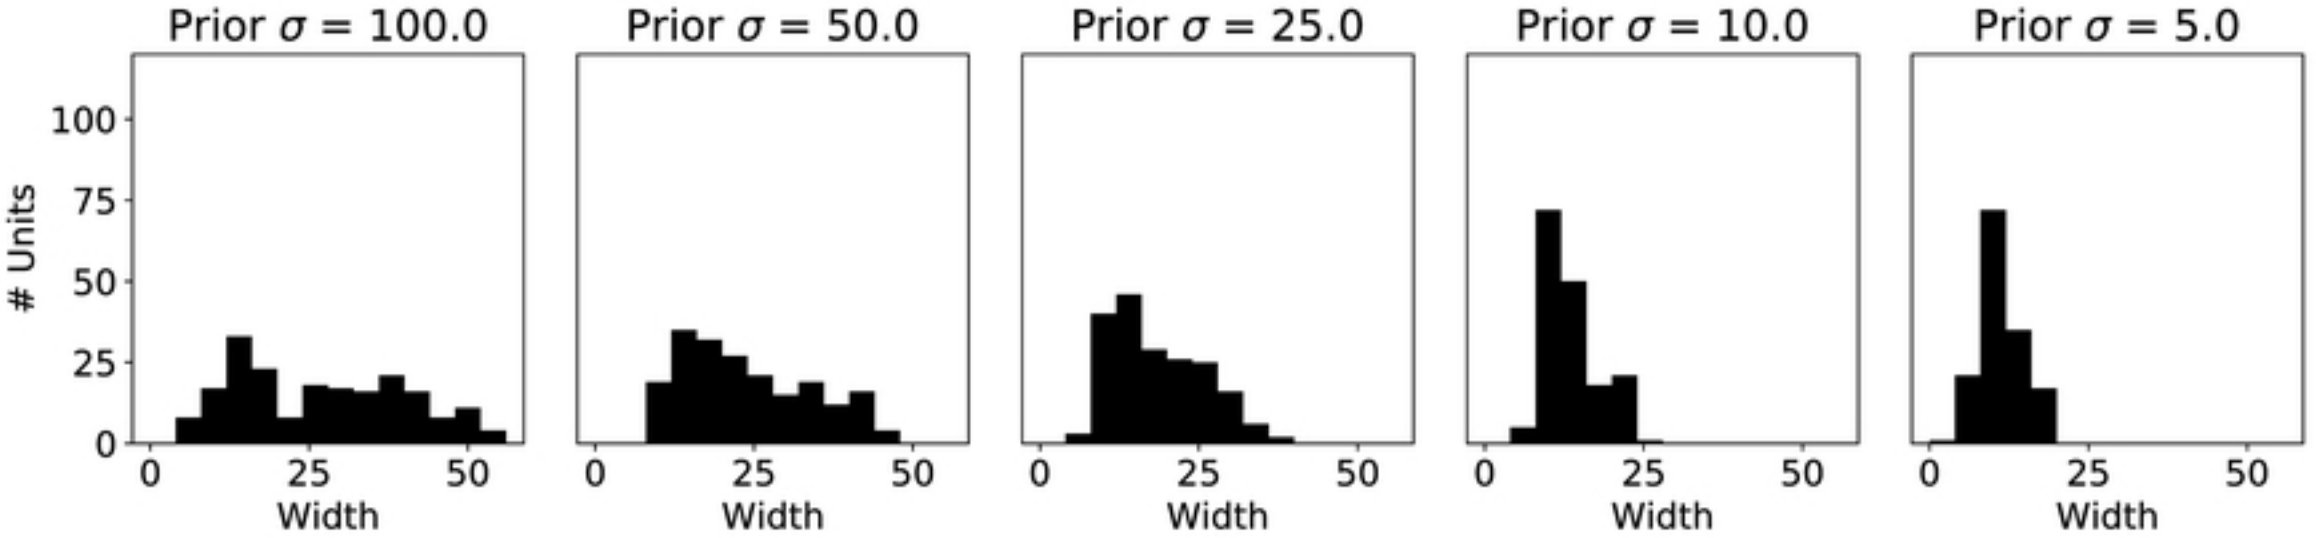
Prior σ = 100.0
Prior σ = 50.0
Prior σ = 25.0
Prior σ = 10.0
Prior σ = 5.0
Units
100
75
50
25
0
0
25
50
Width

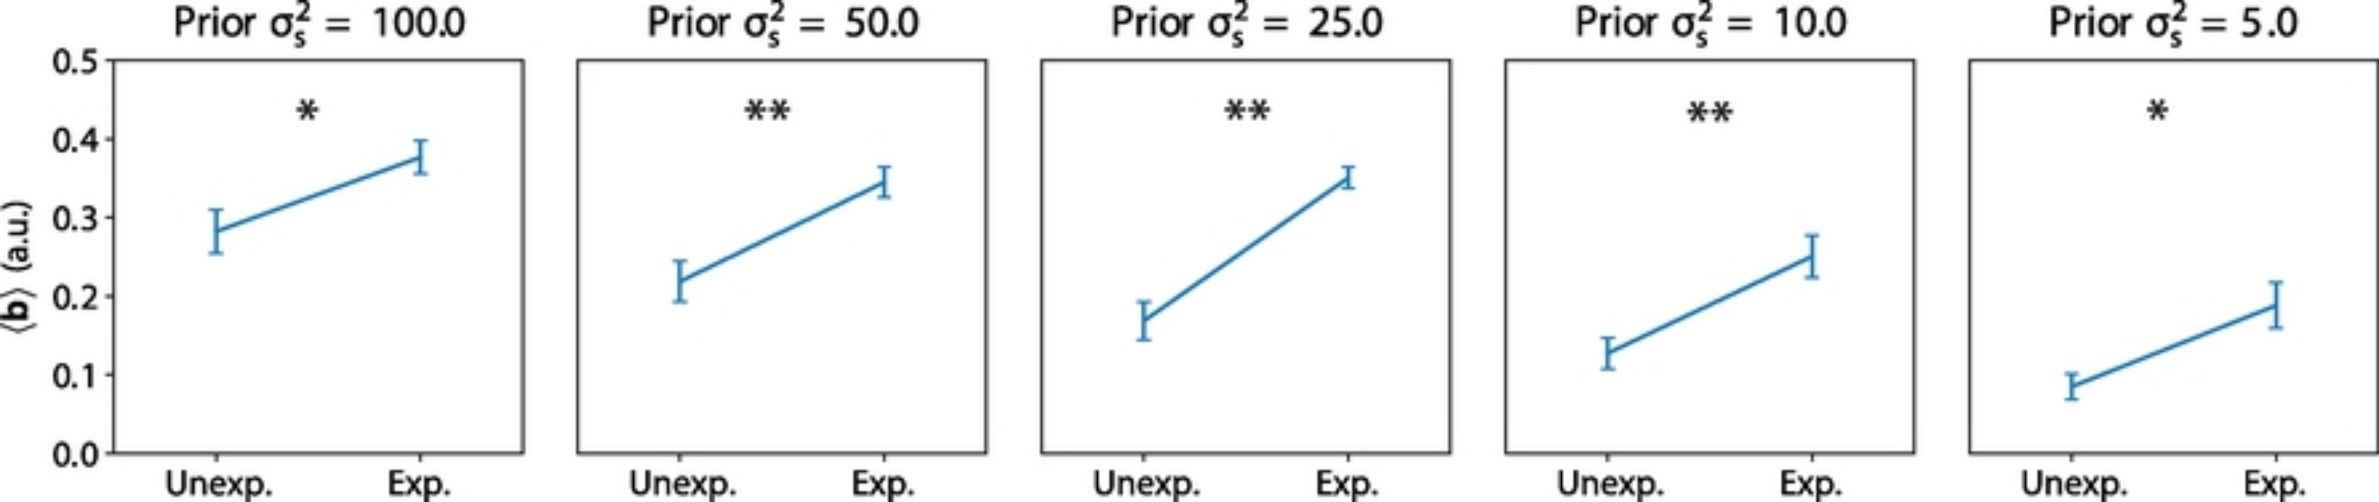
Prior σ_s^2 = 100.0
Prior σ_s^2 = 50.0
Prior σ_s^2 = 25.0
Prior σ_s^2 = 10.0
Prior σ_s^2 = 5.0
⟨b⟩ (a.u.)
0.0
0.1
0.2
0.3
0.4
0.5
*
**
**
**
*
Unexp.
Exp.