

Cortical oscillations implement a backbone for sampling-based computation in spiking neural networks

A. Korcsak-Gorzo*^{2,3,1} M. G. Müller*⁴ A. Baumbach^{1,5} L. Leng¹ O. J. Breitwieser¹
 S. J. van Albada^{2,6} W. Senn⁵ K. Meier¹ R. Legenstein⁴
 M. A. Petrovici^{5,1}

¹ Kirchhoff-Institute for Physics, Heidelberg University, Heidelberg, Germany

² Institute of Neuroscience and Medicine (INM-6) and Institute for Advanced Simulation (IAS-6) and JARA-Institute Brain Structure-Function Relationships (INM-10), Jülich Research Centre, Jülich, Germany

³ RWTH Aachen University, Aachen, Germany

⁴ Institute of Theoretical Computer Science, Graz University of Technology, Graz, Austria

⁵ Department of Physiology, University of Bern, Bern, Switzerland

⁶ Institute of Zoology, University of Cologne, Cologne, Germany

* equal contribution

Abstract

Brains need to deal with an uncertain world. Often, this requires visiting multiple interpretations of the available information or multiple solutions to an encountered problem. This gives rise to the so-called mixing problem: since all of these “valid” states represent powerful attractors, but between themselves can be very dissimilar, switching between such states can be difficult. We propose that cortical oscillations can be effectively used to overcome this challenge. By acting as an effective temperature, background spiking activity modulates exploration. Rhythmic changes induced by cortical oscillations can then be interpreted as a form of simulated tempering. We provide a rigorous mathematical discussion of this link and study some of its phenomenological implications in computer simulations. This identifies a new computational role of cortical oscillations and connects them to various phenomena in the brain, such as sampling-based probabilistic inference, memory replay, multisensory cue combination and place cell flickering.

1 Introduction

The ability to build an internal, predictive model of reality endows an agent with a clear evolutionary benefit. How the mammalian brain accomplishes this feat remains a subject of debate, but the representation of uncertainty certainly plays a role, considering the probabilistic nature of sensory data and uncertainty about past and future events. A good representation of an uncertain reality must allow efficient access to a large variety of plausible beliefs about the environmental state. In the context of sampling-based models, this is usually referred to as the mixing problem and is one of the many facets of the curse of dimensionality.

For complex distributions over sensory data, as is characteristic for natural scenes, the mixing problem is notoriously challenging. In this manuscript, we put forward a hypothesis for how the brain can efficiently overcome this challenge. In doing so, we unify two individually well-studied, but previously unlinked aspects of cortical dynamics under a common normative framework: probabilistic inference and cortical oscillations. In particular, we consider the interpretation of spiking activity in cortex as probabilistic inference via sampling, which has gained ample experimental (Fiser et al., 2010; Berkes et al., 2011; Orbán et al., 2016) and theoretical (Buesing et al., 2011; Pecevski et al., 2011; Probst et al., 2015; Petrovici et al., 2016; Dold et al., 2019) support over the last decade.

Mathematically, these models are closely related to Gibbs sampling, which tends to get stuck in single states of high probability which act as local attractors. We propose that this problem of sampling-based representations can be overcome by firing rate oscillations. Firing rate oscillations over multiple frequency bands are a naturally emerging phenomenon in spiking networks (Brunel, 2000; Izhikevich, 2006; Destexhe, 2009; Muller and Destexhe, 2012) and have been extensively studied in the mammalian brain (Buzsáki and Draguhn, 2004; Buzsáki, 2006). Notably, they appear to play an important role both during awake perception (Başar et al., 2001; Wang, 2010; Klimesch, 2012) and during sleep (Klinzing et al., 2019; Adamantidis et al., 2019), suggesting a fundamental role in cognition and learning.

We propose that the background firing rate of cortical neurons can be interpreted as a (computational) temperature, and can accordingly modify the probability landscape sampled by cortical circuits. If the background activity is oscillatory, the network temperature also changes periodically and phase-dependent stationary distributions emerge. By cyclically alternating between "hot" and "cold" periods, cortical networks can effectively instantiate a tempering schedule, with hot phases corresponding to flat probability distributions in which the network can move freely and cold phases representing the target multimodal distribution. This allows networks to escape from local minima and efficiently sample from challenging distributions characterized by multiple high-probability modes separated by large low-probability volumes of the state space.

With respect to experimental neuroscience, the suggested computational mechanisms relate to various psychophysical and physiological phenomena, from single neurons to behavior. The tempering in our model modulates the gain of the neuronal transfer functions, similarly to the stochastic sampling of a scene through an oscillatory modulation of attentional gain (Klimesch, 2012; Marzecová et al., 2018), particularly through top-down input (Reynolds and Chelazzi, 2004; Larkum et al., 2004). The stochasticity in our model by which stored memories are selectively recalled is mirrored in the randomness of hippocampal replay during sleep that goes beyond the more regular behavior (Stella et al., 2019), or in free memory recall in humans (Katkov et al., 2017). The oscillatory recall that supports cognitive computation in our model can also be related to creative thinking (Lustenberger et al., 2015), to midbrain oscillatory activity during stimulus disambiguation (Pettigrew, 2001), to mind wandering (Compton et al., 2019) and to local sleep (Stålesen Ramfjord et al., 2020).

In this work, we provide an analytical treatment of tempering in spiking networks induced by cortical background oscillations and demonstrate the benefits of this phenomenon in simulations. We explicitly consider both current-based and conductance-based synaptic interactions as well as different network architectures, and discuss links to experimental data. These observations establish a novel connection between multiple observed cortical phenomena, as well as between these experimental findings and normative theoretical models of brain computation.

2 Experiments and results

To understand how cortical oscillations affect computation at the network scale, we first study the behavior of single neurons under variable levels of background activity. For mathematical tractability, we consider current-based leaky integrate-and-fire (LIF) neurons, for which we can derive analytical expressions for relevant statistical properties that affect the neuronal response. In particular, this allows us to precisely define the notion of (computational) temperature, first for individual neurons and later for entire networks (Fig. 1a). We then consider such recurrent networks as models of computation in sensory cortex, performing Bayesian inference via spike-based sampling; their underlying statistical model allows a rigorous interpretation of cortical oscillations as a tempering algorithm. Subsequently, we generalize our models by lifting imposed restrictions on model parameters, while also switching to conductance-based synapses as a more general representation of biological neuronal interaction. We show how our previous findings apply to different cortical assemblies and discuss how our models can be linked to experimental data.

2.1 Single-neuron statistics

Cortical neurons are embedded in a noisy environment. In addition to functional input I_{in} , their many presynaptic partners provide them with an effectively stochastic background (Ecker et al., 2010). This background activity leads to stochastic single-neuron behavior (Shadlen and Newsome, 1998). To understand this behavior, we consider a simple LIF neuron model with current-based input synapses (see *Methods*). The neuron receives a large number of background inputs, with firing rates ν_i and synaptic efficacies w_i . We first consider the free membrane potential u_{free} of this neuron, that is, the membrane potential in the hypothetical case that there is no firing threshold. The steady-state distribution $p(u_{\text{free}})$ is well-described by a Gaussian (Fig. 1b) with moments

$$\mu_u := \mathbb{E}[u_{\text{free}}] = E_1 + \frac{I_{\text{in}} + \sum_i w_i \nu_i \tau_{s,i}}{g_1}, \quad (1)$$

$$\sigma_u^2 := \text{Var}[u_{\text{free}}] = \sum_i \nu_i w_i^2 \frac{\tau_{s,i}^2}{2g_1^2 (\tau_m + \tau_{s,i})}, \quad (2)$$

where E_1 and g_1 are the leak potential and conductance, and τ_m and τ_s are the membrane and synaptic time constants (Petrovici, 2016, Section 4.3). Here, \sum_i runs over all background presynaptic partners. Note that excitatory and inhibitory inputs (defined by the sign of the synaptic weight w_i) can cancel each other out in the mean, but always add up towards the variance of the free membrane potential distribution.

Upon introducing a firing threshold, some portion of the free membrane potential will lie above it, causing the neuron to spike stochastically. The shape of the neuronal response function depends strongly on the characteristic time constant of

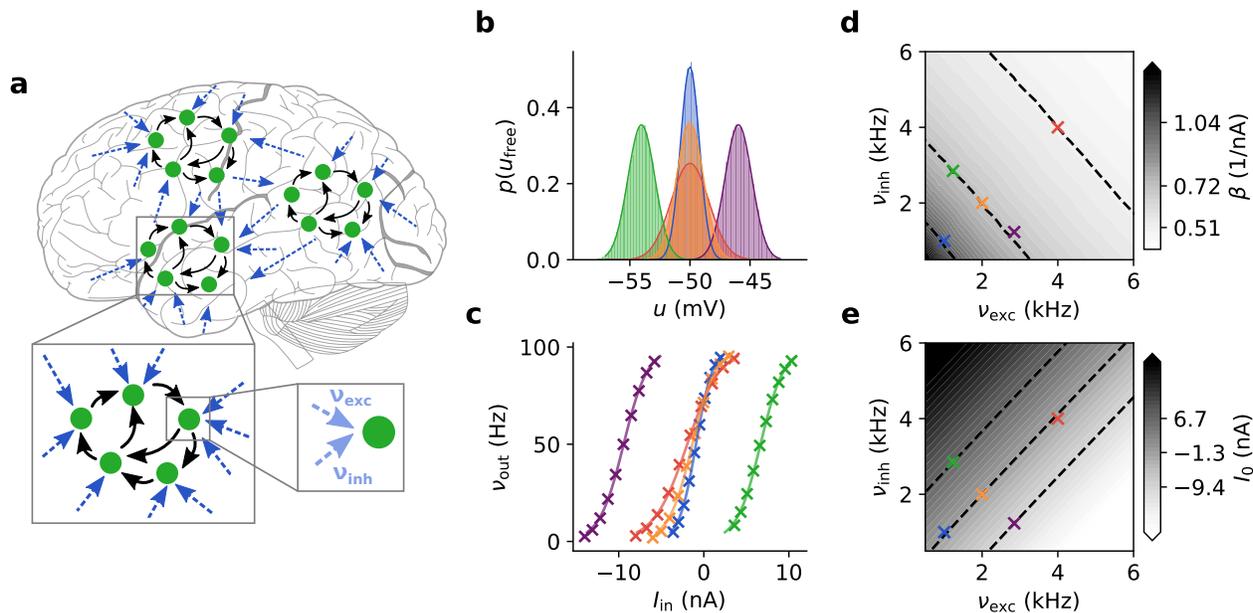


Figure 1: **Response functions of neurons in an ensemble.** (a) Cortical ensemble of networks. The spike input received by a neuron can be partitioned into functional (solid black arrows) and background (dashed dark blue arrows) input. The background can be partitioned into an excitatory and an inhibitory subset (dashed light blue arrows). In the following panels, we consider one such neuron under five different illustrative background regimes, each of which is assigned a specific color. (b) Steady-state free membrane potential distributions. Shaded areas: numerical simulation; solid lines: analytical approximation using Eqs. 1 and 2. Purple, orange, green: same σ_u , different μ_u ; blue, orange, red: same μ_u , different σ_u . (c) Corresponding neuronal response functions. Crosses: numerical simulation; solid lines: logistic fit with Eq. 3. (d) Slope parameter β and (e) offset I_0 of response functions under various background regimes defined by their respective pairs of excitatory and inhibitory input rates ($\nu_{\text{exc}}, \nu_{\text{inh}}$). Dashed isolines indicate configurations of constant slope (cf. Eq. 4) or offset, with specific values given as colorbar ticks. Note the approximate linearity of the contour lines.

the neuronal membrane. Cortical neurons under strong presynaptic bombardment have been shown to operate in a high-conductance regime (Destexhe et al., 2003), which greatly reduces the effective membrane time constant τ_m . Under such conditions, the neuronal response function, that is, its firing rate in response to a constant current I_{in} , (Fig. 1c) can be well approximated by a logistic function (Petrovici et al., 2016):

$$\nu_{\text{out}}(I_{\text{in}}) = \frac{1}{1 + \exp[-\beta(I_{\text{in}} - I_0)]}. \quad (3)$$

Hence, the stochastic response of the neuron is characterized by two parameters, the offset I_0 and the slope β of the sigmoid. Both of these depend on the background activity. The response function can be intuitively understood as the area under the free membrane potential distribution that lies above the firing threshold. Thus, its shape is similar to the integral of $p(u_{\text{free}})$, its offset I_0 has a similar linear dependence on μ_u and its slope parameter β will decrease for increasing σ_u . Their exact dependence on the background rates is shown in Fig. 1d and e. In particular, the relationship between the slope of the response function and the standard deviation of the free membrane potential distribution is well-approximated by a linear function, which allows us, in turn, to establish the relationship between the slope parameter β and the total (i.e., summing over all background presynaptic partners) excitatory and inhibitory background firing rates ν_{exc} and ν_{inh} using Eq. 2:

$$\frac{1}{\beta} \propto \sigma_u \propto \sqrt{w_{\text{exc}}^2 \nu_{\text{exc}} + w_{\text{inh}}^2 \nu_{\text{inh}}}. \quad (4)$$

Below, we show how this relationship also describes the ensemble temperature of a spiking network.

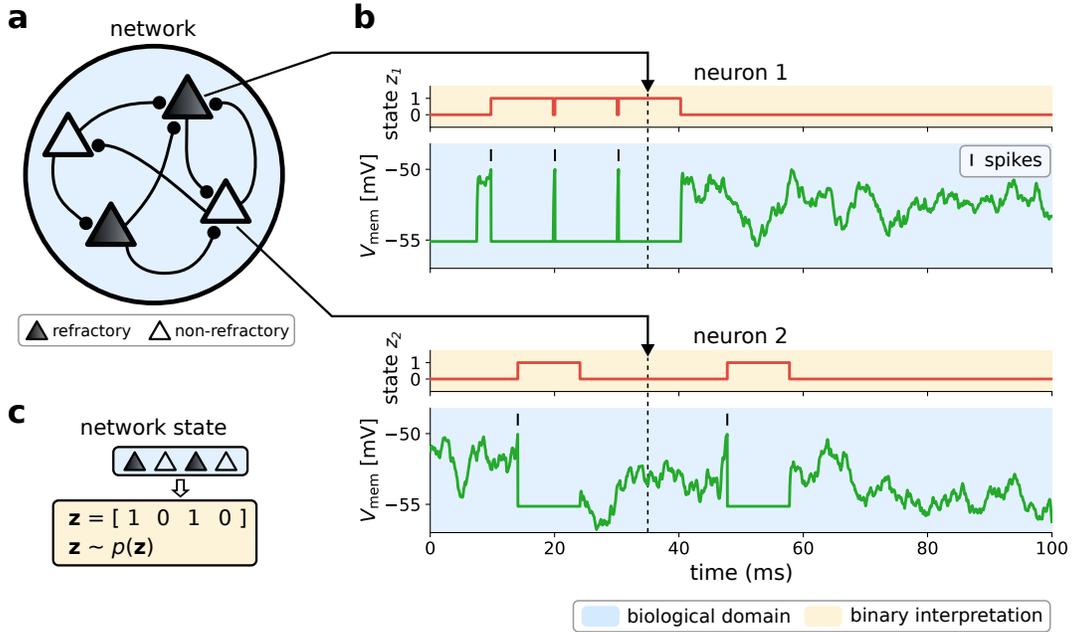


Figure 2: **Spiking network activity as sampling from binary distributions.** The two levels of abstraction are color-coded: blue represents the biological domain, yellow the binary-variable interpretation. **(a)** Network schematic. Each neuron encodes a binary random variable. At any given time, a subset of neurons (dark) have spiked recently and are refractory. **(b)** Relationship between neuron dynamics and sampled states. Membrane traces (lower panels, green) following a random walk due to the stochastic background input interspersed with output spikes followed by reset and refractoriness. The refractory period τ_{ref} following each spike is associated with the state $z = 1$; non-refractory neurons encode $z = 0$ (upper panes, red line). **(c)** The network state $\mathbf{z}(t)$ represents the vector of individual neuron states $z_k(t)$.

