Stochastic inference with spiking neurons in the high-conductance state

Mihai A. Petrovici, ^{1,*} Johannes Bill, ^{1,2} Ilja Bytschok, ¹ Johannes Schemmel, ¹ and Karlheinz Meier ¹ Kirchhoff Institute for Physics, University of Heidelberg, Heidelberg, Germany ² Institute for Theoretical Computer Science, Graz University of Technology, Styria, Austria (Received 11 November 2013; revised manuscript received 2 July 2016; published 20 October 2016)

The highly variable dynamics of neocortical circuits observed *in vivo* have been hypothesized to represent a signature of ongoing stochastic inference but stand in apparent contrast to the deterministic response of neurons measured *in vitro*. Based on a propagation of the membrane autocorrelation across spike bursts, we provide an analytical derivation of the neural activation function that holds for a large parameter space, including the high-conductance state. On this basis, we show how an ensemble of leaky integrate-and-fire neurons with conductance-based synapses embedded in a spiking environment can attain the correct firing statistics for sampling from a well-defined target distribution. For recurrent networks, we examine convergence toward stationarity in computer simulations and demonstrate sample-based Bayesian inference in a mixed graphical model. This points to a new computational role of high-conductance states and establishes a rigorous link between deterministic neuron models and functional stochastic dynamics on the network level.

DOI: 10.1103/PhysRevE.94.042312

I. INTRODUCTION

In responding to environmental stimuli, brains have to make predictions based on incomplete, noisy, and ambiguous data. The recent hypothesis that the brain copes with this challenge by performing Bayesian, rather than logical inference [1–3], has been strengthened by electrophysiological data that identified neural correlates of the involved computations [4,5] and theoretical work on potential spiking network implementations [6–8].

In probabilistic inference, possible values of a quantity are described by a random variable (RV) z_k and all dependencies between RVs are stored in a joint distribution $p(z_1, \ldots, z_K)$. The belief about a set of unobserved RVs $\{z_1, \ldots, z_M\}$ given an observed set of RVs is represented by the posterior distribution $p(z_1, \ldots, z_M | z_{M+1}, \ldots, z_K)$. In particular, the posterior contains information on the most likely conclusion, as well as on all potential alternatives.

With regard to the representation of probability distributions in the brain, theoretical work [2] has argued in favor of sample-based codes, i.e., instead of providing the entire distribution at any point in time, samples $z^{(t)} \sim p(z_1, \dots, z_K)$ are used as a proxy. When modeling large systems, this offers three important advantages. First, approximate solutions can be provided at any time, with increasingly reliable results as the calculation progresses ("anytime computing"). Second, marginalization comes at no cost, as $p(z_k)$ can be determined by simply neglecting the values of all other RVs. Third, some sampling algorithms support a high degree of parallelization with an algorithmic structure that is reminiscent of neural networks [9]. Recently, a theory has been suggested that combines these advantages by implementing Markov chain Monte Carlo (MCMC) sampling in spiking networks of abstract model neurons [7]. In this framework, spike patterns are interpreted as samples of binary RVs as follows [see Fig. 1(a)]:

$$z_k^{(t)} = 1 \Leftrightarrow \text{Neuron } k \text{ fired in } (t - \tau_{\text{on}}, t].$$
 (1)

The duration τ_{on} of the active state following a spike is a free parameter for the timescale on which a neuron affects downstream cells. The neuron model in Ref. [7] is inherently stochastic, with an instantaneous firing probability,

$$r_k(t) = \lim_{\Delta t \to 0} \frac{p(\text{spike in } [t, t + \Delta t))}{\Delta t} = \begin{cases} \frac{\exp(v_k)}{\tau} & \text{if } z_k = 0\\ 0 & \text{if } z_k = 1 \end{cases},$$
(2)

where v_k is an abstract membrane potential [Fig. 1(b)].

In contrast to this idealized stochastic neuron model, *in vitro* experiments have demonstrated the largely deterministic nature of single neurons [10]. Similarly, microscopic models of neural circuits typically rely on deterministic dynamics of their constituents [see Fig. 1(c)]. The aim of this paper is to demonstrate how a network of deterministic neurons in a biologically plausible spiking environment can quantitatively reproduce the stochastic dynamics required for sampling from a well-defined distribution and perform inference given observations. In this letter, we extend our earlier discussion of this approach [11].

We start by calculating the dynamics of a single leaky integrate-and-fire (LIF) neuron in a spiking noisy environment and derive its activation function in the high-conductance state. For this biologically relevant regime, existing analytical descriptions of neuronal response functions (e.g., Refs. [16,17]) are not applicable. Here, we provide an approach

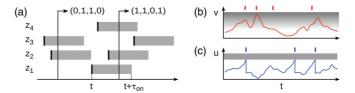


FIG. 1. (a) Spike patterns as samples of a random vector z. The RV z_k is active for duration τ_{on} (gray bar) after each spike. (b) In stochastic neuron models, internal state variables (red) modulate the firing probability (gray). (c) In contrast, deterministic neurons elicit a spike when the membrane potential (blue) reaches a threshold voltage.

^{*}Corresponding author: mpedro@kip.uni-heidelberg.de

based on the propagation of the membrane autocorrelation throughout spike bursts. This establishes an equivalence to the abstract, inherently stochastic units [Eq. (2)]. On the network level, we show how conductance-based synapses approximate the interaction for sampling from a well-defined target distribution. Furthermore, we show that the distribution sampled by LIF networks remains a good approximation of the target distribution even for large networks with strong recurrent interaction. We complement our study with a demonstration of probabilistic inference by implementing a small graphical model for pattern recognition in a recurrent network of LIF neurons.

II. LIF DYNAMICS IN A HIGH-CONDUCTANCE STATE

We consider deterministic LIF neurons defined by $C_m \frac{d}{dt} u_k = g_l(E_l - u_k) + I_k$, with membrane potential u_k , capacitance C_m , leak potential E_l , leak conductance g_l , and input current I_k . When u_k crosses a threshold ϑ from below, a spike is emitted and u_k is reset to ϱ for a refractory period $\tau_{\rm ref}$. We formally partition the total input current I_k into recurrent synaptic input, diffuse synaptic noise, and additional external currents: $I_k = I_k^{\rm rec} + I_k^{\rm noise} + I_k^{\rm ext}$. The currents $I_k^{\rm rec}$ and $I_k^{\rm noise}$ are mediated through synapses and obey $I_k^{\rm syn} = \sum_i g_{ki} \left(E_i^{\rm rev} - u_k \right)$ with reversal potential $E_i^{\rm rev}$ of the ith synapse. The temporal evolution of the conductance g_{ki} is modeled as a low-pass filter on presynaptic spikes: $\frac{d}{dt} g_{ki} = -g_{ki}/\tau_{\rm syn} + \sum_s w_{ki} \delta(t - t_i^s)$, with synaptic time constant $\tau_{\rm syn}$, weight w_{ki} , and spike times t_i^s .

We start by considering a single neuron that receives diffuse synaptic noise I_k^{noise} in the form of Poisson spike trains from its surrounding. The capacity of recurrent networks to produce such noise has been shown in, e.g., Ref. [12]. For the following analysis of individual neurons in a noisy environment, we omit the index k and set $I^{\text{rec}} = 0$.

When a conductance-based LIF neuron receives strong synaptic stimulation, it enters a so-called high-conductance state (HCS, [13]), characterized by accelerated membrane dynamics. For an analytical treatment, it is advantageous to rewrite the membrane dynamics as $\tau_{\rm eff} \frac{d}{dt} u = u_{\rm eff}(t) - u$, such that u decays toward an effective leak potential $u_{\rm eff}$ with an effective time constant $\tau_{\rm eff} = C_m/g_{\rm tot}$ [14] (see Appendix A). The total conductance $g_{\rm tot}$ subsumes both leakage and synaptic conductances. In a high input rate regime, $\tau_{\rm eff} \to 0$ and the effective potential $u_{\rm eff}$ simply becomes a linear transformation of the synaptic noise input (see Appendix B). Using methods similar to Ref. [15], it can be shown that, in this regime, $u_{\rm eff}(t)$ can be described as an Ornstein-Uhlenbeck (OU) process $du = \frac{1}{\tau_{\rm syn}}(\bar{u} - u)dt + \sigma dW$ with parameters (see Appendix C):

$$\bar{u} = \left\{ I^{\text{ext}} + g_l E_l + \sum_i \nu_i w_i E_i^{\text{rev}} \tau_{\text{syn}} \right\} / \langle g_{\text{tot}} \rangle, \quad (3)$$

$$\sigma^2 = \left\{ \sum_{i} \nu_i \left[w_i \left(E_i^{\text{rev}} - \bar{u} \right) \right]^2 \tau_{\text{syn}} \right\} / \langle g_{\text{tot}} \rangle^2. \tag{4}$$

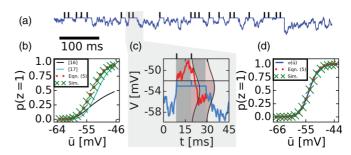


FIG. 2. (a) Membrane potential u(t) and spikes of an LIF neuron in a spiking noisy environment. (b) Prediction of the activation function (red) compared to simulation results (green), as well as to other predictions from literature [16,17]. (c) In a HCS, u (blue) and $u_{\rm eff}$ (red) are nearly identical when the neuron is not refractory. After each refractory period (gray), the predicted distribution of $u_{\rm eff}$ (pink) is used for the propagation in [Eq. (6)]. (d) High-noise regime: theoretical prediction (red) vs. simulation results (green), with a fitted logistic function $\sigma(\bar{u})$ (blue).

III. THE ACTIVATION FUNCTION OF LIF NEURONS IN THE HCS

Similar to the abstract model [7], we define the refractory state of a neuron as z(t) = 1. The mapping of spikes to RV states [Eq. (1)] naturally leads to the concept of an activation function $p(z = 1|\bar{u})$, where $\bar{u} = \langle u_{\text{eff}} \rangle_t$. In the following, we derive a general expression for the activation function of an LIF neuron under Poisson stimulus, which is of particular use in the cortically relevant HCS regime. Figure 2(a) shows an exemplary simulation (membrane potential and spike train) of such a scenario. The activation function in Fig. 2(b) is obtained by sweeping over I^{ext} (see Appendices G and H).

Related setups have already been examined in Refs. [16,17]. However, the methods in Refs. [16,17] are tailored to certain parameter ranges, which, in particular, do not include the HCS regime with refractoriness. This is because Ref. [16] requires $\tau_{\rm syn} \ll \tau_{\rm eff}$, whereas Ref. [17] assumes $\tau_{\rm ref} \ll \tau_{\rm eff}, \tau_{\rm syn}$, which leads to discrepancies in the predicted activation functions [Fig. 2(b)]. The deeper reason for the observed discrepancies is found in the lack of an appropriate propagation of the autocorrelation of $u_{\rm eff}$ through $\tau_{\rm ref}$. Here, we propose a derivation that explicitly includes this propagation and thereby covers a large parameter space, including the cases studied in Refs. [16,17], as well as the HCS.

