
A sampling-based circuit for optimal decision making

Camille E. Rullán Buxó
Center for Neural Science
New York University
New York, NY 10003
ch2880@nyu.edu

Cristina Savin
Center for Neural Science
Center for Data Science
New York University
New York, NY 10003
csavin@nyu.edu

Abstract

Many features of human and animal behavior can be understood in the framework of Bayesian inference and optimal decision making, but the biological substrate of such processes is not fully understood. Neural sampling provides a flexible code for probabilistic inference in high dimensions and explains key features of sensory responses under experimental manipulations of uncertainty. However, since it encodes uncertainty implicitly, across time and neurons, it remains unclear how such representations can be used for decision making. Here we propose a spiking network model that maps neural samples of a task-specific marginal distribution into an instantaneous representation of uncertainty via a procedure inspired by online kernel density estimation, so that its output can be readily used for decision making. Our model is consistent with experimental results at the level of single neurons and populations, and makes predictions for how neural responses and decisions could be modulated by uncertainty and prior biases. More generally, our work brings together conflicting perspectives on probabilistic brain computation.

One of the central questions of perception is how organisms reliably estimate hidden or abstract quantities of interest using noisy and ambiguous sensory information. Almost equally important is representing the reliability of these estimates, especially in complex environments and situations of risk, where the uncertainty associated with a choice may radically change the optimal course of action. From basic functions, such as cue combination or motor control, to higher level cognitive tasks, such as decision making and planning, there is substantial evidence that both humans and animals represent and use uncertainty information to guide their actions [1]. Neural correlates have been identified in a number of regions, including the orbitofrontal cortex [2], the cingulate cortex [3], and the lateral intraparietal area (LIP) [4, 5], but the principles behind how uncertainty is represented in neural circuits to support circuit computations remain hotly debated.

Since behavior has been shown to be Bayes-optimal in many situations [6], the problem of perception can be modelled in the framework of Bayesian inference, whereby an observer combines prior information with current observations according to an internal model to arrive at a posterior estimate of a quantity of interest (the ‘latent variable’). Moreover, natural statistics are strongly non-Gaussian and there is evidence that human subjects use varied, non-Gaussian prior representations to support behavior [7, 8], so observers must be able to perform inference flexibly, efficiently and adaptively. This raises two questions: which neural organizations allow for these kinds of computations, and how are the associated neural representations mapped into behaviorally relevant action plans?

Currently there are several theories for how neural activity may represent probability distributions. One large class of models assume that neural responses encode parameters of underlying posterior distributions; this includes probabilistic population codes [9, 10], their predecessors, kernel density estimators and distributional population codes [11, 12], and most recently, distributed distributional codes (DDC) [13]. A second class of models relies on stochasticity in recurrent circuits to approxi-

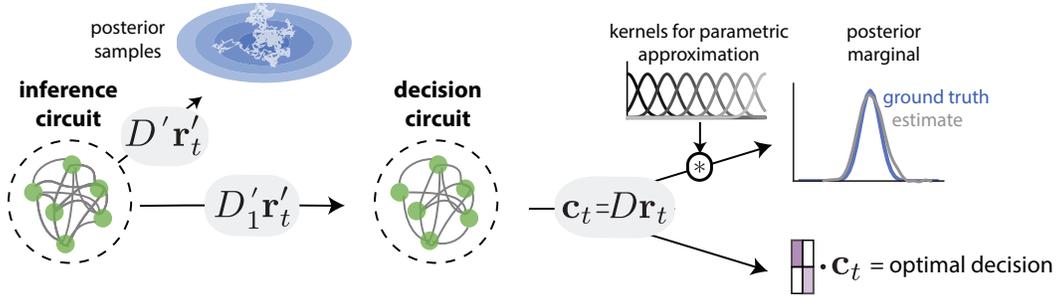


Figure 1: **Schematic of sampling-based optimal decision making in spiking networks.** Approximate inference is performed in a first spiking recurrent network by distributed sampling [22] (example trajectory from joint posterior shown in white). Samples from the task-relevant marginal posterior are read out linearly from this circuit and serve as input to a second network, which integrates them and converts them into a parametric representation. The associated parameters (bank of kernels in gray) are read out linearly from this second circuit and combined with a cost function, reflecting the potentially asymmetric cost of different errors, to generate the final optimal decision.

mately encode probability distributions via sampling [14, 15]. This idea was originally motivated by the practical success of Markov Chain Monte Carlo (MCMC) sampling when performing inference in complex graphical models [14], and had the appeal of being able to flexibly represent complex probability distributions, something that parametric models could not achieve at the time.¹ Sampling has been invoked to explain aspects of perceptual decision making [16], response variability in V1 neurons [17, 18] and structured spontaneous activity in the cortex [19]. Moreover, recent theoretical work has focused on improving the computational efficiency of neural sampling, by increasing sampling speed [20, 21] and improving the robustness of the representation [22]. Nonetheless, sampling-based codes represent uncertainty only implicitly, distributed across time and neurons [22]. How the brain maps neural sampling dynamics into uncertainty-calibrated decisions remains a key open question.

Here, we propose the first spiking circuit for mapping neural samples into appropriate actions. At the core of our idea is the observation that although inference may involve complex high-dimensional posteriors, individual decisions are usually based on low-dimensional marginals of these distributions, which can be represented explicitly using a traditional parametric code. The proposed recurrent network model takes as inputs samples from the task-relevant posterior marginal and integrates them over time in a procedure inspired by online kernel density estimation [23]. We demonstrate the ability of this circuit to perform decision making in several toy examples. We also analyze the spiking activity of the network to look for signatures of uncertainty and probabilistic computation at the level of single neurons and population activity and find that our model recapitulates a variety of empirical observations in neural data. The model also makes predictions for how uncertainty and prior biases affect neural activity and population latent dynamics in this decision circuit.

1 Spiking neural network for sampling-based marginalization

In a Bayesian framework, perceptual decision making involves several key computations. First, given a sensory stimulus, s , one needs to compute the posterior over the latent variables, x , that may have given rise to it, $P(x|s)$. Second, uncertainty about nuisance variables is averaged out, to obtain a marginal distribution over the task-relevant latent x_i , $P(x_i|s)$. Finally, this marginal is combined with a task-specific cost function to yield the final decision. The first two steps are straightforward under a sampling framework [22], but parametric representations are needed for the final step to convert the information represented by samples into a spatially and temporally localized code [9].