2.2 Tempering in spiking networks

In a spiking network, the information conveyed by a neuron at any point in time can be described as binary: the neuron either spikes or it does not. A spike has a twofold effect: it initiates a refractory period and it elicits postsynaptic potentials (PSPs) in postsynaptic partner neurons. We can therefore view the binary state z of a neuron being refractory ($z = 1$) or not ($z = 0$) following a spike as corresponding to the state communicated to its downstream partners, up to scaling by synaptic weights (Fig. 2a). Moreover, due to the background input discussed above, each neuron responds stochastically to inputs from its presynaptic partners. Thus, each neuron can be interpreted as sampling from the conditional distribution $p(z_k = 1 | \mathbf{z}_{\setminus k})$, i.e., the probability of the k -th neuron to be in the state "1" given the states of all other neurons.

In general, the joint distribution sampled by the network cannot be given in closed form. To allow an analytical approach, we begin with a set of assumptions about the network parameters (see *Methods*), but later show that they can be relaxed without affecting the computational network properties discussed here. Under these assumptions, the activity of an LIF network can be interpreted as sampling from a joint Boltzmann distribution

$$p_T(\mathbf{z}) \propto \exp[-E(\mathbf{z})/(k_B T)] , \quad (5)$$

where $E(\mathbf{z}) = -\frac{1}{2} \sum_{kj} W_{kj} z_k z_j - \sum_k B_k z_k$ represents the energy of a particular joint state \mathbf{z} , with W_{kj} denoting effective recurrent synaptic weights and B_k effective individual neuron biases, which subsume biasing effects induced by leak and threshold potentials, unbalanced input etc. For the exact derivation of the Boltzmann equivalence, we refer to Petrovici et al. (2016). Note that the Boltzmann parametrization with unitless weights \mathbf{W} and biases \mathbf{B} is different from the biological domain of synaptic conductances and currents. The (linear) conversion is given in the *Methods*.

The ensemble (Boltzmann) temperature T scales these effective weights and biases multiplicatively, identically to its effect in statistical physics: as the temperature of an ensemble rises, particle interactions (here: synaptic weights) and external fields (here: neuronal biases) become increasingly inconsequential. Since weights and biases can both be interpreted as movements along the I_{in} -axis of the neuronal response function, their simultaneous multiplicative scaling by T is equivalent to a horizontal stretching of the response function, thus allowing us to identify

$$\beta = 1/(k_B T) , \quad (6)$$

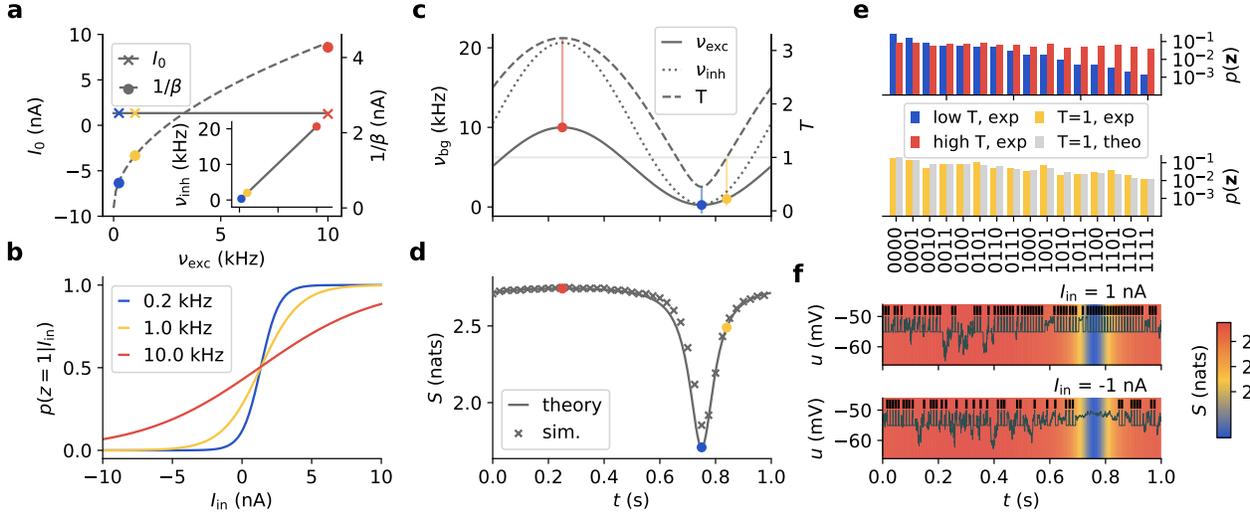


Figure 3: **Effects of oscillatory background activity on sampled distributions.** (a) A linear relationship $\nu_{\text{inh}}(\nu_{\text{exc}})$ (inset) can keep the offset constant (solid line), while only changing the slope of activation functions and thereby the temperature of a network (dashed line). This relationship also reflects the relative strengths of afferent excitatory and inhibitory weights. The three colored markers correspond to the minimum, reference ($T = 1$), and maximum rates used for the background oscillations in (c) and (d). (b) Three example activation functions with constant offset for different background rates. Colors correspond to those in panel a. Here, we emphasize the binary-state interpretation by plotting $p(z = 1|I_{\text{in}}) = \nu_{\text{out}}\tau_{\text{ref}}$ (cf. Fig. 1c). (c) Time course of excitatory and inhibitory background rates (Eq. 8), along with the associated temperature (Eq. 4). (d) Simulated vs. calculated (Eq. 32) entropy course $S(t)$. The slight lag is due to the finite relaxation time constants $\tau_s, \tau_{\text{ref}}$ of the network (Eq. 32 only holds strictly for quasistatic temperature changes). (e) Distributions sampled by a 4-neuron network at the three temperatures marked in (a). States are ordered according to their respective probabilities at low temperature to visually emphasize the effect of tempering. (f) Effect of tempering on individual membrane potentials and spiking activity. The background color represents the corresponding entropy.

again analogous to statistical physics, for a Boltzmann constant k_B that relates the (unitless) reference temperature $T = 1$ to a chosen set of neuron and background parameters via the resulting activation function (here, the unit of k_B is nA). Equation 4 thus establishes an exact relationship between the ensemble temperature and the background firing rates.

The ensemble temperature also controls the entropy of the sampled distribution, which effectively describes the “disorder” of the network and corresponds to the uniformity of the sampled distribution. For higher temperatures, the sampled distribution becomes more uniform and its entropy increases, equalizing state probabilities. This, in turn, promotes mixing: high-probability modes surrounded by low-probability states become flatter, allowing the network to escape such local attractors. Thus, oscillatory background activity can be interpreted as tempering, a periodic cycle of heating and cooling, with hot phases for mixing and cold phases for reading out the most relevant samples of the correct distribution.

In order to study the effect of temperature variations without affecting neuronal offsets, excitatory and inhibitory background rates need to be balanced. Such a balance is also well-documented in vivo (Shu et al., 2003; Haider et al., 2006). Note that this is not simply achieved by setting $w_{\text{exc}}\nu_{\text{exc}}\tau_{\text{exc}} = w_{\text{inh}}\nu_{\text{inh}}\tau_{\text{inh}}$ and thus effectively equalizing the effects of excitation and inhibition; while this would leave μ_u unchanged, it would still affect I_0 (cf. Fig. 1b,c). Still, a balanced regime can be achieved by a linear dependence between firing rates (Fig. 3a,b), following one of the isolines in Fig. 1e, which are well approximated by

$$\nu_{\text{inh}} = \nu_0 + m\nu_{\text{exc}}. \quad (7)$$

The exact parameters ν_0 and m that are necessary for balance depend on synaptic time constants and background weights (see *Methods*). This allows excitation and inhibition to vary synchronously, as observed in vivo. While this approach enables a strict realization of simulated tempering, the achieved effect does not rely strongly on such a balance, as we discuss later.

Since any periodic change in background rates can be approximated by a Fourier series, we consider a simple sinusoidal oscillation as a basis function for modeling cortical oscillations:

$$\nu_{\text{exc}}(t) = \frac{\nu_{\text{max}} - \nu_{\text{min}}}{2} \sin(2\pi f_{\text{osc}}t) + \frac{\nu_{\text{max}} + \nu_{\text{min}}}{2}, \quad (8)$$

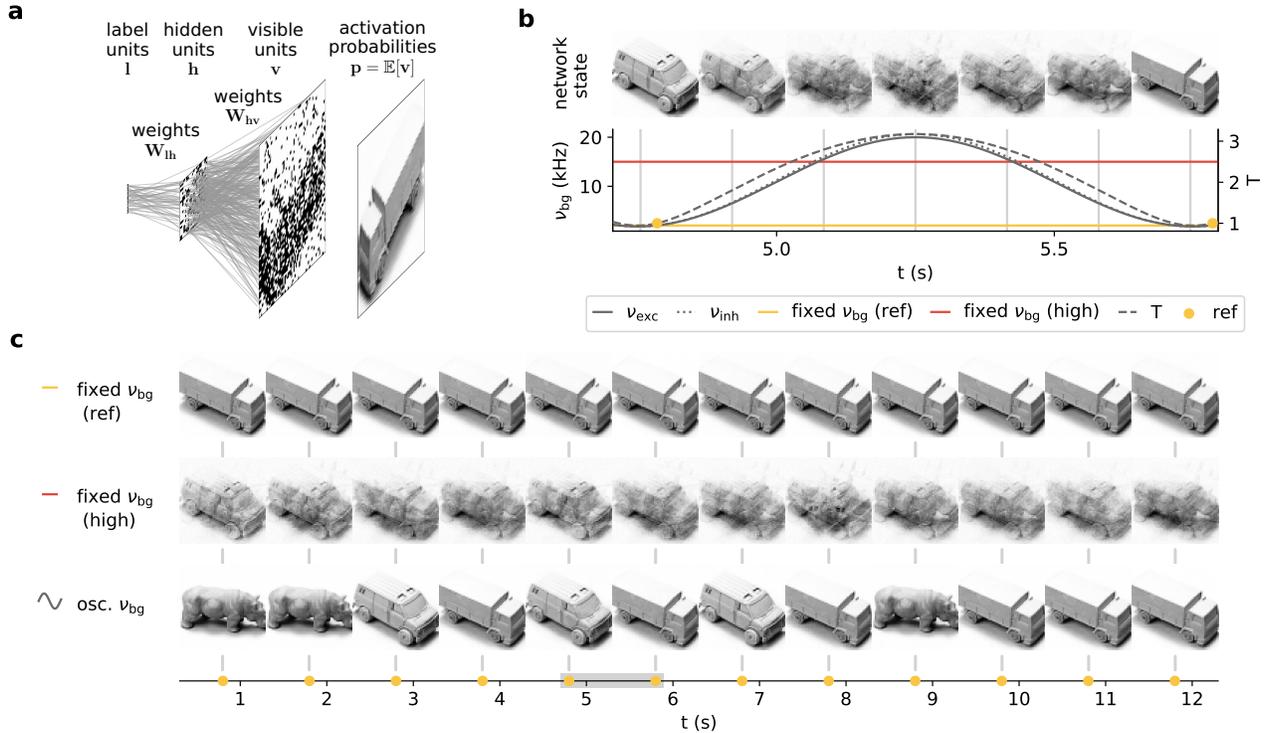


Figure 4: **Background oscillations improve generative properties of spiking sampling networks.** (a) Hierarchical 3-layer (visible v , hidden h and label l) network architecture and example layerwise activity. For a better representation of the visible layer statistics, we consider neuronal activation probabilities $p(v|h)$ rather than samples thereof, to speed up the calculation of averages over (conditional) visible layer states. Here, we show a network trained on images from the NORB dataset. (b) Evolution of the activation probabilities of the visible layer (top) over one period of the background oscillation (bottom). (c) Evolution of the visible layer over multiple periods of the oscillation compared to a network with constant background input at the reference rate (2 kHz, top) and at a high rate (10 kHz, middle), cf. also yellow and red lines in (b). The activation probabilities are shown whenever the reference rate (see panel b) is reached. The gray bar denotes the period shown in (b).

with minimum ν_{\min} , maximum rate ν_{\max} and oscillation frequency f_{osc} , which implicitly also fixes $\nu_{\text{inh}}(t)$, as discussed above. The temperature and resulting gain modulation thus also varies periodically, with the square root of a sine (Fig. 3c). The resulting effects on the sampled distribution are shown in Fig. 3e: at high temperatures, the distribution becomes flat, while at low temperatures, the high-probability maxima become even more pronounced. This is mirrored by the change in entropy (Fig. 3d). In high-temperature/high-entropy states, membrane potentials are extremely noisy, causing neurons to fire randomly and independently, whereas in low-temperature/low-entropy states, membrane potentials are nearly constant and neurons are "frozen" in certain states, firing either persistently or not at all (Fig. 3f).

2.3 Mixing in high-dimensional multimodal data spaces

In the following, we discuss the computational role of background oscillations for spiking networks trained to represent complex distributions over high-dimensional visual data. As a simplified model of cortical visual hierarchy, we consider recurrent layered LIF networks, which we train as simultaneous generative and discriminative models (Fig. 4a). These two forms of computation happen concurrently and bidirectionally: the label neurons classify the state of the visible layer, while the visible neurons adapt their states to produce images that are compatible to the category represented by the label layer. High-dimensional, but well-recognizable visual data confronts such networks with two contradictory challenges. On the one hand, they need to produce good samples, i.e., clean images corresponding to particular categories; this corresponds to the buildup of sharp, high-probability modes separated by large vanishing-probability volumes of the state space that correspond to out-of-distribution samples. On the other hand, they need to be able to switch between different modes in order to fully sample from the target distribution; this is at fundamental odds with the probability landscape described above.