In Fig. 2(a), two modes of firing can be observed: a "bursting" mode, where the effective membrane potential $u_{\rm eff}$ after the refractory period is still above threshold, and a freely evolving mode, where the neuron does not spike again immediately after the refractory period. This is illustrated by the distributions in Fig. 2(c). Our approach relies on the calculation of burst lengths and their associated occurrence probability P_n .

Denoting the relative occurrence of burst lengths n by P_n , the average drift time from ϱ to ϑ between the kth and (k+1)st spike in a burst by $\overline{\tau_k^b}$ and the average duration of the freely evolving mode that follows an n-spike-burst by T_n , we identify the following relation:

$$p(z=1|\bar{u}) = \frac{\sum_{n} P_{n} n \tau_{\text{ref}}}{\sum_{n} P_{n} \left(n \tau_{\text{ref}} + \sum_{k=1}^{n-1} \overline{\tau_{k}^{\text{b}}} + T_{n}\right)}.$$
 (5)

Knowing the OU process governing $u_{\rm eff}$, recursive expressions for P_n and T_n can be derived. These terms have considerable impact on the activation function if $\tau_{\rm ref} \approx \tau_{\rm syn}$. In the limit of strong noise stimuli, we can calculate the average drift time $\overline{\tau_k^{\rm b}}$ in a quasistatic approximation [17] of $u_{\rm eff}$ and assume that $u \approx u_{\rm eff}$ for the freely evolving mode, thus obtaining (see Appendix D):

$$P_{n} = \left(1 - \sum_{i=1}^{n-1} P_{i}\right) \int_{\vartheta}^{\infty} du_{n-1} p(u_{n-1} | u_{n-1} \geqslant \vartheta)$$

$$\times \left[\int_{-\infty}^{\vartheta} du_{n} p(u_{n} | u_{n-1})\right], \qquad (6)$$

$$T_{n} = \int_{\vartheta}^{\infty} du_{n-1} p(u_{n-1} | u_{n-1} \geqslant \vartheta)$$

$$\times \left[\int_{-\infty}^{\vartheta} du_{n} p(u_{n} | u_{n} < \vartheta, u_{n-1}) \langle T(\vartheta, u_{n}) \rangle\right], \qquad (7)$$

$$\overline{\tau_{k}^{b}} = \int_{\vartheta}^{\infty} du_{k} \ln\left(\frac{\varrho - u_{k}}{\vartheta - u_{k}}\right)$$

$$\times \int_{\vartheta}^{\infty} du_{k-1} p(u_{k} | u_{k} \geqslant \vartheta, u_{k-1}). \qquad (8)$$

Figure 2(c) displays an intuitive picture of the integrals in Eqs. (6) and (7). The transfer function $p(u_n|u_{n-1})$ is the Green's function of the OU process at time $t - t_s = \tau_{\text{on}}$ and $\langle T(\vartheta, u_n) \rangle$ denotes the average time needed for the membrane to reach ϑ starting from u_n , which can be given in closed form [18]. The dependency of P_n , T_n , and $\overline{\tau_k^b}$ on the moments of the OU process [Eqs. (3) and (4)] renders Eq. (5) a function of \bar{u} .

So far, we have only considered $\tau_{\rm eff}=0$. To further improve the prediction, we take into account finite values of $\tau_{\rm eff}$ by means of an expansion in $\sqrt{\tau_{\rm eff}/\tau_{\rm syn}}$. Due to the symmetry of the PSP shape in $\tau_{\rm eff}$ and $\tau_{\rm syn}$, this can be done analogously to Refs. [16,17], where the opposite limit of large $\tau_{\rm eff}$ and small $\tau_{\rm syn}$ is used (see Appendix D). A comparison between our prediction of $p(z=1|\bar{u})$ and results from a numerical simulation is shown in Fig. 2(b).

IV. SAMPLING WITH NETWORKS OF LIF NEURONS

We can now reconcile the response of LIF neurons with the inherently stochastic neuron model [Eq. (2)] that requires a logistic activation function for constant potential v: $p(z = 1|v) = \sigma(v) := [1 + \exp(-v)]^{-1}$. In the HCS regime, this logistic activation can be approximated by LIF neurons with high accuracy. Figure 2(d) shows our theoretical prediction and simulation results for the activation function alongside a fitted logistic function. For the translation from the LIF domain to the abstract model [Eq. (2)] we have employed a linear mapping,

$$v = (\bar{u} - \bar{u}^0)/\alpha,\tag{9}$$

with scaling factor α and \bar{u}^0 denoting the potential for which $p(z=1)=\frac{1}{2}$.

We next connect the neurons to form a recurrent network. In addition to noise stimuli, an LIF neuron in a network receives synaptic currents I_k^{rec} from other neurons. For certain connectivity structures, it is possible to predict the target distribution [7,19] of states $z^{(t)}$ that arise from the stochastic

dynamics of the recurrent network. We use the emulation of Boltzmann machines (BMs) as an example case. The joint distribution reads.

$$p_{\mathrm{B}}(z) = Z^{-1} \exp(z^{\mathsf{T}} W z / 2 + z^{\mathsf{T}} b), \tag{10}$$

where W is a symmetric zero-diagonal weight matrix, b is a bias vector, and Z is the normalizing partition function. This probabilistic model underlies state-of-the-art machine learning algorithms for image [20] and speech recognition [21]. It has been shown [7] that a network of abstract neurons (Eq. 2) with linear membrane potentials,

$$v_k = b_k + \sum_{j=1}^K W_{kj} z_j, \tag{11}$$

will sample from the desired target distribution [Eq. (10)]. This finding uses the fact that individual neurons sample from the conditionals $p(z_k = 1|z_{\setminus k}) = \sigma(v_k)$, with $z_{\setminus k} = \{z_j \mid j \neq k\}$, in an MCMC updating scheme.

As shown above, LIF neurons in a spiking noisy environment closely approximate this logistic activation function if the synaptic currents I_k^{rec} shift the mean membrane potential \bar{u}_k according to the linear interaction [Eq. (11)]. Using the linear transformation [Eq. (9)] between v_k and \bar{u}_k , and estimating the effect of a conductance-based synapse of weight w_{kj} , we arrive at the following translation between the abstract and the LIF domain (see Appendix E):

$$b_k = \left(\bar{u}_k^b - \bar{u}_k^0\right) / \alpha,\tag{12}$$

$$W_{kj} = \frac{1}{\alpha C_m} \frac{w_{kj} \left(E_{kj}^{\text{rev}} - \mu \right)}{\frac{1}{\tau_{\text{syn}}} - \frac{1}{\tau_{\text{eff}}}} \left[\frac{1 - e}{e} - \frac{\tau_{\text{eff}}}{\tau_{\text{syn}}} (e^{-\frac{\tau_{\text{syn}}}{\tau_{\text{eff}}}} - 1) \right], \tag{13}$$

where \bar{u}_k^b is the mean potential \bar{u}_k that establishes $p(z_k = 1 | \bar{u}_k = \bar{u}_k^b) = \sigma(b_k)$ in Eq. (5), and E_{kj}^{rev} denotes the reversal potential for synapse w_{kj} . The idea behind Eq. (13) is to match the integrals of individual postsynaptic potentials (PSPs) on v_k and \bar{u}_k . Since the membrane loses any memory following a reset, in contrast to Ref. [7], we use the synaptic conductance as a memory carrier.

The remaining systematic difference to the abstract model lies in the additive—instead of renewing—nature of PSPs elicited by the same presynaptic neuron, which has a noticeable effect in the case of fast consecutive spikes (bursts). Since the membrane potential closely follows the effective potential in the HCS, renewing PSPs can be achieved by using renewing postsynaptic conductances. For this, we have used short-term synaptic depression [22] with a recovery time constant equal to $\tau_{\rm syn}$.

The sampling quality with networks of LIF neurons was examined in computer simulations of BMs, with randomly drawn parameters b_k and W_{kj} [Figs. 3(a) and 3(b)]. The target distribution $p_B(z)$ is approximated by the distribution $p_N(z)$ of network states with $\tau_{on} = 10$ ms. The chosen integration time T = 10 s displays a conservative estimate of the maximum duration a neuronal ensemble will experience stable stimulus conditions in a behaving organism and can thus be expected to sample from a stable target distribution. For this integration

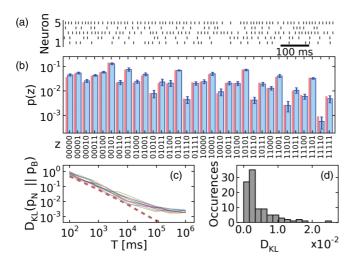


FIG. 3. (a) Spike pattern of a recurrent network of K=5 LIF neurons during sampling from a randomly generated Boltzmann machine. (b) Sampled distribution $p_{\rm N}(z)$ of network states (blue) and target distribution $p_{\rm B}(z)$ (red). $p_{\rm N}(z)$ was estimated from ten 10 -s simulation runs (error bars: std. deviation between runs). (c) $D_{\rm KL}(p_{\rm N}||p_{\rm B})$ as a function of integration time T for 10 trials. The red dotted line shows convergence for the theoretically optimal abstract model [7]. (d) $D_{\rm KL}(p_{\rm N}||p_{\rm B})$ when sampling for $T=10^6$ ms from 100 different randomly generated target distributions.

time, we find that the recurrent network of LIF neurons accurately encodes the target distribution, within the precision imposed by the sample-based representation. The network distribution $p_N(z)$ becomes increasingly more reliable as more samples are considered [Fig. 3(c)]. After a few samples, the network has generated a coarse approximation of $p_B(z)$ that could serve as an "educated guess" in online computation tasks. Only for simulation times T well beyond biologically

relevant timescales do systematic errors in $p_N(z)$ become apparent. The sampling quality holds for a variety of target distributions [Fig. 3(d); see Appendix I for simulation details).

These observations hold for larger-scale networks as well (Fig. 4). For weak coupling, the networks are in an asynchronous irregular state of firing, which enables highly accurate sampling of the target distribution. As synaptic weights increase, the network activity expectedly becomes more synchronous and the sampled distribution overall less accurate, especially in low-probability regions of the state space, but the high-probability modes are always sampled from with high fidelity.

V. DEMONSTRATION OF PROBABILISTIC INFERENCE

We conclude our investigation of sampling in recurrent networks of LIF neurons with an example of Bayesian inference based on incomplete observations. A fully connected BM of K=144 neurons, aligned on a 12×12 grid, was trained as an associative network [23] to store patterns of handwritten digits 0, 3, and 4 in the weights W_{kj} and biases b_k . Each pixel of the image grid was assigned to one network neuron. The resulting joint distribution p(z) displays "prior knowledge" stored by the network.