Our approach uses spike-based distributed sampling for approximate inference and marginalization, then converts the resulting samples into an instantaneous parametric representation that can be used

¹Since recent DDCs are also able to represent complex multidimensional posteriors, the arguments in favor of sampling have moved away from representational flexibility, to its ability to explain many nontrivial features of neural variability in the cortex.

for decision making (Fig.1). For the inference portion of the model, we use the distributed sampling scheme developed by Savin and Deneve [22], implemented using the Poisson version of the encoding model, proposed by Rullán Buxó and Pillow [24]. Briefly, the recurrent circuit embeds Langevin sampling dynamics into a population of spiking neurons so that samples can be read out from the instantaneous neural rates (computed by low-pass filtering the spikes with an exponential kernel), \mathbf{r}'_t via a fixed linear decoder, \mathbf{D}' : $\mathbf{x}_t = \mathbf{D}'\mathbf{r}'_t$ (see Suppl. Info. and Refs. [22, 24] for details)². Given these dynamics, samples from the i -th marginal are obtained by projecting activity along the axis defined by row i of the decoding matrix, \mathbf{D}'_i .

The core of our contribution is the second circuit, which receives inputs $x_{i,t} = \mathbf{D}'_i\mathbf{r}'_t$ from the inference circuit, and constructs a parametric approximation of the associated distribution as a mixture of Gaussian kernels, $\phi_k(\cdot)$, parametrized by kernel weights, \mathbf{c} : $P(x_i) = \sum_k c_k \phi_k(x_i)$. For simplicity, here we assume a collection of K Gaussian basis functions, although alternative choices of kernels are also possible. The parameters \mathbf{c} adapt dynamically with each input sample, following dynamics of the form:

$$\dot{\mathbf{c}}_t = \frac{1}{\alpha}(\phi(x_{it}) - f(\mathbf{c}_{t-1})) \quad (1)$$

where vector $\phi(x_{it})$ concatenates all basis functions evaluated at the current input; function $f(\cdot)$ ensures that the fixed point solution yields a properly normalized distribution. Throughout this paper, we define $f(\bar{\mathbf{c}}_t) = K\langle\phi\rangle\bar{\mathbf{c}}_t = \beta\bar{\mathbf{c}}_t$, such that at the fixed point, the norm is 1, $\sum_k \frac{\phi_k(x_t)}{\beta} = \sum_k c_{k,t} = 1$ (see Suppl. Info. for details). The α term in Eq. 1 can be thought of as a leak, or the discount rate of internal evidence. This term becomes important for studying the dynamics of neural activity and decision making when the marginal posterior is changing, but otherwise does not qualitatively affect our results.

The dynamics implementing the parametric approximation of the marginal are embedded into a spiking neuron population, using the same encoding model used for the first module [24, 25, 27], so that the parameters can themselves be read out from neural activity via another arbitrary, fixed linear decoder, D , of size $K \times N$. The communication between circuits involves a linear map along the relevant axis of the inference circuit decoder, followed by point nonlinearities, with a form determined by the kernel $\phi_k(\cdot)$; this map could be implemented by local dendritic non-linearities.

As a simple proof of concept, we first consider inference in a linear Gaussian graphical model (Fig. 2A; see Suppl. Info. for details and simulation parameters), which results in Gaussian posterior marginals. Before stimulus onset, the inference circuit samples from a broad Gaussian prior centered at zero. Once the stimulus is on ($t = 0.15$ s), the dynamics switch to sampling from a sharper and correlated posterior (Fig. 2B). We take the first dimension x_1 as the decision-relevant latent dimension, with corresponding samples forming the stream of inputs to the decision circuit of size $N = 160$, where they are converted into a parametric representation of the marginal (Fig. 2C). The spiking activity of the neurons in the second circuit reflect these probabilistic computations, both before (spontaneous activity) and after stimulus presentation (Fig. 2D) so that parameters \mathbf{c}_t can be read out from the neural responses (Fig. 2E). They also closely match the ground truth $\sim 150 - 200$ ms after stimulus onset, once the dynamics have stabilized (Fig. 2F).

To explore how this representation could be used for optimal decision making, we consider a slightly more involved scenario, in which the task-relevant marginal is a bimodal distribution (formally, a two-component Gaussian mixture with unequal variance, Fig. 3A), a classic example where different choices of cost functions can have very different optimal decisions.³ For instance the task of estimating x_1 has the posterior mean as the optimum, while reporting the most likely value would use a maximum a posteriori (MAP) estimate. Here, we consider a binary categorization task, which requires reporting the sign of x_1 as an abstract analogue to binary discrimination (for instance, deciding whether or not an oriented grating is tilted leftwards or rightwards relative to the vertical). The optimal decision in this case is to report the side of the decision boundary that has more mass (integrating out the posterior $\int_0^\infty P(x_i|s)dx_i$). For our concrete example, there is slightly more mass

²The decoder weights are pre-specified by the linear dynamical system, consistent with the encoding framework by [25]; however, there exist extensions on how the weights could be learned through local plasticity [26].

³Bimodal posteriors are common in probabilistic accounts of bistable percepts [28, 29]

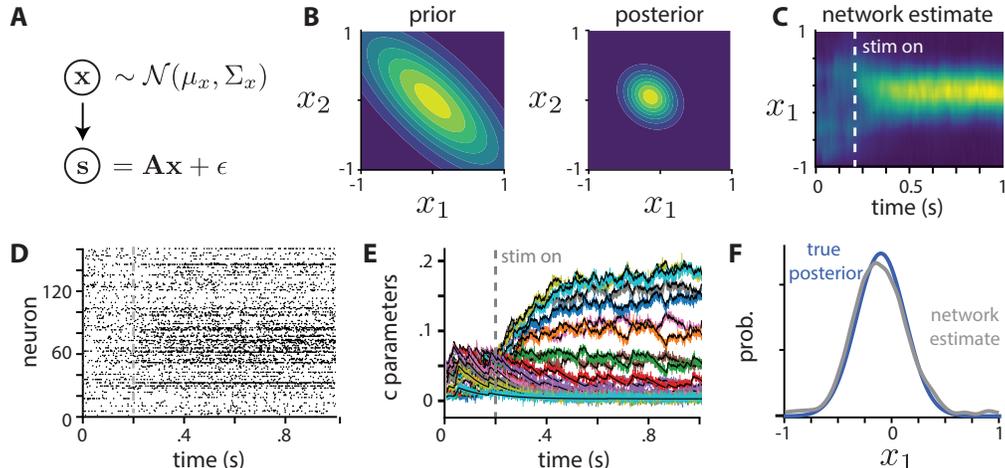


Figure 2: **Decision circuit and proof of concept decision making.** **A.** Graphical model for our toy example: inference in a linear Gaussian model. **B.** An example prior and posterior for two latent variables. **C.** The time evolution of the parametric approximation of the marginal for x_1 of the prior (before stimulus onset), and the posterior (afterward). **D.** Spiking activity encoding this information. **E.** Decoded parameters, c ; ground truth in black, network estimates in color. **F.** Comparison between estimate and ground truth posterior marginal at the end the trial.