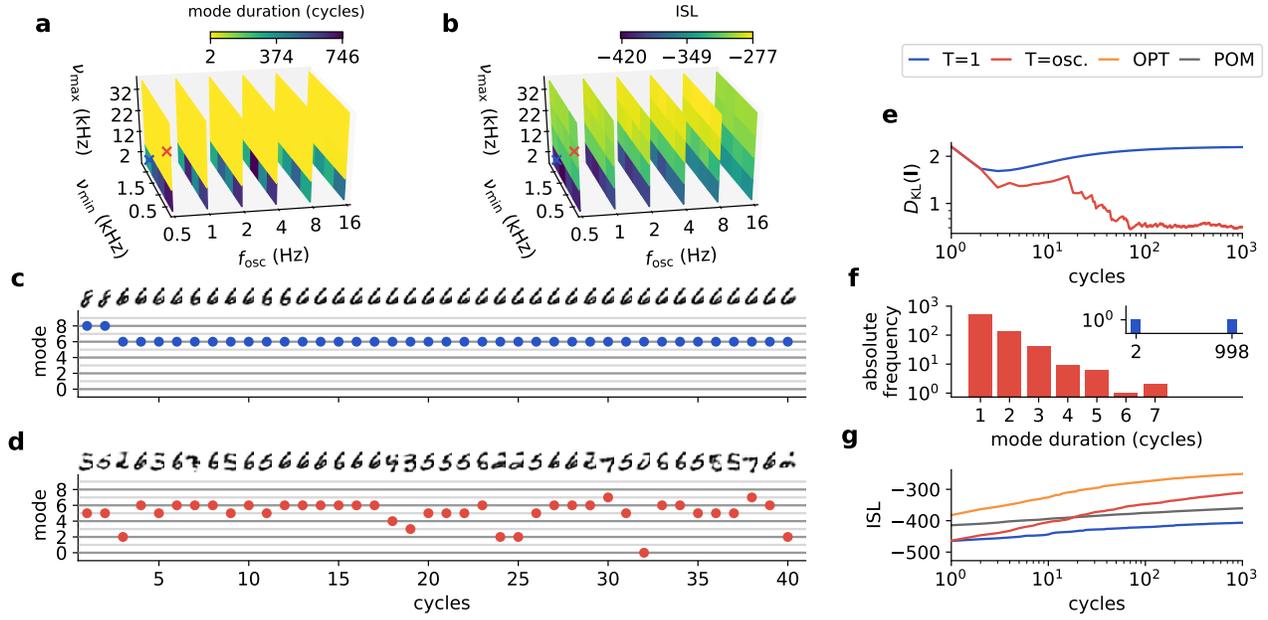


Figure 5: **Parameter dependence of tempering effectiveness.** (a-b) Generative properties, i.e. mode duration and indirect sampling likelihood (ISL) of a spiking sampling network trained on the MNIST dataset for systematically varied parameters of the background oscillation cycle. Values represent averages over 10 runs of 10^4 samples. The following panels show results using the parameter set marked with a red/blue cross in (a) and (b). Networks receiving oscillatory background are shown in red, constant background ($\nu_{max} = \nu_{min}$) at $T = 1$ in blue. (c-d) Comparison of visible and label layer (mode) activity between constant and oscillatory background. Cortical oscillations are seen to have several beneficial effects over a large volume of the studied parameter space: (e) a more uniform mode distribution, as seen in a lower Kullback-Leibler divergence of the label layer activation with respect to the uniform distribution; (f) shorter mode durations than under constant input (inset) and (g) an ISL that converges to a higher value. Optimal sampling (OPT, orange) and the product of marginals (POM, gray) shown as upper and lower boundary of the ISL. Note that the ISL prefers images that are more similar to the entire test set, which is why the "average MNIST images" described by the POM have a higher ISL than the unimodal distribution induced by constant background at $T = 1$.

This so-called mixing problem is well-known and quasi-ubiquitous for sampling models.

One solution to this problem was proposed by Marinari and Parisi (1992) in the context of Markov-chain Monte Carlo sampling for Ising models, which is intimately related to our form of spike-based sampling in both dynamics and sampled distribution (Petrovici et al., 2016). This simulated tempering method describes a cyclic heating and cooling schedule reminiscent of the periodic temperature variation induced by cortical oscillations discussed above (Eq. 8). In-between readouts at the reference temperature, a temporary rise in temperature flattens the probability landscape, allowing the network to escape from local attractors. Thus, Eqs. 4, 5 and 8 establish a rigorous analogy between simulated tempering and cortical oscillations, which thereby take on the computational role of enabling mixing in challenging real-world scenarios.

To evaluate these effects, we considered two example scenarios based on well-studied visual datasets: NORB (Huang and LeCun, 2004) and MNIST (LeCun and Cortes, 2010). Network training was done using a variant of wake-sleep learning (Salakhutdinov, 2010), a contrastive Hebbian scheme inspired by biological phenomenology and widely used for sampling models (see, in particular, Leng et al., 2018). A background rate of $\nu_{exc} = \nu_{inh} = 2$ kHz was chosen as reference, implicitly defining the reference temperature $T = 1$.

For the network trained on NORB, the effect of oscillatory background clearly instantiates the desired tempering schedule (Fig. 4b): at low background rates, the network produces sharp images, while being able to switch modes in the high-rate phases. Note especially how the network enters a superposition of several "clean" states at higher background rates. Figure 4c further illustrates the effects of different background regimes. Under the constant reference background, networks produce sharp images, but mix poorly. Increasing the constant background to a higher rate promotes mixing, but at the cost of blurring the generated images. Tempering through background oscillations effectively combines these two regimes, allowing a better sampling of the target distribution at phases where the reference temperature is reached.

The effectiveness of this tempering schedule depends on the parameters of the background oscillations: ν_{min} , ν_{max} and

f_{osc} . In particular, the frequency f_{osc} plays an important role, as it represents a trade-off between exploration and exploitation of the network’s state space. Low frequencies guarantee that the network has time to relax towards its momentary stationary distribution p_T , with $f_{\text{osc}} \rightarrow 0$ representing the quasistatic limit. This also enables perfect sampling from the target distribution at $T = 1$, as the network loses memory of previous states occupied at higher temperatures. However, lower oscillation frequencies come at the cost of slower sampling, as they increase the time between consecutive readouts. Furthermore, frequencies significantly lower than 0.1 Hz are rarely observed in vivo (Buzsáki, 2006). In the following, we study the behavior of spiking sampling networks under different background oscillation regimes.

Two important quality criteria for any sampling network are its mixing speed and sample fidelity. In principle, Eq. 5 allows an analytical evaluation of these properties, but in practice this is unfeasible for high-dimensional distributions. We therefore use a sample-based measure, the indirect sampling likelihood (ISL, see Breuleux et al., 2010). The ISL accumulates fidelity values for all generated samples, assigning high values if they are similar to images in the test set and low values otherwise. Additionally, the rate at which the ISL increases over time implicitly represents a measure of the mixing speed. We further use the distribution of times between label switches as a more explicit measure of mixing times for different image categories.

Using an LIF network configured for sampling from MNIST data (see *Methods* for detailed parameter settings), we studied tempering under a range of biologically plausible regimes, with total (population) background rates varying between 0.5 and 30 kHz and oscillation frequencies ranging from the alpha range to the first slow-wave band (Buzsáki and Draguhn, 2004). The most important prerequisite for effective tempering is the maximum background rate, as the temperature between readouts has to be high enough for frequent mode switches. For our networks, this required input rates above 10 kHz in the high-temperature phases (Fig. 5a,b). On the other hand, the minimum background rates in the cold phases have a much smaller influence. In terms of oscillation frequency, the fastest oscillations were already sufficient to significantly improve mixing in the studied range.

In general, effective tempering is achieved over a wide range of oscillation parameters (yellow and light green areas) covering all studied frequency bands. Overall, the best performance was achieved in the slow-wave regime (Fig. 5c-g), which would be particularly useful for the replay and consolidation of varied memories during slow-wave sleep (Diekelmann and Born, 2010; Klinzing et al., 2019).

2.4 Impact of conductance-based synaptic input

Up to this point, we have used a mathematically tractable model providing an exact link between background input and network behavior. This link leads to a clear interpretation of background oscillations as a schedule of temperature changes within networks of neurons. We now move on to show that these conceptual results hold over a wider range of neuron and synapse models with different parameters. We therefore relax our assumptions and consider more general network architectures with a range of physiological parameter settings. In particular, we consider conductance-based synaptic interactions, which are known to be a good description of the behavior of biological neurons. As before, we first investigate the properties of individual neurons before moving on to network-level effects.

2.4.1 Background input scenarios

The behavior of neurons with conductance-based synapses under synaptic bombardment in general differs from the current-based case. In particular, the variance of the membrane potential does not grow monotonically with the rate of the background input (Fig. 6, see also Eq. 18 in *Methods*). As this growth was shown above to underlie the increase of the temperature, this raises the question whether the same, simple relationship between background input rate and sampling temperature is present in the conductance-based case. However, neuron response functions can still depend monotonically on the input rates, as the relative impact of inputs is also weakened by an increasing total conductance. In the following, we show that as a result, the influence of background input rates on the effective temperature using conductance-based neurons matches the influence in the current-based case.

To this end, we investigate four different scenarios that reflect the diversity and complexity of conductance-based interactions. These scenarios differ along a number of degrees of freedom: different synaptic conductances for background input, affecting the variance of u , as well as multiple ways of changing the excitatory and inhibitory input rates.

As for current-based neurons (Eq. 7), we assume that the excitatory and inhibitory input rates ν_{exc} and ν_{inh} depend on an abstract parameter α that linearly regulates the level of background activity via

$$\nu_{\text{exc}} = \alpha \nu_{\text{exc},1} + \nu_{\text{exc},0} \quad \text{and} \quad \nu_{\text{inh}} = \alpha \nu_{\text{inh},1} + \nu_{\text{inh},0} \quad (9)$$

where $\nu_{\text{exc},1}$ and $\nu_{\text{inh},1}$ are some excitatory and inhibitory base rates, and $\nu_{\text{exc},0}$ and $\nu_{\text{inh},0}$ are rate offsets. Note that these rates are always positive as we only consider $\alpha > 0$.

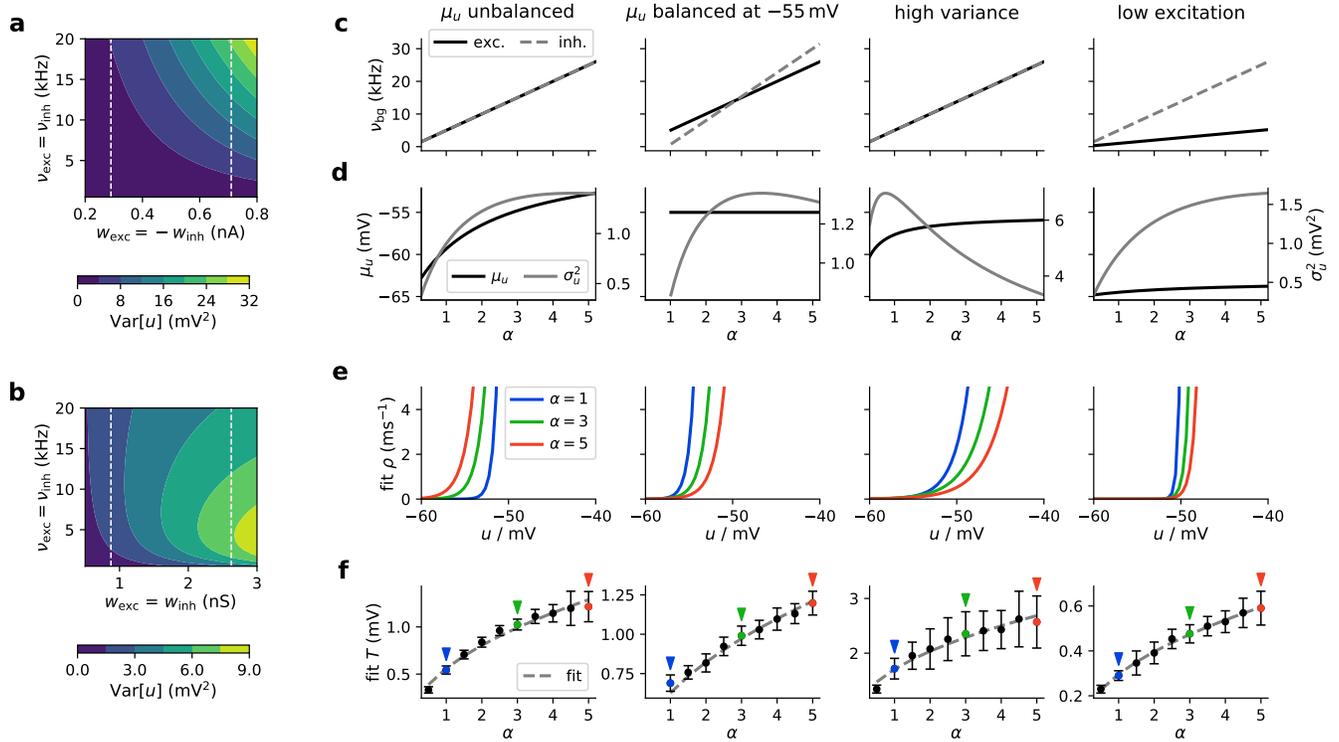


Figure 6: Background input rate sets the sampling temperature for conductance-based LIF neurons. (a, b) Variance of the free membrane potential for current-based and conductance-based synaptic interactions as a function of background rate and strength under balanced input rates and weights. Note that $\text{Var}[u]$ is always monotonic in ν_{exc} for current-based synapses, which does not hold for the conductance-based case (see the change of $\text{Var}[u]$ along the dashed white lines). The following panels show four different biologically interesting scenarios of conductance-based background input (ordered in columns). From left to right: μ_u unbalanced (one of many possible settings in which μ_u increases), μ_u balanced at -55 mV (with constant mean membrane potential close to the firing threshold), high variance (same as the first scenario but with larger synaptic conductances, resulting in a higher variance of the membrane potential), and low excitation (with lower increase in excitation compared to inhibition). (c) Excitatory and inhibitory background input rates as a function of the scaling value α . Note that for the μ_u balanced at -55 mV scenario the balance can only be achieved for $\alpha \geq 1$. (d) Mean and variance of free membrane potential resulting from background input scaled by α . (e) Result of fitting the stochastic model (see Eq. 10) for different values of α . The parameters of ρ (including the temperature T) were obtained by averaging the results of $N = 50$ independent fits. (f) Temperature values resulting from fitting. Dots denote mean, whiskers standard deviation over $N = 50$ independent runs (see *Methods* for details). Importantly, the temperature increases monotonically with α in all cases. The dashed gray line shows fits according to the theoretical prediction (Eq. 4), closely matching mean values. The colored means (also highlighted by colored arrows) indicate the values used for the firing intensity plots in panel (e). The soft threshold values u_T resulting from fitting are given in Fig. 6 Supp. 1b.