The probabilistic model was augmented by adding real-valued input channels for each pixel, associated with random variables $y_k \in \mathbb{R}$, $1 \le k \le K$. The resulting generative model p(y, z) has the structure shown in Fig. 5(a) and connects the latent network variables z_k to observable inputs y_k by means of likelihood functions $p(y_k|z_k)$, which we have chosen to be Gaussian with unit variance [Fig. 5(b)]. The likelihoods $p(y_k|z_k)$ tend to align the network state with the observation, i.e., $z_k = 1$ for $y_k > 0$, while the prior p(z) reconciles the observations with knowledge on consistent activation patterns z. The task for the network is to calculate and represent

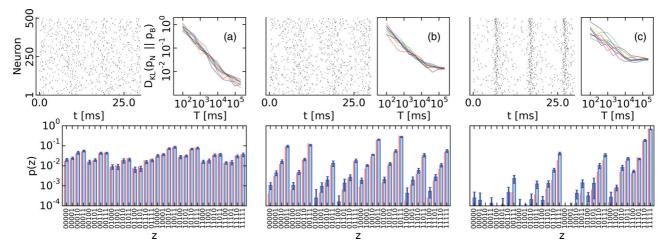


FIG. 4. Sampling in large networks. Spike raster (top left), joint distribution (bottom), and KL divergence (top right) as in Figs. 3(a)–3(c), respectively. Since the target distribution cannot be computed analytically for 500 RVs, we define the Gibbs-sampling estimate after 10^6 steps as a reference distribution. For the joint, the sampled distribution over 5 RVs (out of 500) obtained from the LIF network (after $T_{\rm sim}=10^4$ ms) is plotted alongside the Gibbs estimate. Error bars indicate standard deviation between 10 simulation runs. In all panels, biases were drawn from a β distribution: $b_k \sim 1.2[\mathcal{B}(0.5,0.5)-0.5]$. (a) Small weights: $W_{kj} \sim 0.6[\mathcal{B}(0.5,0.5)-0.5]$. (b) Moderate weights: $W_{kj} \sim 1.2[\mathcal{B}(0.5,0.5)-0.5]$. Systematic deviations between the sampled distribution and the reference manifest mainly in low-probability modes and are due to the difference in PSP shapes between the abstract model and LIF neurons.

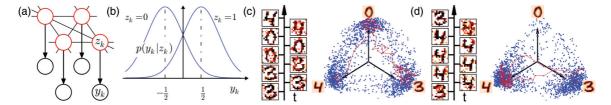


FIG. 5. (a) Graphical model used for the probabilistic inference task. (b) A Gaussian likelihood model provides input to the sampling neurons. (c) Two-dimensional projection of states $z^{(t)}$ when sampling from the prior p(z), which was trained to store hand-written digits 0, 3, and 4. Solid line: network trajectory over 200 ms. Color maps: Marginals of z_k averaged over 25 ms. The time arrow covers the duration of the red trajectory and consecutive snapshots are 25 ms apart. (d) As described in (c) when provided with incomplete input y that is incompatible with digit 0 and ambiguous with respect to digits 3 and 4.

the posterior distribution according to Bayes' rule: $p(z|y) \propto p(z)p(y|z)$. A short derivation (see Appendix F) shows that the posterior p(z|y) is a BM for any input y with the following abstract membrane potential:

$$v_k = b_k + y_k + \sum_j W_{kj} z_j.$$
 (14)

In the LIF domain, the sum $b_k + y_k$ is equivalent to an effective bias [Eq. (I1)] and corresponds to an external current $I_k^{\text{ext}} = I_k^b + I_k^y$ that shifts \bar{u}_k appropriately. Thus, a neuron receives synaptic input from recurrent connections and noise sources, as well as an external current, i.e., $I_k = I_k^{\text{rec}} + I_k^{\text{noise}} + I_k^{\text{ext}}$.

In case of $I_k^y = 0 \ \forall k$, the network samples from the prior distribution p(z) = p(z|y = 0). A two-dimensional projection of network states $z^{(t)} \sim p(z)$ is shown in Fig. 5(c). The sampled distribution has three distinct modes that correspond to the three handwritten digits stored in the recurrent weight matrix. A closer look at the network trajectory reveals that the system stays in one mode ("digit") for some duration, traverses the state space, and then samples from a different mode of the distribution.

A typical inference scenario with incomplete observations is shown in Fig. 5(d). Four input channels at the center were picked to inject positive currents $I_k^y > 0$ to the network while all other inputs remained uninformative. Positive currents I_k^y were chosen such that the observation y appeared incompatible with digit 0, and remained ambiguous with respect to digits 3 and 4 (see Appendix J). In accordance with Bayes' rule, the resulting bimodal posterior distribution p(z|y) has a suppressed 0 mode, but preserves the 3 and 4 modes.

VI. DISCUSSION

We have shown how recurrent networks of conductance-based neurons in a spiking noisy environment can perform probabilistic inference through sampling from a well-defined posterior distribution. Our approach extends Bayesian spiking network implementations to deterministic neuron models widely used in computational neuroscience. We have provided an analytical derivation of the bursty firing response of LIF neurons under Poisson bombardment, which holds for a wide range of parameter regimes (both high and low ratios of $\tau_{\rm syn}/\tau_{\rm eff}$, as well as both high and low $\tau_{\rm ref}$). Our approach is based on the propagation of the membrane autocorrelation throughout bursts and can thereby provide a prediction of the

activation function in a regime where existing approaches [16,17] do not hold (see also Fig. 6 in Appendix K). We have further shown how high-frequency spiking inputs that could be provided by the surrounding network can lead to fast membrane dynamics, which enable individual LIF neurons to correctly encode conditional distributions given information from their presynaptic partners. Thereby, our derivation also identifies a potential functional role of biologically observed high-conductance states and synaptic memory within a Bayesian framework of brain computation.

For mathematical tractability, simplifying modeling assumptions had to be made. The neuron model uses an absolute refractory time τ_{ref} , which matches the activation time constant τ_{on} , and neglects any gradual recovery effects. On the network level, we have assumed statistically independent noise sources and instantaneous axonal transmission. One important difference to cortical structure is the requirement of a symmetric connectivity matrix. The precise symmetry is a consequence of all neurons sharing the same parameters and can be relaxed as neurons become diverse $[\alpha \mapsto \alpha_k \text{ in Eq. (13)}].$ Furthermore, while Dale's principle is known to not hold universally [24], negative coupling in large networks between otherwise excitatory neurons could be, in principle, introduced through populations of inhibitory interneurons. LIF PSPs differ from the theoretically optimal rectangular shape, which could impair convergence to the target distribution outside of the high noise regime [25]. However, computer simulations indicate that in many biologically relevant scenarios the above approximations are not critical (see also Fig. 7 in Appendix L). In particular, the sampling properties of LIF networks remain preserved as network size and interaction strength increase (Fig. 4). Embedding LIF sampling in cortical-size networks appears feasible in light of our results, but remains a matter for future work.

For neuroscientific modeling, our analysis of LIF neurons can be readily transferred to other neuron [26] and synapse models. Beyond neuroscience, the ability to perform probabilistic inference with deterministic neurons displays a promising computing paradigm for neuromorphic hardware systems, which typically implement physical models of integrate-and-fire neurons [27,28]. The distributed nature of the proposed LIF sampling networks allows us to exploit the inherent parallelism of neuromorphic architectures and fosters their application to online data evaluation and robotics. In this context, our results have already provided the basis for the implementation of Bayesian networks [29] and learning [30].

ACKNOWLEDGMENTS

We thank W. Maass for his essential support, as well as S. Habenschuss, M. Brixner, and the three anonymous reviewers for their helpful comments. This research was supported by EU Grants No. 269921 (BrainScaleS), No. 237955 (FACETS-ITN), No. 604102 (Human Brain Project), the Austrian Science Fund FWF Grant No. I753-N23 (PNEUMA), and the Manfred Stärk Foundation.

M.A.P. and J.B. contributed equally to this work.

APPENDIX A: CONDUCTANCE-BASED LIF NEURON MODEL

We restate the set of equations that govern the conductancebased LIF model. For the membrane potential, we have

$$C_m \frac{du}{dt} = g_l(E_l - u) + I, \tag{A1}$$

with capacitance C_m , membrane potential u_k , leak potential E_l , leak conductance g_l , and input current I. The input current I can be formally partitioned into recurrent synaptic inputs, diffuse synaptic noise, and additional "external" currents, i.e.,

$$I = I^{\text{rec}} + I^{\text{noise}} + I^{\text{ext}}.$$
 (A2)

 $I^{\rm ext}$ may represent external current stimuli, average synaptic stimulus currents, or changes in the leak mechanism (i.e., changes in g_l or E_l). The total synaptic current $I^{\rm syn} = I^{\rm rec} + I^{\rm noise}$ obeys the equation

$$I^{\text{syn}} = \sum_{\text{syn}\,i} g_i^{\text{syn}} \left(E_i^{\text{rev}} - u \right),\tag{A3}$$

where $g_i^{\rm syn}$ represents the conductance at the *i*th synapse and $E_i^{\rm rev}$ the corresponding reversal potential. The synaptic conductance obeys the ODE

$$\frac{dg_i^{\text{syn}}}{dt} = -\frac{g_i^{\text{syn}}}{\tau_{\text{syn}}} + \sum_{\text{snk s}} w_i \delta(t - t_s), \tag{A4}$$

with synaptic time constant τ_{syn} and weight w_i . The sum runs over all presynaptic spikes s. The solution to this equation is a superposition of exponentials:

$$g_i^{\text{syn}} = \sum_{\text{spk } s} w_i \Theta(t - t_s) \exp\left(-\frac{t - t_s}{\tau_{\text{syn}}}\right). \tag{A5}$$

Putting all of the above together, we obtain the full ODE for the membrane potential:

$$C_m \frac{du}{dt} = g_l(E_l - u) + \sum_i \sum_{\text{spk } s} w_i \Theta(t - t_s)$$

$$\times \exp\left(-\frac{t - t_s}{\tau_{\text{syn}}}\right) \left(E_i^{\text{rev}} - u\right) + I^{\text{ext}}. \quad (A6)$$

We can now divide the right-hand side of [Eq. (A6)] by $g_{\text{tot}} = g_l + \sum_i g_i^{\text{syn}}$ and rearrange the terms in order to obtain

$$\tau_{\rm eff} \frac{du}{dt} = u_{\rm eff} - u,\tag{A7}$$

with a new effective membrane time constant,

$$\tau_{\rm eff} = \frac{C_m}{g_{\rm tot}},\tag{A8}$$

and effective leak potential,

$$u_{\text{eff}}(t) = \frac{g_l E_l + \sum_i g_i^{\text{syn}}(t) E_i^{\text{rev}} + I^{\text{ext}}}{g_{\text{tot}}(t)}.$$
 (A9)

This transformation is routinely used in studies of conductance-based neurons (see, e.g., Ref. [14]). Here, we have made the time dependencies explicit, since the following HCS approximation will serve to eliminate t in the denominator.