on the negative side, such that with equal costs the optimal decision would be to report ‘L’.⁴ We also define an asymmetric cost scenario where the cost of mistakenly reporting a negative value is higher than the alternative so the optimal decision changes to ‘R’.

The dynamics of the decision circuit are shown in Fig. 3B. In this example, the inference circuit happens to start its trajectory in the rightward mixture component; since Langevin dynamics are slow at crossing regions of low probability [20, 21], initially the decision circuit only ‘sees’ samples from one component, but eventually both components are explored and the parametric marginal converges to the ground truth. In parallel, a cost-weighted linear projection of the corresponding neural responses tracks the evolution of the estimated cost associated to the two options (Fig. 3C). Indeed, for equal costs option ‘L’ starts with a high value, as it looks improbable under the initial approximate posterior, with the estimate converging to the correct low value with more samples; comparing the two options yields the optimal decision, ‘L’. Similarly, when high costs bias the decision away from this option, the network labels ‘R’ as the preferred option. Any other choice of cost function can be accommodated by changing the linear readout from the decision circuit. The same is true for estimation tasks based on the posterior mean, which we simulated to reproduce a reduction in the slope of a psychometric function with increasing uncertainty (Fig. 3D) and a shift in due to prior biases (Fig. 3E).

2 Neural signatures of sampling-based decision making

The inference and decision circuits share a lot of their underlying structure. Both encode information using the same spiking code; both their responses are driven by samples from the same posterior, directly in the inference circuit, and indirectly — via a linear projection — in the decision circuit. Are the two computations distinguishable at the level of the neural activity?

To investigate the neural implications of our model, we reproduced the experiment used to probe neural correlates of distributed sampling from [22]. Briefly, we assumed that the first circuit performs inference in response to 9 distinct test stimuli, under the assumption of a linear Gaussian graphical model. At the level of the decision circuit, this setup yields corresponding Gaussian marginals with equal variance and means that are evenly spaced on an interval from -1 to 1 . We probe the responses

⁴Note that codes that only consider the mean and variance of the posterior or ones that only represent the mode (MAP) would yield suboptimal decisions in this example; the correct answer requires knowledge of the full distribution as provided by sampling.

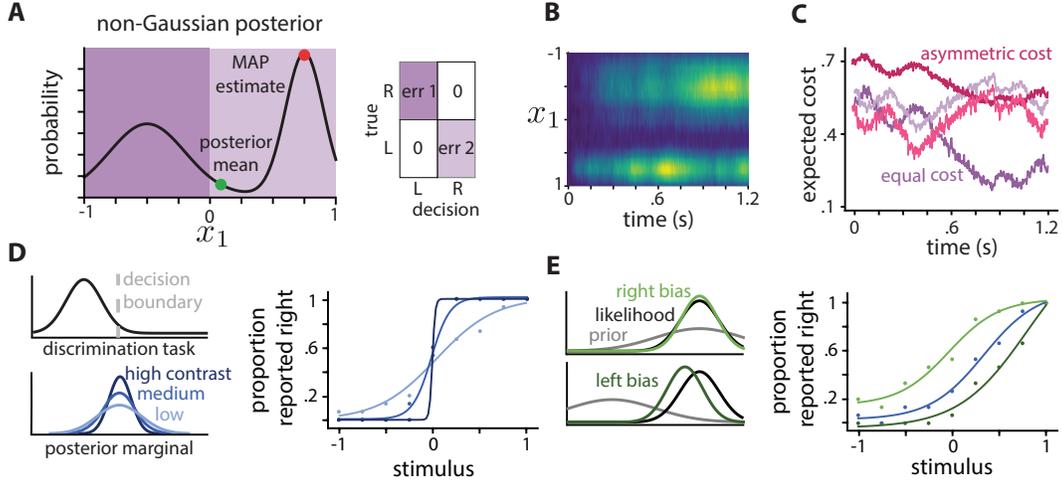


Figure 3: **Toy examples of decision making.** **A.** A binary decision when the relevant marginal is a two-component Gaussian mixture, so the mean (green) and the mode (red) of the distribution differ. The goal is to optimally report the sign of x_1 , taking into account any asymmetry in the cost structure. **B.** Evolution of the network’s marginal estimate; inference circuit dynamics start in the right posterior mode. **C.** Weighting output by the cost function generates dynamic estimates of the expected cost of the two decisions (L/R) in the case of equal costs (light purple = R, dark purple = L) or asymmetric cost (light pink = R, dark pink = L). **D.** Posterior mean estimation under uncertainty. Stimuli were selected from one of nine equally spaced measurement distributions with varying uncertainty (left). Psychometric curves showing the proportion of rightward judgements as a function of stimulus location and uncertainty (right). **E.** Same as D but for inference with prior-induced biases.

of the decision circuit to these stimuli and analyze the resulting spiking activity of the circuit the same way an experimentalist would analyze cortical data. We furthermore quantify how these responses change with uncertainty and prior biases. All figures shown are for the decision circuit; equivalent results for the inference circuit can be found in [22].

Signatures of uncertainty in single neurons

We quantified a range of response statistics that are commonly used to characterize spiking activity in the cortex using simulated data from the model. For each stimulus condition, we further varied posterior uncertainty between three levels, from low to high uncertainty (analogous to changing from high to low contrast in V1 experiments) (Fig. 4A).