We consider $\alpha \in [0.5, 5]$, which results in background excitatory and inhibitory input rates ≤ 30 kHz given the base and offset rate values for the different scenarios (see below). These values are similar to the rate range used in the current-based simulations above and represent a reasonable assumption for cortical neurons as they typically have a large number of presynaptic partners (Megias et al., 2001). Using the input rates and the synaptic parameters, it is possible to calculate the mean μ_u and variance σ_u^2 as a function of the background input scaling factor α ((Petrovici, 2016, Section 6.5); Zerlaut et al. (2018)).

The four different scenarios (Fig. 6, see *Methods*) considered here are:

- μ_u *unbalanced*: A general case of realistic synaptic conductances not tuned to any specific regime with equal excitatory and inhibitory background rates ($\nu_{\text{exc},1} = \nu_{\text{inh},1} = 5$ kHz and $\nu_{\text{exc},0} = \nu_{\text{inh},0} = 0$ kHz). Although $\nu_{\text{exc}} = \nu_{\text{inh}}$ in this case (Fig. 6a, first column), excitatory and inhibitory inputs are not balanced as the synaptic time constants differ (see *Methods*). As a result, both μ_u and σ_u^2 increase over the considered range of α (Fig. 6b, first column).
- μ_u *balanced at -55 mV*: This scenario mimics the regime of cortical up-states, where neurons have membrane potentials close to the firing threshold (here: $u_{\text{th}} = -50$ mV). Balancing neurons in this fashion, i.e., keeping μ_u at -55 mV regardless of the value of α , can be achieved with specific choices of $\nu_{\text{inh},1}$ and $\nu_{\text{inh},0}$ (i.e., $\nu_{\text{inh}} \neq \nu_{\text{exc}}$ here, see Fig. 6a, second column, and *Methods*). Note that balancing here requires some minimal background input level. Here, the variance σ_u^2 first increases before reaching a peak and slowly decreasing again (Fig. 6b, second column).
- *High variance*: Larger synaptic conductances give rise to a different regime, which is approximately balanced (small change of μ_u for changing α), but differs from the previous scenarios in two significant ways: first, the overall variance is much higher, and second, the variance decreases with α (Fig. 6b, third column). This scenario uses $\nu_{\text{exc}} = \nu_{\text{inh}}$ (as the first scenario).
- *Low excitation*: It is unclear whether in the brain excitatory and inhibitory input levels are similar, in particular within oscillations, and it has previously been suggested that oscillations mostly affect inhibition (Sohal and Hasselmo, 1998a). We therefore also consider a scenario in which the inhibitory rates increase much more strongly than the excitation (Fig. 6c, last column). This results in a marginal effect on the mean free membrane potential, while the variance increases with α (Fig. 6d, last column).

All parameters for the different scenarios are given in *Methods*.

2.4.2 Fitting stochastic models to quantify temperature changes

To gain an understanding of the stochasticity induced by background activity, we fit stochastic neuron models to data produced by LIF models with background input. This method allows quantifying behavior changes regardless of the precise input conditions and neuron parameters, thus making it possible to describe the sampling temperature even when an analytical treatment is not possible.

To this end, we used the fitting method proposed by Jolivet et al. (2006). The stochastic model is identical to the LIF neuron model except for a stochastic firing criterion with instantaneous firing intensity

$$\rho(u) = \frac{1}{\Delta t} \exp\left(\frac{u - u_T}{T}\right), \quad (10)$$

where T is the temperature, u_T is the soft threshold (i.e., the value of u where the firing intensity reaches $1/\Delta t$), and Δt is the resolution of the discrete-time simulation. Exponential escape rate functions have been shown to account for neural data (Jolivet et al., 2006) and also match theoretical models of sampling with spiking neurons (Buesing et al., 2011).

We performed this fitting procedure for different values of α to examine how the model changes based on the background input. For every fit, we used a number of presynaptic spike trains to excite the LIF neuron (see *Methods*) and recorded the firing times. The resulting firing rates varied markedly (Fig. 6 Supp. 1a), highlighting the different operating regimes induced by different levels of background input. We fit exponential firing intensities to the data (Fig. 6e). The temperature values resulting from the fitting procedure are shown in Fig. 6f.

We found that in all four scenarios, T increases with α , even as the variance of the membrane potential plateaus or decreases. This confirms the effect of increasing background input rates on the temperature even in conductance-based conditions with diverse properties. We have shown above that T grows with $\sqrt{\alpha}$ in the current-based case. A fit of the mean temperatures resulting from the fitting procedure shows that this relationship also describes the change of the temperature very well in the conductance-based case (Fig. 6f, dashed gray lines).

These results confirm the role of background rates as an effective ensemble temperature in conductance-based networks, and this method can be used to quantify the temperature changes. Depending on the parameters of the background activity, the covered temperature range can vary significantly. Nevertheless, as show in the following, these temperature changes have important functional consequences at the network level.

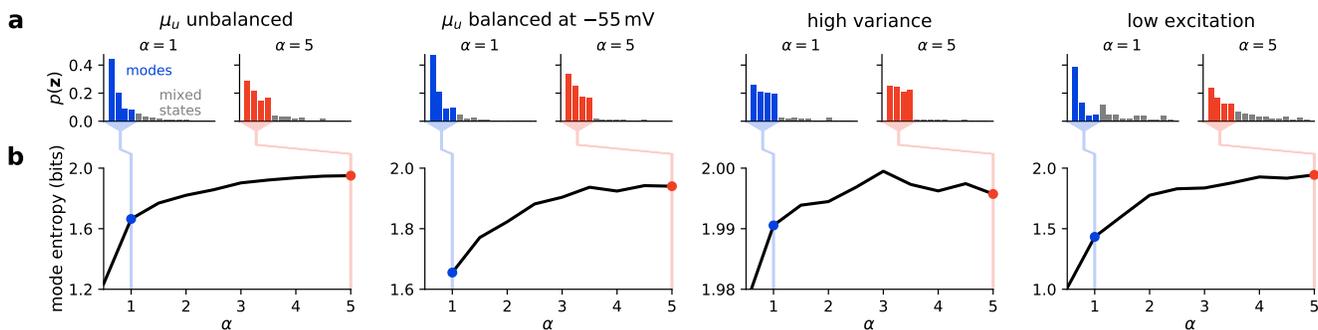


Figure 7: **Effect of temperature variations in networks of conductance-based LIF neurons.** (a) State probabilities using a simple winner-take-all network for low ($\alpha = 1$) and high ($\alpha = 5$) temperatures (see *Methods*). Modes (states with exactly one neuron active) are shown in color, mixed states are shown in gray. At high temperatures, the distribution generally becomes flatter. (b) Entropy of modes at each value of α (in bits). As the mode distribution becomes flatter (i.e., closer to a uniform distribution), the entropy increases. Note that while the temperature increases in all scenarios (Fig. 6f), the values and changes of the entropy can vary dramatically.

2.5 Entropy in networks with conductance-based synapses

To confirm that these single-neuron effects result in similar effects on the network level, we investigated a simple network of conductance-based LIF neurons. The network had a winner-take-all (WTA) structure with four distinct modes with unequal probability (see *Methods*) and received background input with varying α for each scenario. We found that in every scenario, the mode distribution becomes more uniform for high levels of background input (Fig. 7a). To quantify the changes, we calculated the entropy of the mode probabilities (Fig. 7b). In each of the scenarios, the entropy increases with α , indicating an increase of the sampling temperature. In the high-variance case, the effect is small but shows the same trend (larger entropy for larger α), which is surprising as the variance of u here decreases as α is increased. This can be explained as follows: even as the variance decreases, the overall synaptic conductance evoked by background input grows. Therefore, as α increases, the effect of the background input grows stronger relative to the input from the recurrent network connections, thus leading to more equal responses.

This experiment confirms that changes in the background activity of conductance-based networks give rise to the same qualitative phenomena as in the current-based case. We next show the relevance of this effect in a behaviorally relevant sampling task.

2.6 Background oscillations and behaviorally relevant sampling tasks

As in networks of current-based neurons, we expect that background oscillations structure computations into distinct phases when using conductance-based networks, which we investigate next. In contrast to the current-based sampling experiments, we do not restrict ourselves to the precise conditions required for unbiased sampling, and instead consider the more general case of arbitrary parameters.

The link between activity levels and sampling temperature described previously suggests that brain networks alternate between sampling at high temperatures, allowing rapid traversing of the state space for good mixing, and sampling at low temperatures, promoting convergence to states of high probability. We investigated this effect using a stimulus disambiguation task, in which a network was required to find coherent interpretations for conflicting inputs across three different sensory modalities (Fig. 8a). Each sensory modality (vertical columns in Fig. 8a) was represented by a group of three neuronal assemblies, with each of these assemblies encoding one of three possible interpretations of the input. To represent mutually exclusive interpretations, the assemblies representing each sensory modality (boxes in Fig. 8a) had inhibitory lateral connections, instantiating a winner-take-all (WTA) network. Assemblies representing the same interpretation across different sensory modalities were set up as mutually and recurrently excitatory (solid lines in Fig. 8a). Thus, when such a triplet of assemblies across sensory modalities was active, the network encoded a coherent interpretation of the input. As one assembly encoding each interpretation received a small bias input (Fig. 8a), none of the three competing interpretations of the input was favored, leading to sensory ambiguity. To correctly represent such an ambiguous situation, the network is expected to sample all three interpretations.

Viewed as a sampling task, this encodes a distribution with three high-probability states, in each of which all assemblies encoding a single, coherent interpretation are active (while all other assemblies are silent). This triple will then inhibit other assemblies due to the WTA structures employed for each sensory modality. We say that the network has found a valid

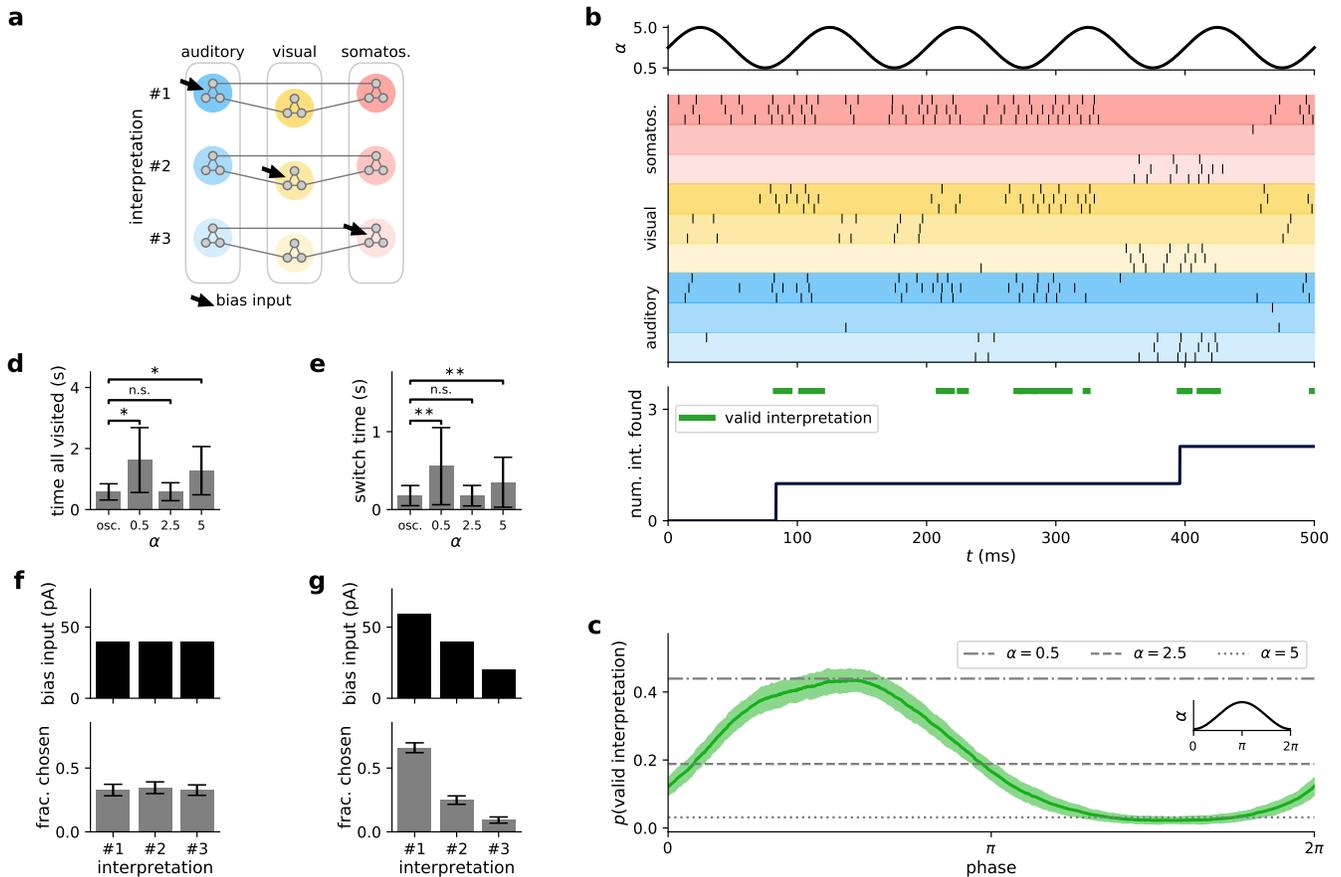


Figure 8: Background oscillations structure computations into sampling episodes in conductance-based networks. **(a)** Setup of stimulus disambiguation task. Assemblies (large colored circles) encode interpretation of stimulus spanning three different sensory modalities (auditory, visual, somatosensory). For each modality, every interpretation is encoded by the activity of a small assembly (different interpretations are represented by different color hues). One assembly per sensory modality receives a bias input, resulting in ambiguous input. Simultaneous activity of connected assemblies encoding the same interpretation across all modalities encodes selection of an interpretation (i.e., a solution). **(b)** Network activity with oscillating background input. Top: background activity scaling by α . Middle: spike raster plot of network activity (color coding of assemblies as in panel a). Bottom: number of solutions, (i.e., interpretations) visited so far. Green bars show times in which the network state encodes a solution. **(c)** Probability of network state encoding a solution depending on the phase of the oscillating background input. Gray lines show solution probabilities in networks without oscillating background activity for different activity levels $\alpha \in \{0.5, 2.5, 5\}$. Inset shows background activity phase for reference. **(d)** Mean time until the network has visited all three solutions for oscillating and constant background activity (over $N = 100$ runs for each case). Bars and whiskers show mean and standard deviation, respectively. Significance was calculated with Wilcoxon rank-sum test with $* \hat{=} p < 10^{-10}$ and n.s. $\hat{=} p > 0.05$. **(e)** Time between choosing distinct solutions (see *Methods* for details). Plot as in panel d with $** \hat{=} p < 10^{-100}$. **(f)** For constant bias inputs (top, cf. panel a), solutions are chosen equally often (bottom, means and standard deviations). **(g)** For different bias inputs (top), the probability of choosing the corresponding solution matches the input values (bottom). Plot as in panel f.

interpretation if one linked assembly triple is active ($> 50\%$ of neurons per assembly fired within the last 10 ms) while all other assemblies remain silent ($< 50\%$ of neurons fired; see *Methods* for details). As the recurrent connectivity within each assembly is rather strong, the network tends to lock into one such state, making mixing difficult. However, as generally no interpretation is preferred over the others, the goal of the sampling process is to visit all solutions (with visitation frequencies corresponding to their relative biases) in a reasonable amount of time.

