Stochasticity from function — Why the Bayesian brain may need no noise

Dominik Dold [a,b,*,1], Ilja Bytschok [a,1], Akos F. Kungl [a,b], Andreas Baumbach [a],
Oliver Breitwieser [a], Walter Senn [b], Johannes Schemmel [a], Karlheinz Meier [a],
Mihai A. Petrovici [a,b,*,1]

[a] *Kirchhoff-Institute for Physics, Heidelberg University, Im Neuenheimer Feld 227, D-69120 Heidelberg, Germany*
[b] *Department of Physiology, University of Bern, Bühlplatz 5, CH-3012 Bern, Switzerland*

ARTICLE INFO

ABSTRACT

An increasing body of evidence suggests that the trial-to-trial variability of spiking activity in the brain is not mere noise, but rather the reflection of a sampling-based encoding scheme for probabilistic computing. Since the precise statistical properties of neural activity are important in this context, many models assume an ad-hoc source of well-behaved, explicit noise, either on the input or on the output side of single neuron dynamics, most often assuming an independent Poisson process in either case. However, these assumptions are somewhat problematic: neighboring neurons tend to share receptive fields, rendering both their input and their output correlated; at the same time, neurons are known to behave largely deterministically, as a function of their membrane potential and conductance. We suggest that spiking neural networks may have no need for noise to perform sampling-based Bayesian inference. We study analytically the effect of auto- and cross-correlations in functional Bayesian spiking networks and demonstrate how their effect translates to synaptic interaction strengths, rendering them controllable through synaptic plasticity. This allows even small ensembles of interconnected deterministic spiking networks to simultaneously and co-dependently shape their output activity through learning, enabling them to perform complex Bayesian computation without any need for noise, which we demonstrate in silico, both in classical simulation and in neuromorphic emulation. These results close a gap between the abstract models and the biology of functionally Bayesian spiking networks, effectively reducing the architectural constraints imposed on physical neural substrates required to perform probabilistic computing, be they biological or artificial.

1. Introduction

An ubiquitous feature of in-vivo neural responses is their stochastic nature (Arieli, Sterkin, Grinvald, & Aertsen, 1996; Azouz & Gray, 1999; Henry, Bishop, Tupper, & Dreher, 1973; Schiller, Finlay, & Volman, 1976; Snowden, Treue, & Andersen, 1992; Vogels, Spileers, & Orban, 1989). The clear presence of this variability has spawned many functional interpretations, with the Bayesian-brain hypothesis arguably being the most notable example (Brascamp, Van Ee, Noest, Jacobs, & van den Berg, 2006; Deco, Rolls, & Romo, 2009; Fiser, Berkes, Orbán, & Lengyel, 2010; Körding & Wolpert, 2004; Maass, 2016; Rao, Olshausen, & Lewicki, 2002). Under this assumption, the activity of a neural network

is interpreted as representing an underlying (prior) probability distribution, with sensory data providing the evidence needed to constrain this distribution to a (posterior) shape that most accurately represents the possible states of the environment given the limited available knowledge about it.

Neural network models have evolved to reproduce this kind of neuronal response variability by introducing noise-generating mechanisms, be they extrinsic, such as Poisson input (Brunel, 2000; Fourcaud & Brunel, 2002; Gerstner, Kistler, Naud, & Paninski, 2014; Stein, 1967) or fluctuating currents (Maass & Zador, 1998; Moreno-Bote, 2014; Neftci, Pedroni, Joshi, Al-Shedivat, & Cauwenberghs, 2016; Smetters & Zador, 1996; Steinmetz, Manwani, Koch, London, & Segev, 2000; Yarom & Hounsgaard, 2011), or intrinsic, such as stochastic firing (Chichilnisky, 2001; Dayan et al., 2003; Gerstner & Kistler, 2002; Ostojic & Brunel, 2011; Plesser & Gerstner, 2000; Stevens & Zador, 1996) or membrane fluctuations (Aitchison & Lengyel, 2016; Orbán, Berkes, Fiser, & Lengyel, 2016).

---

* Corresponding authors at: Kirchhoff-Institute for Physics, Heidelberg University, Im Neuenheimer Feld 227, D-69120 Heidelberg, Germany.
E-mail addresses: dodo@kip.uni-heidelberg.de (D. Dold), petrovici@pyl.unibe.ch (M.A. Petrovici).
[1] Authors with equal contributions.

However, while representing, to some degree, reasonable approximations, none of the commonly used sources of stochasticity is fully compatible with biological constraints. Contrary to the independent white noise assumption, neuronal inputs are both auto- and cross-correlated to a significant degree (Averbeck, Latham, & Pouget, 2006; Deger, Helias, Boucsein, & Rotter, 2012; Fiser, Chiu, & Weliky, 2004; Nelson, Salin, Munk, Arzi, & Bullier, 1992; Rosenbaum, Tchumatchenko, & Moreno-Bote, 2014; Salinas & Sejnowski, 2001; Segev et al., 2002), with obvious consequences for a network's output statistics (Moreno-Bote, Renart, & Parga, 2008). At the same time, the assumption of intrinsic neuronal stochasticity is at odds with experimental evidence of neurons being largely deterministic units (Mainen & Sejnowski, 1995; Rauch, La Camera, Lüscher, Senn, & Fusi, 2003; Zador, 1998). Although synaptic transmissions from individual release sites are stochastic, averaged across multiple sites, contacts and connections, they largely average out (Markram, Muller, Ramaswamy, Reimann, Abdellah, Sanchez, et al., 2015). Therefore, it remains an interesting question how cortical networks that use stochastic activity as a means to perform probabilistic inference can realistically attain such apparent randomness in the first place.

We address this question within the normative framework of sampling-based Bayesian computation (Buesing, Bill, Nessler, & Maass, 2011; Leng et al., 2018; Orbán et al., 2016; Pecevski, Buesing, & Maass, 2011; Petrovici, Bill, Bytschok, Schemmel, & Meier, 2016; Probst et al., 2015), in which the spiking activity of neurons is interpreted as Markov Chain Monte Carlo sampling from an underlying distribution over a high-dimensional binary state space. In contrast to other work on deterministic chaos in functional spiking networks, done mostly in the context of reservoir computing (e.g., Monteforte & Wolf, 2012; Pyle & Rosenbaum, 2017), we provide a stringent connection to the spike-based representation and computation of probabilities, as well as the synaptic plasticity required for learning the above. We demonstrate how an ensemble of dynamically fully deterministic, but functionally probabilistic networks, can learn a connectivity pattern that enables probabilistic computation with a degree of precision that matches the one attainable with idealized, perfectly stochastic components. The key element of this construction is self-consistency, in that all input activity seen by a neuron is the result of output activity of other neurons that fulfill a functional role in their respective subnetworks. The present work supports probabilistic computation in light of experimental evidence from biology and suggests a resource-efficient implementation of stochastic computing by completely removing the need for any form of explicit noise.

2. Methods

2.1. Neuron model and simulation details

We consider deterministic Leaky Integrate-and-Fire (LIF) neurons with conductance-based synapses and dynamics described by

$$C_{\mathrm{m}} \frac{\mathrm{d}u_k}{\mathrm{d}t} = g_{\mathrm{l}}\left(E_{\mathrm{l}} - u_k\right) + \sum_{x \in \{\mathrm{e},\mathrm{i}\}} g_{k,x}^{\mathrm{syn}}(E_x^{\mathrm{rev}} - u_k), \qquad (1)$$

$$g_{k,x}^{\mathrm{syn}}(t) = \sum_{\text{synapses } j} \sum_{\text{spikes } s} w_{kj}\, \theta(t - t_s)\, \exp\left(-\frac{t - t_s}{\tau^{\mathrm{syn}}}\right), \qquad (2)$$

$$u_k(t_s) \geq \vartheta \Rightarrow u_k(t \in (t_s, t_s + \tau_{\mathrm{ref}}]) = \varrho, \qquad (3)$$

with membrane capacitance C_{m}, leak conductance g_{l}, leak potential E_{l}, excitatory and inhibitory reversal potentials $E_{\mathrm{e/i}}^{\mathrm{rev}}$ and conductances $g_{k,\mathrm{e/i}}^{\mathrm{syn}}$, synaptic strength w_{kj}, synaptic time constant τ^{syn} and $\theta(t)$ the Heaviside step function. For $g_{k,\mathrm{e/i}}^{\mathrm{syn}}$, the first

sum covers all synaptic connections projecting to neuron k. A neuron spikes at time t_s when its membrane potential crosses the threshold ϑ, after which it becomes refractory. During the refractory period τ_{ref}, the membrane potential is clamped to the reset potential ϱ. We have chosen the above model because it provides a computationally tractable abstraction of neurosynaptic dynamics (Rauch et al., 2003), but our general conclusions are not restricted to these specific dynamics.

We further use the short-term plasticity mechanism described in Fuhrmann, Segev, Markram, and Tsodyks (2002) to modulate synaptic interaction strengths with an adaptive factor $U_{\mathrm{SE}} \times R(t)$, where the time-dependence is given by[2]

$$\frac{\mathrm{d}R}{\mathrm{d}t} = \frac{1 - R}{\tau_{\mathrm{rec}}} - U_{\mathrm{SE}} R \delta(t - t_s), \qquad U_{\mathrm{SE}}, R \in [0, 1], \qquad (4)$$

where $\delta(t)$ is the Dirac delta function, t_s denotes the time of a presynaptic spike, which depletes the reservoir R by a fraction U_{SE}, and τ_{rec} is the time scale on which the reservoir R recovers. This enables a better control over the inter-neuron interaction, as well as over the mixing properties of our networks (Leng et al., 2018).

Background input, such as spikes from a Poisson source, enters Eq. (1) as synaptic input, but without short-term plasticity (as in Petrovici et al. (2016)) to facilitate the mathematical analysis (see Supporting information for more details).

All simulations were performed with the network specification language PyNN 0.8 (Davison et al., 2009) and the spiking neural network simulator NEST 2.4.2 (Gewaltig & Diesmann, 2007).

2.2. Sampling framework

As a model of probabilistic inference in networks of spiking neurons, we adopt the framework introduced in Petrovici et al. (2016), Probst et al. (2015). There, the neuronal output becomes stochastic due to a high-frequency bombardment of excitatory and inhibitory Poisson stimuli (Fig. 1A), elevating neurons into a high-conductance state (HCS) (Destexhe, Rudolph, & Paré, 2003; Kumar, Schrader, Aertsen, & Rotter, 2008), where they attain a high reaction speed due to a reduced effective membrane time constant. Under these conditions, a neuron's response (or activation) function becomes approximately logistic and can be represented as $\varphi(\mu) = (1 + \exp(-(\mu - u_0)/\alpha))^{-1}$ with inverse slope α and inflection point u_0. Together with the mean free membrane potential μ and the mean effective membrane time constant τ_{eff} (see Eqs. (17b) and (20c)), the scaling parameters α and u_0 are used to translate the weight matrix $\mathbf{W}$ and bias vector $\mathbf{b}$ of a target Boltzmann distribution $p_{\mathbf{z}}^{\mathrm{target}} = p(\mathbf{z}) \propto \exp\left(\frac{1}{2}\mathbf{z}^T\mathbf{W}\mathbf{z} + \mathbf{z}^T\mathbf{b}\right)$ with binary random variables $\mathbf{z} \in \{0, 1\}^n$ to synaptic weights and leak potentials in a sampling spiking network (SSN):

$$w_{kj} = \frac{\alpha W_{kj} C_{\mathrm{m}} \frac{\tau_{\mathrm{ref}}}{\tau^{\mathrm{syn}}} \left(1 - \frac{\tau^{\mathrm{syn}}}{\tau_{\mathrm{eff}}}\right)\left(E_{kj}^{\mathrm{rev}} - \mu\right)^{-1}}{\left[\tau^{\mathrm{syn}}\left(e^{-\frac{\tau_{\mathrm{ref}}}{\tau^{\mathrm{syn}}}} - 1\right) - \tau_{\mathrm{eff}}\left(e^{-\frac{\tau_{\mathrm{ref}}}{\tau_{\mathrm{eff}}}} - 1\right)\right]}, \qquad (5)$$

$$\mathbf{E_l} = \frac{\tau_{\mathrm{m}}}{\tau_{\mathrm{eff}}}(\alpha\mathbf{b} + u_0) - \sum_{x \in \{\mathrm{e},\mathrm{i}\}} \frac{\langle g_x^{\mathrm{syn}}\rangle}{g_{\mathrm{l}}} E_x^{\mathrm{rev}}, \qquad (6)$$

where w_{kj} is the synaptic weight from neuron j to neuron k, $\mathbf{E_l}$ a vector containing the leak potentials of all neurons, $\mathbf{b}$ the corresponding bias vector, $E_{kj}^{\mathrm{rev}} \in \{E_{\mathrm{e}}^{\mathrm{rev}}, E_{\mathrm{i}}^{\mathrm{rev}}\}$, depending on the nature of the respective synapse, and $\tau_{\mathrm{m}} = \frac{C_{\mathrm{m}}}{g_{\mathrm{l}}}$ (see Eq. (68) to Eq. (73) for a derivation). This translation effectively enables sampling from $p_{\mathbf{z}}^{\mathrm{target}}$, where a refractory neuron is considered to represent the state $z_k = 1$ (see Fig. 1B,C).