Measuring the tuning function of individual neurons revealed that they decrease their peak firing and increase their width as uncertainty increases (Fig.4B). This modulation differs from the behavior of the inference module [22] or from data in early visual cortex [30]. The difference can be easily understood in the context of our unified coding scheme. In both modules, individual neurons are responsible for particular directions in signal space (here signal refers to either x_t or c_t), defined by the corresponding columns of the decoding matrix. They will increase their firing whenever the underlying signal is in their preferred region of stimulus space. In the inference module, this means that the average firing rate of the neuron will reflect the posterior mean. In the decision module, neurons inherit their stimulus selectivity from their preferred c axis. As uncertainty increases, their preferred c_k will reduce its amplitude, triggering a corresponding reduction in the neuron’s firing for a preferred stimulus. Likewise, c_k s far away from the stimulus, which were originally inactive, will start to participate in the signal; their corresponding neurons will also increase their firing rates, which leads to a widening of their tuning function. Putting together the different effects, as a function of similarity of the neuron’s preferred stimulus with the presented stimulus, we find reductions in firing rates for the preferred stimulus and an increase in firing away from the preferred stimulus, with the magnitude of the effects scaled by uncertainty (Fig. 4C).

In contrast, the effects of uncertainty on variability were consistent across circuits (see corresponding figure in Ref.[22]), with mean Fano factors close to 1 in the low uncertainty conditions and systematic

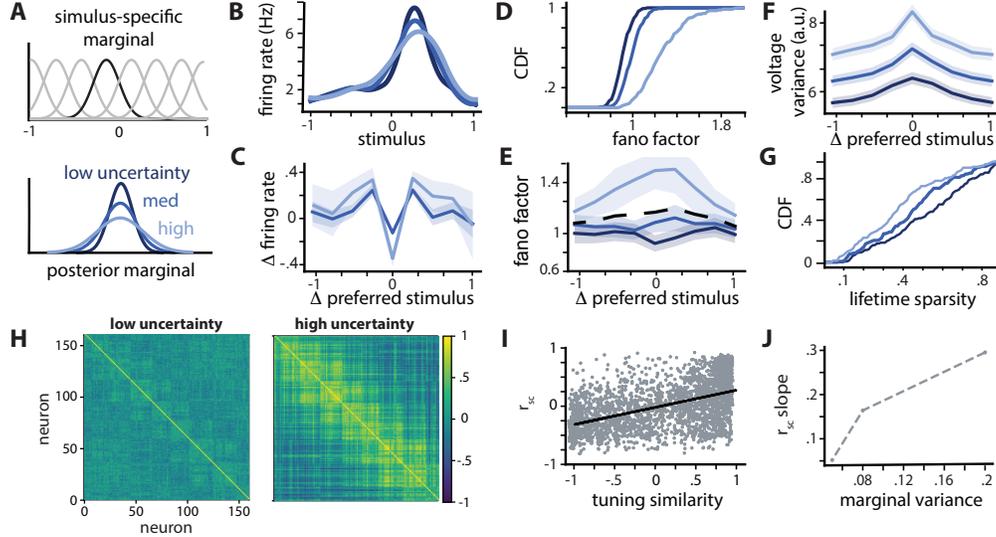


Figure 4: **Modulation of single neuron responses by uncertainty.** **A.** Simulated experiment with a range of stimuli that induce marginals with different means and the same variance; the uncertainty manipulations change their width. **B.** Tuning curves for a representative example neuron as uncertainty changes (with spline smoothing). **C.** Fractional changes in firing rate relative to the high certainty condition as a function of the distance between the stimulus and the neuron’s preferred stimulus. **D.** Fano factor cumulative distributions. **E.** Average Fano factors and **F.** Voltage variance as a function of the stimulus similarity to the neuron’s preferred stimulus. Black dashed line in E. shows average across all conditions. **G.** Lifetime sparsity cumulative distributions. **H.** Noise correlations between neuron pairs; neurons sorted by stimulus preference. **I.** Noise correlations as a function of tuning similarity; linear regression fit in black. **J.** Slope of regression line as a function of marginal uncertainty. All summary statistics were measured using 20 trials per stimulus condition.

rightward shifts of the Fano factor distribution with increasing uncertainty (Fig. 4D). This property can again be understood in the context of the encoding scheme: the variability of individual neurons can be decomposed into Poisson variability due to the encoding scheme [24] and signal variability. The first is largely unchanged across manipulations, but the increase in uncertainty leads to proportional increases in the signal — directly in the case of the inference circuit and indirectly in the decision module, since the variance of an average estimate also changes with the variance of the random variable being averages. Similar Fano factor increases are also a robust feature of experimental data [31, 17, 32].

Recent work also showed that average Fano factors over a range of uncertainties were roughly constant across stimulus orientations, both in a sampling-based model of V1 activity [17] and in data from awake macaque V1 [32]. While our model recapitulates this behavior on average (Fig. 4E, black dashed line), Fano factors start to increase at the preferred orientation for high uncertainty stimuli, a novel prediction for future data analyses. We found a similar pattern in the voltage variance (see detailed model description in Suppl.Info.), which increases with uncertainty and peaks at the preferred orientation (Fig. 4F), in agreement with previous experimental results [17, 33] as well as previous distributed sampling theory [34]. We also found that the sparsity of neural responses decreases with uncertainty, similar to that in the inference module (Fig. 4G).

Lastly, we analyzed the properties of pairwise correlations between neurons in the same setup. When visualizing the matrix of noise correlations between neurons (measured by Pearson correlations between their instantaneous firing rates over the course of a trial), with neurons indexed by their preferred stimulus, we found its structure to be modulated by uncertainty, with correlations expanding across wider stimulus ranges in high uncertainty conditions (Fig. 4H). While this observation is less intuitively explainable at the mechanistic level, it is likely to reflect the expansion of the active c_k range, possibly paired with the competitive process that ensures the normalization of the approximate marginal distribution. When organizing the noise correlations by the tuning preference of the two neurons, we found that similarly tuned neurons have higher correlations than neurons with different

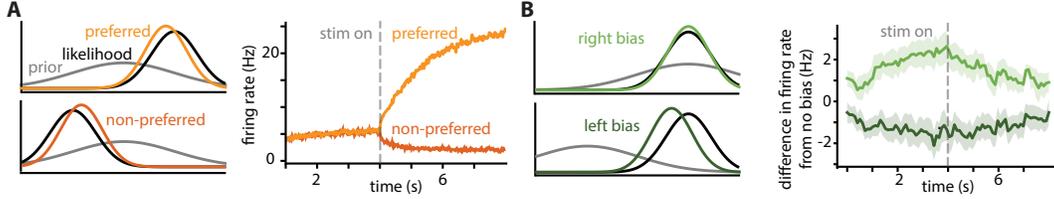


Figure 5: **Signatures of internal evidence integration.** Distributions in color represent posterior marginal. **A.** Left: a neuron’s perspective on probabilistic decision making comparing a scenario when the posterior mean is close to its preferred stimulus (top) vs. the posterior attributing high probability to non-preferred stimuli (bottom). Right: Evolution of average firing rates over time for the two scenarios. **B.** Prior manipulations (left) and the change in average neuron activity relative to no bias, for the two conditions (right).

tuning preferences (Fig. 4I), as seen experimentally [35]. Moreover, this dependence increases in strength with uncertainty (Fig. 4J), as previously reported for the inference circuit [22], leading to a robust prediction for this kind of sampling based computation, which can be tested experimentally.