We compared the behavior of the network for oscillating α in the same frequency bands as above ($\alpha \in [0.5, 5]$, i.e., total background rates in $[2.5, 25]$ kHz) with a constant-background scenario. To allow a fair comparison of oscillating to constant background for both low and high background activity, we tuned the synaptic parameters so the variance of the network firing rates was minimal for different α values (see *Methods* for details).

Fig. 8b shows network activity over the first 500 ms of a simulation run. The network can clearly jump between attractors and sample different valid interpretations of the ambiguous input. To test whether oscillating background has an advantage over constant input, we performed $N = 100$ simulations lasting 20 s each and calculated the probability for the network state to encode a valid interpretation at any point in time. Fig. 8c shows that background oscillations structure sampling-based computations by defining times when good solutions can be read out from the network. We find that at certain phases, the network with background oscillations has a much higher probability of encoding a valid interpretation compared to medium or high level constant background input. Moreover, this probability itself oscillates at the same frequency as the background, but with a certain phase lag that depends on the network parameters. In contrast, constant background input produces constant valid-state probabilities, with minimum and maximum values corresponding to those achieved with background oscillations. However, constant background also implies a trade-off between spending time in valid states and being able to mix between these. For example, networks receiving low background input produce the same valid-state probability as the maximum value achieved with oscillations, but tend to converge to one solution and stay there for a long time (see Fig. 8 Supp. 1), thus exhibiting much worse mixing behavior.

We quantified mixing by measuring (i) the time it took the network to visit each solution at least once (Fig. 8d), and (ii) the time it took on average to move from one solution to another (Fig. 8e). On both measures, two regimes achieve comparatively high performance: oscillatory background or constant background at a well-tuned intermediate activity level. However, constant background that is high enough to also facilitate mixing represents a computationally unreliable regime: solutions are found at random points in time and are comparatively ephemeral (cf. Fig. 8 Supp. 1c,d).

In contrast, oscillatory background has a significantly higher probability of producing long-lived valid solution at well-defined readout phases of the oscillation. Cortical oscillations thus provide explicit temporal structure to sampling-based computations in spiking neural networks. This structure provides good solutions with high probability while inheriting the good mixing properties of high-temperature networks.

For these experiments, we used an equal bias input for each of the three interpretations (Fig. 8f, top). Thus, each interpretation should occur equally often. We verified this for oscillating background input (Fig. 8f, bottom), where it is indeed the case. In reality, bias input may occur at any level, and the resulting frequency of visiting interpretation states changes accordingly. We repeated the analysis ($N = 100$ simulations lasting 20 s) for this case, and found that the visitation levels correspond to the level of bias input (Fig. 8g). This highlights a further advantage of background-oscillation-induced tempering. Since constant background rates can only produce constant temperature, increasing their base level to promote mixing necessarily skews the relative strength of individual attractors by equalizing them (cf. Figs. 3 and 7). Cortical oscillations, on the other hand, preserve the relative dominance of the different modes in the readout phases.

2.7 Constraining sampling models with experimental data

So far, we have shown that sampling networks benefit from temperature oscillations in a variety of parameter regimes. To understand the operating regime of cortex, it is important to constrain models by experimental data. In this final section, we provide an example of how such links can be established.

We previously discussed how the probability of a valid interpretation changes over the phase of the background input oscillation. One possible way to link the sampling models to experimental data is by considering how the changes within one cycle match recordings from the brain. One study that touches upon this question is Jezek et al. (2011), which investigated place-cell flickering in rat hippocampus in relation to theta oscillations. The place cells encode the current belief on the spatial location of the animal based on visual cues. Even in unambiguous situations, their activity is not static, but occasionally switches to alternative interpretations for one or a few theta cycles. Typically, the recorded activity matches only one interpretation, but over brief periods mixed interpretations are present. Jezek et al. highlight that these mixed interpretations are more likely in the first half of the cycle.

We investigated how this relates to the behavior of a simple model of such flickering activity while changing the model parameters along one of the most important degrees of freedom: the ratio of inhibitory to excitatory background input. We constructed a simple circuit model with two assemblies encoding correct and incorrect interpretations of the spatial context (see *Methods*). The model parameters were chosen such that the flickering was similar to the data shown by Jezek et al. (2011) in the absence of cue switching (i.e., occasional switches to the incorrect interpretation, see Fig. 9a). We then varied the ratio of the synaptic conductances of inhibitory and excitatory background input to change the mean input ratio $\mathbb{E}[g_{\text{inh}}/g_{\text{exc}}] \in [0.75, 3]$. This has a strong influence on the behavior of individual neurons (Fig. 9b), resulting also in a change of the firing rate of the network in relation to the background oscillations (Fig. 9c): for low inhibition, the background input acts mostly excitatorily, thus, high background input results in large network input. For high ratios of $\mathbb{E}[g_{\text{inh}}/g_{\text{exc}}]$, the converse holds: large background input activity inhibits the network, and the network is most active in low input phases. The transition between these two regimes occurs around $\mathbb{E}[g_{\text{inh}}/g_{\text{exc}}] = 1.8$ (Fig. 9c). Fitting stochastic models for each value of $\mathbb{E}[g_{\text{inh}}/g_{\text{exc}}]$ showed that at this value, the activation functions are aligned at 0.5 for all values of α , corresponding

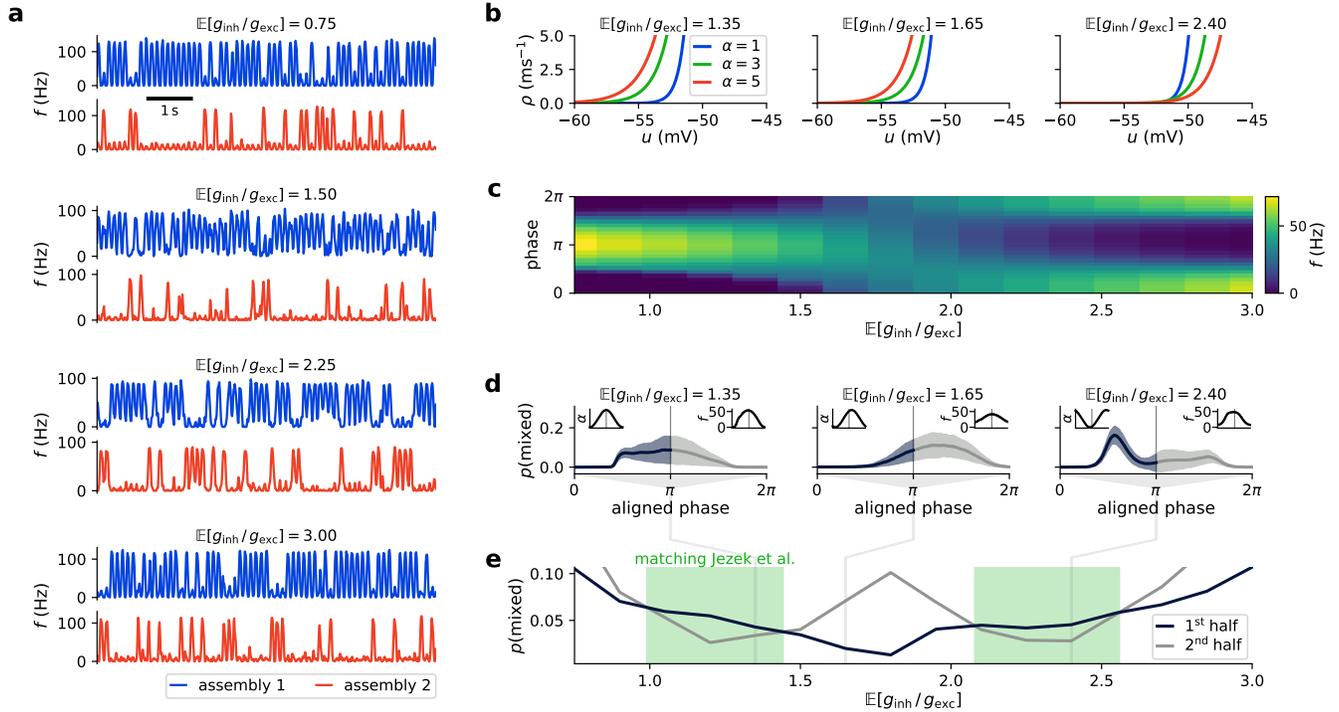


Figure 9: Relating model features to experimental data. (a) Activity in a simple model of flickering place-cell activity in the hippocampus (see Fig.3c in Jezek et al., 2011). Assembly firing rates encoding correct and incorrect interpretations of context are shown in blue and red, respectively. The model behavior holds as the ratio of the mean inhibitory to excitatory background input is varied (plot shows $\mathbb{E}[g_{inh}/g_{exc}] \in \{0.75, 1.5, 2.25, 3\}$). (b) Firing intensity behavior of individual neurons determined by fitting stochastic models (as in Fig. 6e) drastically changes as the mean background input ratio $\mathbb{E}[g_{inh}/g_{exc}]$ is increased. (c) As $\mathbb{E}[g_{inh}/g_{exc}]$ changes, the minimum firing rate shifts along the background input phase. The background input cycles are defined here as having their onset time (phase zero) when the input rate is at its minimum, i.e. the maximum input occurs at phase π . (d) Probability of mixed interpretations during a cycle. Solid line shows mean, shaded area standard deviation over $N = 100$ network runs of 50 s each. Insets show background input (left) and network firing rate (right, cf. panel b). To match Jezek et al. (2011), the phase is reordered for every value of $\mathbb{E}[g_{inh}/g_{exc}]$ so that phase zero is aligned with the minimum firing rate in the network (see right inset). (e) Mixed state probability in first and second half of network activity cycle as the mean input ratio is varied. Green shading shows areas in which the mixed probability is larger in the first half of the cycle, matching the results of Jezek et al. (2011).

to unbiased sampling (cf. Fig. 3b). This shows that, similar to the current-based case, this regime can be achieved by adjusting the balance of excitation and inhibition in the conductance-based case.

We then calculated the probability of mixed interpretations within background input cycles. To match the procedure of Jezek et al. (2011), we defined the phase in relation to the network activity recorded in the model (Fig. 9c). Jezek et al. (2011) segregated the network activity into cycles using the recorded firing rate such that the minimum firing rate corresponds to phase zero. We reordered the phase at every value of $\mathbb{E}[g_{inh}/g_{exc}]$ accordingly (Fig. 9d, insets). We then analyzed the probability of mixed states in the first and second half of the cycle (as Jezek et al., 2011) at every value of $\mathbb{E}[g_{inh}/g_{exc}]$ (Fig. 9d). When analyzing how the mean probability changes over the range of conductance ratios, we found that there were two regions matching the situation described by Jezek et al. (2011). Interestingly, the conductance ratio of corresponding to unbiased sampling at $\mathbb{E}[g_{inh}/g_{exc}] = 1.8$ does not fall into either range (Fig. 9e). This raises the question whether synapses that mediate the effect of background input on cortical assemblies are optimally tuned towards achieving a balanced regime – i.e., unbiased temperature changes – or whether it is computationally useful for some cortical functions to tune the effect of oscillatory activity towards being explicitly biased.

We emphasize that this experiment only provides a first example of linking sampling models to experimental data. Clearly, more data is needed to fully constrain sampling models of cortex. In particular, experiments establishing more direct links between neuronal properties and sampling model features such as probability distributions or sampling temperatures represent a necessary prerequisite for a quantitative understanding of sampling computations in the brain.

3 Discussion

Oscillatory activity is a naturally emerging phenomenon in spiking neuronal networks. As it is well-known that background input increases the variability of neuronal firing, oscillatory background implies oscillatory variability. In the context of ensemble theory, this creates an immediate link to the notion of temperature. We have shown that the level of background input determines the sampling temperature in networks of LIF neurons and demonstrated that this effect leads to functional advantages in sampling networks when oscillatory background input is present. This holds in the case of current-based synaptic interactions, for which we have presented an analytical treatment of cortical oscillations as tempering, as well as for conductance-based synaptic interactions, where we have studied a broad range of physiologically relevant parameters in computer simulations. Our results suggest that the ubiquity of oscillations in human and animal brains provides a clear benefit for behaviorally relevant computations, which is elucidated by considering the analogy to simulated tempering.

Related theoretical work Our considerations rest on the assumption that for fixed parameters, spiking networks sample from a stationary distribution. This has been shown to hold under only mild constraints for a large class of neuron and network models in Habenschuss et al. (2013). They also showed that in the presence of periodic input, a phase-specific stationary distribution exists which is influenced by the network parameters and the properties of the inputs. The existence of such a distribution naturally leads to questions about its specific nature, given specific ensemble dynamics such as those arising in networks of connected LIF neurons, and its functional properties for cortical computation. In this work, we have shown that the phase-dependent component is a temperature scaling of a Boltzmann distribution, with periodic background alternating between its exploration and its exploitation.

An alternative way of promoting mixing was proposed by Leng et al. (2018). There, synaptic short-term plasticity was shown to weaken local attractors. This mechanism has a similar effect, but is different from a change in temperature. Since this form of plasticity only affects active synapses, it only suppresses active local modes rather than flattening the entire distribution. These dynamics ensure that local modes can be abandoned quickly, as synapses can be weakened significantly by only a few spikes, but they come at the cost of changing the sampled distribution. In contrast, cortical oscillations induce a well-defined temporal structure that promotes an undistorted readout. For mathematical tractability, we first considered LIF neurons with current-based synaptic interactions and network structures that are easily amenable to contrastive Hebbian training. We then showed that our results hold for a larger class of biological settings by considering conductance-based synapses and competing neural assemblies. This suggests that the computational role we propose for cortical oscillations is generalizable to a diverse set of cortical structures and their associated functions. Indeed, it has already been observed that oscillations appear to have a similar function throughout the cortex (Lundqvist et al., 2020).