APPENDIX B: THE HIGH-CONDUCTANCE STATE

In a first approximation, assuming a rapidly firing Poisson background $(\nu_{\text{syn}} \to \infty)$, the total average conductance can become arbitrarily large $(\langle g_{\text{tot}} \rangle \to \infty)$, causing the membrane potential to follow the effective potential nearly instantaneously $(\langle \tau_{\text{eff}} \rangle \to 0)$. Equation (A9) can then be rewritten as

$$u \approx u_{\text{eff}} = \frac{g_l E_l + \sum_i \langle g_i^{\text{syn}} \rangle E_i^{\text{rev}} + \sum_i \Delta g_i^{\text{syn}} E_i^{\text{rev}} + I^{\text{ext}}}{\langle g_{\text{tot}} \rangle + \sum_i \Delta g_i^{\text{syn}}},$$
(B1)

where

$$\Delta g_i^{\text{syn}} = g_i^{\text{syn}} - \langle g_i^{\text{syn}} \rangle \tag{B2}$$

denotes the fluctuations of the synaptic conductances.

For a single Poisson source with rate v_i connected to the neuron by a synapse with weight w_i and time constant $\tau_{\rm syn}$, the conductance course can be seen as a sum of independent random variables, each of them representing the conductance change caused by a single spike. In the limit of large v_i , the central limit theorem guarantees the convergence of the conductance distribution to a Gaussian, with moments given by

$$\langle g_i^{\text{syn}} \rangle = \sum_{\text{spk } s} \left\langle w_i \Theta(t - t_s) \exp\left(-\frac{t - t_s}{\tau_{\text{syn}}}\right) \right\rangle$$

$$= \lim_{T \to \infty} \frac{\langle N \rangle}{T} w_i \int_0^T \exp\left(-\frac{t}{\tau_{\text{syn}}}\right) dt$$

$$= w_i v_i \tau_{\text{syn}}, \tag{B3}$$

and

$$\operatorname{Var}\left[g_{i}^{\operatorname{syn}}\right] = \sum_{\operatorname{spk} s} \operatorname{Var}\left[w_{i}\Theta(t - t_{s}) \exp\left(-\frac{t - t_{s}}{\tau_{\operatorname{syn}}}\right)\right]$$

$$= \lim_{T \to \infty} \langle N \rangle \left\{ \left\langle \left[w_{i}\Theta(t - t_{s}) \exp\left(-\frac{t - t_{s}}{\tau_{\operatorname{syn}}}\right)\right]^{2}\right\rangle + \left\langle \left[w_{i}\Theta(t - t_{s}) \exp\left(-\frac{t - t_{s}}{\tau_{\operatorname{syn}}}\right)\right]^{2}\right\}$$

$$= \lim_{T \to \infty} v_{i} T \left\{ \frac{1}{T} w_{i}^{2} \int_{0}^{T} \exp\left(-2\frac{t}{\tau_{\operatorname{syn}}}\right) dt - \frac{1}{T^{2}} \left[\int_{0}^{T} \exp\left(-\frac{t}{\tau_{\operatorname{syn}}}\right) dt\right] \right\}$$

$$= \frac{1}{2} w_{i}^{2} v_{i} \tau_{\operatorname{syn}}, \tag{B4}$$

so the relative fluctuations of g^{syn} are of the order

$$\frac{\sqrt{\text{Var}[g_i^{\text{syn}}]}}{\langle g_i^{\text{syn}} \rangle} = \sqrt{\frac{1}{2 \nu_i \tau_{\text{syn}}}}$$
(B5)

and vanish in the limit of large firing rates.

This warrants an expansion of Eq. (B1) in Δg_i^{syn} , $\forall i$. Considering only the first-order term, we obtain

$$u(t) = \frac{I^{\text{ext}} + g_l E_l + \sum_i g_i^{\text{syn}}(t) E_i^{\text{rev}}}{\langle g_{\text{tot}} \rangle},$$
 (B6)

which renders u simply a linear transformation of the synaptic noise current $J^{\text{syn}} = \sum_{i} g_{i}^{\text{syn}} E_{i}^{\text{rev}}$.

APPENDIX C: DERIVATION OF THE EQUIVALENCE TO AN OU PROCESS

From Eq. (A4), we can find that the synaptic noise J^{syn} obeys the first-order inhomogenous ODE,

$$\frac{dJ^{\text{syn}}}{dt} = -\frac{J^{\text{syn}}}{\tau_{\text{syn}}} + \sum_{\text{syn } i} \sum_{\text{spk } s} \Delta J_i^{\text{syn}} \delta(t - t_s), \qquad (C1)$$

where $\Delta J_i^{\text{syn}} = w_i E_i^{\text{rev}}$. This equation is highly reminiscent of the ODE that defines the OU process:

$$dx(t) = \theta[\mu - x(t)]dt + \sigma dW(t). \tag{C2}$$

It is well-known that the PDF of the OU process,

$$f(x,t|x_0) = \sqrt{\frac{\theta}{\pi\sigma^2(1 - e^{-2\theta t})}}$$

$$\times \exp\left\{\frac{-\theta}{\sigma^2} \left[\frac{(x - \mu + (\mu - x_0)e^{-\theta t})^2}{1 - e^{-2\theta t}}\right]\right\},$$
(C3)

is the unique solution of the Fokker-Planck equation,

$$\frac{1}{\theta} \frac{\partial f(x,t)}{\partial t} = \frac{\partial}{\partial x} [(x - \mu)f] + \frac{\sigma^2}{2\theta} \frac{\partial^2 f}{\partial x^2}, \quad (C4)$$

with starting condition $x_0 := x(t = 0)$. In the following, we prove that, under certain assumptions, the distribution of the synaptic input J^{syn} obeys the same Fokker-Planck equation. To this end, we follow an approach similar to Ref. [15].

Consider the PDF of the synaptic input $f(J^{\text{syn}},t)$. We can use the Chapman-Kolmogorov equation to describe its evolution after a short time interval Δt as an integral over all possible intermediate states J':

$$f(J^{\text{syn}}, t + \Delta t) = \int_{-\infty}^{\infty} f(J^{\text{syn}}, t + \Delta t | J', t) f(J', t) dJ'. \quad (C5)$$

For a small enough Δt , the probability of the occurrence of multiple spikes within Δt can be neglected. As incoming spikes are assumed to be generated by Poisson processes, the probability of a single spike occurring in Δt is $\Delta t \sum_i \nu_i$. By summing over the two possible histories of $J^{\rm syn}$ within Δt (either a single incoming spike or no spike at all), we can use Eq. (C1) to find

$$f(J^{\text{syn}}, t + \Delta t | J') = \left[1 - \Delta t \sum_{i} \nu_{i}\right] \delta \left[J^{\text{syn}} - J' \exp\left(-\frac{\Delta t}{\tau_{\text{syn}}}\right)\right] + \Delta t \sum_{i} \nu_{i} \delta \left[J^{\text{syn}} - \left(J' + \Delta J_{i}^{\text{syn}}\right) \exp\left(-\frac{\Delta t}{\tau_{\text{syn}}}\right)\right], \quad (C6)$$

where v_i represents the afferent firing frequency at the ith synapse. Plugging this into Eq. (C5) and integrating over J' yields

$$f(J^{\text{syn}}, t + \Delta t) = \left(1 - \Delta t \sum_{i} \nu_{i}\right) \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right) f\left[J^{\text{syn}} \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right), t\right] + \Delta t \sum_{i} \nu_{i} \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right) f\left[J^{\text{syn}} \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right) - \Delta J_{i}^{\text{syn}}, t\right].$$
(C7)

We can now expand $f(x,t + \Delta t)$ up to first order in Δt ,

$$f(J^{\text{syn}}, t + \Delta t) \approx f(J^{\text{syn}}, t) + \left. \frac{\partial f(J^{\text{syn}}, t + \Delta t)}{\partial \Delta t} \right|_{\Delta t = 0} \Delta t,$$
 (C8)

and rearrange the terms to obtain

$$\frac{f(J^{\text{syn}}, t + \Delta t) - f(J^{\text{syn}}, t)}{\Delta t} = \frac{\partial f(J^{\text{syn}}, t + \Delta t)}{\partial \Delta t} \Big|_{\Delta t = 0}$$

$$= \left\{ -\sum_{i} \nu_{i} \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right) f \left[J^{\text{syn}} \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right), t \right] + \left(1 - \Delta t \sum_{i} \nu_{i}\right) \frac{1}{\tau_{\text{syn}}} \left\{ \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right) f \left[J^{\text{syn}} \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right), t \right] + \exp\left(2\frac{\Delta t}{\tau_{\text{syn}}}\right) J^{\text{syn}} \frac{\partial f \left[J^{\text{syn}} \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right), t \right]}{\partial J^{\text{syn}} \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right)} \right\} + \sum_{i} \nu_{i} \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right) f \left[J^{\text{syn}} \exp\left(\frac{\Delta t}{\tau_{\text{syn}}}\right) - \Delta J^{\text{syn}}_{i}, t \right] + (\cdots) \Delta t \right\}_{\Delta t = 0}.$$
(C9)

By taking the limit $\Delta t \rightarrow 0$, we obtain

$$\frac{\partial f(J^{\text{syn}},t)}{\partial t} = \frac{1}{\tau_{\text{syn}}} \frac{\partial}{\partial J^{\text{syn}}} [J^{\text{syn}} f(J^{\text{syn}},t)] + \sum_{i} \nu_{i} [f(J^{\text{syn}} - \Delta J_{i}^{\text{syn}},t) - f(J^{\text{syn}},t)]. \tag{C10}$$

In the limit of small synaptic weights (i.e., $\Delta J_i^{\text{syn}} \to 0$), we can expand the second term on the right-hand side up to the second order in ΔJ_i^{syn} . This yields, after some rearrangement,

$$\frac{\partial f(J^{\text{syn}},t)}{\partial t} = \frac{1}{\tau_{\text{syn}}} \frac{\partial}{\partial J^{\text{syn}}} \left[\left(J^{\text{syn}} - \sum_{i} \nu_{i} \Delta J_{i}^{\text{syn}} \tau_{\text{syn}} \right) f(J^{\text{syn}},t) \right] + \frac{\sum_{i} \nu_{i} \Delta J_{i}^{\text{syn}2}}{2} \frac{\partial^{2} f(J^{\text{syn}},t)}{\partial J^{\text{syn}2}}, \tag{C11}$$

which is the exact equivalent of the Fokker-Planck equation of the OU process [Eq. (C4)]. Since u(t) is only a linear transformation of $J^{\text{syn}}(t)$, it can also be approximated by an OU process in the limit of large input frequencies and small synaptic weights, with Eqs. (B6) and (C11) giving the specific time constant, mean value, and variance:

$$\theta = \frac{1}{\tau_{\text{syn}}},\tag{C12}$$

$$\mu = \frac{I^{\text{ext}} + g_l E_l + \sum_i \nu_i w_i E_i^{\text{rev}} \tau_{\text{syn}}}{\langle g_{\text{tot}} \rangle}, \quad (C13)$$

$$\frac{\sigma^2}{2} = \frac{\sum_i \nu_i \left[w_i \left(E_i^{\text{rev}} - \mu \right) \right]^2 \tau_{\text{syn}}}{2 \langle g_{\text{tot}} \rangle^2}.$$
 (C14)

We conclude this section with two important notes. First, for the above methodology to be generally applicable, we must be able to take the limit $\Delta J_i^{\rm syn} \to 0$ for arbitrary first and second moments of $f(J^{\rm syn},t)$ without modifying them. This is possible if at least one excitatory and one inhibitory input is present, which then give us two degrees of freedom with a proper choice of $v^{\rm exc} \to \infty$ and $v^{\rm inh} \to \infty$. Second, all higher moments (3 and above) need to vanish in the abovementioned limit. This has been shown to also be the case under the above conditions [31].