Signatures of internal evidence integration So far, we have investigated the neural responses of the second module through the lens of statistics commonly used to characterize sensory neuron activity. Here, we analyze the same responses as one would for a decision circuit, such as LIP.

In particular, we focused on the unfolding of single neuron responses over time after stimulus onset (Fig. 5A), segregating the neurons into subsets whose tuning aligns with the presented stimulus (light pink) or not (dark pink). We found that neurons tuned to the stimulus ramp up their responses over time, as a reflection of the approximate posterior accumulating more posterior samples and thus increasing the corresponding c_k values. In contrast, neurons tuned to other stimuli decrease their responses, as the posterior density associated to that stimulus region decreases. The amplitude of these ramps depends on the overall uncertainty level, such that firing rates saturate at higher levels for higher certainty stimuli when the stimulus is in the preferred direction, similar to LIP responses found by Roitman and Shadlen [36] during sensory evidence integration. The key distinction is that here we are integrating *internal* evidence (a stream of samples) rather than external evidence (a stream of noisy stimuli). Or, put another way, the model we’re proposing computes a distribution over the decision variable and not a point estimate of the decision variable itself, as is the case in traditional evidence integration. What has been taken to be a key signature of evidence integration (firing rates increasing as more evidence is shown over time) is also consistent with our probabilistic coding scheme, despite the different computational goals.

The second manipulation we consider is one of the prior, inspired by an experiment by Rao and collaborators [37], in which the prior mean shifts rightwards or leftwards relative to a ‘no bias’ condition in which the prior is uninformative (Fig. 5B). After isolating the subset of neurons activated in the no bias condition, we compared their responses to the ‘right bias’ condition (light green) in which the prior aligns with the likelihood slightly sharpening the posterior without significantly changing its mean, and ‘left bias’ (dark green), which shifts the posterior away from the monitored neural population. Before the stimulus is presented, the neurons increase their activity when encoding the right bias, when the c_k s that these neurons are tuned to have high values, and similarly decrease their firing for the left bias. However, after the stimulus is turned on, this discrepancy slowly decreases: as the likelihood is strongly informative, the differences between posteriors are much less prominent than the differences between the two priors, so as more samples are integrated, the difference in responses are reduced so as to reflect the relatively minor difference between the likelihood and both posteriors. Again, LIP is known to show similar responses, where the effect of prior biases diminishes as the circuit accumulates more sensory evidence [37].

Signatures of probabilistic computation at the population level

The encoding of information in our model is naturally distributed across entire spiking neural networks and allows us to find robust signatures of the underlying probabilistic computation in the form of neural trajectories in low-dimensional manifolds [22]. If the map from neural responses to the

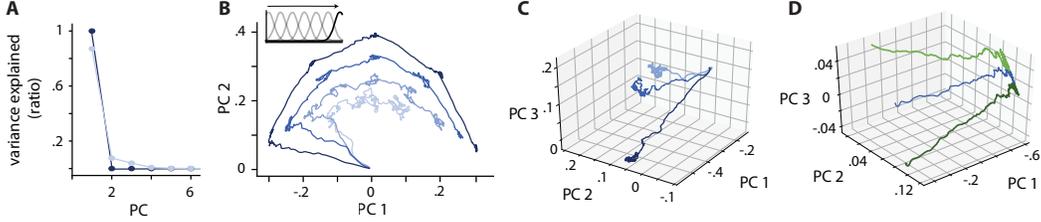


Figure 6: **Latent population dynamics.** **A.** Variance explained by increasing size of latent manifold for low (dark blue) and high (light blue) uncertainty; estimates use trial-average PCA for responses to 9 probe stimuli. **B.** Latent dynamics of the decision circuit in response to sequentially presenting the test stimuli (inset), for different uncertainty levels (colors). **C.** Latent dynamics of three presentations of the same stimulus, for different uncertainty conditions. **D.** Latent dynamics for three presentations of the same stimulus with an uninformative prior (blue) or left/right prior biases (green); initial conditions differ by prior, the likelihood is the same across conditions.

embedding manifold is easy to understand for the sampling circuit (the manifold maps one-to-one to the latent axes of the encoded posterior and the circuit dynamics to MCMC samples from it), this is not necessarily the case for the decision circuit. In particular the effective dimensionality of this circuit may change depending on uncertainty level. The intuition for this is simple: if the test stimuli are relatively sparse and the posteriors sharp, then a fraction of the parameters c_k s will not be recruited by any of the test trials. These ‘collapsed’ axes reduce the effective dimensionality of the dynamics; increasing uncertainty recruits more of these parameters, increasing effective dimensionality. Paired with the fact that c_k s are strongly correlated, dynamics appear almost invariably one dimensional in our toy examples (Fig. 6A), as seen in LIP data [38].

We can look at the low dimensional embedding of these dynamics for signatures of both uncertainty and prior biases. For instance, we simulated a scenario where the model is presented with the full sequence of test stimuli (from left to right) at different uncertainty levels, and traced the corresponding decision circuit trajectories in the space of their first two principal components (Fig. 6B), similar to [39]. This revealed structured trajectories whose geometry directly reflects posterior uncertainty. We further probed the effect of uncertainty on single stimuli (Fig. 6C). Here, the trajectories start at the same point (a location defined by the prior distribution) and gradually separate out to reflect different posteriors (Fig. 6C). The stimulus with the highest uncertainty (light blue) is closer to the prior, which nicely mirrors the observation that with limited evidence a Bayesian observer depends more on prior knowledge. The same pattern was seen in monkeys performing a time interval reproduction task [40]. The variance around the specific fixed point also depends on uncertainty, with more variability for higher uncertainty, analogous to (and inherited from) the inference circuit.