---

[2] In Fuhrmann et al. (2002) the postsynaptic response only scales with $R(t)$, whereas here we scale it with $U_{\mathrm{SE}} \times R(t)$.
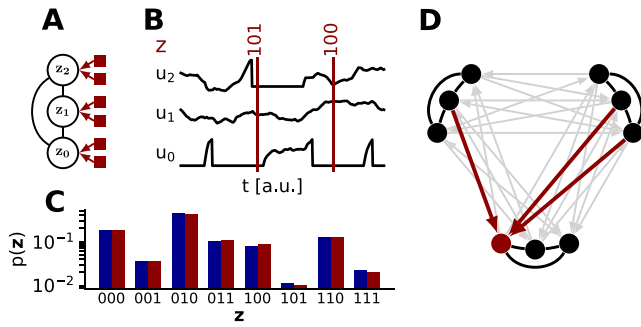
Fig. 1. Sampling spiking networks (SSNs) with and without explicit noise. **(A)** Schematic of a sampling spiking network, where each neuron (circles) encodes a binary random variable $z_i \in \{0, 1\}$. In the original model, neurons were rendered effectively stochastic by adding external Poisson sources of high-frequency balanced noise (red boxes). **(B)** A neuron represents the state $z_k = 1$ when refractory and $z_k = 0$ otherwise. **(C)** The dynamics of neurons in an SSN can be described as sampling (red bars) from a target distribution (blue bars). **(D)** Instead of using Poisson processes as a source of explicit noise, we replace the Poisson input with spikes coming from other networks performing spike-based probabilistic inference by creating a sparse, asymmetric connectivity matrix between several SSNs. For instance, the red neuron receives not only information-carrying spikes from its home network (black lines), but also spikes from the other two SSNs as background (red arrows), and in turn projects back towards these networks. Other such background connections are indicated in light gray. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2.3. Measures of network performance

To assess how well a sampling spiking network (SSN) samples from its target distribution, we use the Kullback–Leibler divergence (Kullback & Leibler, 1951)

$$D_{KL}\left(p^{net} \parallel p^{target}\right) = \sum_{\mathbf{z}} p_{\mathbf{z}}^{net} \ln\left(\frac{p_{\mathbf{z}}^{net}}{p_{\mathbf{z}}^{target}}\right), \tag{7}$$

which is a measure for the similarity between the sampled distribution p^{net} and the target distribution p^{target}. For inference tasks, we determine the network's classification rate on a subset of the used data set which was put aside during training. Furthermore, generative properties of SSNs are investigated either by letting the network complete partially occluded examples from the data set or by letting it generate new examples.

2.4. Learning algorithm

Networks were trained with a Hebbian wake-sleep algorithm

$$\Delta \mathbf{W}_{ij} = \eta \left[p_{z_i=1, z_j=1}^{target} - p_{z_i=1, z_j=1}^{net} \right], \tag{8}$$

$$\Delta \mathbf{b}_i = \eta \left[p_{z_i=1}^{target} - p_{z_i=1}^{net} \right], \tag{9}$$

which minimizes the $D_{KL}\left(p^{net} \parallel p^{target}\right)$ (Ackley, Hinton, & Sejnowski, 1985). η is a learning rate (see Supporting information for used hyperparameters). For high-dimensional datasets (e.g. handwritten letters and digits), Boltzmann machines were trained with the CAST algorithm (Salakhutdinov, 2010), a variant of wake-sleep with a tempering scheme, and then translated to SSN parameters with Eqs. (5) and (6) instead of training the SSNs directly to reduce simulation time.

2.5. Experiments and calculations

Details to all experiments as well as additional figures and captions to videos can be found in the Supporting information. Detailed calculations are presented at the end of the main text.

3. Results

We approach the problem of externally-induced stochasticity incrementally. Throughout the remainder of the manuscript, we discern between background input, which is provided by other functional networks, and explicit noise, for which we use the conventional assumption of Poisson spike trains. We start by analyzing the effect of correlated background on the performance of SSNs. We then demonstrate how the effects of both auto- and cross-correlated background can be mitigated by Hebbian plasticity. This ultimately enables us to train a fully deterministic network of networks to perform different inference tasks without requiring any form of explicit noise. This is first shown for larger ensembles of small networks, each of which receives its own target distribution, which allows a straightforward quantitative assessment of their sampling performance $D_{KL}\left(p^{net} \parallel p^{target}\right)$. We study the behavior of such ensembles both in computer simulations and on mixed-signal neuromorphic hardware. Finally, we demonstrate the capability of our approach for truly functional, larger-scale networks, trained on higher-dimensional visual data.

3.1. Background autocorrelations

Unlike ideal Poisson sources, single spiking neurons produce autocorrelated spike trains, with the shape of the autocorrelation function (ACF) depending on their refractory time τ_{ref} and mean spike frequency $\bar{r} = p(z = 1)\tau_{ref}^{-1}$. For higher output rates, spike trains become increasingly dominated by bursts, i.e., sequences of equidistant spikes with an interspike interval (ISI) of ISI $\approx \tau_{ref}$. These fixed structures also remain in a population, since the population autocorrelation is equal to the averaged ACFs of the individual spike trains.

We investigated the effect of such autocorrelations on the output statistics of SSNs by replacing the Poisson input in the ideal model with spikes coming from other SSNs. As opposed to Poisson noise, the autocorrelation $C(S_x, S_x, \Delta) = \frac{\langle S_x(t)S_x(t+\Delta)\rangle - \langle S_x\rangle^2}{\text{Var}(S_x)}$ of the SSN-generated (excitatory or inhibitory) background S_x, $x \in \{e, i\}$ (Fig. 2B) is non-singular and influences the free membrane potential (FMP) distribution (Fig. 2C) and thereby activation function (Fig. 2D) of individual sampling neurons. With increasing firing rates (controlled by the bias of the neurons in the background SSNs), the number of significant peaks in the ACF increases as well (see Eq. (54)):

$$C(S_x, S_x, n\tau_{ref}) \approx \sum_{k=1}^{\infty} e^{k \ln \bar{p}} \delta\left([n - k]\tau_{ref}\right), \tag{10}$$

where $\bar{p}$ is the probability for a burst to start. This regularity in the background input manifests itself in a reduced width σ' of the FMP distribution (see Eq. (30))

$$f(u_i^{free}) \sim \mathcal{N}(\mu' = \mu, \sigma' = \sqrt{\beta}\sigma) \tag{11}$$

with a scaling factor $\sqrt{\beta}$ that depends on the ACF, which in turn translates to a steeper activation function (see Eqs. (36) and (37))

$$p(z_i = 1) \approx \int_{\vartheta}^{\infty} f(u)du \approx \varphi(\mu)\Big|_{u_0'=u_0, \alpha'=\sqrt{\beta}\alpha}, \tag{12}$$

with inflection point u_0' and inverse slope α'. Thus, autocorrelations in the background input lead to a reduced width of the FMP distribution and hence to a steeper activation function compared to the one obtained using uncorrelated Poisson input. For a better intuition, we used an approximation of the activation function of LIF neurons, but the argument also holds for the exact expression derived in Petrovici et al. (2016), as verified by simulations (Fig. 2D).
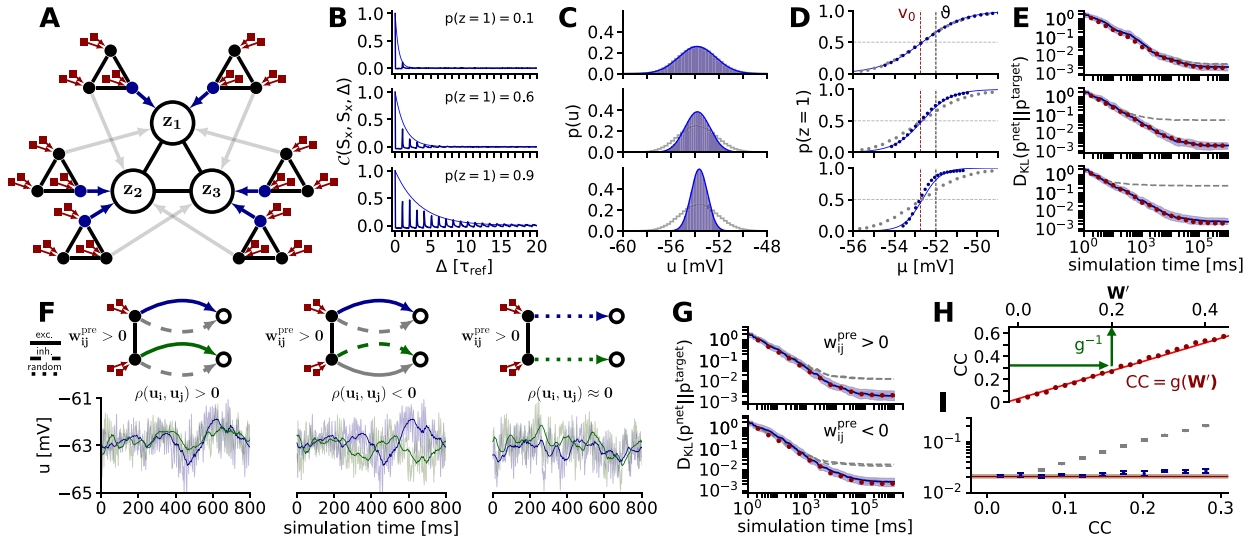
Fig. 2. Effect of correlated background on SSN dynamics and compensation through reparametrization. **(A)** Feedforward replacement of Poisson noise by spiking activity from other SSNs. In this illustration, the principal SSN consists of three neurons receiving background input only from other functional SSNs that sample from their own predetermined target distribution. For clarity, only two out of a total of [260, 50, 34] (top to bottom in (B)) background SSNs per neuron are shown here. By modifying the background connectivity (gray and blue arrows) the amount of cross-correlation in the background input can be controlled. At this stage, the background SSNs still receive Poisson input (red boxes). **(B)** By appropriate parametrization of the background SSNs, we adjust the mean spike frequency of the background neurons (blue) to study the effect of background autocorrelations $C(S_x, S_x, \Delta)$. Higher firing probabilities increase the chance of evoking bursts, which induce background autocorrelations for the neurons in the principal SSN at multiples of τ_{ref} (dark blue: simulation results; light blue: $e^{k \ln \tilde{p}}$ with $k = \frac{\Delta}{\tau_{\text{ref}}}$, see Eq. (10)). **(C)** Background autocorrelation narrows the FMP distribution of neurons in the principal SSN: simulation (blue bars) and the theoretical prediction (Eq. (11), blue line) vs. background Poisson noise of the same rate (gray). Background intensities correspond to (B). **(D)** Single-neuron activation functions corresponding to (B,C) and the theoretical prediction (Eq. (12), blue line). For autocorrelated noise, the slope of the response curve changes, but the inflection point (with $p(z = 1) = 0.5$) is conserved. **(E)** Kullback–Leibler divergence $D_{\text{KL}}\left(p^{\text{net}} \parallel p^{\text{target}}\right)$ (median and range between the first and third quartile) for the three cases shown in (B,C,D) after sampling from 50 different target distributions with 10 different random seeds for the 3-neuron network depicted in (A). Appropriate reparametrization can fully cancel out the effect of background autocorrelations (blue). The according results without reparametrization (gray) and with Poisson input (red) are also shown. **(F)** A pair of interconnected neurons in a background SSN generates correlated noise, as given by Eq. (13). The effect of cross-correlated background on a pair of target neurons depends on the nature of synaptic projections from the background to the principal SSN. Here, we depict the case where their interaction w_{ij}^{pre} is excitatory; the inhibitory case is a mirror image thereof. Left: If forward projections are of the same type, postsynaptic potentials will be positively correlated. Middle: Different synapse types in the forward projection only change the sign of the postsynaptic potential correlations. Right: For many background inputs with mixed connectivity patterns, correlations can average out to zero even when all input correlations have the same sign. **(G)** Same experiment as in (E), with background connection statistics adjusted to compensate for input cross-correlations. The uncompensated cases from (F, left) and (F, middle) are shown in gray. **(H)** Correlation-cancelling reparametrization in the principal SSN. By transforming the state space from $\mathbf{z} \in \{0, 1\}^n$ to $\mathbf{z}' \in \{-1, 1\}$, input correlations attain the same functional effect as synaptic weights (Eq. (15)); simulation results given as red dots, linear fit as red line. Weight rescaling followed by a transformation back into the $\mathbf{z} \in \{0, 1\}^n$ state space, shown in green (which affects both weights and biases) can therefore alleviate the effects of correlated background. **(I)** Similar experiment as in (E) for a network with ten neurons, with parameters adjusted to compensate for input cross-correlations. As in the case of autocorrelated background, cross-correlations can be cancelled out by appropriate reparametrization. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Apart from the above effect, the background autocorrelations do not affect neuron properties that depend linearly on the synaptic noise input, such as the mean FMP and the inflection point of the activation function (equivalent to zero bias). Therefore, the effect of the background autocorrelations can be functionally reversed by rescaling the functional (from other neurons in the principal SSN) afferent synaptic weights by a factor equal to the ratio between the new and the original slope α'/α (Eqs. (5) and (6)), as shown in Fig. 2E.