A similar function of brain rhythms related to slower oscillations was proposed by Sohal and Hasselmo (1998b), who suggested on theoretical grounds that during the hippocampal theta cycle, modulation of GABA_B synapses performs a process similar to simulated annealing in a model of population dynamics. Such a mechanism was shown to be advantageous for sequence disambiguation (Sohal and Hasselmo, 1998a). In this work, we propose that temperature control takes place on the level of individual neurons via input regardless of the synapse type. Thus, the mechanism we propose for incorporating such an annealing in neural networks has a much more general scope. Savin et al. (2014) showed the benefits of rhythmic changes of neuron excitability in a model of probabilistic memory recall, resulting in a similar kind of annealing as in our model. Our work shows how such a schedule of excitability changes arises in spiking neural networks via background input, thus suggesting an implementation of this mechanism on the cellular level.

Related experimental work and model predictions Across the entire spectrum of cortical rhythms, individual components of these oscillations are characterized by their frequency and amplitude. We have shown that an effective tempering schedule can be achieved across a wide range of frequencies and amplitudes, roughly corresponding to the range lying between slow and alpha waves (Buzsáki and Draguhn, 2004). At higher frequencies, internal network dynamics cannot react quickly enough to the changes in temperature and the sampling quality deteriorates. However, this soft upper limit depends on the dominant time constants of individual neurons and synapses, which in our simulations were on the order of 10 ms. For faster dynamics, as often observed in vivo (τ_m in Koch et al. (1996), τ_s in Häusser and Roth (1997); Bartos et al. (2001), τ_{ref} in McCormick et al. (1985); Steriade et al. (1998)), correspondingly faster oscillations can be accommodated; for example, oscillations in the gamma band could be employed by ensembles with synaptic time constants and refractory times in the order of a few milliseconds, as discussed in recent sampling-based modeling approaches (Aitchison and Lengyel, 2016; Echeveste et al., 2020). Thus, this form of tempering can be exploited both for inference in the awake state, where oscillations are typically fast, and during sleep, for functions such as memory retrieval and consolidation (Klimesch, 2012; Klinzing et al., 2019; Adamantidis et al., 2019).

The oscillation frequency, and thus the rate of change in temperature, carries another subtle effect. For slow waves, the effect of a single transition from maximum to minimum temperature is similar to simulated annealing (Kirkpatrick et al.,

1983). As the network effectively has more time to relax towards its corresponding thermodynamic equilibrium, it will, at least statistically, tend towards the global minimum energy state. On the other hand, faster oscillations are more akin to tempered transitions (Neal, 1996; Salakhutdinov, 2009), where low-temperature states are more diverse. Indeed, the extreme scenario of quenching (extremely rapid cooling) could be implemented by switches between synchronized cortical up and down states (Destexhe, 2009; Haider et al., 2006; Shu et al., 2003). Thus, different oscillatory phenomena in the brain can shift the focus from finding a small set of maximum probability modes to finding a larger range of relevant modes. Similarly, the oscillation amplitude can also control the effective breadth of the exploration space, with larger maximum rates promoting larger jumps between more dissimilar network states.

The benefits of cortical oscillations also extend to other facets of Bayesian inference. For example, when the state distribution is constrained by partial observations, such as in our cue disambiguation task, tempering helps explore the conditional distribution and find multiple ways to solve this pattern completion problem. Similarly, this can help with finding multiple solutions to a given problem, such as assigning multiple categories to particular input patterns. Importantly, this also highlights the potential benefits of background oscillations during learning (see also Capone et al. (2019)), where exploration plays an essential role.

Our results demonstrate that oscillations provide an additional benefit to improved mixing: they serve as a reference for reading out computational results, reducing the amount of data requiring processing, and facilitating the temporal organization of neural computations. Furthermore, they can also serve as a means of input filtering, increasing susceptibility to coherent stimuli (Grothe et al., 2012). In general, it is well known that information encoding in relation to a background oscillation can be found in the brain, for example in the hippocampus (O’Keefe and Recce, 1993), where place cells convey information by firing earlier or later in relation to the theta rhythm. A similar form of coding takes place in our models, as the network distribution changes during each cycle of the background input.

Furthermore, cyclic background input results in the network generating a stream of candidate solutions, with one such state arriving in each cycle. This leads to a form of computing in discrete steps, as computations are structured into episodes defined by background oscillations. A similar type of structured computation has been suggested to take place in monkey and human visual brains during the processing of visual inputs (Buschman and Miller, 2010). These experiments showed that shifts in attention were aligned to beta-band oscillations, and every shift took place within a single cycle. In our model, we find similar shifts of the state taking place within each cycle as the temperature decreases.

Temperature changes from oscillations predict that the time course of the network state variability is coupled to the oscillation phase, as we have shown in our model. This suggests that a similar coupling could be found in sampling-based computations in the brain. Jezek et al. (2011) have given a hint that this can indeed be the case in hippocampal circuits by showing that ambiguous interpretations of the network input are more likely in the first half of the theta cycle. We have shown how a simple model can reproduce these findings. However, this data is rather coarse, and our results also suggest that a similar behavior can emerge in multiple operating regimes (cf. Fig. 9e). Thus, more detailed experimental data are required to adequately constrain sampling models based on background oscillations.

In particular, experimental data could elucidate the question whether cortical networks are tuned to an unbiased sampling regime. In our model, achieving unbiased sampling requires some tuning of neuron parameters and of the time course of the background oscillation. While our analysis based on a simple model of the results of Jezek et al. (2011) suggested that such a tuning might not be present, a model more closely matching biological networks (e.g., by incorporating additional features such as short-term plasticity, neuronal adaptation, and more specific inhibition) would help to either corroborate this finding or provide more insight into how such a tuning might be achieved in brain networks. In general, the balance between excitation and inhibition is of renewed interest in this context, as it connects directly to experimental data. Individual neurons or neuron populations can, for example, use unbalanced rates to implement biases for their associated random variables. Moreover, we expect that different networks tune their background inputs to different balances, depending on which biases are beneficial for their respective tasks.

Applications Recent years have seen an increasing interest in using spike-based computation on specialized hardware to perform energy-efficient computations (Roy et al., 2019). This has spurred efforts to develop models which allow efficient learning and inference with spiking neural networks. Some of these platforms explicitly exploit the stochasticity of their components for computation (Al-Shedivat et al., 2015; Sengupta et al., 2016). By offering a mechanism for modulating neuronal stochasticity, oscillatory background can enhance computation in stochastic neuromorphic networks, for example in generative spiking models (Dold et al., 2019; Kungl et al., 2019).

The implicit time window for reading out correct solutions suggests that such oscillations may also improve the performance of networks used for constraint satisfaction problems (Jonke et al., 2016; Binas et al., 2016; Fonseca Guerra and Furber, 2017). These are solved by shaping the stationary distribution of the network so that solution states have a high probability. However, it is not clear at any given point in time whether the current state is a solution candidate or a transitional state. In contrast, in an oscillation-driven tempering schedule, it is known that solutions are likely at low-temperature phases.

Overall, the parallels with a variety of empirical phenomena and the advantages for spike-based sampling demonstrated here make neuronal oscillations a likely mechanism for supporting stochastic computations in the brain and a useful tool for fulfilling this same function in biologically inspired neural networks.

4 Methods

4.1 Neuron models

Current-based LIF model

The membrane potential u of a current-based leaky integrate-and-fire (LIF) neuron evolves according to

$$C_m \frac{du}{dt} = g_l(E_l - u) + I(t) , \quad (11)$$

with membrane capacitance C_m , leak potential E_l and leak conductance g_l . The resulting membrane time constant is $\tau_m = C_m/g_l$. When the membrane voltage u reaches a threshold value v_{th} from below, a spike is emitted and the membrane potential is fixed to a reset value v_{reset} for the refractory time τ_{ref} (see Table 1). The input current $I(t)$ is a sum of synaptic currents

$$I(t) = I_{rec}(t) + I_{in}(t) + I_{bg}(t) , \quad (12)$$

where we distinguish between functional input I_{rec} , synaptic background input I_{bg} and any other form of bias input I_{in} (see Fig. 1a). Assuming exponential synaptic kernels, the input current obeys

$$\frac{dI}{dt} = \frac{I_{in} - I}{\tau_s} + \sum_j w_j S_j(t) , \quad (13)$$

where w_j and τ_s respectively denote the synaptic weight and time constant. The sum goes over all presynaptic spike sources j , including both background and recurrent input, with the corresponding spike trains $S_j(t) = \sum_f \delta(t - t_j^{(f)})$, where $t_j^{(f)}$ denotes the f th spike time of spike source j .

Stochastic background synaptic input is modeled as Poissonian. Without loss of generality, we endow each neuron with a single excitatory and a single inhibitory Poisson source characterized by rates ν_{exc} and ν_{inh} and corresponding connection strengths w_{exc} and w_{inh} .

The resulting distribution of the free membrane potential u_{free} (no spiking, $v_{th} \rightarrow \infty$) is well described by a Gaussian with moments given by Eqs. 1 and 2 (for more details see (Petrovici, 2016, Section 4.3)). In general, more background input, originating from either larger weights $w_{exc}, |w_{inh}|$ or higher frequencies ν_{exc}, ν_{inh} , increases the variance.

The resulting neuronal response function can be calculated from this distribution using a recursive approach (Petrovici et al., 2016). In the high-conductance state (Destexhe et al., 2003), the membrane time constant becomes small, leading to a more symmetric response function, which is well-approximated by a logistic function (Eq. 3). In the interpretation of spiking neurons as binary random variables, the neuronal response becomes an expression for the conditional probability of a neuron to be in state "1" given the states of its presynaptic partners $p(z_k = 1 | \mathbf{z}_{\setminus k})$. Neuron parameters are given in Table 1.

Conductance-based LIF model

We performed additional experiments with conductance-based models to investigate the behavior in this more biologically realistic case. In this model, $u(t)$ evolves according to

$$C_m \frac{du}{dt} = -g_l(u - E_l) - g_{exc}(u - E_{exc}) - g_{inh}(u - E_{inh}) , \quad (14)$$

where $g_{exc}(t)$ and $g_{inh}(t)$ are the excitatory and inhibitory conductances at time t , and E_{exc} and E_{inh} are the excitatory and inhibitory reversal potentials, respectively. Just like in the current-based case, synaptic kernels are modeled as exponential:

$$\frac{dg_{exc}}{dt} = -\frac{g_{exc}}{\tau_{exc}} + \sum_{j \in \text{PRE}_{exc}} w_j S_j(t) \quad \frac{dg_{inh}}{dt} = -\frac{g_{inh}}{\tau_{inh}} + \sum_{j \in \text{PRE}_i} w_j S_j(t) \quad (15)$$

where the sums run over the sets of excitatory and inhibitory presynaptic spike sources, w_j is the quantal synaptic conductance of the synapse with the presynaptic neuron j , τ_{exc} and τ_{inh} are the time constants of excitatory and inhibitory

model	C_m pF	g_l nS	E_l mV	E_{exc} mV	E_{inh} mV	τ_{exc} ms	τ_{inh} ms	v_{th} mV	v_{reset} mV	τ_{ref} ms
current-based	200	2000	-50			10	10	-50	-55.1	10
conductance-based	250	25	-65	0	-80	2	3	-50	-65	3

Table 1: Neuron parameters.

synapses, respectively, and $S_j(t) = \sum_f \delta(t - t_j^{(f)})$ is the spike train of the presynaptic neuron j . The spiking mechanism is equivalent to the current-based case. The neuron parameters are given in Table 1.

Unlike in the current-based case, the variance of the free membrane potential has a non-monotonic dependence on background rates ν , becoming inversely proportional with ν for intense background input. This is a consequence of the decreased effective membrane time constant

$$\mathbb{E}[\tau_{eff}] = \frac{C_m}{g_l + \mathbb{E}[g_{exc}] + \mathbb{E}[g_{inh}]} = \frac{C_m}{g_l + w_{exc}\tau_{exc}\nu_{exc} + w_{inh}\tau_{inh}\nu_{inh}} \quad (16)$$

which also decreases the amplitude of the induced PSP. The resulting membrane potential distribution is still a Gaussian with moments ((Petrovici, 2016, Section 4.3); Zerlaut et al. (2018)):

$$\mu_u = \mathbb{E}[u]_{COBA} = \frac{g_l E_l + \sum_{x \in \{exc, inh\}} w_x \nu_x \tau_x E_x}{g_l + \sum_{x \in \{exc, inh\}} w_x \nu_x \tau_x} \quad \text{and} \quad (17)$$

$$\sigma_u^2 = \text{Var}[u]_{COBA} = \sum_{x \in \{exc, inh\}} \frac{\nu_x w_x^2 (E_x - \mathbb{E}[u])^2}{2(\mathbb{E}[\tau_{eff}] + \tau_x)} \left(\frac{\mathbb{E}[\tau_{eff}] \tau_x}{C_m} \right)^2. \quad (18)$$

The variance depends non-monotonically on the input rates as both $\mathbb{E}[\tau_{eff}]$ and $\mathbb{E}[u]$ depend on ν_x .

In the low-input limit ($g_l \gg g_e, g_i$), τ_{eff} is largely independent of the synaptic conductance and the generated membrane distribution $p(u)$ behaves similarly to the current-based case. In particular, the variance increases with both increasing input rates ν_{exc} and ν_{inh} and increasing synaptic weights w_{exc} and w_{inh} (see Eq. 2). In the high-conductance limit, i.e., $g_l \ll g_e, g_i$ and thereby $\mathbb{E}[\tau_{eff}] \rightarrow 0$, the variance becomes largely independent of the synaptic strengths and inversely proportional to the input rates ν_{exc} and ν_{inh} (Petrovici, 2016, Section 4.3):

$$\lim_{\sum_x \nu_x \rightarrow \infty} \text{Var}[u]_{COBA} = \frac{\sum_x w_x^2 \nu_x \tau_x (E_x - \mathbb{E}[u])^2}{(\sum_x w_x \nu_x \tau_x)^2} \propto \frac{1}{\sum_x \nu_x \tau_x}. \quad (19)$$

However, functional synaptic input also has a decreasing effect with increasing background conductance. The result of these two opposing phenomena is that the effect of increasing background rates on the neuronal response function in the conductance-based case matches the one for current-based neurons (see Eq. 3, Eq. 10 and Fig. 6f).