Finally we also considered a scenario where the likelihood is fixed, but the prior varies from leftward biased (dark green) to uninformative (blue) and rightward biased (light green). In this scenario, all three trajectories start in different regions of phase space (reflecting their respective priors) but quickly converge to roughly the same region of phase space (since the likelihood is narrow and dominates posterior estimates), with slight offsets that reflect the bias. Overall, we found that the decision circuit dynamics are low-dimensional and finely structured to reflect aspects of probabilistic computation, in a very different way from the sampling trajectory that drives them.

Inference under changing conditions

The online nature of the marginalization procedure implemented by the decision circuit allows it to integrate samples from any posterior, static or dynamic, within a time window specified by parameter α . This will lead to dependencies over time in the outputs of this circuit within the time horizon of integration, with possibly measurable behavioral signatures. Although the current implementation of the distributed sampling network only considers inference with static stimuli, we can get a sense for how the coding scheme is affected by stimulus dynamics by simulating a slightly more complex setup, with a succession of two static stimuli.

Figure 7A shows a simple simulation of a changing distribution, where the sensory input switches halfway through the simulation such that the mean of the marginal posterior changes. The decoded

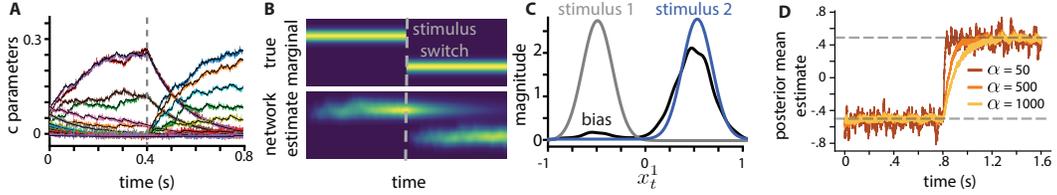


Figure 7: **Switching between two static stimuli.** **A.** Dynamics of decoded parameters \mathbf{c} when the stimulus switches 400ms (dashed line) into the simulated trial. **B.** The evolution of the decoded marginal over time. **C.** The decoded marginal distribution (black) 400ms after the stimulus switch. **D.** Decoded posterior mean estimate over time, for different values of integration time constant α . Dashed lines show ground truth for the two stimuli.

marginal distribution shows lingering traces of the previous stimulus, with the approximate marginal exhibiting bimodality (Fig.7B) that still persists 400ms after the switch (Fig.7C). Our model predicts that the strength of this bias should be related to the uncertainty of the preceding stimuli, and the similarity between the stimuli (formally, the distance between their corresponding posterior means). The temporal extent of this bias is controlled by the leak, or evidence discount rate, of the marginal posterior (Fig.7D). This is reminiscent of the history dependence of choice biases seen in sequential decision making [41, 42]. It may even be beneficial for dynamic inference, since exponential discounting of past evidence can be used to approximate Bayes-optimal inference in non-stationary environments [43]. It leads however to an interesting conundrum: on the one side, a precise representation of the marginal requires a large α so as to integrate as many samples as possible; on the other side, α should be small if we are to minimize sequential biases. The optimal trade-off between the two is dictated by the time scale of changes in the the sensory stimuli. Indeed, discount rates are known to reflect the expected rate of change of the statistics of the environment in a similar manner [44]. We expect that the integration time constant α is fixed within a given experimental context, but may adapt over time to reflect the time constants of environmental stability. The α parameter can in principle be estimated independently at the level of behavior and in neural data, which could be used to generate experimental predictions of the theory.

3 Discussion

Although behavioral evidence supports the idea that humans and animals use uncertainty to guide nearly optimal behavior (but see [45–47] for counter-arguments), the neural underpinnings of probabilistic computation remain controversial [15, 6, 48]. Contrary to the tradition of viewing parametric and sampling-based codes as mutually exclusive, here we use them both as useful data structures, needed at different computational stages for approximately optimal decision making. First, sensory circuits use sampling to represent the joint statistics over many features and to marginalize out nuisance variables, which is hard with parametric representations [49]. Second, in decision areas, parametric codes provide a compact, quasi-instantaneous representation of the marginals which can be easily combined with costs to yield approximately optimal decisions. Thus, our model provides a computationally well-justified reconciliation between competing probabilistic neural codes, which is complementary to recent attempts at unifying parametric and sampling-based models through the lens of Bayesian encoding and decoding [50].

Our decision circuit could in principle be adapted to represent a two-dimensional marginal posterior using two-dimensional kernel functions. The corresponding circuit size would need to grow linearly in the number of parameters, \bar{c} , and in principle exponentially with the number of marginal posterior dimensions. However, higher dimensional marginals may not be necessary: one should be able to include the decision variable explicitly as part of sampling-based hierarchical inference, making the relevant marginals always 1-D. These marginals could, in principle, change to adapt to different tasks using any of the biologically plausible proposals for dynamic information routing between neural circuits [51–53].

Our model involves two stages of computation and relies on the same spike-based framework [24] to encode the corresponding probabilistic quantities. This shared encoding means that many classic measures of responses in single cell and neuron pairs are preserved across processing stages and

largely consistent with those reported for early sensory cortical responses. Furthermore, the points of difference match experimental points of contention between parametric and sampling-based codes (e.g. the modulation of tuning functions by uncertainty) [9, 22]. Finally, although the idea that neural computation relies on coordinated activity between neurons is central to the construction of both circuits, they significantly diverge in the properties of their latent trajectories. Specifically, the two circuits have different slow dynamics that reflect the different underlying computations, leading to diverging experimental predictions. Uncertainty manifests mainly in the entropy of latent trajectories for sampling [22], whereas in the decision circuit it affects amplitude of mostly deterministic trajectories. These different signatures point to approaches based on latent dynamical systems models [54, 55] as a way to potentially disambiguate between different stages of probabilistic computation.