3.2. Background cross-correlations

In addition to being autocorrelated, background input to pairs of neurons can be cross-correlated as well, due to either shared inputs or synaptic connections between the neurons that generate said background. These background cross-correlations can manifest themselves in a modified cross-correlation between the outputs of neurons, thereby distorting the distribution sampled by an SSN.

However, depending on the number and nature of presynaptic background sources, background cross-correlations may cancel out to a significant degree. The correlation coefficient (CC) of the FMPs of two neurons fed by correlated noise amounts to

(see Eq. (59))

$$\rho(u_i^{\text{free}}, u_j^{\text{free}}) \propto \sum_{l,m} w_{il} w_{jm} \left(E_{il}^{\text{rev}} - \mu_i\right)\left(E_{jm}^{\text{rev}} - \mu_j\right) \tag{13}$$

$$\cdot \int d\Delta\, \lambda_{li,mj}\, C\left(S_{l,i}, S_{m,j}, \Delta\right) \widetilde{C}\left(\kappa, \kappa, \Delta\right),$$

where l sums over all background spike trains $S_{l,i}$ projecting to neuron i and m sums over all background spike trains $S_{m,j}$ projecting to neuron j. $\widetilde{C}(\kappa, \kappa, \Delta)$ is the unnormalized autocorrelation function of the postsynaptic potential (PSP) kernel κ, i.e., $\widetilde{C}(\kappa, \kappa, \Delta) = \langle \kappa(t)\kappa(t + \Delta)\rangle$, and $C\left(S_{l,i}, S_{m,j}, \Delta\right)$ the cross-correlation function of the background inputs. $\lambda_{li,mj}$ is given by $\lambda_{li,mj} = \sqrt{\text{Var}\left(S_{l,i}\right)\text{Var}\left(S_{m,j}\right)}$. The background cross-correlation is gated into the cross-correlation of FMPs by the nature of the respective synaptic connections: if the two neurons connect to the cross-correlated inputs by synapses of different type (one excitatory, one inhibitory), the sign of the CC is switched (Fig. 2F). However, individual contributions to the FMP CC also depend on the difference of the mean free membrane potential and the reversal potentials, so the gating of cross-correlations is not symmetric for excitatory and inhibitory synapses. Nevertheless, if the connectivity statistics (in-degree and synaptic weights) from the background sources to an SSN are chosen appropriately

and enough presynaptic partners are available, the total pairwise cross-correlation between neurons in an SSN cancels out on average, leaving the sampling performance unimpaired (Fig. 2G). Note that this way of reducing cross-correlations is independent of the underlying weight distribution of the networks providing the background; the required cross-wiring of functional networks could therefore, in principle, be encoded genetically and does not need to be learned. Furthermore, a very simple cross-wiring rule, i.e., independently and randomly determined connections, already suffices to accomplish low background cross-correlations and therefore reach a good sampling performance.

Whereas this method is guaranteed to work in an artificial setting, further analysis is needed to assess its compatibility with the cortical connectome with respect to connectivity statistics or synaptic weight distributions. However, even if cortical architecture prevents a clean implementation of this decorrelation mechanism, SSNs can themselves compensate for residual background cross-correlations by modifying their parameters, similar to the autocorrelation compensation discussed above.

To demonstrate this ability, we need to switch from the natural state space of neurons $\mathbf{z} \in \{0, 1\}^N$ to the more symmetric space $\mathbf{z}' \in \{-1, 1\}^N$.[3] By requiring $p(\mathbf{z}') \overset{!}{=} p(\mathbf{z})$ to conserve state probabilities (and thereby also correlations), the desired change of state variables $\mathbf{z}' = 2\mathbf{z} - \mathbf{1}$ can be achieved with a linear parameter transformation (see Eqs. (66) and (67)):

$$\mathbf{W}' = \frac{1}{4}\mathbf{W} \text{ and } \mathbf{b}' = \frac{1}{2}\mathbf{b} + \frac{1}{4}\sum_i \mathrm{col}_i\mathbf{W}. \qquad (14)$$

In the $\{-1, 1\}^N$ state space, both synaptic connections w'_{ij} and background cross-correlations $\rho(S_i, S_j)$ shift probability mass between mixed states $(z_i, z_j) = \pm(1, -1)$ and aligned states $(z_i, z_j) = \pm(1, 1)$ (see Supporting information, Fig. S1). Therefore, by adjusting $\mathbf{b}$ and $\mathbf{W}$, it is possible to find a $\mathbf{W}'$ (Fig. 2H) that precisely conserves the desired correlation structure between neurons:

$$w'_{ij} = g^{-1}[\rho(S_i, S_j)] \approx \frac{\rho(S_i, S_j) - g_0}{g_1}, \qquad (15)$$

with constants g_0 and g_1 (Fig. 2I). Therefore, when an SSN learns a target distribution from data, background cross-correlations are equivalent to an offset in the initial network parameters and are automatically compensated during training.

For now, we can conclude that the activity of SSNs constitutes a sufficient source of stochasticity for other SSNs, since all effects that follow from replacing Poisson noise in an SSN with functional output from other SSNs (which at this point still receive explicit noise) can be compensated by appropriate parameter adjustments. These are important preliminary conclusions for the next sections, where we show how all noise can be eliminated in an ensemble of interconnected SSNs endowed with synaptic plasticity without significant penalty to their respective functional performance.

3.3. Sampling without explicit noise in large ensembles

We initialized an ensemble of 100 6-neuron SSNs with an inter-network connectivity of $\epsilon = 0.1$ and random synaptic weights. As opposed to the previous experiments, none of the neurons in the ensemble receive explicit Poisson input and the activity of the ensemble itself acts as a source of stochasticity instead, as depicted in Fig. 1D. No external input is needed to kick-start network activity, as some neurons spike spontaneously,

due to the random initialization of parameters (see Fig. 3A). The existence of inhibitory weights disrupts the initial regularity, initiating the sampling process. Ongoing learning (Eqs. (8) and (9)) shapes the sampled distributions towards their respective targets (Fig. 3B), the parameters of which were drawn randomly (see Supporting information). Our ensemble achieved a sampling performance (median D_{KL}) of $1.06^{+0.27}_{-0.40} \times 10^{-3}$, which is similar to the median performance of an idealized setup (independent, Poisson-driven SSNs as in Petrovici et al. (2016)) of $1.05^{+0.15}_{-0.35} \times 10^{-3}$ (errors are given by the first and third quartile). To put the above D_{KL} values in perspective, we compare the sampled and target distributions of one of the SSNs in the ensemble at various stages of learning (Fig. 3C). Thus, despite the fully deterministic nature of the system, the network dynamics and achieved performance after training is essentially indistinguishable from that of networks harnessing explicit noise for the representation of probability. Instead of training ensembles, they can also be set up by translating the parameters of the target distributions to neurosynaptic parameters directly, as discussed in the previous section (see Supporting information, Fig. S2).

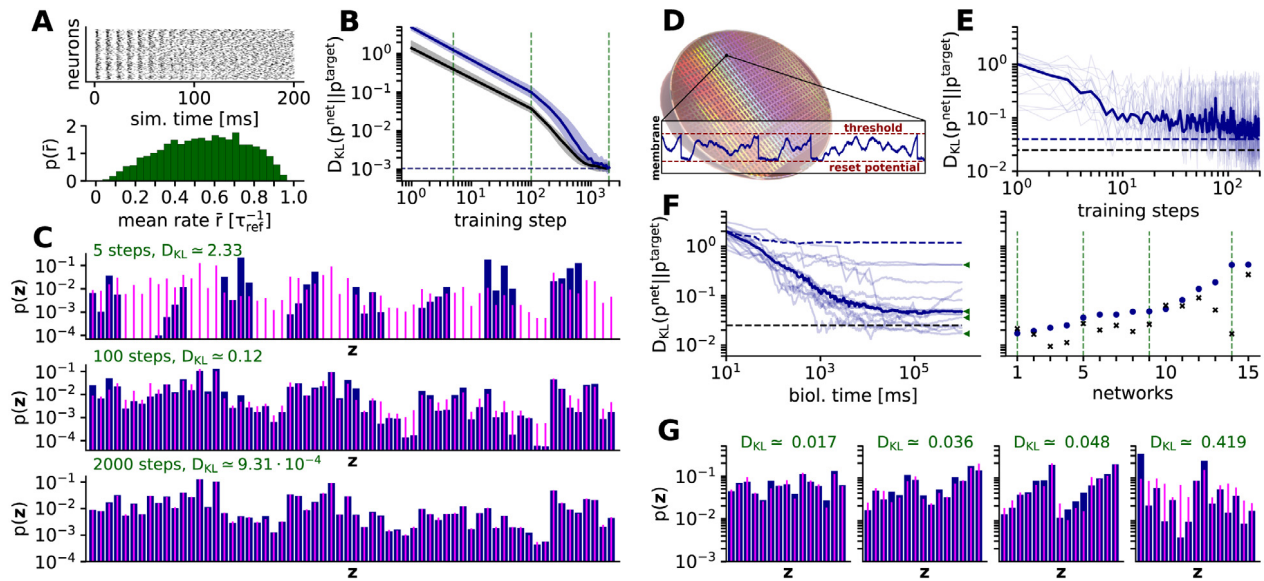### 3.4. Implementation on a neuromorphic substrate

To test the robustness of our results, we studied an implementation of noise-free sampling on an artificial neural substrate, which incorporates unreliable components and is therefore significantly more difficult to control. For this, we used the BrainScaleS system (Schemmel, Fieres, & Meier, 2008), a mixed-signal neuromorphic platform with analog neurosynaptic dynamics and digital inter-neuron communication (Fig. 3D, see also Supporting information, Fig. S3). A major advantage of this implementation is the emulation speedup of 10^4 with respect to biological real-time; however, for clarity, we shall continue using biological time units instead of actual emulation time.

The additional challenge for our neuronal ensemble is to cope with the natural variability of the substrate, caused mainly by fixed-pattern noise, or with other limitations such as a finite weight resolution (4 bits) or spike loss, which can all be substantial (Petrovici et al., 2014; Schmitt et al., 2017). It is important to note that the ability to function when embedded in an imperfect substrate with significant deviations from an idealized model represents a necessary prerequisite for viable theories of biological neural function.

We emulated an ensemble of 15 4-neuron SSNs, with an inter-SSN connectivity of $\epsilon = 0.2$ and with randomly drawn target distributions (see Supporting information). The biases were provided by additional bias neurons and adjusted during learning via the synaptic weights between bias and sampling neurons, along with the synapses within the SSNs, using the same learning rule as before (Eqs. (8) and (9)). After 200 training steps, the ensemble reached a median D_{KL} of $3.99^{+1.27}_{-1.15} \cdot 10^{-2}$ (errors given by the distance to the first and third quartile) compared to $1.18^{+0.47}_{-0.55}$ before training (Fig. 3E). As a point of reference, we also considered the idealized case by training the same set of SSNs without interconnections and with every neuron receiving external Poisson noise generated from the host computer, reaching a D_{KL} of $2.49^{+3.18}_{-0.71} \cdot 10^{-2}$.

This relatively small performance loss of the noise-free ensemble compared to the ideal case confirms the theoretical predictions and simulation results. Importantly, this was achieved with only a rather small ensemble, demonstrating that large numbers of neurons are not needed for realizing this computational paradigm.

In Fig. 3F and Video S1, we show the sampling dynamics of all emulated SSNs after learning. While most SSNs are able to approximate their target distributions well, some sampled

---

[3] The $z = 0$ state for a silent neuron is arguably more natural, because it has no effect on its postsynaptic partners during this state. In contrast, $z \in \{-1, 1\}$ would, for example, imply efferent excitation upon spiking and constant efferent inhibition otherwise.