APPENDIX D: DERIVATION OF THE ACTIVATION FUNCTION

We can distinguish between two firing modes of the neuron. The first mode can be classified as "burst spiking" and occurs when multiple spikes occur in rapid succession with an expected ISI of $\langle \Delta t_k \rangle = \tau_{\rm ref} + \langle t_k^{s+1} - t_k^s \rangle_s = \tau_{\rm ref} + \overline{\tau_k^b}$, where $\overline{\tau_k^b}$ represents the average drift time from the reset to the threshold potential following the kth refractory period within a burst. In this case, for each spike within a burst,

$$u_{\text{eff}}(t_s) \geqslant \vartheta,$$
 (D1)

and also, for all but the last spike,

$$u_{\rm eff}(t_s + \tau_{\rm ref}) \geqslant \vartheta.$$
 (D2)

The second mode appears between such bursts, where the membrane potential evolves freely in the subthreshold regime. If we define, just like in the abstract model, that the *k*th neuron is in the state $z_k = 1$ for a duration $\tau_{\rm on} = \tau_{\rm ref}$ following a spike,

we can write

$$p(z=1) = \frac{\sum_{n} P_{n} n \tau_{\text{ref}}}{\sum_{n} P_{n} \left(n \tau_{\text{ref}} + \sum_{k=1}^{n-1} \overline{\tau_{k}^{b}} + T_{n}\right)},$$
 (D3)

where P_n represents the distribution of burst lengths (conditioned on the existence of the first spike) and T_n is the mean time interval between the end of a burst (i.e., the endpoint of its last refractory period) and the next spike. The variables P_n , T_n , and $\overline{\tau_k^b}$ depend on all neuron and noise parameters, but for calculating the activation function (Fig. 2 in the main manuscript and Fig. 6 in the Appendix), we only vary \bar{u} .

We can now calculate both P_n and T_n iteratively. The idea behind this approach is to propagate the membrane potential PDF from spike to spike within a burst and cut off the irrelevant parts for a particular burst length n. We denote the spike times within a burst of length n by t_0, \ldots, t_{n-1} and the endpoint of such a burst by $t_n := t_{n-1} + \tau_{\text{ref}}$. For brevity, we also use $u_i := u(t_i)$. Assuming a first spike at some time t_0 ($u_0 := u(t_0) = \vartheta$), a "burst" of length n = 1 requires a subthreshold free membrane potential after the first refractory period ($u_1 := u(t_0 + \tau_{\text{ref}}) < \vartheta$). This occurs with probability

$$P_{1} := p(u_{1} < \vartheta | u_{0} = \vartheta)$$

$$= \underbrace{\int_{-\infty}^{\vartheta} du_{1} p(u_{1} | u_{0} = \vartheta)}_{\mathcal{T}_{1}}, \tag{D4}$$

where $p(u_{i+1}|u_i) := f(u, \tau_{ref}|u_i)$, which was defined in Eq. (C3). On average, the neuron then stays in the subthreshold regime for a period equal to the mean first passage time from u_1 to ϑ , so the mean duration of the time interval until the onset of the next burst can be expressed as

$$T_1 = \int_{-\infty}^{\vartheta} du_1 p(u_1 | u_0 = \vartheta) \langle T(\vartheta, u_1) \rangle.$$
 (D5)

The first-passage time problem of the OU process has often been discussed [32]. While no closed-form expression for the distribution of first-passage times $T(b,a) = \inf t \ge 0$: x(t) = b|x(0) = a is known, its moments can be computed analytically [18]. In particular, the mean first passage time reads

$$\langle T(b,a)\rangle$$

$$= \frac{\theta}{\sigma} \sqrt{\frac{\pi}{2}} \int_{a}^{b} dx \exp\left[\frac{(x-\mu)^{2}}{2\sigma^{2}}\right] \left[1 + \operatorname{erf}\left(\frac{x-\mu}{\sqrt{2}\sigma}\right)\right]. \quad (D6)$$

A burst of n=2 spikes can only occur when the effective membrane potential lies above the spiking threshold $(u_1 \ge \vartheta)$ after the first refractory period and below after the second $(u_2 < \vartheta)$. This makes P_2 and T_2 recursive functions of P_1 :

$$P_{2} = p(u_{2} < \vartheta, u_{1} \geqslant \vartheta | u_{0} = \vartheta)$$

$$= p(u_{1} \geqslant \vartheta | u_{0} = \vartheta) \ p(u_{2} < \vartheta | u_{1} \geqslant \vartheta, u_{0} = \vartheta)$$

$$\stackrel{[\text{Eq.}(D4)]}{=} \underbrace{(1 - \mathcal{I}_{1})}_{=1 - P_{1}} \underbrace{\int_{\vartheta}^{\vartheta} du_{1} p(u_{1} | u_{1} \geqslant \vartheta)}_{\mathcal{I}_{2}} \underbrace{\int_{-\infty}^{\vartheta} du_{2} p(u_{2} | u_{1})}_{\mathcal{I}_{2}}$$

$$T_{2} = \int_{\vartheta}^{\infty} du_{1} p(u_{1} | u_{1} \geqslant \vartheta)$$

$$(D7)$$

where $p(u_i|u_i \ge \vartheta)$ is a shorthand notation for $p(u_i|u_i \ge \vartheta, u_{i-1} \ge \vartheta, \dots, u_1 \ge \vartheta, u_0 = \vartheta)$. In particular, this represents a renormalization of the PDF of the effective membrane potential to values above the spiking threshold after i refractory periods.

 $\times \left[\int_{-\infty}^{\vartheta} du_2 p(u_2|u_2 > \vartheta, u_1) \langle T(u_2, \vartheta) \rangle \right],$

(D8)

We can now continue this recursion up to an arbitrary burst length and write

$$P_{n} = p(u_{n} < \vartheta, u_{n-1} \geqslant \vartheta, \dots, u_{1} \geqslant \vartheta | u_{0} = \vartheta)$$

$$= p(u_{1} \geqslant \vartheta | u_{0} = \vartheta)$$

$$\times p(u_{n} < \vartheta, u_{n-1} \geqslant \vartheta, \dots, u_{2} \geqslant \vartheta | u_{1} \geqslant \vartheta, u_{0} = \vartheta)$$

$$\stackrel{[Eq.(D4)]}{=} (1 - \mathcal{I}_{1}) \ p(u_{2} \geqslant \vartheta | u_{1} \geqslant \vartheta, u_{0} = \vartheta)$$

$$\times p(u_{n} < \vartheta, u_{n-1} \geqslant \vartheta, \dots, u_{3} \geqslant \vartheta |$$

$$\times |u_{2} \geqslant \vartheta, u_{1} \geqslant \vartheta, u_{0} = \vartheta)$$

$$\stackrel{[Eq.(D7)]}{=} (1 - \mathcal{I}_{1})(1 - \mathcal{I}_{2}) \ p(u_{3} \geqslant \vartheta | u_{2} \geqslant \vartheta, u_{1} \geqslant \vartheta, u_{0} = \vartheta)$$

$$\times p(u_{n} < \vartheta, u_{n-1} \geqslant \vartheta, \dots, u_{4} \geqslant \vartheta |$$

$$\times |u_{3} \geqslant \vartheta, \dots, u_{1} \geqslant \vartheta, u_{0} = \vartheta)$$

$$= \prod_{i=1}^{n-1} (1 - \mathcal{I}_{i}) \ p(u_{n} < \vartheta | u_{n-1} \geqslant \vartheta, \dots, u_{1} \geqslant \vartheta, u_{0} = \vartheta)$$

$$(D9)$$

$$= \left(1 - \sum_{i=1}^{n-1} P_i\right)$$

$$\times \underbrace{\int_{\vartheta}^{\infty} du_{n-1} p(u_{n-1}|u_{n-1} \geq \vartheta) \left[\int_{-\infty}^{\vartheta} du_n p(u_n|u_{n-1})\right]}_{\mathcal{I}_n}$$
(D10)

$$T_{n} = \int_{\vartheta}^{\infty} du_{n-1} p(u_{n-1}|u_{n-1}|\geqslant \vartheta)$$

$$\times \left[\int_{-\infty}^{\vartheta} du_{n} p(u_{n}|u_{n} < \vartheta, u_{n-1}) \langle T(u_{n}, \vartheta) \rangle \right]. \quad (D11)$$

The transition from a product to a sum between Eqs. (D9) and (D10) requires the identity

$$\prod_{i=1}^{n-1} (1 - \mathcal{I}_i) = 1 - \sum_{i=1}^{n-1} P_i,$$
 (D12)

which can be easily shown by induction from $P_n = \mathcal{I}_n \prod_{i=1}^{n-1} (1 - \mathcal{I}_i)$ [Eq. (D9)] and $P_1 = \mathcal{I}_1$ [Eq. (D4)]. Since $\lim_{n \to \infty} P_n = 0$, one can stop the recursion at some small enough P_n .

What remains to be calculated is the average time to threshold $\overline{\tau_k^b}$ within a burst that follows the kth refractory period. Since we assume a HCS, we are looking at a regime in which $\tau_{\rm eff} \ll \tau_{\rm syn}$. Therefore, we can assume $u_{\rm eff}$ to be approximately unchanged during the short time interval $\overline{\tau_k^b}$ (adiabatic approximation, see also Ref. [17]). For a fixed u_k , the jump time can be easily calculated from Eq. (A7):

$$\tau_k^{\mathsf{b}}(u_k) = \ln\left(\frac{\varrho - u_k}{\vartheta - u_k}\right).$$
(D13)

The average jump time can then be obtained by integrating over all suprathreshold values of u_k , which in turn have probabilities that follow from integrating over all suprathreshold values of u_{k-1} :

$$\overline{\tau_k^b} = \int_{\vartheta}^{\infty} du_k \ln\left(\frac{\varrho - u_k}{\vartheta - u_k}\right) \int_{\vartheta}^{\infty} du_{k-1} p(u_k | u_k > \vartheta, u_{k-1}).$$
(D14)

With Eqs. (D3), (D10), (D11), (D14), and (D6), one could now predict the activation function of an LIF unit in an extreme high-noise regime ($\tau_{eff} \rightarrow 0$). We can, however, generalize our approach by taking the finite nature of the effective time constant into account.