Our decision circuit can be thought of as a probabilistic form of evidence integration akin to that proposed by Boerlin and Deneve [25], with the ‘evidence’ given by samples from the network’s internal model rather than noisy external inputs. This close analogy prompts the idea that from the perspective of a brain trying to act in the world—or our decision circuit—there may be little point in making a distinction between internal and external noise [56]. This significantly complicates the experimental validation of neural sampling in a way that can only be addressed by further theoretical work. To explore such complex scenarios, a natural next step is to expand the first module from inference with static stimuli to inference based on continuous stimulus streams with nontrivial temporal correlations. This could in principle be achieved by using particle filtering as underlying sampling dynamics [57], with particles embedded in the same spiking recurrent circuit (as previously demonstrated for multiple parallel chains sampling from the same static posterior [22]). Exploring the interaction between such inference dynamics and our probabilistic decision circuit could reveal new neural and behavioral signatures of sampling-based probabilistic computation in the brain.

Acknowledgments and Disclosure of Funding

We thank Edoardo Balzani, Pedro Herrero-Vidal and Colin Bredenberg for helpful discussions and feedback on earlier versions of this manuscript. CRB is supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE1839302. CS is supported by National Institute of Mental Health Award 1R01MH125571-01, by the National Science Foundation under NSF Award No.1922658 and a Google faculty award.

References

- [1] Wei Ji Ma and Mehrdad Jazayeri. Neural coding of uncertainty and probability. *Annual Review of Neuroscience*, 37(1), 2014.
- [2] Adam Kepecs, Naoshige Uchida, Hatim A. Zariwala, and Zachary F. Mainen. Neural correlates, computation and behavioural impact of decision confidence. *Nature*, 455(7210):227–231, September 2008.
- [3] E. R. Stern, R. Gonzalez, R. C. Welsh, and S. F. Taylor. Updating beliefs for a decision: Neural correlates of uncertainty and underconfidence. *Journal of Neuroscience*, 30(23):8032–8041, June 2010.
- [4] R. Kiani and M. N. Shadlen. Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, 324(5928):759–764, May 2009.
- [5] Tianming Yang and Michael N. Shadlen. Probabilistic reasoning by neurons. *Nature*, 447(7148):1075–1080, June 2007.
- [6] Alexandre Pouget, Jeffrey M Beck, Wei Ji Ma, and Peter E Latham. Probabilistic brains: knowns and unknowns. *Nature Neuroscience*, 16(9):1170–1178, September 2013.
- [7] Neil M.T. Houlshby, Ferenc Huszár, Mohammad M. Ghassemi, Gergő Orbán, Daniel M. Wolpert, and Máté Lengyel. Cognitive tomography reveals complex, task-independent mental representations. *Current Biology*, 23(21):2169–2175, November 2013.

- [8] Thomas L. Griffiths and Joshua B. Tenenbaum. Optimal predictions in everyday cognition. *Psychological Science*, 17(9):767–773, September 2006.
- [9] Jeffrey M. Beck, Wei Ji Ma, Roozbeh Kiani, Tim Hanks, Anne K. Churchland, Jamie Roitman, Michael N. Shadlen, Peter E. Latham, and Alexandre Pouget. Probabilistic Population Codes for Bayesian Decision Making. *Neuron*, 60(6):1142–1152, December 2008.
- [10] Wei Ji Ma, Jeffrey M Beck, Peter E Latham, and Alexandre Pouget. Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11):1432–1438, November 2006.
- [11] Richard S. Zemel, Peter Dayan, and Alexandre Pouget. Probabilistic interpretation of population codes. *Neural Computation*, 10(2):403–430, 1998.
- [12] Richard S. Zemel and Peter Dayan. Distributional population codes and multiple motion models. In *Proceedings of the 11th International Conference on Neural Information Processing Systems*, NIPS’98, page 174–180, Cambridge, MA, USA, 1998. MIT Press.
- [13] Eszter Vértes and Maneesh Sahani. Flexible and accurate inference and learning for deep generative models. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- [14] Patrik Hoyer and Aapo Hyvärinen. Interpreting neural response variability as monte carlo sampling of the posterior. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems*, volume 15. MIT Press, 2003.
- [15] József Fiser, Pietro Berkes, Gergő Orbán, and Máté Lengyel. Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3):119–130, March 2010.
- [16] Ralf M. Haefner, Pietro Berkes, and József Fiser. Perceptual Decision-Making as Probabilistic Inference by Neural Sampling. *Neuron*, 90(3):649–660, May 2016.
- [17] Gergő Orbán, Pietro Berkes, József Fiser, and Máté Lengyel. Neural variability and sampling-based probabilistic representations in the visual cortex. *Neuron*, 92(2):530–543, October 2016.
- [18] Rodrigo Echeveste, Laurence Aitchison, Guillaume Hennequin, and Máté Lengyel. Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference. *Nature neuroscience*, 23(9):1138–1149, 2020.
- [19] P. Berkes, G. Orban, M. Lengyel, and J. Fiser. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013):83–87, January 2011.
- [20] Guillaume Hennequin, Laurence Aitchison, and Mate Lengyel. Fast sampling-based inference in balanced neuronal networks. In *Advances in Neural Information Processing Systems 27*, pages 2240–2248. 2014.
- [21] Laurence Aitchison and Máté Lengyel. The Hamiltonian brain: Efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLOS Computational Biology*, 12(12):1–24, 12 2016.
- [22] Cristina Savin and Sophie Denève. Spatio-temporal representations of uncertainty in spiking neural networks. In *Advances in Neural Information Processing Systems 27*, pages 2024–2032. 2014.
- [23] Matej Kristan, Aleš Leonardis, and Danijel Skočaj. Multivariate online kernel density estimation with gaussian kernels. *Pattern Recognition*, 44(10):2630–2642, 2011.
- [24] Camille E. Rullán Buxó and Jonathan W. Pillow. Poisson balanced spiking networks. *PLOS Computational Biology*, 16(11):e1008261, November 2020.
- [25] Martin Boerlin, Christian K. Machens, and Sophie Denève. Predictive coding of dynamical variables in balanced spiking networks. *PLoS Computational Biology*, 9(11):e1003258, November 2013.