Fig. 3. Sampling without explicit noise from a set of predefined target distributions in software (A–C) and on a neuromorphic substrate (D–G). **(A)** (top) Temporal evolution of spiking activity in an ensemble of 100 interconnected 6-neuron SSNs with no source of explicit noise. An initial burst of regular activity caused by neurons with a strong enough positive bias quickly transitions to asynchronous irregular activity due to inhibitory synapses. (bottom) Distribution of mean neuronal firing rates of the ensembles shown in (B, C) after training. **(B)** Median sampling quality of the above ensemble during learning, for a test sampling period of 10^6 ms. At the end of the learning phase, the sampling quality of individual networks in the ensemble (blue) is on par with the one obtained in the theoretically ideal case of independent networks with Poisson background (black). Error bars given over 5 simulation runs with different random seeds. **(C)** Illustration of a single target distribution (magenta) and corresponding sampled distribution (blue) of a network in the ensemble at several stages of the learning process (dashed green lines in (B)). **(D)** Photograph of a wafer from the BrainScaleS neuromorphic system used in (E), (F) and (G) before post-processing (i.e., adding additional structures like buses on top), which would mask the underlying modular structure. Blue: exemplary membrane trace of an analog neuron receiving Poisson noise. **(E)** Performance of an ensemble consisting of 15 4-neuron SSNs with no external noise during learning on the neuromorphic substrate, shown in light blue for each SSN and with the median shown in dark blue. The large fluctuations compared to (B) are a signature of the natural variability of the substrate's analog components. The dashed blue line represents the best achieved median performance at $D_{KL}\left(p^{net} \parallel p^{target}\right) = 3.99 \times 10^{-2}$. For comparison, we also plot the optimal median performance for the theoretically ideal case of independent, Poisson-driven SSNs emulated on the same substrate, which lies at $D_{KL}\left(p^{net} \parallel p^{target}\right) = 2.49 \times 10^{-2}$ (dashed black line). **(F)** Left: Demonstration of sampling in the neuromorphic ensemble of SSNs after 200 training steps. Individual networks in light blue, median performance in dark blue. Dashed blue line: median performance before training. Dashed black line: median performance of ideal networks, as in (E). Right: Best achieved performance, after 100 s of bio time (10 ms of hardware time) for all SSNs in the ensemble depicted as blue dots (sorted from lowest to highest D_{KL}). For comparison, the same is plotted as black crosses for their ideal counterparts. **(G)** Sampled (blue) and target (magenta) distributions of four of the 15 SSNs. The selection is marked in (F) with green triangles (left) and vertical green dashed lines (right). Since we made no particular selection of hardware neurons according to their behavior, hardware defects have a significant impact on a small subset of the SSNs. Despite these imperfections, a majority of SSNs perform close to the best value permitted by the limited weight resolution (4 bits) of the substrate. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

distributions are significantly skewed (Fig. 3G). This is caused by a small subset of dysfunctional neurons, which we have not discarded beforehand, in order to avoid an implausibly fine-tuned use-case of the neuromorphic substrate. These effects become less significant in larger networks trained on data instead of predefined distributions, where learning can naturally cope with such outliers by assigning them smaller output weights. Nevertheless, these results demonstrate the feasibility of self-sustained Bayesian computation through sampling in physical neural substrates, without the need for any source of explicit noise. Importantly, and in contrast to other approaches (Jordan, Petrovici, Breitwieser, Schemmel, Meier, Diesmann, & Tetzlaff, 2017), every neuron in the ensemble plays a functional role, with no neuronal real-estate being dedicated to the production of (pseudo-)randomness.
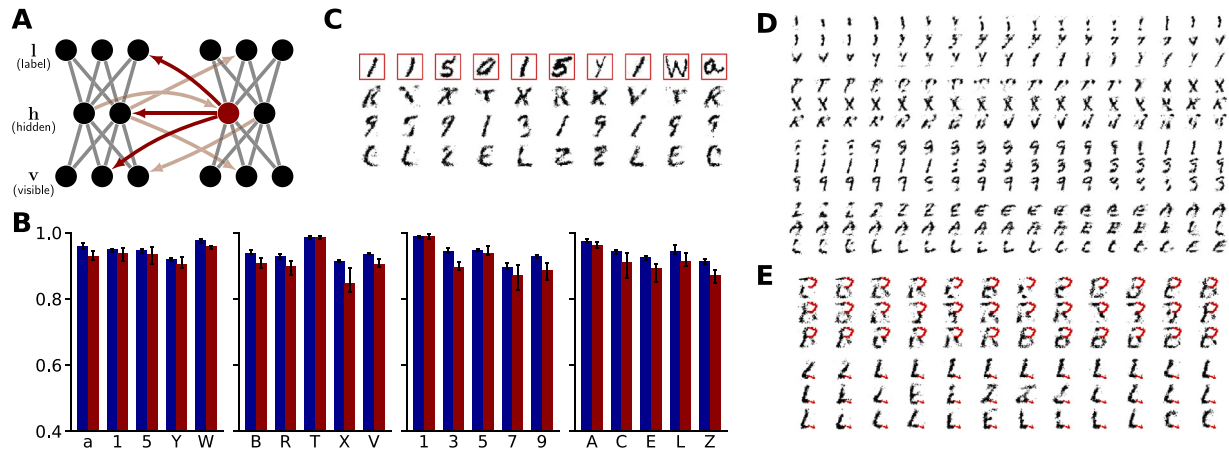
3.5. Ensembles of hierarchical SSNs

When endowed with appropriate learning rules, hierarchical spiking networks can be efficiently trained on high-dimensional visual data (Kheradpisheh, Ganjtabesh, Thorpe, & Masquelier, 2018; Lee, Delbruck, & Pfeiffer, 2016; Leng et al., 2018; Petrovici et al., 2017; Schmitt et al., 2017; Zenke & Ganguli, 2018). Such hierarchical networks are characterized by the presence of several layers, with connections between consecutive layers, but no lateral connections within the layers themselves. When both feedforward and feedback connections are present, such networks are able to both classify and generate images that are similar to those used during training.

In these networks, information processing in both directions is Bayesian in nature. Bottom-up propagation of information enables an estimation of the conditional probability of a particular label to fit the input data. Additionally, top-down propagation of neural activity allows generating a subset of patterns in the visible layer conditioned on incomplete or partially occluded visual stimulus. When no input is presented, such networks will produce patterns similar to those enforced during training ("dreaming"). In general, the exploration of a multimodal solution space in generative models is facilitated by some noise-generating mechanism. We demonstrate how even a small interconnected set of hierarchical SSNs can perform these computations self-sufficiently, without any source of explicit noise.

We used an ensemble of four 3-layer hierarchical SSNs trained on a subset of the EMNIST dataset (Cohen, Afshar, Tapson, & van Schaik, 2017), an extended version of the widely used MNIST dataset (LeCun, Bottou, Bengio, & Haffner, 1998) that includes digits as well as capital and lower-case letters. All SSNs had the same structure, with 784 visible units, 200 hidden units and 5 label units (Fig. 4A). To emulate the presence of networks with different functionality, we trained each of them on a separate subset of the data. (To combine sampling in space with sampling in time, multiple networks can also be trained on the same data,

Fig. 4. Bayesian inference on visual input. **(A)** Illustration of the connectivity between two hierarchical SSNs in the simulated ensemble. Each SSN had a visible layer **v**, a hidden **h** and a label layer **l**. Neurons in the same layer of an SSN were not interconnected. Each neuron in an SSN received only activity from the hidden layers of other SSNs as background (no sources of explicit noise). **(B)** An ensemble of four such SSNs (red) was trained to perform generative and discriminative tasks on visual data from the EMNIST dataset. We used the classification rate of restricted Boltzmann machines trained with the same hyperparameters as a benchmark (blue). Error bars are given (on blue) over 10 test runs and (on red) over 10 ensemble realizations with different random seeds. **(C)** Illustration of a scenario where one of the four SSNs (red boxes) received visual input for classification (B). At the same time, the other SSNs continuously generated images from their respective learned distributions. **(D)** Pattern generation and mixing during unconstrained dreaming. Here, we show the activity of the visible layer of all four networks from (B), each spanning three rows. Time evolves from left to right. For further illustrations of the sampling process in the ensemble of hierarchical SSNs, see Supporting information, Figs. S4 and S5 and Video S2. **(E)** Pattern completion and rivalry for two instances of incomplete visual stimulus. The stimulus consisted of the top right and bottom right quadrant of the visible layer, respectively. In the first run, we clamped the top arc of a "B" compatible with either a "B" or an "R" (top three rows, red), in the second run we chose the bottom line of an "L" compatible with an "L", an "E", a "Z" or a "C" (bottom three rows, red). An ensemble of SSNs performs Bayesian inference by implicitly evaluating the conditional distribution of the unstimulated visible neurons, which manifests itself here as sampling from all image classes compatible with the ambiguous stimulus (see also Supporting information, Fig. S6). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

see Supporting information Fig. S5 and Video S2.) Since training the spiking ensemble directly was computationally prohibitive, we trained four Boltzmann machines on the respective datasets and then translated the resulting parameters to neurosynaptic parameters of the ensemble using the methods described earlier in the manuscript (see also Supporting information, Fig. S2).

To test the discriminative properties of the SSNs in the ensemble, one was stimulated with visual input, while the remaining three were left to freely sample from their underlying distribution. We measured a median classification rate of $91.5^{+3.6}_{-3.0}\%$ with errors given by the distance to the first and third quartile, which is close to the $94.0^{+2.1}_{-1.5}\%$ achieved by the idealized reference setup provided by the abstract Boltzmann machines (Fig. 4B). At the same time, all other SSNs remained capable of generating recognizable images (Fig. 4C). It is expected that direct training and a larger number of SSNs in the ensemble would further improve the results, but a functioning translation from the abstract to the biological domain already underpins the soundness of the underlying theory.

Without visual stimulus, all SSNs sampled freely, generating images similar to those on which they were trained (Fig. 4D). Without any source of explicit noise, the SSNs were capable to mix between the relevant modes (images belonging to all classes) of their respective underlying distributions, which is a hallmark of a good generative model. We further extended these results to an ensemble trained on the full MNIST dataset, reaching a similar generative performance for all networks (see Supporting information Fig. S5 and Video S2).

To test the pattern completion capabilities of the SSNs in the ensemble, we stimulated them with incomplete and ambiguous visual data (Fig. 4E). Under these conditions, SSNs only produced images compatible with the stimulus, alternating between different image classes, in a display of pattern rivalry. As in the case of free dreaming, the key mechanism facilitating this form of exploration was provided by the functional activity of other neurons in the ensemble.

4. Discussion

Based on our findings, we argue that sampling-based Bayesian computation can be implemented in fully deterministic ensembles of spiking networks without requiring any explicit noise-generating mechanism. Our approach has a firm theoretical foundation in the theory of sampling spiking neural networks, upon which we formulate a rigorous analysis of network dynamics and learning in the presence or absence of noise.

While in biology various explicit sources of noise exist (Branco & Staras, 2009; Faisal, Selen, & Wolpert, 2008; White, Rubinstein, & Kay, 2000), these forms of stochasticity are either too weak (in case of ion channels) or too high-dimensional for efficient exploration (in the case of stochastic synaptic transmission, as used for, e.g., reinforcement learning (Seung, 2003)). Furthermore, a rigorous mathematical framework for neural sampling with stochastic synapses is still lacking. On the other hand, in the case of population codes, neuronal population noise can be highly correlated, affecting information processing by, e.g., inducing systematic sampling biases (Averbeck et al., 2006).

In our proposed framework, each network in an ensemble plays a dual role: while fulfilling its assigned function within its home subnetwork, it also provides its peers with the spiking background necessary for stochastic search within their respective solution spaces. This enables a self-consistent and parsimonious implementation of neural sampling, by allowing all neurons to take on a functional role and not dedicating any resources purely to the production of background stochasticity. The underlying idea lies in adapting neuro-synaptic parameters by (contrastive) Hebbian learning to compensate for auto- and cross-correlations induced by interactions between the functional networks in the ensemble. Importantly, we show that this does not rely on the presence of a large number of independent presynaptic partners for each neuron, as often assumed by models of cortical computation that use Poisson noise (see, e.g., Xie and Seung (2004)). Instead, only a small number of ensembles is

necessary to implement noise-free Bayesian sampling. This becomes particularly relevant for the development of neuromorphic platforms by eliminating the computational footprint imposed by the generation and distribution of explicit noise, thereby reducing power consumption and bandwidth constraints.

For simplicity, we chose networks of similar size in our simulations. However, the presented results are not contingent on network sizes in the ensemble and largely independent of the particular functionality (underlying distribution) of each SSN. Their applicability to scenarios where different SSNs learn to represent different data is particularly relevant for cortical computation, where weakly interconnected areas or modules are responsible for distinct functions (Bertolero, Yeo, & D'Esposito, 2015; Bullmore & Sporns, 2009; Chen, He, Rosa-Neto, Germann, & Evans, 2008; Meunier, Lambiotte, & Bullmore, 2010; Song, Sjöström, Reigl, Nelson, & Chklovskii, 2005). Importantly, these ideas scale naturally to larger ensembles and larger SSNs. Since each neuron only needs a small number of presynaptic partners from the ensemble, larger networks lead to a sparser interconnectivity between SSNs in the ensemble and hence soften structural constraints. Preliminary simulations show that the principle of using functional output as noise can even be applied to connections *within* a single SSN, eliminating the artificial separation between network and ensemble connections (see Fig. S7 and Video S3 in the Supporting information).

Even though we have used a simplified neuron model in our simulations to reduce computation time and facilitate the mathematical analysis, we expect the core underlying principles to generalize. This is evidenced by our results on neuromorphic hardware, where the dynamics of individual neurons and synapses differ significantly from the mathematical model. Such an ability to compute with unreliable components represents a particularly appealing feature in the context of both biology and emerging nanoscale technologies.