4.2 Spike response of sampling neurons

In the experiments underlying Fig. 1, we connect a current-based sampling neuron with one excitatory and one inhibitory Poisson source with a weights w_{exc} and w_{inh} , where $w_{exc} = -w_{inh}$, and vary the corresponding firing rates ν_{exc} and ν_{inh} . Background input parameters are listed in Table 3. We can freely choose the mapping of background rates to the Boltzmann temperature. For simplicity, we chose $T = 1$ in the lower range of physiological values, such that the readout can happen at low points in the oscillation cycle:

$$T = 1 \iff \nu_{exc} = \nu_{inh} = 2 \text{ kHz} \quad (20)$$

which results in a slope of $\beta = 1.39 \text{ nA}^{-1}$ and an offset of $I_0 = 1.34 \text{ nA}$ (see Fig. 1d,e). Since shifting the offset implies a change of the neuron's bias, we only have one degree of freedom when changing the temperature T . We, again arbitrarily, choose the excitatory rate ν_{exc} . The relationship $\nu_{inh} = h(\nu_{exc})$ is then found by interpolating the measured response functions from Fig. 1e.

In practice, this function can be approximated with the linear fit in Eq. 7 with $\nu_0 = -0.13 \text{ kHz}$, $m = 1.04$ and the coefficient of determination $r^2 = 0.99998$.

The five explicitly marked background configurations shown in Fig. 1b-e are given in Table 2.

color	ν_{exc} kHz	ν_{inh} kHz
blue	1.000	1.000
orange	2.000	2.000
red	4.000	4.000
green	2.848	1.232
purple	1.232	2.848

Table 2: Background parameters for the colored response functions and membrane potential distributions in Fig. 1.

4.3 Temperature as a function of background rates

The relationship between our temperature definition and the background rates can be approximated by linking the probability density function of the membrane potential to the derivative of the logistic response function. In the diffusion approximation, the free membrane potential distribution is Gaussian:

$$f(u; \mu_u, \sigma_u) = \frac{1}{\sqrt{2\pi}\sigma_u} \exp\left(-\frac{(u - \mu_u)^2}{2\sigma_u^2}\right). \quad (21)$$

In the high-conductance state, the cumulative distribution function (CDF) has a very similar shape to the (logistic) response function (Eq. 3). In particular, they have approximately the same derivative at their inflection point (for details see Petrovici et al. (2016)). With the parameter transformation $u_{\text{in}} = I_{\text{in}}/g_l$ and $\beta = \beta_u g_l$, where β_u is the slope in the potential domain, the response function reads:

$$\nu_{\text{out}}(u_{\text{in}}) = \frac{1}{1 + \exp(-\beta u_{\text{in}}/g_l)}. \quad (22)$$

The slope of the CDF at its inflection point is

$$\partial_u F|_{u=0} = f|_{u=0} = \frac{1}{\sqrt{2\pi}\sigma_u}, \quad (23)$$

whereas for the activation function it is

$$\partial_{u_{\text{in}}} \nu_{\text{out}}|_{u_{\text{in}}=0} = \frac{\beta \exp(-\beta u_{\text{in}}/g_l)}{g_l (1 + \exp(-\beta u_{\text{in}}/g_l))^2}|_{u_{\text{in}}=0} = \frac{\beta}{4g_l}. \quad (24)$$

Equating the two creates a direct correspondence between the inverse temperature β and the width σ_u of the free membrane potential distribution:

$$\beta(\sigma_u) \approx \frac{4g_l}{\sqrt{2\pi}\sigma_u}. \quad (25)$$

In our case, with $w_{\text{exc}} = w_{\text{inh}}$, $\tau_{\text{exc}} = \tau_{\text{inh}}$ and $1/k_B = \beta_{\text{ref}}$, plugging in the expression for σ_u from Eq. 2, the temperature in Eq. 6 is given by

$$T = \frac{\beta_{\text{ref}}}{\beta} = \sqrt{\frac{\nu_{\text{exc}} - \nu_{\text{inh}}}{\nu_{\text{exc,ref}} - \nu_{\text{inh,ref}}}}. \quad (26)$$

4.4 Entropy of spiking sampling networks

Networks of current-based LIF neurons can sample, to a very good approximation, from binary Boltzmann distributions

$$p(\mathbf{z}) = \frac{1}{Z} \exp\left(\frac{-E(\mathbf{z})}{k_B T}\right) \quad (27)$$

with energy function

$$E(\mathbf{z}) = -\frac{1}{2} \sum_{k,j} W_{kj} z_k z_j - \sum_k B_k z_k. \quad (28)$$

where \mathbf{W} is a symmetric zero-diagonal matrix and \mathbf{B} a bias vector (Petrovici et al., 2016). The associated neuronal response function represents a conditional state probability and reads

$$p(z_k = 1 | \mathbf{z}_{\setminus k}) = \frac{1}{1 + \exp\left(-\sum_j W_{kj} z_j - B_k\right)}. \quad (29)$$

The synaptic strength w_{kj} and input current $I_{in,k}$ in the equivalent LIF network can be related to the Boltzmann parameters W_{kj} and B_k via the slope of the response function β (cf. Eq. 3):

$$w_{kj} = \frac{W_{kj}}{\beta} \frac{g_l (\tau_s - \tau_m)}{\tau_s (1 - \exp(-1)) - \tau_m \left(1 - \exp\left(\frac{\tau_{ref}}{\tau_m}\right)\right)}, \quad (30)$$

$$I_{in,k} = \frac{B_k}{\beta} + I_0. \quad (31)$$

Biases are implemented via a shift of the leak potential E_l . Tempering is implemented by changing $\nu_{exc}, \nu_{inh} = f(\nu_{exc})$ via Eq. 7 according to a sinusoidal schedule with parameters shown in Table 3. The entropy is given by

$$S(p_T) = \sum_{\mathbf{z}} -p_T(\mathbf{z}) \log p_T(\mathbf{z}). \quad (32)$$

Depending on the base of the logarithm, the unit of S is either nats or bits.

Parameter choice of the Boltzmann distribution For the entropy scaling in Fig. 3, we use a 4-neuron network with random weights and biases distributed according to

$$\hat{W}_{kj} \propto \mathcal{N}(0.0, 0.5) \quad W_{kj} = \frac{\hat{W}_{kj} + \hat{W}_{jk}}{2} \quad B_k \propto \mathcal{N}(0.0, 0.5), \quad (33)$$

where $\mathcal{N}(\mu, \sigma)$ is the normal distribution with mean μ and standard deviation σ . The third and fourth neuron’s bias is set to ± 1 to ensure one leak-over-threshold and one leak-below-threshold neuron for Fig. 3e.

4.5 Image generation examples: NORB and MNIST

The layer sizes of our hierarchical networks are given in Table 3.

NORB In order to reduce the pixelation in Fig. 4a we do not plot the visible state $\mathbf{v} \in \{0, 1\}^{3600}$ directly but instead show the activation probability $p(\mathbf{v})$ that is imprinted by the instantaneous state of the hidden layer:

$$p_{T=1}(\mathbf{v} | \mathbf{h}) = \frac{1}{1 + \exp\left(-\mathbf{W}_{vh} \mathbf{h} - \mathbf{B}_v\right)}. \quad (34)$$

The temperature schedule of the oscillating background case can be found in Table 3. For the static background input we use the reference configuration (Eq. 20) and retrieve samples every $1/f_{osc} = 1$ s in order to get an equal-time comparison.

MNIST In Fig. 5 we use a similar network structure to the one in Fig. 4, with parameters from Leng et al. (2018). Background configurations are varied according to Eq. 7 as before and sine parameters are given in Table 3.

Kullback-Leibler Divergence The Kullback-Leibler divergence is a standard measure of the discrepancy between two probability distributions. Intuitively, it measures how many bits are wasted when encoding a distribution Q according to the optimal encoding for distribution P . For a discrete probability distribution P with respect to another Q , this divergence is defined as:

$$D_{KL}(P || Q) = \sum_i P(i) \log \left(\frac{P(i)}{Q(i)} \right). \quad (35)$$

Note that Q must be strictly positive, whereas P may have states with zero probabilities associated to it.

network unit	number of neurons	$\nu_{\text{exc},\text{min}}$ kHz	$\nu_{\text{exc},\text{max}}$ kHz	$\nu_{\text{inh},\text{min}}$ Hz	$\nu_{\text{inh},\text{max}}$ nA	f_{osc} nA	w_{exc}	w_{inh}
response function (see Fig. 1)	1	0.5 - 30	0.5 - 30	0.4 - 31	0.4 - 31	const.	0.5	-0.5
entropy (see Fig. 3)	4	0.25	10	0.1	10.3	1	1.0	-0.5
NORB (see Fig. 4)	(3600, 500, 10)	0.5	20	0.4	20.6	1	0.5	-0.5
MNIST (see Fig. 5)	(784, 400, 10)	0.5	22	0.4	22.7	$\frac{2}{3}$	0.5	-0.5

Table 3: Hierarchical network parameters.

Indirect sampling likelihood We quantitatively evaluate how well the samples generated by our networks reflect the target distribution by calculating the indirect sampling likelihood (ISL) described in Breuleux et al. (2010). The ISL measures the similarity between the generated samples and samples from the dataset that were not shown during training (test set). Each test sample \mathbf{y}_j and generated sample \mathbf{x}_i is a d -dimensional binary vector, whereby each \mathbf{x}_i is given by the instantaneous visible layer activity $\mathbf{v} \in \{0, 1\}^d$.

For retrieving the ISL, a density model \mathcal{P} is trained on N generated samples and the likelihood of each test sample under \mathcal{P} calculated. For d -dimensional binary vectors, a non-parametric kernel density estimator is suitable:

$$\mathcal{P}(\mathbf{y}) = \frac{1}{N} \sum_{i=1}^N \prod_{j=1}^d \gamma^{1_{\mathbf{y}_j = \mathbf{x}_{i,j}}} (1 - \gamma)^{1_{\mathbf{y}_j \neq \mathbf{x}_{i,j}}} , \quad (36)$$

which is essentially a mixture model representing the \mathbf{x}_i . The hyperparameter $\gamma \in [0.5, 1)$ determines, how much the empirical distribution over \mathbf{x}_i is smoothed out. We set $\gamma = 0.95$ following the setting from Breuleux et al. (2010). The two exponents denote identity functions that compare an individual test to a generated sample and count the identical and different pixels respectively. Intuitively, the ISL penalizes each test sample that is far from the generated samples.

In Fig. 5g, we plot $\log \mathcal{P}(\mathbf{y})$ averaged over all test samples versus the number of samples. This time course reveals how many main modes of the target distribution are covered and how fast. Note that the ISL does not necessarily evaluate how diverse the network output is, but rather how well the test set is covered – repetitive samples would yield a high ISL when compared to a not very diverse test set.

An upper bound of the ISL for our experiment is given by the output of an optimal sampler (OPT) and the lower bound by the product of marginals (POM) (see Fig. 5g). The optimal sampler draws randomly, without replacement, from a pool of 10^5 images that were generated with Adaptive Simulated Tempering (AST) (Salakhutdinov, 2010), a complex algorithm that is constructed for optimal mixing properties. The product-of-marginals sampler generates examples by independently sampling each vector component from its respective intensity distribution over the training set. Hence, the marginal probability distribution for each component is preserved and correlations between components, i.e., the overall structure, discarded.

One known drawback of the ISL is that it does not represent an accurate reflection of human perceptual judgment of image quality (Borji, 2019). Therefore, we additionally checkrf the sampling quality by eye, and evaluated the activation probability of the visible layer as shown in Fig. 5c,d. Based on this, we picked a point on the $f_{\text{osc}} = 0.5$ Hz plane with an intermediate ISL value for display in Fig. 5c-f.

Mode duration We calculate the mode duration as the average time between two mode switches, where the current mode is defined as the most active label unit. This reflects the network’s own interpretation of its current visible state \mathbf{v} and as such requires the network to be self-consistent. In practice, we did not find significant deviations (see Fig. 5c and d) from this assumption. Due to computational constraints, we only simulated 1000 s in a single run and averaged over multiple simulations for improved statistics. As such, the mode duration of 998 s in Fig. 5f is an artifact reflecting inability to mix.

4.6 Conductance-based input scenarios

Figure 6a-b show the different behaviors of $\text{Var}[u]$ using the neuron parameters given in Table 1. For the conductance-based models, we investigated four scenarios for conductance-based background input chosen to cover a range of behaviors of $\mathbb{E}[u]$ and $\text{Var}[u]$ for increasing background input frequencies ν_{exc} and ν_{inh} (see above and *main text*). This is accomplished by different ways of changing the rates ν_{exc} and ν_{inh} with the scaling parameter α (see Eq. 9) and different values of the background input weights w_{exc} and w_{inh} . All parameters are given in Table 4.

Using Eq. 17, we can balance the membrane potential at a target value of $\hat{u} = -55$ mV by choosing $\nu_{\text{inh},1}$ and $\nu_{\text{inh},0}$ (see

scenario	parameter	$\nu_{\text{exc},1}$	$\nu_{\text{exc},0}$	$\nu_{\text{inh},1}$	$\nu_{\text{inh},0}$	w_{exc}	w_{inh}	$w_{\text{stim,max}}$
	unit	kHz	kHz	kHz	kHz	nS	nS	nS
μ_u unbalanced		5	0	5	0	0.5	0.5	30
μ_u balanced at -55 mV		5	0	7.3	-6.7	0.5	0.5	70
high variance		5	0	5	0	2.5	3.75	70
low excitation		1	0	5	0	1	3.75	70

Table 4: Background input parameters for the different conductance-based input scenarios.

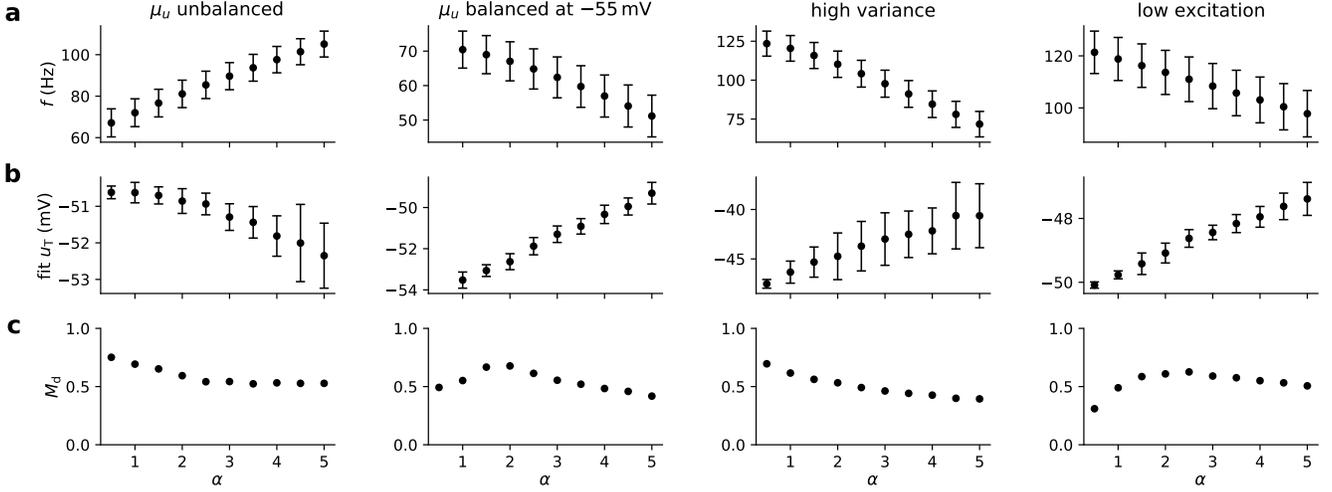


Figure 6 Supp. 1: **Details of stochastic model fitting.** (a) Firing rate in response to stimulus used for fitting. Dots denote mean, fliers standard deviation over $N = 50$ independent runs. (b) Soft threshold values resulting from fitting. (c) Quality of fit as described by M_d criterion (see *Methods*), measuring match of firing intensities of LIF model and stochastic models resulting from fitting. The fit models reproduce spiking behavior reasonably well ($M_d = 1$ indicates a perfect match, $M_d = 0$ indicates no overlap of PSTHs).