- [26] Wieland Brendel, Ralph Bourdoukan, Pietro Verzechi, Christian K. Machens, and Sophie Denève. Learning to represent signals spike by spike. *PLOS Computational Biology*, 16(3):e1007692, March 2020.
- [27] Sophie Denève and Christian K Machens. Efficient codes and balanced networks. *Nature Neuroscience*, 19(3):375–382, February 2016.
- [28] David Reichert, Peggy Series, and Amos J Storkey. Neuronal adaptation for sampling-based probabilistic inference in perceptual bistability. In *Advances in Neural Information Processing Systems*, volume 24, 2011.
- [29] Cristina Savin, Peter Dayan, and Máté Lengyel. Optimal recall from bounded metaplastic synapses: Predicting functional adaptations in hippocampal area CA3. *PLOS Computational Biology*, 10(2):1–22, 02 2014.
- [30] Todd W. Troyer, Anton E. Krukowski, Nicholas J. Priebe, and Kenneth D. Miller. Contrast-invariant orientation tuning in cat visual cortex: Thalamocortical input tuning and correlation-based intracortical connectivity. *Journal of Neuroscience*, 18(15):5908–5927, 1998.
- [31] Mark M Churchland, Byron M Yu, John P Cunningham, Leo P Sugrue, Marlene R Cohen, Greg S Corrado, William T Newsome, Andrew M Clark, Paymon Hosseini, Benjamin B Scott, David C Bradley, Matthew A Smith, Adam Kohn, J Anthony Movshon, Katherine M Armstrong, Tirin Moore, Steve W Chang, Lawrence H Snyder, Stephen G Lisberger, Nicholas J Priebe, Ian M Finn, David Ferster, Stephen I Ryu, Gopal Santhanam, Maneesh Sahani, and Krishna V Shenoy. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*, 13(3):369–378, February 2010.
- [32] Alexander S. Ecker, Philipp Berens, Georgios A. Keliris, Matthias Bethge, Nikos K. Logothetis, and Andreas S. Tolias. Decorrelated neuronal firing in cortical microcircuits. *Science*, 327(5965):584–587, 2010.
- [33] Ian M. Finn, Nicholas J. Priebe, and David Ferster. The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron*, 54(1):137–152, April 2007.
- [34] Edoardo Balzani and Cristina Savin. Population-level signatures of probabilistic computation. In *COSYNE 2019*, 2019.
- [35] Matthew A. Smith and Adam Kohn. Spatial and temporal scales of neuronal correlation in primary visual cortex. *Journal of Neuroscience*, 28(48):12591–12603, 2008.
- [36] Jamie D. Roitman and Michael N. Shadlen. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience*, 22(21):9475–9489, November 2002.
- [37] V. Rao, G. C. DeAngelis, and L. H. Snyder. Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas. *Journal of Neuroscience*, 32(29):10063–10074, July 2012.
- [38] Surya Ganguli, James W. Bisley, Jamie D. Roitman, Michael N. Shadlen, Michael E. Goldberg, and Kenneth D. Miller. One-dimensional dynamics of attention and decision making in LIP. *Neuron*, 58(1):15–25, April 2008.
- [39] Gouki Okazawa, Long Sha, and Roozbeh Kiani. Linear integration of sensory evidence over space and time underlies face categorization. *Cell (in press)*, November 2020.
- [40] Hansem Sohn, Devika Narain, Nicolas Meirhaeghe, and Mehrdad Jazayeri. Bayesian computation through cortical latent dynamics. *Neuron*, 103(5):934–947.e5, September 2019.
- [41] Daniel E. Acuña and Paul Schrater. Structure learning in human sequential decision-making. *PLOS Computational Biology*, 6(12):1–12, 12 2010.
- [42] Anne E Urai, Jan Willem De Gee, Konstantinos Tsetsos, and Tobias H Donner. Choice history biases subsequent evidence accumulation. *Elife*, 8:e46331, 2019.

- [43] Angela J Yu and Jonathan D Cohen. Sequential effects: Superstition or rational behavior? In *Advances in Neural Information Processing Systems 21*, pages 1873–1880. 2009.
- [44] Christopher M Glaze, Joseph W Kable, and Joshua I Gold. Normative evidence accumulation in unpredictable environments. *eLife*, 4, August 2015.
- [45] Jeffrey S Bowers and Colin J Davis. Bayesian just-so stories in psychology and neuroscience. *Psychological bulletin*, 138(3):389, 2012.
- [46] Dobromir Rahnev and Rachel N Denison. Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, 41, 2018.
- [47] A Emin Orhan and Wei Ji Ma. Efficient probabilistic inference in generic neural networks trained with non-probabilistic feedback. *Nature communications*, 8(1):1–14, 2017.
- [48] Jeff Beck, Ralf Haefner, Xaq Pitkow, Cristina Savin, and Eszter Vertes. Competing theories of probabilistic computations in the brain. <https://openreview.net/forum?id=uZMMO2obl50>. Accessed: 2021-05-28.
- [49] J. M. Beck, P. E. Latham, and A. Pouget. Marginalization in neural circuits with divisive normalization. *Journal of Neuroscience*, 31(43):15310–15319, October 2011.
- [50] Richard D. Lange, Sabyasachi Shivkumar, Ankani Chattoraj, and Ralf M. Haefner. Bayesian encoding and decoding as distinct perspectives on neural coding. *bioarxiv*, October 2020.
- [51] David C Van Essen, Charles H Anderson, and Bruno A Olshausen. Dynamic routing strategies in sensory, motor, and cognitive processing. *Large-scale neuronal theories of the brain*, pages 271–299, 1994.
- [52] Christopher H Donahue and Daeyeol Lee. Dynamic routing of task-relevant signals for decision making in dorsolateral prefrontal cortex. *Nature neuroscience*, 18(2):295–301, 2015.
- [53] Guangyu Robert Yang, John D Murray, and Xiao-Jing Wang. A dendritic disinhibitory circuit mechanism for pathway-specific gating. *Nature communications*, 7(1):1–14, 2016.
- [54] John P Cunningham and Byron M Yu. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500–1509, August 2014.
- [55] Jakob H Macke, Lars Busing, John P. Cunningham, Byron M. Yu, Krishna V. Shenoy, and Maneesh Sahani. Empirical models of spiking in neural populations. *Advances in Neural Information Processing Systems 24*, 2011.
- [56] Jeffrey M. Beck, Wei Ji Ma, Xaq Pitkow, Peter E. Latham, and Alexandre Pouget. Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability. *Neuron*, 74(1):30–39, April 2012.
- [57] Anna Kutschireiter, Simone Carlo Surace, Henning Sprekeler, and Jean-Pascal Pfister. Non-linear Bayesian filtering and learning: a neuronal dynamics for perception. *Scientific Reports*, 7(1):8722, December 2017.