Finally, the suggested noise-free Bayesian brain reconciles the debate on spatial versus temporal sampling (Ma, Beck, Latham, & Pouget, 2006; Orbán et al., 2016). In fact, the networks of spiking neurons that provide each other with virtual noise may be arranged in parallel sensory streams. An ambiguous stimulus will trigger different representations on each level of these streams, forming a hierarchy of probabilistic population codes. While these population codes learn to cover the full sensory distribution in space, they will also generate samples of the sensory distribution in time (see Fig. S5 in the Supporting information). Attention may select the most likely representation, while suppressing the representations in the other streams. Analogously, possible actions may be represented in parallel motor streams during planning and a motor decision may select the one to be performed. When recording in premotor cortex, such a selection causes a noise reduction (Churchland, Byron, Ryu, Santhanam, & Shenoy, 2006), that we suggest is effectively the signature of choosing the most probable action in a Bayesian sense.

5. Conclusion

From a generic Bayesian perspective, cortical networks can be viewed as generators of target distributions. To enable such computation, models assume neurons to possess sources of perfect, well-behaved noise — an assumption that is both impractical and at odds with biology. We showed how local plasticity in an ensemble of spiking networks allows them to co-shape their activity towards a set of well-defined targets, while reciprocally using the very same activity as a source of (pseudo-)stochasticity. This enables purely deterministic networks to simultaneously learn a variety of tasks, completely removing the need for true randomness. While reconciling the sampling hypothesis with the deterministic nature of single neurons, this also offers an efficient blueprint for in-silico implementations of sampling-based inference.

6. Calculations

6.1. Free membrane potential distribution with colored noise

In the high-conductance state (HCS), it can be shown that the temporal evolution of the free membrane potential (FMP) of an LIF neuron stimulated by balanced Poisson inputs is equivalent to an Ornstein–Uhlenbeck (OU) process with the following Green's function (Petrovici, 2016):

$$f(u, t|u_0) = \sqrt{\frac{1}{2\pi\sigma^2(1 - e^{-2\theta t})}}$$
$$\cdot \exp\left(-\frac{1}{2\sigma^2}\frac{(u - \mu + (\mu - u_0)e^{-\theta t})^2}{1 - e^{-2\theta t}}\right). \tag{16}$$

with

$$\theta = \frac{1}{\tau^{\text{syn}}}, \tag{17a}$$

$$\mu = \frac{g_l E_l + \sum_{k \in \{e,i\}} \nu_k w_k E_k^{\text{rev}} \tau^{\text{syn}}}{\langle g^{\text{tot}} \rangle}, \tag{17b}$$

$$\sigma^2 = \frac{\sum_{k \in \{e,i\}} \nu_k w_k^2 (E_k^{\text{rev}} - \mu)^2 \tau^{\text{syn}}}{\langle g^{\text{tot}} \rangle^2}, \tag{17c}$$

$$\langle g^{\text{tot}} \rangle = g_l + \sum_{k \in \{e,i\}} w_k \nu_k \tau_k^{\text{syn}}, \tag{17d}$$

where ν_k are the noise frequencies, w_k the noise weights and we dropped the index notation u_k^{free} used in previous sections for convenience. The stationary FMP distribution is then given by a Gaussian (Gerstner & Kistler, 2002; Petrovici, 2016):

$$f(u) = \sqrt{\frac{1}{2\pi\sigma^2}} \exp\left(-\frac{(u - \mu)^2}{2\sigma^2}\right). \tag{18}$$

Replacing the white noise $\eta(t)$ in the OU process, defined by $\langle\eta\rangle = $ const. and $\langle\eta(t)\eta(t')\rangle = \nu\delta(t - t') + \nu^2$ (Gerstner & Kistler, 2002), with (Gaussian) colored noise η_c, defined by $\langle\eta_c\rangle = $ const. and $\langle\eta_c(t)\eta_c(t')\rangle = \gamma(t - t')$ (Häunggi & Jung, 1994) where $\gamma(t - t')$ is a function that does not vanish for $t - t' \neq 0$, the stationary solution of the FMP distribution is still given by a Gaussian with mean μ' and width σ' (Cáceres, 1999; Häunggi & Jung, 1994). Since the noise correlations only appear when calculating higher-order moments of the FMP, the mean value of the FMP distribution remains unchanged $\mu' = \mu$. However, the variance $\sigma'^2 = \langle(u(t) - \langle u(t)\rangle)^2\rangle$ of the stationary FMP distribution changes due to the correlations, as discussed in the next section.

6.2. Width of free membrane potential distribution

In the HCS, the FMP can be approximated analytically as (Petrovici, 2016)

$$u(t) = u_0 + \sum_{k \in \{e,i\}} \sum_{\text{spikes } s} \Lambda_k \Theta(t - t_s)$$
$$\cdot \left[\exp\left(-\frac{t - t_s}{\tau_k^{\text{syn}}}\right) - \exp\left(-\frac{t - t_s}{\langle\tau_{\text{eff}}\rangle}\right)\right], \tag{19}$$

with

$$u_0 = \frac{g_l E_l + (\langle g^{\text{tot}}\rangle - g_l)\mu}{\langle g^{\text{tot}}\rangle}, \tag{20a}$$

$$\Lambda_k = \frac{\tau_k^{\text{syn}} w_k (E_k^{\text{rev}} - \mu)}{\langle g^{\text{tot}}\rangle (\tau_k^{\text{syn}} - \langle\tau_{\text{eff}}\rangle)}, \tag{20b}$$

$$\langle\tau_{\text{eff}}\rangle = \frac{C_m}{\langle g^{\text{tot}}\rangle}. \tag{20c}$$

By explicitly writing the excitatory and inhibitory noise spike trains as $S_{e/i}(t') = \sum_{\text{spikes } s} \delta(t' - t_s)$, this can be rewritten to

$$u(t) = u_0 + \sum_{k \in \{e,i\}} \Lambda_k \int dt' S_k(t') \Theta(t - t')$$

$$\cdot \left[\exp\left(-\frac{t-t'}{\tau_k^{\text{syn}}}\right) - \exp\left(-\frac{t-t'}{\langle\tau_{\text{eff}}\rangle}\right) \right] \tag{21a}$$

$$= u_0 + \Lambda_e(S_e * \kappa_e)(t) + \Lambda_i(S_i * \kappa_i)(t) \tag{21b}$$

$$= u_0 + [(\Lambda_e S_e + \Lambda_i S_i) * \kappa](t), \tag{21c}$$

where $*$ denotes the convolution operator and with

$$\kappa_{e/i}(t) = \Theta(t) \left[\exp\left(-\frac{t}{\tau_{e/i}^{\text{syn}}}\right) - \exp\left(-\frac{t}{\langle\tau_{\text{eff}}\rangle}\right) \right]. \tag{22}$$

For simplicity, we assume $\tau_e^{\text{syn}} = \tau_i^{\text{syn}}$. The width of the FMP distribution can now be calculated as

$$\langle (u(t) - \langle u(t) \rangle)^2 \rangle = \langle u(t)^2 \rangle - \langle u(t) \rangle^2 \tag{23a}$$

$$= \langle [u_0 + (S_{\text{tot}} * \kappa)(t)]^2 \rangle - \langle u_0 + (S_{\text{tot}} * \kappa)(t) \rangle^2 \tag{23b}$$

$$= \langle [(S_{\text{tot}} * \kappa)(t)]^2 \rangle - \langle (S_{\text{tot}} * \kappa)(t) \rangle^2 \tag{23c}$$

$$= \langle [(S_{\text{tot}} * \kappa)(t) - \langle (S_{\text{tot}} * \kappa)(t) \rangle]^2 \rangle, \tag{23d}$$

where the average is calculated over t and $S_{\text{tot}}(t) = \Lambda_e S_e(t) + \Lambda_i S_i(t)$. Since the average is an integral over t, i.e. $\langle(\cdot)\rangle \to \lim_{T\to\infty} \frac{1}{T} \int_{-T/2}^{T/2}(\cdot) \, dt$, we can use the identity $\int(f * g)(t) \, dt = (\int f(t) \, dt)(\int g(t) \, dt)$, that is $\langle f * g \rangle = \langle f \rangle \int g(t) \, dt = \langle f \rangle * g$ in the limit of $T \to \infty$, to arrive at the following solution:

$$\langle (u(t) - \langle u(t) \rangle)^2 \rangle$$
$$= \langle [(S_{\text{tot}}(t) - \langle S_{\text{tot}}(t) \rangle) * \kappa(t)]^2 \rangle. \tag{24}$$

More generally, we obtain with a similar calculation the autocorrelation function (ACF) of the FMP:

$$\langle \bar{u}(t) \bar{u}(t + \Delta) \rangle$$
$$= \langle ((\bar{S}_{\text{tot}} * \kappa)(t))((\bar{S}_{\text{tot}} * \kappa)(t + \Delta)) \rangle, \tag{25}$$

with $\bar{x}(t) = x(t) - \langle x(t) \rangle$ and by using $\langle \bar{u}(t) \bar{u}(t + \Delta) \rangle = \langle u(t) u(t + \Delta) \rangle - \langle u(t) \rangle^2$. This can be further simplified by applying the Wiener–Khintchine theorem (Khintchine, 1934; Wiener, 1930), which states that $\lim_{T\to\infty} \langle x(t) x(t + \Delta) \rangle_T = \mathcal{F}^{-1}(|\mathcal{F}(x)|^2)(\Delta)$ with $\langle(\cdot)\rangle_T \to \int_{-T/2}^{T/2}(\cdot) \, dt$ (due to $\int x(t) x(t + \Delta) \, dt = (x(t) * x(-t))(\Delta)$). Thus, for the limit $T \to \infty$, we can rewrite this as

$$\langle ((\bar{S}_{\text{tot}} * \kappa)(t))((\bar{S}_{\text{tot}} * \kappa)(t + \Delta)) \rangle$$

$$= \lim_{T\to\infty} \frac{1}{T} \mathcal{F}^{-1}(|\mathcal{F}(\bar{S}_{\text{tot}} * \kappa)|^2)(\Delta) \tag{26a}$$

$$= \lim_{T\to\infty} \frac{1}{T} \mathcal{F}^{-1}(|\mathcal{F}(\bar{S}_{\text{tot}})\mathcal{F}(\kappa)|^2)(\Delta) \tag{26b}$$

$$= \lim_{T\to\infty} \frac{1}{T} \left(\mathcal{F}^{-1}(|\mathcal{F}(\bar{S}_{\text{tot}})|^2) \right.$$

$$\left. * \mathcal{F}^{-1}(|\mathcal{F}(\kappa)|^2) \right)(\Delta), \tag{26c}$$

and by applying the Wiener–Khintchine theorem again in reverse

$$\langle \bar{u}(t) \bar{u}(t + \Delta) \rangle = \left(\lim_{T\to\infty} \frac{1}{T} \langle [\bar{S}_{\text{tot}}(t)][\bar{S}_{\text{tot}}(t + \Delta')] \rangle_T \right.$$

$$\left. * \langle [\kappa(t)][\kappa(t + \Delta')] \rangle_\infty \right)(\Delta), \tag{27}$$

where the variance of the FMP distribution is given for $\Delta = 0$. Thus, the unnormalized ACF of the FMP can be calculated by convolving the unnormalized ACF of the background spike trains (S_{tot}) and the PSP shape (κ). In case of independent excitatory and inhibitory Poisson noise (i.e., $\langle \bar{S}(t)\bar{S}(t') \rangle = \nu \delta(t - t')$), we get

$$\langle [\bar{S}_{\text{tot}}(t)][\bar{S}_{\text{tot}}(t + \Delta')] \rangle = \Lambda_e^2 \langle \bar{S}_e(t) \bar{S}_e(t + \Delta') \rangle$$

$$+ \Lambda_i^2 \langle \bar{S}_i(t) \bar{S}_i(t + \Delta') \rangle \tag{28a}$$

$$= \sum_{k \in \{e,i\}} \Lambda_k^2 \nu_k \delta(\Delta') \tag{28b}$$

and therefore

$$\text{Var}(u) = \Big(\sum_{k \in \{e,i\}} \Lambda_k^2 \nu_k \delta(\Delta')$$

$$* \langle [\kappa(t)][\kappa(t + \Delta')] \rangle_\infty \Big)(\Delta = 0) \tag{29a}$$

$$= \sum_{k \in \{e,i\}} \Lambda_k^2 \nu_k \langle \kappa^2(t) \rangle \tag{29b}$$

$$= \sum_{k \in \{e,i\}} \Lambda_k^2 \nu_k \int_0^\infty \kappa^2(t) \, dt, \tag{29c}$$

which agrees with the result given in Petrovici (2016). If the noise spike trains are generated by processes with refractory periods, the absence of spikes between refractory periods leads to negative contributions in the ACF of the noise spike trains. This leads to a reduced value of the variance of the FMP and hence, also to a reduced width of the FMP distribution. The factor $\sqrt{\beta}$ by which the width of the FMP distribution (Eqn. (11)) changes due to the introduction of colored background noise is given by

$$\beta = \frac{\sigma_{\text{colored}}^2}{\sigma_{\text{Poisson}}^2} \tag{30}$$

$$= \frac{\int d\Delta \ \langle \bar{S}_{\text{tot}}(t) \bar{S}_{\text{tot}}(t + \Delta) \rangle \cdot \int dt \ \kappa(t) \kappa(t + \Delta)}{\sum_{k \in \{e,i\}} \Lambda_k^2 \nu_k \int_0^\infty \kappa^2(t) \, dt}. \tag{31}$$

For the simplified case of a Poisson process with refractory period, one can show that $\int d\Delta \ \langle \bar{S}_{\text{tot}}(t) \bar{S}_{\text{tot}}(t + \Delta) \rangle$ has a reduced value compared to a Poisson process without refractory period (Gerstner & Kistler, 2002), leading to $\beta \leq 1$. Even though we do not show this here for neuron-generated spike trains, the two cases are similar enough that $\beta \leq 1$ can be assumed to apply in this case as well.