If we go back to Eq. (A7) and leave $\tau_{\rm eff} = C/\langle g_{\rm tot} \rangle$ small but finite, we can still perform all the remaining approximations, but are required to modify Eq. (B6):

$$\tau_{\text{eff}}\dot{u}(t) = \frac{I^{\text{ext}} + g_l E_l}{\langle g_{\text{tot}} \rangle} + \frac{J^{\text{syn}}(t)}{\langle g_{\text{tot}} \rangle} - u(t). \tag{D15}$$

Together with Eq. (C1), we now have a system of first-order ODEs, which can be solved analytically by standard techniques (variation of constants). The PSPs are then no longer a linear transformation of the exponentially shaped PSCs, but rather α -shaped (more precisely, a difference of exponentials):

$$u_s(t) = \Theta(t - t_s) A \frac{\left(e^{-\frac{t - t_s}{\tau_{\text{eff}}}} - e^{-\frac{t - t_s}{\tau_{\text{syn}}}}\right)}{\tau_{\text{eff}} - \tau_{\text{syn}}}, \tag{D16}$$

with $A = \frac{w_i(E_i^{\rm rev} - \langle u_{\rm eff} \rangle) \tau_{\rm syn_i}}{\langle g_{\rm tot} \rangle}$. This shape causes a lower PSP peak than in the case of exponential PSPs, decreasing the overall width of the membrane potential distribution. Intuitively speaking, this results in a horizontal shift and compression of the activation function.

More recently, analytical treatments of these phenomena have been proposed [33]. In these approaches, large membrane time constants (equivalent to a long $\tau_{\rm eff}$) and small synaptic time constants are usually considered. However, Eq. (D16) is symmetric in $\tau_{\rm eff}$ and $\tau_{\rm syn}$, so the same argument applies to our case as well, but the two time constants need to be switched. It

is, for example, possible to correct the first passage time from the reset to the threshold potential by using an expansion in $\sqrt{\tau'/\tau}$ (with τ' and τ being the smaller and the larger of the two time constants, respectively) [16]:

$$\langle T(\vartheta, u) \rangle = \tau \sqrt{\pi} \int_{\frac{u-\mu}{\sigma}}^{\frac{\vartheta_{\text{eff}} - \mu}{\sigma}} dx \exp(x^2) [\text{erf}(x) + 1], \quad (D17)$$

with μ and σ^2 the first two moments of the free membrane potential distribution and an effective threshold,

$$\vartheta_{\rm eff} \approx \vartheta - \zeta \left(\frac{1}{2}\right) \sqrt{\frac{\tau'}{2\tau}} \sigma,$$
 (D18)

in which ζ denotes the Riemann zeta function. In our particular case, the expansion is done in $\sqrt{\tau_{\rm eff}/\tau_{\rm syn}}$, so $\tau' = \tau_{\rm eff}$ and $\tau = \tau_{\rm syn}$. With this approximation, we assume that u converges from ρ to $u_{\rm eff}$ in negligible time after it is released from the refractory state. Afterwards, its convergence to $u_{\rm eff}$ is determined by Eq. (A7). Note how Eq. (D17) is equivalent to a change of the integration variable and limits in the original equation [Eq. (D6)] for the first passage time.

APPENDIX E: TRANSLATION OF SYNAPTIC WEIGHTS

A sufficient condition for a single neuron to sample from the correct conditional distribution is given by its activation function:

$$p(z_k = 1|z_{\setminus k}) = \sigma(v_k). \tag{E1}$$

The relationship between the abstract model and the LIF implementation is defined by the lateral dilation α and relative offset \bar{u}_k^0 of the LIF activation function:

$$p(z_k = 1 | z_{\setminus k}) = \sigma\left(\frac{\bar{u}_k - \bar{u}_k^0}{\alpha}\right). \tag{E2}$$

The parameters α and \bar{u}_k^0 can be determined by fitting Eq. (E2) either to simulation results or to the theoretical prediction [Eq. (D3)]. As a consequence, also synaptic weights need to be rescaled by the factor α . Additionally, the difference in PSP shapes needs to be taken into account.

We choose a translation rule in which we set the LIF synaptic weights w_{ij} such that the area under a PSP [Eq. (D16)] during the refractory state of the corresponding afferent neuron (i.e., for a duration τ_{ref}) is equal to W_{ij} τ_{ref} α :

$$W_{kj}\tau_{\text{ref}}\alpha = \int_{0}^{\tau_{\text{ref}}} \frac{w_{kj} \left(E_{kj}^{\text{rev}} - \langle u_{\text{eff}} \rangle\right) \tau_{\text{syn}}}{\langle g_{\text{tot}} \rangle}$$

$$\times \frac{\exp\left(-\frac{t-t_{s}}{\tau_{\text{eff}}}\right) - \exp\left(-\frac{t-t_{s}}{\tau_{\text{syn}}}\right)}{\tau_{\text{eff}} - \tau_{\text{syn}}} dt$$

$$= \frac{w_{kj}\tau_{\text{syn}}}{\langle g_{\text{tot}} \rangle} \frac{\left(E_{kj}^{\text{rev}} - \mu\right)}{\tau_{\text{eff}} - \tau_{\text{syn}}}$$

$$\times \left[\tau_{\text{syn}} \left(e^{-\frac{\tau_{\text{ref}}}{\tau_{\text{syn}}}} - 1\right) - \tau_{\text{eff}} \left(e^{-\frac{\tau_{\text{ref}}}{\tau_{\text{eff}}}} - 1\right)\right]. \quad (E3)$$

By setting $\tau_{ref} = \tau_{syn}$, we obtain the mapping between the abstract and LIF synaptic weight domains:

$$W_{kj} = \frac{1}{\alpha C_m} \frac{w_{kj} \left(E_{kj}^{\text{rev}} - \mu\right)}{1 - \frac{\tau_{\text{syn}}}{\tau_{\text{eff}}}}$$
$$\times \left[\tau_{\text{syn}} (e^{-1} - 1) - \tau_{\text{eff}} (e^{-\frac{\tau_{\text{syn}}}{\tau_{\text{eff}}}} - 1)\right]. \tag{E4}$$

Additionally, depressing short-term plasticity [22] has been applied to attenuate the amplitudes of consecutively arriving α -shaped PSPs from a network neuron and emulate renewing synapses. In particular, within the Tsodyks-Markram short-term plasticity model [22], the synaptic efficacy parameter and recovery time constant have been chosen as $U_{\rm SE}=1$ and $\tau_{\rm rec}=\tau_{\rm syn}$, respectively.

Analogously to the weights, the biases can be determined from the condition $\sigma(b_k) = \sigma([\bar{u}_k^b - \bar{u}_k^0]/\alpha)$ in the absence of recurrent activity, i.e., $I_k^{\rm rec} = 0$.

APPENDIX F: PROBABILISTIC MODEL FOR DEMONSTRATION OF INFERENCE

We define a joint model p(y,z) over real-valued input nodes $y = (y_1, \ldots, y_K)$ and binary latent variables $z = (z_1, \ldots, z_K)$ as sketched in Fig. 4(a) of the main manuscript. The real-valued variables y_k encode intensities of the input pixels. The latent variables z_k correspond to neurons in the network. Our aim is to demonstrate that the network can sample from the posterior distribution p(z|y).

For this example, we have chosen a particularly simple likelihood function, namely a Gaussian emission model with variance $\sigma^2=1$ and mean values $\mu=\pm\frac{1}{2}$. From the graphical model we identify the structure of the joint distribution,

$$p(y,z) = p(z) \prod_{k=1}^{K} p(y_k|z_k),$$
 (F1)

where p(z) is a Boltzmann distribution and the likelihood is defined by

$$p(y_k|z_k) = \mathcal{N}(y_k; \mu = 1/2)^{z_k} \mathcal{N}(y_k; \mu = -1/2)^{1-z_k},$$
 (F2)

which is equivalent to

$$\log p(y_k|z_k) = z_k [\log \mathcal{N}(y_k; 1/2) - \log \mathcal{N}(y_k; -1/2)] + \log \mathcal{N}(y_k; -1/2)$$

= $z_k y_k + \log \mathcal{N}(y_k; -1/2)$, (F3)

using normal distributions

$$\mathcal{N}(y\mu[;\mu]) := \exp[-(y-\mu)^2/2]/\sqrt{2\pi},$$
 (F4)

with unit variance. The posterior of this model reads

$$p(z|y) = \frac{p(y,z)}{p(y)}$$

$$= \exp\left(\frac{1}{2}z^{\mathsf{T}}Wz + z^{\mathsf{T}}(b+y)\right) / \text{Norm}, \quad (F5)$$

where the normalizing constant depends on the input y, but is independent of the network variables and thus defines a Boltzmann distribution over z for any y. In particular, we identify the abstract membrane potential

TABLE I. Neuron parameters used for the simulations in the main manuscript.

C_m	0.1 nF	Membrane capacitance
g_l	5 nS	Leak conductance
E_l	-65 mV	Leak potential
ρ	-53 mV	Reset potential
$E_{ m exc}^{ m rev}$	0 mV	Excitatory reversal potential
$E_{ m inh}^{ m rev}$	$-90~\mathrm{mV}$	Inhibitory reversal potential
ϑ	-52 mV	Threshold voltage
$ au_{ m syn}$	10 ms	Synaptic time constant
$ au_{ m ref}$	10 ms	Refractory time constant

 $v_k = b_k + y_k + \sum_j W_{kj} z_j$ for sampling from p(z|y) by means of a spiking network.

APPENDIX G: SIMULATION PARAMETERS

All simulations have been performed with the NEURON simulation package [34] and the PyNN API [35], with a time step of dt = 0.01 ms. For the LIF neuron, we have chosen the following parameters (Table I; compare with, e.g., Ref. [36] for parameters fitted to experimental data):

Synaptic noise was implemented as bombardment by inhibitory and excitatory Poisson stimuli with rates $v_{inh} =$ $v_{\rm exc} = 5000 \, \rm Hz$. The excitatory synaptic weight for the noise stimuli was set to $w_{\rm exc}^{\rm noise} = 0.0035 \,\mu{\rm S}$. The inhibitory weight $w_{\rm inh}^{\rm noise}$ was adjusted as to yield $p(z_k=1)\approx 0.5$ with no current stimulus present. For above parameters, this happens at an average free membrane potential of $\bar{u} = -55 \,\mathrm{mV}$. This determines $w_{\rm inh}^{\rm noise}$ according to

$$\left| \frac{E_{\text{inh}}^{\text{rev}} - \bar{u}}{E_{\text{exc}}^{\text{rev}} - \bar{u}} \right| = \frac{w_{\text{exc}}^{\text{noise}}}{w_{\text{inh}}^{\text{noise}}}.$$
 (G1)

APPENDIX H: THE ACTIVATION FUNCTION OF LIF NEURONS IN A SPIKING NOISY ENVIRONMENT (FIG. 2 IN MAIN MANUSCRIPT)

In order to sweep through the activation function, the external current I^{ext} was varied. However, in order to facilitate a comparison with the logistic activation function of the abstract model, we have represented p(z = 1) as a function of \bar{u} instead. The latter is equivalent to the mean μ of the corresponding Ornstein-Uhlenbeck process, with Eq. (C13) allowing a direct translation between I^{ext} and \bar{u} .

The abscissa values in Fig. 2(d) in the main manuscript represent averages of the free membrane potential obtained from 10 simulation runs with a total duration of $T_{\rm sim}=100\,{\rm s}$ and firing threshold θ set to $E_{\rm exc}^{\rm rev}=0\,{\rm mV}$. The deviations from the theoretical prediction [Eq. (D3)] are smaller than the size of the symbols, therefore no error bars are shown.