In the next section, we will show that the factor β can be used to rescale the inverse slope of the activation function to transform the activation function of a neuron receiving white noise to the activation function of a neuron receiving equivalent (in frequency and weights), but colored noise. That is, the rescaling of the FMP distribution width due to the autocorrelated background noise translates into a rescaled inverse slope of the activation function.

6.3. Approximate inverse slope of LIF activation function

As stated earlier, the FMP of an LIF neuron in the HCS is described by an OU process with a Gaussian stationary FMP distribution (both for white and colored background noise). As a first approximation, we can define the activation function as the probability of the neuron having a FMP above threshold (see Eq. (18))

$$p(z_i = 1) \approx \int_\vartheta^\infty f(u) du \tag{32}$$

$$= \int_\vartheta^\infty \sqrt{\frac{1}{2\pi\sigma^2}} \exp\left(-\frac{(u-\mu)^2}{2\sigma^2}\right) du \tag{33}$$

$$= \frac{1}{2} \left(1 - \text{erf}\left(\frac{\vartheta - \mu}{\sqrt{2}\sigma}\right) \right). \tag{34}$$

Even though this is only an approximation (as we are neglecting the effect of the reset), the error function is already similar to the logistic activation function observed in simulations (Petrovici et al., 2016).

The inverse slope of a logistic activation function is defined at the inflection point, i.e.,

$$\alpha^{-1} = \frac{\mathrm{d}}{\mathrm{d}\mu}\varphi\left(\frac{\mu - u_0}{\alpha}\right)\bigg|_{\mu = u_0}. \tag{35}$$

By calculating the inverse slope via the activation function derived from the FMP distribution, we get

$$\alpha^{-1} = \frac{\mathrm{d}}{\mathrm{d}\mu}p(z_i = 1)\bigg|_{\mu = \vartheta}, \tag{36}$$

$$= \sqrt{\frac{1}{2\pi\sigma^2}}, \tag{37}$$

from which it follows that the inverse slope α is proportional to the width of the FMP distribution σ. Thus, rescaling the variance of the FMP distribution by a factor β leads, approximately, to a rescaling of the inverse slope of the activation function $\alpha' = \sqrt{\beta}\alpha$.

6.4. Origin of side-peaks in the noise autocorrelation function

For high rates, the spike train generated by an LIF neuron in the HCS shows regular patterns of interspike intervals which are roughly equal to the absolute refractory period. This occurs (i) due to the refractory period introducing regularity for higher rates, since ISI's $< \tau_{\mathrm{ref}}$ are not allowed and the maximum firing rate of the LIF neuron is bounded by $\frac{1}{\tau_{\mathrm{ref}}}$, and (ii) due to an LIF neurons's tendency to spike consecutively when the effective membrane potential

$$u_{\mathrm{eff}}(t) = \frac{g_l E_l + \sum_{k \in \{e,i\}} g_k^{\mathrm{syn}}(t) E_k^{\mathrm{rev}}}{g^{\mathrm{tot}}(t)}, \tag{38}$$

$$\tau_{\mathrm{eff}}\dot{u} = u_{\mathrm{eff}} - u, \tag{39}$$

is suprathreshold after the refractory period (Petrovici, 2016). The probability of a consecutive spike after the neuron has spiked once at time t is given by (under the assumption of the HCS)

$$p_1 = p(\text{spike at } t + \tau_{\mathrm{ref}}| \text{ first spike at } t) \tag{40}$$

$$= \int_{\vartheta}^{\infty} \mathrm{d}u_{t+\tau_{\mathrm{ref}}} f(u_{t+\tau_{\mathrm{ref}}}, \ \tau_{\mathrm{ref}} \mid u_t = \vartheta),$$

due to the effective membrane potential following an Ornstein–Uhlenbeck process (whereas the FMP is a low-pass filter thereof, however with a very low time constant τ_{eff}), see Eq. (16). The probability to spike again after the second refractory period is then given by

$$p_2 = p(\text{spike at } t + 2\tau_{\mathrm{ref}}| \text{ spike at } t + \tau_{\mathrm{ref}}, \ t) \tag{41}$$

$$= \frac{\int_{\vartheta}^{\infty} \int_{\vartheta}^{\infty} \mathrm{d}u_2 \mathrm{d}u_1 \, f(u_2, \tau_{\mathrm{ref}} \mid u_1) f(u_1, \tau_{\mathrm{ref}} \mid u_0 = \vartheta)}{\int_{\vartheta}^{\infty} \mathrm{d}u_1 \, f(u_1, \tau_{\mathrm{ref}} \mid u_0 = \vartheta)},$$

with $u_n = u_{t+n\tau_{\mathrm{ref}}}$, or in general after $n - 1$ spikes

$$p_n = p(\text{spike at } t + n\tau_{\mathrm{ref}} \mid \text{spike at } t + (n-1)\tau_{\mathrm{ref}}, \ldots, \ t)$$

$$= \int_{\vartheta}^{\infty} \mathrm{d}u_{n-1} \, f^{(n-1)}(u_{n-1}), \tag{42}$$

$$f^n(u_n) = \frac{\int_{\vartheta}^{\infty} \mathrm{d}u_{n-1} \, f(u_n, \tau_{\mathrm{ref}} \mid u_{n-1}) f^{(n-1)}(u_{n-1})}{\int_{\vartheta}^{\infty} \mathrm{d}u_{n-1} \, f^{(n-1)}(u_{n-1})}, \tag{43}$$

for $n > 1$ and $f^1(u_1) = f(u_1, \tau_{\mathrm{ref}} \mid u_0 = \vartheta)$. The probability to observe n spikes in such a sequence is then given by

$$P_n = \prod_{i=1}^{n-1} p_i, \tag{44}$$

and the probability to find a burst of length n (i.e., the burst ends)

$$p(\text{burst of length n}) = P_n \cdot (1 - p_n). \tag{45}$$

With this, one can calculate the average length of the occurring bursts $\sum_{i=1}^{\infty} i \cdot p(\text{burst of length i})$, from which we can already see how the occurrence of bursts depends on the mean activity of the neuron. A simple solution can be found for the special case of $\tau^{\mathrm{syn}} \ll \tau_{\mathrm{ref}}$, since then the effective membrane potential distribution has already converged to the stationary distribution after every refractory period, i.e., $f(u_n, \tau_{\mathrm{ref}} \mid u_{n-1}) = f(u_n)$ and hence

$$p_n = p(\text{spike at } t + n\tau_{\mathrm{ref}}| \text{ spike at } t + (n - 1)\tau_{\mathrm{ref}}, \ldots, \ t)$$

$$= \int_{\vartheta}^{\infty} \mathrm{d}u f(u) = \bar{p} \tag{46}$$

for all n. Thus, for this special case the average burst length can be expressed as

$$\sum_{i=1}^{\infty} i \cdot p(\text{burst of length i}) = \sum_{i=1}^{\infty} i \cdot \bar{p}^{i-1}(1 - \bar{p}), \tag{47}$$

$$= \frac{1}{1 - \bar{p}}. \tag{48}$$

By changing the mean membrane potential (e.g., by adjusting the leak potential or adding an external (bias) current), the probability of consecutive spikes $\bar{p}$ can be directly adjusted and hence, also the average length of bursts. Since these bursts are fixed structures with interspike intervals equal to the refractory period, they translate into side-peaks at multiples of the refractory period in the spike train ACF, as we demonstrate below.

The ACF of the spike train S is given by

$$C(S, S, \Delta) = \frac{\langle S_t S_{t+\Delta} \rangle - \langle S \rangle^2}{\mathrm{Var}(S)}, \tag{49}$$

where the first term of the numerator is $\langle S_t S_{t+\Delta} \rangle = p(\text{spike at } t + \Delta, \text{spike at } t)$ (notation as in Eqs. (21a) and (23a)). This term can be expressed as

$$p(\text{spike at } t + \Delta, \text{spike at } t)$$

$$= p(\text{spike at } t + \Delta \mid \text{spike at } t) \cdot p(\text{spike at } t), \tag{50}$$

$$= p(\text{spike at } t + \Delta \mid \text{spike at } t) \cdot \langle S \rangle, \tag{51}$$

where we assumed that the first spike starts the burst at a random time t. Therefore, in order to calculate the ACF, we have to calculate the probability that a spike occurs at time $t + \Delta$ given that the neuron spikes at time t. This has to include every possible combination of spikes during this interval. In the following, we argue that at multiples of the refractory period, the main contribution to the ACF comes from bursts.

- First, for $\Delta < \tau_{\mathrm{ref}}$, the term $p(\text{spike at } t + \Delta \mid \text{spike at } t)$ in Eq. (50) vanishes since the neuron is refractory and cannot spike during this interval. Thus, the ACF becomes negative as only the term $-\frac{\langle S \rangle^2}{\mathrm{Var}(S)}$ in Eq. (49) remains, where both numerator and denominator are positive.
- For $\Delta = \tau_{\mathrm{ref}}$, a spike can only occur when the neuron bursts with probability $p_1 = \int_{\vartheta}^{\infty} \mathrm{d}u_{t+\tau_{\mathrm{ref}}} f(u_{t+\tau_{\mathrm{ref}}}, \ \tau_{\mathrm{ref}} \mid u_t = \vartheta)$, where we assumed for simplicity that the first spike starts the burst spontaneously.

- Since for $\tau_{\text{ref}} < \Delta < 2\tau_{\text{ref}}$, the neuron did not burst with probability $1-p_1$, it is possible to find a spike in this interval, leading again to negative, but diminished, values in the ACF.
- For $\Delta = 2\tau_{\text{ref}}$, we now have two ways to observe spikes at t and $t + 2\tau_{\text{ref}}$: (i) The spikes are part of a burst of length 2 or (ii) there was no intermediate spike and the spikes have an ISI of $2\tau_{\text{ref}}$. Since for larger rates, having large ISIs that are exact multiples of τ_{ref} is unlikely, we can neglect the contribution of (ii).
- If we go further to $\Delta = n\tau_{\text{ref}}$, we get even more additional terms including bursts of length $< n$. However, these terms can again be neglected as, compared to having a burst of length n, it is rather unlikely to get a burst pattern with missing intermediate spikes, i.e., having partial bursts which are a multiple of τ_{ref} apart.
- Finally, for $\Delta \to \infty$, we have $\langle S_t S_{t+\Delta}\rangle - \langle S\rangle^2 = \langle S_t\rangle \langle S_{t+\Delta}\rangle - \langle S\rangle^2 = \langle S\rangle \langle S\rangle - \langle S\rangle^2 = 0$ and the ACF (Eq. (49)) vanishes.

Consequently, we can approximate the ACF at multiples of the refractory period by calculating the probability of finding a burst of n spikes (Eq. (43)):

$$C(S, S, n\tau_{\text{ref}}) \approx \sum_{k=1}^{\infty} P_{k+1}\delta\big([n-k]\tau_{\text{ref}}\big), \tag{52}$$

and for the special case of $\tau^{\text{syn}} \ll \tau_{\text{ref}}$

$$C(S, S, n\tau_{\text{ref}}) \approx \sum_{k=1}^{\infty} \bar{p}^k \delta\big([n-k]\tau_{\text{ref}}\big) \tag{53}$$

$$= \sum_{k=1}^{\infty} e^{k \ln \bar{p}}\delta\big([n-k]\tau_{\text{ref}}\big). \tag{54}$$

Hence, since increasing the mean rate (or bias) of the neuron leads to an increase in $\bar{p}$ and thus to a reduced decay constant $\ln \bar{p}$, more significant side-peaks emerge.

For $\tau^{\text{syn}} \approx \tau_{\text{ref}}$, the effective membrane distribution is not yet stationary and therefore, this approximation does not hold. To arrive at the exact solution, one would have to repeat the above calculation for all possible spike time combinations, leading to a recursive integral (Gerstner & Kistler, 2002). Furthermore, one would also need to take into account the situation where the first spike is itself part of a burst, i.e., is not the first spike in the burst. To circumvent a more tedious calculation, we use an approximation which is in between the two cases $\tau^{\text{syn}} \ll \tau_{\text{ref}}$ and $\tau^{\text{syn}} \approx \tau_{\text{ref}}$: we use $\bar{p} = \int_{\vartheta}^{\infty} \mathrm{d}u f(u, \tau_{\text{ref}} \mid \vartheta)$, which provides a reasonable approximation for short bursts.

6.5. Cross-correlation of free membrane potentials receiving correlated input

Similarly to the ACF of the membrane potential, one can calculate the crosscorrelation function of the FMPs of two neurons receiving correlated noise input. First, the membrane potentials are given by

$$u_1 = u_0^1 + S_{\text{tot},1} * \kappa, \tag{55}$$

$$u_2 = u_0^2 + S_{\text{tot},2} * \kappa. \tag{56}$$

The covariance function can be written as

$$\langle \bar{u}_1(t)\bar{u}_2(t + \Delta)\rangle$$

$$= \langle u_1(t)u_2(t + \Delta)\rangle - \langle u_1(t)\rangle \langle u_2(t)\rangle \tag{57a}$$

$$= \langle (S_{\text{tot},1} * \kappa)(t)(S_{\text{tot},2} * \kappa)(t + \Delta)\rangle$$

$$\quad - \langle (S_{\text{tot},1} * \kappa)(t)\rangle \langle (S_{\text{tot},2} * \kappa)(t)\rangle \tag{57b}$$

$$= \cdots$$

$$= \Big(\lim_{T \to \infty} \frac{1}{T} \langle \bar{S}_{\text{tot},1}(t)\bar{S}_{\text{tot},2}(t + \Delta')\rangle_T$$

$$\quad * \langle \kappa(t)\kappa(t + \Delta')\rangle_{\infty}\Big)(\Delta), \tag{57c}$$

with $\bar{u} = u - \langle u\rangle$, from which we obtain the crosscorrelation function by normalizing with the product of standard deviations of u_1 and u_2 (for notation, see Eq. (27)). The term containing the input correlation coefficient is $\langle \bar{S}_{\text{tot},1}(t)\bar{S}_{\text{tot},2}(t + \Delta')\rangle$. Plugging in the spike trains, we get four crosscorrelation terms

$$\langle \bar{S}_{\text{tot},1}(t)\bar{S}_{\text{tot},2}(t + \Delta')\rangle$$

$$= \sum_{l,m\in\{e,i\}} \Lambda_{l,1}\Lambda_{m,2} \langle \bar{S}_{l,1}(t)\bar{S}_{m,2}(t + \Delta')\rangle. \tag{58}$$

Since excitatory as well as inhibitory noise inputs are randomly drawn from the same pool of neurons, we can assume that $\langle \bar{S}_{l,1}(t)\bar{S}_{m,2}(t + \Delta')\rangle$ is approximately equal for all combinations of synapse types when averaging over enough inputs, regardless of the underlying correlation structure/distribution of the noise pool. The first term, however, depends on the synapse types since the Λ-terms (Eq. (20b)) contain the distance between reversal potentials and mean FMP:

$$\langle \bar{u}_1(t)\bar{u}_2(t + \Delta)\rangle$$

$$= \zeta_1\zeta_2 \sum_{l,m\in\{e,i\}} w_l w_m \big(E_l^{\text{rev}} - \mu_1\big)\big(E_m^{\text{rev}} - \mu_2\big)$$

$$\cdot \left[\langle \bar{S}_{l,1}(t)\bar{S}_{m,2}(t + \Delta')\rangle * \langle \kappa(t)\kappa(t + \Delta')\rangle\right](\Delta), \tag{59}$$

with constants $\zeta_i = \frac{\tau^{\text{syn}}}{\langle g^{\text{tot}}_i\rangle}\big(\tau^{\text{syn}} - \langle \tau_{\text{eff}}^i\rangle\big)$. The cross-correlation vanishes when, after summing over many inputs, the following identities hold:

$$\langle \Lambda_{e,1}\Lambda_{e,2}\rangle_{\text{inputs}} = -\langle \Lambda_{e,1}\Lambda_{i,2}\rangle_{\text{inputs}}, \tag{60a}$$

$$\langle \Lambda_{i,1}\Lambda_{i,2}\rangle_{\text{inputs}} = -\langle \Lambda_{i,1}\Lambda_{e,2}\rangle_{\text{inputs}}, \tag{60b}$$

where $\langle\langle(\cdot)\rangle\rangle$ is an average over all inputs, i.e., all neurons that provide noise.

While not relevant for our simulations, it is worth noting that the excitatory and inhibitory weights with which each neuron contributes its spike trains can be randomly drawn from non-identical distributions. By enforcing the following correlation between the noise weights of both neurons, one can introduce a skew into the weight distribution which compensates for the differing distance to the reversal potentials:

$$(E_{\text{rev}}^{e,1} - \mu_1)(E_{\text{rev}}^{e,2} - \mu_2)\langle w_e^1 w_e^2\rangle_{\text{inputs}}$$

$$= -(E_{\text{rev}}^{e,1} - \mu_1)(E_{\text{rev}}^{i,2} - \mu_2)\langle w_e^1 w_i^2\rangle_{\text{inputs}} \tag{61}$$

A simple procedure to accomplish this is the following: First, we draw the absolute weights w^1 and w^2 from an arbitrary distribution and assign synapse types randomly with probabilities $p_{e/i}$ afterwards. If w^2 is excitatory, we multiply w^1 by $\frac{|E_{\text{rev}}^{i,2}-\mu_2|}{p_e|E_{\text{rev}}^{i,2}-\mu_2|+p_i|E_{\text{rev}}^{e,2}-\mu_2|}$, otherwise by $\frac{|E_{\text{rev}}^{e,2}-\mu_2|}{p_e|E_{\text{rev}}^{i,2}-\mu_2|+p_i|E_{\text{rev}}^{e,2}-\mu_2|}$. This way, $\langle w^1\rangle$ remains unchanged and the resulting weights suffice Eq. (61).

6.6. State space switch from {0,1} to {-1,1}

To switch from the state space $\mathbf{z} \in \{0, 1\}$ to $\mathbf{z}' \in \{-1, 1\}$ while conserving the state probabilities (i.e., $p(\mathbf{z}) = p(\mathbf{z}')$) one has to adequately transform the distribution parameters $\mathbf{W}$ and $\mathbf{b}$. Since the distributions are of the form $p(\mathbf{z}) = \exp\big(\mathbf{z}^T\mathbf{W}\mathbf{z} + \mathbf{z}^T\mathbf{b}\big)$, this is equivalent to requiring that the energy $E(\mathbf{z}) = \mathbf{z}^T\mathbf{W}\mathbf{z} + \mathbf{z}^T\mathbf{b}$ of each state remains, up to a constant, unchanged.

First, we can write the energy of a state $\mathbf{z}'$ and use the transformation $\mathbf{z}' = 2\mathbf{z} - 1$ to get

$$E(\mathbf{z}') = \frac{1}{2} \sum_{i,j} \mathbf{z}'_i \mathbf{W}'_{ij} \mathbf{z}'_j + \sum_i \mathbf{z}'_i \mathbf{b}'_i \tag{62a}$$

$$= \frac{1}{2} \left(4 \sum_{i,j} \mathbf{z}_i \mathbf{W}'_{ij} \mathbf{z}_j - 2 \sum_{i,j} \mathbf{z}_i \mathbf{W}'_{ij} - 2 \sum_{i,j} \mathbf{W}'_{ij} \mathbf{z}_j \right.$$

$$\left. + \sum_{i,j} \mathbf{W}'_{ij} \right) - \sum_i \mathbf{b}'_i + 2 \sum_i \mathbf{z}_i \mathbf{b}'_i \tag{62b}$$

$$= \frac{1}{2} \sum_{i,j} \mathbf{z}_i 4 \mathbf{W}'_{ij} \mathbf{z}_j + \sum_i \mathbf{z}_i \left(2 \mathbf{b}'_i \right.$$

$$\left. - 2 \sum_j \mathbf{W}'_{ij} \right) + C , \tag{62c}$$

where C is a constant $C = \frac{1}{2} \sum_{i,j} \mathbf{W}'_{ij} - \sum_i \mathbf{b}'_i$ and we used the fact that $\mathbf{W}'_{ij}$ is symmetric. Since constant terms in the energy leave the probability distribution invariant, we can simply compare $E(\mathbf{z}')$ and $E(\mathbf{z})$

$$E(\mathbf{z}) = \frac{1}{2} \sum_{i,j} \mathbf{z}_i^T \mathbf{W}_{ij} \mathbf{z}_j + \sum_i \mathbf{z}_i^T \mathbf{b}_i , \tag{63}$$

and extract the correct parameter transformation:

$$\mathbf{W}_{ij} = 4 \mathbf{W}'_{ij} , \tag{64}$$

$$\mathbf{b}_i = 2 \mathbf{b}'_i - 2 \sum_j \mathbf{W}'_{ij} . \tag{65}$$

From this, we can also calculate the inverse transformation rule for $\mathbf{z} = \frac{1}{2}(\mathbf{z}' + 1)$:

$$\mathbf{W}'_{ij} = \frac{1}{4} \mathbf{W}_{ij} , \tag{66}$$

$$\mathbf{b}'_i = \frac{1}{2} \mathbf{b}_i + \frac{1}{4} \sum_j \mathbf{W}_{ij} . \tag{67}$$

6.7. Translation from Boltzmann to neurosynaptic parameters

As discussed in the methods section, following Petrovici et al. (2016), the activation function of LIF neurons in the HCS is approximately logistic and can be written as

$$p(\mathbf{z}_k = 1 \mid \mathbf{z}_{/k}) = \varphi(\mu)$$

$$= (1 + \exp(-(\mu - u_0)/\alpha))^{-1} , \tag{68}$$

where $\mathbf{z}_{/k}$ is the state vector of all other neurons except the k'th one and μ the mean membrane potential (Eq. (17b)). u_0 and α are the inflection point and the inverse slope, respectively. Furthermore, the conditional probability $p(\mathbf{z}_k = 1 \mid \mathbf{z}_{/k})$ of a Boltzmann distribution over binary random variables $\mathbf{z}_k$, i.e., $p(\mathbf{z}) \propto \exp\left(\frac{1}{2}\mathbf{z}^T\mathbf{W}\mathbf{z} + \mathbf{z}^T\mathbf{b}\right)$, is given by

$$p(\mathbf{z}_k = 1 \mid \mathbf{z}_{/k})$$

$$= \left(1 + \exp\left(-\sum_j W_{kj}\mathbf{z}_j - b_k\right) \right)^{-1} , \tag{69}$$

with symmetric weight matrix $\mathbf{W}$, $W_{ii} = 0 \,\forall i$, and biases $\mathbf{b}$. These equations allow a translation from the parameters of a Boltzmann distribution (b_i, W_{ij}) to parameters of LIF neurons and their synapses (E_l, w_{ij}), such that the state dynamic of the network approximates sampling from the target Boltzmann distribution.

First, the biases $\mathbf{b}$ can be mapped to leak potentials E_l (or external currents) by requiring that, for $\mathbf{W} = 0$ (that is, no synaptic input from other neurons), the activity of each neuron equals the conditional probability of the target Boltzmann distribution

$$(1 + \exp(-(\mu - u_0)/\alpha))^{-1} \overset{!}{=} (1 + \exp(-b_k))^{-1} , \tag{70}$$

leading to the translation rule

$$\mathbf{E}_l = \frac{\tau_m}{\tau_{\text{eff}}}(\alpha \mathbf{b} + u_0) - \sum_{x \in \{e,i\}} \frac{\langle g_x^{\text{syn}} \rangle}{g_l} E_x^{\text{rev}} . \tag{71}$$

To map Boltzmann weights W_{ij} to synaptic weights w_{ij}, we first have to rescale the W_{ij}, as done for the biases in Eq. (71). However, leaky integrator neurons have non-rectangular PSPs, so their interaction strength is modulated over time. This is different from the interaction in Boltzmann machines, where the PSP shape is rectangular (Glauber dynamics). Nevertheless, we can derive a heuristic translation rule by requiring that the mean interaction during the refractory period of the presynaptic neuron is the same in both cases, i.e.,

$$\int_0^{\tau_{\text{ref}}} dt \, PSP(t) \overset{!}{=} \int_0^{\tau_{\text{ref}}} dt \, \alpha W_{ij} \tag{72a}$$

$$= \alpha W_{ij} \tau_{\text{ref}} , \tag{72b}$$

where $PSP(t)$ is given by Eq. (19). From this, we get the translation rule for synaptic weights:

$$w_{kj} = \frac{\alpha W_{kj} C_m \frac{\tau_{\text{ref}}}{\tau^{\text{syn}}} \left(1 - \frac{\tau^{\text{syn}}}{\tau_{\text{eff}}}\right) \left(E_{kj}^{\text{rev}} - \mu\right)^{-1}}{\left[\tau^{\text{syn}}\left(e^{-\frac{\tau_{\text{ref}}}{\tau^{\text{syn}}}} - 1\right) - \tau_{\text{eff}}\left(e^{-\frac{\tau_{\text{ref}}}{\tau_{\text{eff}}}} - 1\right)\right]} . \tag{73}$$

Acknowledgments

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.neunet.2019.08.002.

References

Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9(1), 147–169.

Aitchison, L., & Lengyel, M. (2016). The Hamiltonian brain: efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLoS Computational Biology*, 12(12), e1005186.

Arieli, A., Sterkin, A., Grinvald, A., & Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science*, 273(5283), 1868.

Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5), 358.