The ordinate values and standard errors were calculated from the simulated spike train data according to

$$p(z=1) = \frac{1}{N} \sum_{i=1}^{N} p_i,$$
 (H1)

$$p(z=1) = \frac{1}{N} \sum_{i=1}^{N} p_i,$$
 (H1)
$$s = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N} [p_i - p(z=1)]^2},$$
 (H2)

with $p_i = \frac{N_i^{\rm spk} \tau_{\rm on}}{T_{\rm sim}}$ being the fraction of time spent in z=1 and N_i^{spk} representing the total number of spikes in the *i*th out of N=10 performed simulations. Since the respective standard errors of the mean are smaller than the size of the symbols, no error bars are shown.

APPENDIX I: SAMPLING VIA RECURRENT NETWORKS OF LIF NEURONS (FIGS. 3 AND 4 IN MAIN MANUSCRIPT)

The simulated network consists of K = 5 neurons with a synaptic weight matrix W and a bias vector b (both in the Boltzmann domain). All entries were drawn from a β distribution $\mathcal{B}(0.5,0.5)$ and mapped linearly to the interval [-0.6,0.6]. More specifically, $b_k, W_{kj} \sim 1.2[\mathcal{B}(0.5, 0.5) - 0.5]$. The parameters and mapping of the beta distribution were chosen with the intent of generating diverse distributions, spanning multiple orders of magnitude. The bias b_k , defined in the Boltzmann domain, determines the probability $p(z_k = 1 | z_{\setminus k} = \mathbf{0})$ for neuron k. In the LIF domain, the probability $p(z_k = 1 | z_{\setminus k} = \mathbf{0}) = 0.5$ corresponds to the mean free membrane potential \bar{u}_k^0 . Then, a nonzero bias can be described in the LIF domain as a linear shift from \bar{u}_k^0 to a mean membrane potential \bar{u}_k^b . This yields the linear transformation

$$b_k = (\bar{u}_k^b - \bar{u}_k^0)/\alpha, \tag{I1}$$

where α represents the scaling factor between the two domains. Both quantities \bar{u}_k^0 and α can be determined from the predicted activation function of a single LIF unit. The first quantity constitutes the inflection point of the activation function [at $p(z_k = 1 | z_{\setminus k} = \mathbf{0}) = 0.5$], the latter follows from the slope of

By computing \bar{u}_k^b , we can map any bias b_k of a single unit of the Boltzmann machine onto a yet unconnected LIF neuron. In simulations, \bar{u}_k^b was established by injecting a temporally constant external current I_{ext}^k according to

$$I_k^{\text{ext}} = \left(\alpha b_k + \bar{u}_k^0\right) \langle g_{\text{tot}} \rangle - g_l E_l - \sum_i \nu_i w_i^{\text{noise}} E_i^{\text{rev}} \tau_{\text{syn}}. \quad (12)$$

In order to achieve sampling network dynamics in the LIF domain faithful to those displayed by an equivalent Boltzmann machine, the Boltzmann weight matrix W was translated into LIF network weights w_{ij} according to Eq. (E4). Thus, superposing PSPs saturate the membrane potential, approximating the constant amplitude of a PSP in the abstract neuron model.

For Fig. 3(b) in the main manuscript, this setup of a random Boltzmann machine was simulated N = 10 times with different random seeds for the Poisson background for a duration of $T_{\text{sim}} = 10 \text{ s}$. The red bars show the analytically computed target joint distribution $p_{\rm B}(z)$. The blue bars depict the network distribution $p_N(z)$, calculated from the firing activity of the simulated LIF network set up to match $p_{\rm B}(z)$. The means and error bars have been calculated as in Eqs. (H1) and (H2), respectively.

The above simulations were repeated with a significantly longer duration in order to study systematic deviations due to the LIF implementation. Figure 3(c) in the main manuscript shows the distance between the target distribution $p_{\rm B}(z)$ and its LIF network representation $p_{\rm N}(z)$ in form of the Kullback-Leibler divergence,

$$D_{KL}(p_N||p_B) = \langle \log[p_N(z) / p_B(z)] \rangle_{p_N(z)}. \tag{I3}$$

This estimate has been taken for one set of parameters (W, b) for ten independent trials (thin lines) in an LIF network at integration times $T: 0 \le T \le T_{\rm sim} = 10^6$ ms. The red dashed line displays the averaged $D_{\rm KL}(p_{\rm N}||p_{\rm B})$ for the abstract network model with identical parameters (W, b). The decrease of $D_{\rm KL}(p_{\rm N}||p_{\rm B})$ for longer integration times indicates the increasing precision of the sampling network over time. Eventually, the $D_{\rm KL}(p_{\rm N}||p_{\rm B})$ converges for the LIF network to a nonzero value, reflecting small systematic errors. Figure 3(d) in the main manuscript shows the distribution of $D_{\rm KL}(p_{\rm N}||p_{\rm B})$ values for 100 randomly drawn Boltzmann machines emulated by LIF networks, evaluated from a single run of $T_{\rm sim}=10^6$ ms each.

Figure 4 was generated in the same way as Fig. 3, except for the network size and weight distribution.

APPENDIX J: DEMONSTRATION OF PROBABILISTIC INFERENCE (FIG. 5 IN MAIN MANUSCRIPT)

We trained a fully visible, fully connected Boltzmann machine to store three handwritten digits (0, 3, 4) that were taken from the MNIST data set [37] and were scaled down to 12×12 pixels. The pixel intensities of these patterns (ranging from 0 to 1) were linearly scaled to an activation between 0.05 and 0.95, defining the target statistics $\langle z_k \rangle_T$ and $\langle z_k z_j \rangle_T$ for the Boltzmann machine. Then, external currents I_k^b and synaptic weights w_{kj} were optimized via the update rules $\Delta I_k^b \propto \langle z_k \rangle_T - z_k$ and $\Delta w_{kj} \propto \langle z_k z_j \rangle_T - z_k z_j$, with samples z obtained through sampling from an LIF network set up with synaptic noise and neuron parameters as described above.

The recurrent connections, defined by w_{kj} , induce synaptic currents $I_k^{\rm rec}$ in addition to the noise currents $I_k^{\rm noise}$. Additionally, each neuron's mean effective membrane potential Eq. (C13) is shifted by external currents of the trained quantities I_k^b as well as input currents I_k^y . The latter encode observations y, which are defined in the context of the probabilistic model described in AVI. Hence, the total received current of neuron k in the network amounts to

$$I_k = I_k^{\text{rec}} + I_k^{\text{noise}} + I_k^b + I_k^y.$$
 (J1)

The resulting samples $z^{(t)}$ displayed in subplot Figs. 5(c) and 5(d) were taken for 4000 ms, after a burn-in time that ensured that the network had converged to its equilibrium distribution.

1. Figure 5(c), main manuscript: Sampling from the prior

From the probabilistic model it follows that p(z|y=0) = p(z). This means that the LIF network will sample from the prior when $I_k^y = 0$, $\forall k$. The figure makes use of two projections of network states $z^{(t)}$ to illustrate the sampled distribution:

(1) Star plot: In order to illustrate the 12×12 -dimensional network states $z^{(t)}$, a two-dimensional linear projection in a

star plot has been chosen (blue dots). The axes indicate the three basis vectors \mathbf{B} , representing pixel intensities $\langle z_k \rangle_T$ of the digits (0, 3, 4),

$$\langle z_k \rangle_T^{034} = (\boldsymbol{B}^0, \boldsymbol{B}^3, \boldsymbol{B}^4)^\mathsf{T}, \tag{J2}$$

with a total intensity normalization $||\boldsymbol{B}^{i}|| = \sqrt{\sum_{j} |B_{j}^{i}|^{2}} = 1$.

The network states $z^{(t)}$ acquired from the simulation are projected onto this basis:

$$z^{034}(t) = (\mathbf{B}^0 z^{(t)}, \mathbf{B}^3 z^{(t)}, \mathbf{B}^4 z^{(t)})^{\mathsf{T}}.$$
 (J3)

This three-dimensional vector is then projected onto a twodimensional plane in coordinates

$$z^{\text{proj}}(t) = \begin{pmatrix} \sin\left(\phi_{\text{B}}^{0}\right) & \sin\left(\phi_{\text{B}}^{3}\right) & \sin\left(\phi_{\text{B}}^{4}\right) \\ \cos\left(\phi_{\text{B}}^{0}\right) & \cos\left(\phi_{\text{B}}^{3}\right) & \cos\left(\phi_{\text{B}}^{4}\right) \end{pmatrix} z^{034}(t), \quad (\text{J4})$$

with $(\phi_B^0,\phi_B^3,\phi_B^4)=(0,\frac{2\pi}{3},\frac{4\pi}{3})$ indicating the directions of the normalized basis vectors.

These linear projections $z^{\text{proj}}(t)$ of network states $z^{(t)}$ are used to illustrate similarity of states as distance in the two-dimensional plane. In Fig. 5(c) in the main manuscript, a network evolution time of 4000 ms is shown, samples $z^{(t)}$ taken every 2 ms; i.e., in total 2000 projected states are displayed. We ensured in longer simulations that the total simulation runtime is sufficient to represent the distribution under the mixing of the Markov chain.

The significant clustering of the network states around the directions of the arrows indicates the proximity to the target states (0, 3, 4) for a majority of time.

The transitions of the network states $z^{(t)}$ are depicted as a red trajectory in the star plot. This trajectory connects 100 projected network states within a time interval of 200 ms, demonstrating the time evolution of $z^{\text{proj}}(t)$.

(2) Snapshots: The squared color maps on the time axis display the averaged pixel intensity of the network in a time window 25 ms:

$$\bar{z}_k(t) = \frac{1}{25 \,\mathrm{ms}} \int_{t-12.5 \,\mathrm{ms}}^{t+12.5 \,\mathrm{ms}} z_k^{(t')} dt'.$$
 (J5)

The equidistant time intervals between these snapshots were taken every 25 ms within the time frame of the above mentioned red trajectory of 200 ms.

In summary, the network spends most of the time in distinct modes (digits) and only little time in "blurred" states that could not be clearly assigned to one digit. Furthermore, it spends an approximately equal amount of time in each mode, indicating that the prior p(z) does not favor one of the three digits.