Azouz, R., & Gray, C. M. (1999). Cellular mechanisms contributing to response variability of cortical neurons in vivo. *Journal of Neuroscience*, 19(6), 2209–2223.

Bertolero, M. A., Yeo, B. T., & D'Esposito, M. (2015). The modular and integrative functional architecture of the human brain. *Proceedings of the National Academy of Sciences*, 112(49), E6798–E6807.

Branco, T., & Staras, K. (2009). The probability of neurotransmitter release: variability and feedback control at single synapses. *Nature Reviews Neuroscience*, 10(5), 373.

Brascamp, J. W., Van Ee, R., Noest, A. J., Jacobs, R. H., & van den Berg, A. V. (2006). The time course of binocular rivalry reveals a fundamental role of noise. *Journal of Vision*, 6(11), 8.

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8(3), 183–208.

Buesing, L., Bill, J., Nessler, B., & Maass, W. (2011). Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLOS Computational Biology*, 7(11), e1002211.

Bullmore, E., & Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3), 186.

Cáceres, M. O. (1999). Harmonic potential driven by long-range correlated noise. *Physical Review E*, 60(5), 5208.

Chen, Z. J., He, Y., Rosa-Neto, P., Germann, J., & Evans, A. C. (2008). Revealing modular architecture of human brain structural networks by using cortical thickness from MRI. *Cerebral Cortex*, 18(10), 2374–2381.

Chichilnisky, E. (2001). A simple white noise analysis of neuronal light responses. *Network. Computation in Neural Systems*, 12(2), 199–213.

Churchland, M. M., Byron, M. Y., Ryu, S. I., Santhanam, G., & Shenoy, K. V. (2006). Neural variability in premotor cortex provides a signature of motor preparation. *Journal of Neuroscience*, 26(14), 3697–3712.

Cohen, G., Afshar, S., Tapson, J., & van Schaik, A. (2017). EMNIST: an extension of MNIST to handwritten letters. arXiv preprint arXiv:1702.05373.

Davison, A., Brderle, D., Eppler, J., Kremkow, J., Muller, E., Pecevski, D., Perrinet, L., & Yger, P. (2009). Pynn: a common interface for neuronal network simulators. *Frontiers in Neuroinformatics*, 2, 11. http://dx.doi.org/10.3389/neuro.11.011.2008, https://www.frontiersin.org/article/10.3389/neuro.11.011.2008.

Dayan, P., et al. (2003). Theoretical neuroscience: computational and mathematical modeling of neural systems. *Journal of Cognitive Neuroscience*, 15(1), 154–155.

Deco, G., Rolls, E. T., & Romo, R. (2009). Stochastic dynamics as a principle of brain function. *Progress in Neurobiology*, 88(1), 1–16.

Deger, M., Helias, M., Boucsein, C., & Rotter, S. (2012). Statistical properties of superimposed stationary spike trains. *Journal of Computational Neuroscience*, 32(3), 443–463.

Destexhe, A., Rudolph, M., & Paré, D. (2003). The high-conductance state of neocortical neurons in vivo. *Nature Reviews Neuroscience*, 4(9), 739.

Faisal, A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4), 292.

Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3), 119–130.

Fiser, J., Chiu, C., & Weliky, M. (2004). Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature*, 431(7008), 573.

Fourcaud, N., & Brunel, N. (2002). Dynamics of the firing probability of noisy integrate-and-fire neurons. *Neural Computation*, 14(9), 2057–2110.

Fuhrmann, G., Segev, I., Markram, H., & Tsodyks, M. (2002). Coding of temporal information by activity-dependent synapses. *Journal of Neurophysiology*, 87(1), 140–148.

Gerstner, W., & Kistler, W. M. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press.

Gerstner, W., Kistler, W. M., Naud, R., & Paninski, L. (2014). *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press.

Gewaltig, M.-O., & Diesmann, M. (2007). Nest (neural simulation tool). *Scholarpedia*, 2(4), 1430.

Häunggi, P., & Jung, P. (1994). Colored noise in dynamical systems. *Advances in Chemical Physics*, 89, 239–326.

Henry, G., Bishop, P., Tupper, R., & Dreher, B. (1973). Orientation specificity and response variability of cells in the striate cortex. *Vision Research*, 13(9), 1771–1779.

Jordan, J., Petrovici, M. A., Breitwieser, O., Schemmel, J., Meier, K., Diesmann, M., & Tetzlaff, T. (2017). Stochastic neural computation without noise. arXiv preprint arXiv:1710.04931.

Kheradpisheh, S. R., Ganjtabesh, M., Thorpe, S. J., & Masquelier, T. (2018). STDP-Based spiking deep convolutional neural networks for object recognition. *Neural Networks*, 99, 56–67.

Khintchine, A. (1934). Korrelationstheorie der stationären stochastischen Prozesse. *Mathematische Annalen*, 109(1), 604–615.

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247.

Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1), 79–86.

Kumar, A., Schrader, S., Aertsen, A., & Rotter, S. (2008). The high-conductance state of cortical networks. *Neural Computation*, 20(1), 1–43.

LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324.

Lee, J. H., Delbruck, T., & Pfeiffer, M. (2016). Training deep spiking neural networks using backpropagation. *Frontiers in Neuroscience*, 10, 508.

Leng, L., Martel, R., Breitwieser, O., Bytschok, I., Senn, W., Schemmel, J., Meier, K., & Petrovici, M. A. (2018). Spiking neurons with short-term synaptic plasticity form superior generative networks. *Scientific Reports*, 8(1), 10651.

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11), 1432.

Maass, W. (2016). Searching for principles of brain computation. *Current Opinion in Behavioral Sciences*, 11, 81–92.

Maass, W., & Zador, A. M. (1998). Dynamic stochastic synapses as computational units. In *Advances in neural information processing systems* (pp. 194–200).

Mainen, Z. F., & Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268(5216), 1503.

Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanchez, C. A., et al. (2015). Reconstruction and simulation of neocortical microcircuitry. *Cell*, 163(2), 456–492.

Meunier, D., Lambiotte, R., & Bullmore, E. T. (2010). Modular and hierarchically modular organization of brain networks. *Frontiers in Neuroscience*, 4, 200.

Monteforte, M., & Wolf, F. (2012). Dynamic flux tubes form reservoirs of stability in neuronal circuits. *Physical Review X*, 2(4), 041007.

Moreno-Bote, R. (2014). Poisson-like spiking in circuits with probabilistic synapses. *PLoS Computational Biology*, 10(7), e1003522.

Moreno-Bote, R., Renart, A., & Parga, N. (2008). Theory of input spike auto-and cross-correlations and their effect on the response of spiking neurons. *Neural Computation*, 20(7), 1651–1705.

Neftci, E. O., Pedroni, B. U., Joshi, S., Al-Shedivat, M., & Cauwenberghs, G. (2016). Stochastic synapses enable efficient brain-inspired learning machines. *Frontiers in Neuroscience*, 10, 241.

Nelson, J., Salin, P., Munk, M.-J., Arzi, M., & Bullier, J. (1992). Spatial and temporal coherence in cortico-cortical connections: a cross-correlation study in areas 17 and 18 in the cat. *Visual Neuroscience*, 9(1), 21–37.

Orbán, G., Berkes, P., Fiser, J., & Lengyel, M. (2016). Neural variability and sampling-based probabilistic representations in the visual cortex. *Neuron*, 92(2), 530–543.

Ostojic, S., & Brunel, N. (2011). From spiking neuron models to linear-nonlinear models. *PLoS Computational Biology*, 7(1), e1001056.

Pecevski, D., Buesing, L., & Maass, W. (2011). Probabilistic inference in general graphical models through sampling in stochastic networks of spiking neurons. *PLoS Computational Biology*, 7(12), e1002294.

Petrovici, M. A. (2016). *Form versus function: theory and models for neuronal substrates*. Springer.

Petrovici, M. A., Bill, J., Bytschok, I., Schemmel, J., & Meier, K. (2016). Stochastic inference with spiking neurons in the high-conductance state. *Physical Review E*, 94, 042312. http://dx.doi.org/10.1103/PhysRevE.94.042312, http://link.aps.org/doi/10.1103/PhysRevE.94.042312.

Petrovici, M. A., Schmitt, S., Klähn, J., Stöckel, D., Schroeder, A., Bellec, G., et al. (2017). Pattern representation and recognition with accelerated analog neuromorphic systems. In *Circuits and systems IEEE international symposium on* (pp. 1–4). IEEE.

Petrovici, M. A., Vogginger, B., Müller, P., Breitwieser, O., Lundqvist, M., Muller, L., et al. (2014). Characterization and compensation of network-level anomalies in mixed-signal neuromorphic modeling platforms. *PloS One*, 9(10), e108590.

Plesser, H. E., & Gerstner, W. (2000). Noise in integrate-and-fire neurons: from stochastic input to escape rates. *Neural Computation*, 12(2), 367–384.

Probst, D., Petrovici, M. A., Bytschok, I., Bill, J., Pecevski, D., Schemmel, J., et al. (2015). Probabilistic inference in discrete spaces can be implemented into networks of LIF neurons. *Frontiers in Computational Neuroscience*, 9.

Pyle, R., & Rosenbaum, R. (2017). Spatiotemporal dynamics and reliable computations in recurrent spiking neural networks. *Physical Review Letters*, 118(1), 018103.

Rao, R. P., Olshausen, B. A., & Lewicki, M. S. (2002). *Probabilistic models of the brain: Perception and neural function*. MIT press.

Rauch, A., La Camera, G., Lüscher, H.-R., Senn, W., & Fusi, S. (2003). Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo–like input currents. *Journal of Neurophysiology*, 90(3), 1598–1612.

Rosenbaum, R., Tchumatchenko, T., & Moreno-Bote, R. (2014). Correlated neuronal activity and its relationship to coding, dynamics and network architecture. *Frontiers in Computational Neuroscience*, 8, 102.

Salakhutdinov, R. Learning deep Boltzmann machines using adaptive MCMC. In *Proceedings of the 27th international conference on machine learning* (pp. 943–950).

Salinas, E., & Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, 2(8), 539.

Schemmel, J., Fieres, J., & Meier, K. (2008). Wafer-scale integration of analog neural networks. In *Neural networks, 2008. IJCNN 2008.(IEEE world congress on computational intelligence). IEEE international joint conference on* (pp. 431–438). IEEE.

Schiller, P. H., Finlay, B. L., & Volman, S. F. (1976). Short-term response variability of monkey striate neurons. *Brain Research*, 105(2), 347–349.

Schmitt, S., Klähn, J., Bellec, G., Grübl, A., Guettler, M., Hartel, A., et al. (2017). Neuromorphic hardware in the loop: training a deep spiking network on the brainscales wafer-scale system. In *Neural networks international joint conference on* (pp. 2227–2234). IEEE.

Segev, R., Benveniste, M., Hulata, E., Cohen, N., Palevski, A., Kapon, E., et al. (2002). Long term behavior of lithographically prepared in vitro neuronal networks. *Physical Review Letters*, 88(11), 118102.

Seung, H. (2003). Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron*, 40(6), 1063–1073.

Smetters, D., & Zador, A. (1996). Synaptic transmission: noisy synapses and noisy neurons. *Current Biology*, 6(10), 1217–1218.

Snowden, R. J., Treue, S., & Andersen, R. A. (1992). The response of neurons in areas v1 and mt of the alert rhesus monkey to moving random dot patterns. *Experimental Brain Research*, *88*(2), 389–400.

Song, S., Sjöström, P. J., Reigl, M., Nelson, S., & Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology*, *3*(3), e68.

Stein, R. B. (1967). Some models of neuronal variability. *Biophysical Journal*, *7*(1), 37–68.

Steinmetz, P. N., Manwani, A., Koch, C., London, M., & Segev, I. (2000). Subthreshold voltage noise due to channel fluctuations in active neuronal membranes. *Journal of Computational Neuroscience*, *9*(2), 133–148.

Stevens, C. F., & Zador, A. M. (1996). When is an integrate-and-fire neuron like a Poisson neuron? In *Advances in neural information processing systems* (pp. 103–109).

Vogels, R., Spileers, W., & Orban, G. A. (1989). The response variability of striate cortical neurons in the behaving monkey. *Experimental Brain Research*, *77*(2), 432–436.

White, J. A., Rubinstein, J. T., & Kay, A. R. (2000). Channel noise in neurons. *Trends in Neurosciences*, *23*(3), 131–137.

Wiener, N. (1930). Generalized harmonic analysis. *Acta Mathematica*, *55*(1), 117–258.

Xie, X., & Seung, H. (2004). Learning in neural networks by reinforcement of irregular spiking. *Physical Review E*, *69*(4), 041909.

Yarom, Y., & Hounsgaard, J. (2011). Voltage fluctuations in neurons: signal or noise? *Physiological Reviews*, *91*(3), 917–929.

Zador, A. (1998). Impact of synaptic unreliability on the information transmitted by spiking neurons. *Journal of Neurophysiology*, *79*(3), 1219–1229.

Zenke, F., & Ganguli, S. (2018). Superspike: supervised learning in multilayer spiking neural networks. *Neural Computation*, *30*(6), 1514–1541.