2. Figure 5(d), main manuscript: Sampling from the posterior

Incomplete and ambiguous input y was provided to the network by setting four of the inputs different from zero: $y_k \neq 0$ for $k \in \mathcal{I}$ with \mathcal{I} denoting an index set and $|\mathcal{I}| = 4$. These inputs were chosen at the center of the image where digits 3 and 4 both have black pixels, while digit 0 is white. More precisely, $\mathcal{I} = \{77,78,79,80\}$ when pixels are indexed row-wise starting from the top-left corner. In the LIF implementation, positive values y_k correspond to positive currents I_k^y , as specified in

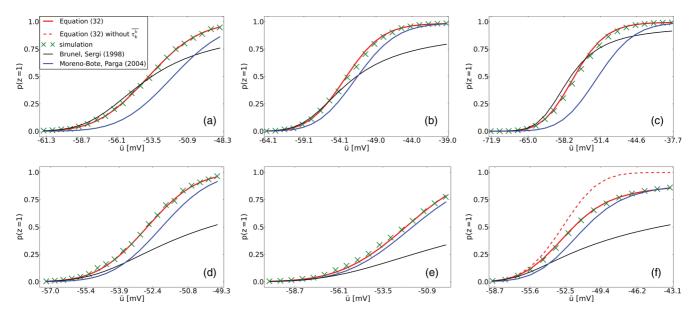


FIG. 6. Comparison of our prediction of the activation function [Eq. (D3)] to simulation data, as well as to the predictions given by Refs. [16] and [17], for several parameter sets. (a) Standard parameter set as given in Table I. This figure shows the same curves as Fig. 2(b) in the main manuscript. (b) Same as described for (a) but with quadrupled membrane capacitance C_m and one quarter of the leak conductance g_l , input rates $v_{\text{exc,inh}}$ and weights $w_{\text{exc,inh}}$. This parameter set is identical to the one used for the top left corner of Fig. 7 and has the effect of slowing the membrane, i.e., increasing τ_{eff} by a factor of 16. (c) Same as described for (a) but with a decreased synaptic time constant $\tau_{\text{syn}} = 3$ ms. The prediction from Ref. [16] is improved, since the correlations in the pre- and post-refractory effective membrane potential are smaller in this scenario. (d) Same as described for (a) but with an increased synaptic time constant $\tau_{\text{syn}} = 30$ ms. The prediction from Ref. [16] deteriorates due to the longer-range membrane potential autocorrelation. Conversely, the prediction from Ref. [17] improves, since the refractory time becomes less important. (e) Same as described for (a) but with a very short refractory time $\tau_{\text{ref}} = 1$ ms. Here, we enter the parameter range where Ref. [17] provides good predictions. (f) Same as described for (a) but with the input rates $v_{\text{exc,inh}}$ and weights $w_{\text{exc,inh}}$ decreased by a factor of 10, thereby slowing the membrane considerably (imperfect HCS). Additionally, we have chosen a large reset-to-threshold distance of $\vartheta - \varrho = 10$ mV. In this scenario, the $\overline{\tau_k^b}$ -term in Eq. (D3) becomes dominant and the activation function departs from the logistic shape that it has in the HCS.

Eq. (I2). We set $I_k^y = 0.831$ nA for the four nonzero inputs, inducing an effective bias $b_k + y_k$ by injecting a total current $I_k^b + I_k^y$. For the plot, the same projections of network states were used as in Fig. 5(c) in the main manuscript.

Under the incomplete input, the network spends only little time in the 0-mode or "blurred" states, while the equilibrium distribution exhibits two distinct modes in the 3- and 4-directions. Thus, the posterior reflects both the almost certain conclusion that "the input is not a zero" and the uncertainty that "the input could either be a three or a four." In particular, the network response is well-suited for further processing (e.g., by other cortical populations or in a technical application). For instance, the network states could be integrated by a linear classifier to recognize the digit class.

APPENDIX K: PREDICTION OF THE ACTIVATION FUNCTION FOR VARIOUS PARAMETER SETS

We have used Eq. (D3) for predicting the activation functions of LIF neurons in the HCS regime. In this regime, the refractory time τ_{ref} and the synaptic time constant τ_{syn} become the dominant time constants:

$$\tau_{\rm eff} \ll \tau_{\rm ref} \approx \tau_{\rm syn}.$$
 (K1)

So far, we have compared our prediction to simulation data, as well as to two other predictions by Refs. [16] and [17] in Fig. 2(b) of the main manuscript. Here, we depart from these assumptions and show in Fig. 6 that our prediction holds for several different parameter sets. In particular, the predictions from Refs. [16] and [17] are only valid when either synaptic time constants [Fig. 6(c)] or refractory times [Fig. 6(e)] become shorter.

APPENDIX L: VALIDITY OF THE LIF SAMPLING FRAMEWORK FOR VARIOUS PARAMETER SETS

For our various network simulations (Figs. 3–5 in the main manuscript) we have used a set of biologically plausible parameters from Ref. [36] (see also Appendix G: Simulation parameters). In particular, we have used high input noise rates and weights in order to achieve the HCS, as well as a small distance between firing threshold and reset, in order to reduce $\overline{\tau_k^b}$, during which a neuron falsely encodes the state z=0. Furthermore, we have assumed $\tau_{\rm ref}=\tau_{\rm syn}$. Figure 7 shows that these are by no means strict constraints.

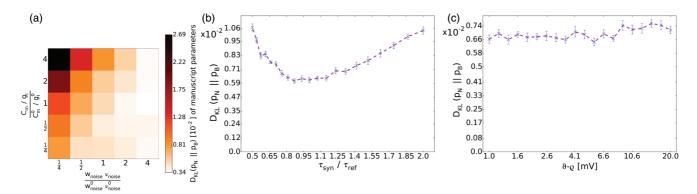


FIG. 7. Sampling with LIF neurons over a broad range of relevant model parameters. All plots depict the Kullback-Leibler divergence between the sampled distribution (with a certain set of parameters) and the target distribution, which is identical to the one used in Fig. 3(b) in the main manuscript. (a) Sweep over neuron size and leakage, as well as background input parameters. The axes represent multiplicative scaling values for four parameters: background (Poisson) synaptic weights $w_{\text{exc,inh}}$ and firing rates $v_{\text{exc,inh}}$ on the abscisa, neuron capacitance C_m and leak conductance g_l on the ordinate. The parameter values used throughout the main manuscript (see also Appendix G: Simulation parameters) therefore have the coordinates (1,1). The network simulation run times were chosen as $T_{\text{sim}} = 10^6$ ms. As expected, the large neuron and weak noise scenario (top left square) does not permit accurate sampling, as the activation function is no longer logistic [see also Fig. 6(f)]. In general, the plot shows that good sampling quality can be achieved for any neuron capacitance and leak as long as the background noise is strong enough (HCS). (b) Sweep over the ratio of the synaptic and refractory time constants. The best sampling performance is, indeed, achieved for $\tau_{\text{syn}} \approx \tau_{\text{ref}}$, but the network still produces good approximations of the target distribution when the two time constants are not precisely identical. All D_{KL} data points result from 20 simulations with run times of $T_{\text{sim}} = 10^5$ ms each. The error bars represent the standard error of the mean. (c) Effect of an increased distance from reset to threshold potential. An increase in $\vartheta - \varrho$ causes a gradual decay of the sampling quality, since the membrane requires additional time to reach a suprathreshold u_{eff} when the refractory period is over. This can, in principle, be accommodated by defining a larger time window τ_{on} during which the neuron is considered to encode the state z = 1. Neverth

- [1] K. Körding and D. Wolpert, Nature **427**, 244 (2004).
- [2] J. Fiser, P. Berkes, G. Orbán, and M. Lengyel, Trends Cogn. Sci. 14, 119 (2010).
- [3] K. Friston, J. Mattout, and J. Kilner, Biol. Cybern. 104, 137 (2011).
- [4] T. Yang and M. N. Shadlen, Nature 447, 1075 (2007).
- [5] P. Berkes, G. Orbán, M. Lengyel, and J. Fiser, Science **331**, 83
- [6] S. Deneve, Neural Comput. 20, 91 (2008).
- [7] L. Buesing, J. Bill, B. Nessler, and W. Maass, PLoS Comput. Biol. **7**, e1002211 (2011).
- [8] R. P. Rao, Advances in Neural Information Processing Systems (MIT Press, Cambridge, MA, 2004), pp. 1113–1120.
- [9] P. O. Hoyer and A. Hyvarinen, Adv. Neural Info. Process. Syst., 293 (2003).
- [10] Z. F. Mainen and T. J. Sejnowski, Science 268, 1503 (1995).
- [11] M. A. Petrovici, J. Bill, I. Bytschok, J. Schemmel, and K. Meier, arXiv:1311.3211.
- [12] A. Destexhe, J. Comp. Neurosci. 27, 493 (2009).
- [13] A. Destexhe, M. Rudolph, and D. Pare, Nature Rev. Neurosci. 4, 739 (2003).
- [14] M. J. Richardson and W. Gerstner, Neural Comput. 17, 923 (2005).
- [15] L. M. Ricciardi and L. Sacerdote, Biol. Cybern. 35, 1 (1979).
- [16] N. Brunel and S. Sergi, J. Theor. Biol. 195, 87 (1998).
- [17] R. Moreno-Bote and N. Parga, Phys. Rev. Lett. 92, 028102 (2004).
- [18] L. M. Ricciardi and S. Sato, J. Appl. Probab. 25, 43 (1988).
- [19] D. Pecevski, L. Buesing, and W. Maass, PLoS Comput. Biol. 7, e1002294 (2011).

- [20] R. Salakhutdinov and G. E. Hinton, *International Conference on Artificial Intelligence and Statistics* (2009), pp. 448–455.
- [21] A.-R. Mohamed, G. E. Dahl, and G. Hinton, IEEE Trans. Audio Speech Lang. Proc. 20, 14 (2012).
- [22] M. Tsodyks and H. Markram, Proc. Natl. Acad. Sci. USA 94, 719 (1997).
- [23] G. Hinton, P. Dayan, B. Frey, and R. Neal, Science 268, 1158 (1995).
- [24] D. Sulzer and S. Rayport, Amino Acids 19, 45 (2000).
- [25] W. Gerstner and J. L. van Hemmen, Netw., Comput. Neural Syst. 3, 139 (1992).
- [26] R. Brette and W. Gerstner, J. Neurophysiol. 94, 3637 (2005).
- [27] S. Mitra, S. Fusi, and G. Indiveri, IEEE Trans. Biomed. Circuits Syst. 3, 32 (2009).
- [28] M. A. Petrovici, B. Vogginger, P. Müller, O. Breitwieser *et al.*, PLoS One **9**, e108590 (2014).
- [29] D. Probst, M. A. Petrovici, I. Bytschok *et al.*, Front. Comput. Neurosci. 9, 13 (2015).
- [30] E. Neftci, S. Das, B. Pedroni, K. Kreutz-Delgado, and G. Cauwenberghs, Front. Neurosci. 7, 272 (2014).
- [31] P. Lánský, Phys. Rev. E 55, 2040 (1997).
- [32] M. U. Thomas, J. Appl. Probab. 12, 600 (1975).
- [33] A. N. Burkitt, Biol. Cybern. 95, 97 (2006).
- [34] M. L. Hines and N. T. Carnevale, *The NEURON Book* (Cambridge University Press, Cambridge, UK, 2006).
- [35] A. P. Davison, D. Brüderle, J. Eppler, J. Kremkow, E. Muller, D. Pecevski, L. Perrinet, and P. Yger, Front. Neuroinform. 2, 11 (2009).
- [36] R. Naud, N. Marcille, C. Clopath, and W. Gerstner, Biol. Cybern. 99, 335 (2008).
- [37] Y. LeCun and C. Cortes (1998).