Eq. 9 and Table 4) so ν_{inh} changes as

$$\nu_{\text{inh}} = \nu_{\text{exc}} \frac{\tau_{\text{exc}} w_{\text{exc}} (E_{\text{exc}} - \hat{u})}{\tau_{\text{inh}} w_{\text{inh}} (\hat{u} - E_{\text{inh}})} + \frac{g_1 (E_1 - \hat{u})}{\tau_{\text{inh}} w_{\text{inh}} (\hat{u} - E_{\text{inh}})} . \quad (37)$$

4.7 Fitting of stochastic models

To assess the effect of background input on individual neurons' temperature in the conductance-based case, we generated data using realistic inputs to conductance-based LIF neurons and used the fitting method described by Jolivet et al. (2006) to obtain a stochastic model.

The stochastic model is identical to the LIF neurons (i.e., membrane potential generation, refractoriness after spike) except for the deterministic spike generation mechanism, which is replaced by a stochastic spike criterion using an instantaneous firing intensity of

$$\rho(t) = \frac{1}{\Delta t} \exp\left(\frac{u(t) - u_T}{T}\right) \quad (38)$$

where T and u_T are parameters (temperature and soft threshold) obtained from the fitting method. Spikes are drawn from a Poisson process with this instantaneous intensity. In our discrete-time simulations, we calculate the probability of a spike within each simulation time step Δt , which is

$$\Pr(\text{spike in } [t, t + \Delta t] | u(t)) = 1 - \exp(-\rho(t)\Delta t) , \quad (39)$$

and draw spikes accordingly.

To fit this model to data recorded from the simulated conductance-based LIF neurons, we estimated the spiking probability given the membrane potential u using a stimulus consisting of 100 inputs (80% excitatory), each firing according to a Poisson process with $f_{\text{stim}} = 5$ Hz. Each input had a synaptic weight drawn from a uniform distribution in $[0, w_{\text{stim,max}}]$ where $w_{\text{stim,max}}$ is a maximum conductance value (see Table 4) adjusted for each scenario so the LIF neuron fires at a reasonable rate (i.e., $50 \text{ Hz} < f < 150 \text{ Hz}$, see Fig. 6 Supp. 1). This stimulus was presented to the deterministic LIF models while recording the output spike times. An identical spike train was then presented to a passive version of the stochastic neurons (no firing mechanism, i.e., its membrane potential is the free membrane potential) which was also reset at every spike of the deterministic model. Binned histograms of u of the passive model at all times and at spike times of the original model allow estimating the firing probability $p(\text{spike}|u)$ (see Jolivet et al., 2006). To fit the model, we insert Eq. 38 into Eq. 39 and reformulate the result to get

$$\frac{u - u_T}{T} = \log \left(-\log \left(1 - p(\text{spike}|u) \right) \right), \quad (40)$$

where we perform linear regression on the right-hand side to get values of T and u_T . The shape of $p(\text{spike}|u)$ is roughly bell-shaped. We found that the best fits result from using only the values of $p(\text{spike}|u)$ from $u < \arg \max_u p(\text{spike}|u)$ for fitting as the values above the peak show increasing background.

We repeated this fitting $N = 50$ times for simulation durations of 100 s and report the mean and standard deviation of the resulting values of u_T and T (Fig. 6d and Fig. 6 Supp. 1b).

The stochastic models were evaluated by simulating 1000 deterministic and 1000 stochastic versions of the model for 1 s using a new stimulus. From these runs, the time-varying firing intensities $\nu_{\text{LIF}}(t)$ and $\nu_{\text{fit}}(t)$ were estimated. A criterion for the quality of fit between the fitted and original models was calculated as

$$M_d = \frac{2 \int \nu_{\text{LIF}}(t) \nu_{\text{fit}}(t) dt}{\int \nu_{\text{LIF}}^2(t) dt + \int \nu_{\text{fit}}^2(t) dt} \quad (41)$$

for every model. This similarity criterion, which determines how well the firing intensities match, is inspired by Mensi et al. (2011).¹ We found that the stochastic models were generally capable of reproducing the LIF behavior reasonably well (i.e., $M_d > 0.5$ for most models, see Fig. 8 Supp. 1c). The quality of the fit generally decreases when the variance is high (all models in the high variance scenario, large α in the remaining scenario). We used the results from fitting only for elucidating the temperature effect in networks of LIF neurons with background input, so any errors resulting from imperfect fits did not carry over to the experiments showing the functional advantages of background oscillations (where we again used LIF neurons and spiking background input for our simulations).

4.8 Illustrative sampling task for conductance-based networks

To illustrate the effect of the background input activity on the sampling behavior of an SNN with conductance-based LIF neurons (Fig. 7), we used a simple winner-take-all (WTA) network. The model consisted of 4 neurons receiving bias input by injecting currents of amplitudes $I_{\text{in}} = [40, 60, 80, 40]$ pA, respectively. Each neuron additionally received input from an external neuron (Poisson spiking at $f = 75$ Hz, synaptic conductance $w_{\text{in}} = w_{\text{stim,max}}$, see above and Table 4). Neurons had inhibitory lateral connections (conductance $w_{\text{inh}} = 3w_{\text{in}}$).

We ran this network for 100 s. From the spikes of each neuron, we computed network states by setting the state z_j of each neuron j to 1 if the neuron fired within the last 10 ms and otherwise to 0 (see Berkes et al., 2011). This allowed us to estimate the fraction of time the network spent in each state (Fig. 7a). Note that here, the duration of $z_j = 1$ after a spike does not coincide with the refractory period or the synaptic time constant.

We used only the modes (states in which one neuron was exclusively active, i.e., $z_j = 1$ for some $j = j_0$ and $z_k = 0$ for all $k \neq j_0$) to compute the entropy (Fig. 7b).

4.9 Stimulus disambiguation task

We illustrate the advantage of oscillatory background input using a biologically relevant task with several distinct modes which are far apart in the state space, thus making mixing hard. The circuit consisted of 3 winner-take-all (WTA) groups (Fig. 8a). Every group contained 3 assemblies, each formed by 3 neurons with strong recurrent connectivity (all neuron pairs bidirectionally connected) and lateral inhibition (bidirectional connections between all neuron pairs that are not part of the same assembly). Between groups, the n^{th} assemblies were bidirectionally linked: in every assembly, 2 (different) neurons

¹The M_d criterion as stated by Mensi et al. (2011, Eq. 16) seems to contain an error, therefore, it is slightly adapted here so its properties match those discussed by Mensi et al., i.e., a value of 1 indicates a perfect match, while a value of 0 indicates no match (e.g., if $\nu_{\text{fit}}(t) \equiv 0$).

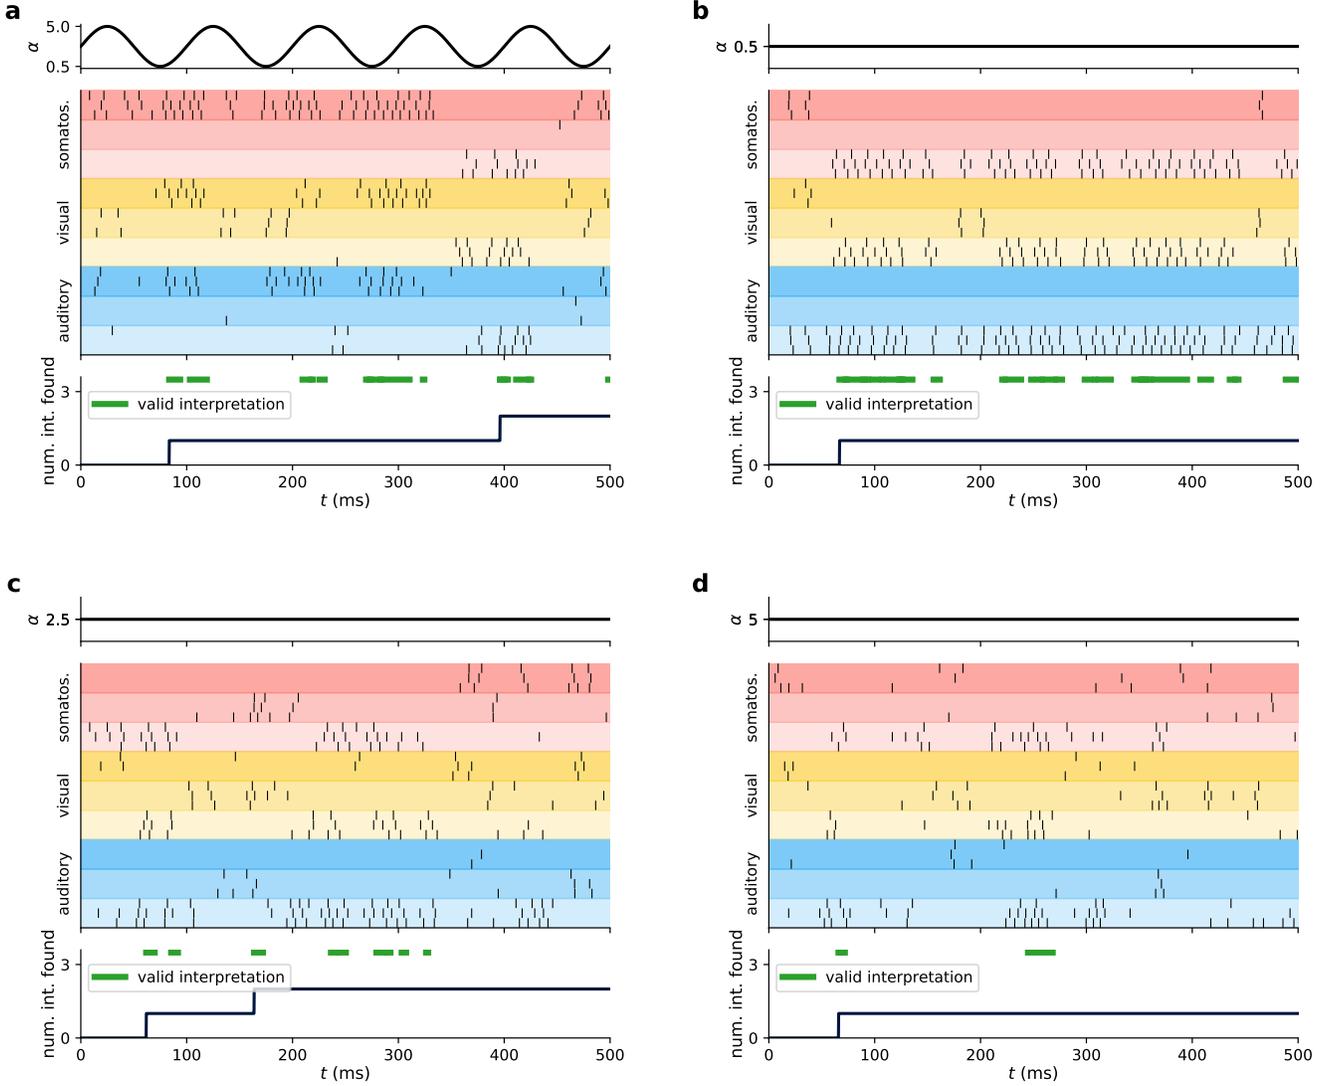


Figure 8 Supp. 1: **Activity in conductance-based network sampling task.** (a) Sample activity for oscillating background input (see Fig. 8b for details). (b-d) Sample activity for constant background input with $\alpha \in \{0.5, 2.5, 5\}$.

were connected to the other 2 assemblies (see Fig. 8a). One assembly in each group received additional bias input (Fig. 8f, g). This results in a stimulus disambiguation task with conflicting input. The synapse parameters are given in Table 5.

We found that using the parameters of each of the scenarios led to complete silence in the network for either low or high background input. To obtain network activity for all values of α , we adjusted the parameters of the background input. Starting from the μ_u unbalanced scenario, we increased the inhibitory background input conductance while keeping the excitatory conductance constant. The variance of the network firing rate (for oscillating background input) decreased until a ratio of $w_{inh}/w_{exc} = 1.35$, after which it increased again. We chose the value at the minimum (i.e., $w_{inh} = 0.675$). This results in network activity also for low and high constant background input, which allows us to compare the effect of background oscillations to constant input without favoring the former. As there are no input units exciting the network and the background input to each neuron is insufficient to evoke network activity, it was necessary to inject a current into each neuron so neurons did not remain silent. At $w_{inh} = 0.675$, we used $I_{in} = 350$ pA, which resulted in a mean firing rate of $f \approx 17$ Hz (with oscillating background input).

The circuit defines a sampling problem with 3 modes, i.e., interpretations. Network states were defined to encode the n^{th} interpretation if the n^{th} assembly in each group was simultaneously active while all other assemblies remained inactive. An assembly was regarded as active at every point in time if 50% of its neurons (i.e., at least 2 of the 3 neurons within an assembly) fired within the last 10 ms, otherwise, it was regarded as inactive. This definition allows one to characterize the network state at each time step of the discrete-time simulation as a state either encoding a particular interpretation or none.

connection unit	g nS	E_{syn} mV	syn. delay ms
between assemblies	17	0	2
within assemblies	8.5	0	2
inhibitory	17	-80	0.1

Table 5: Connection parameters for conductance-based stimulus disambiguation experiments.

Background activity was provided to the network via Poisson sources. Each neuron received independent background input, with rates scaled by α as in previous experiments (as in the μ_u unbalanced scenario, see the first column in Fig. 6c-f and Table 4). The scaling factor α was sinusoidally modulated over time, with $\alpha(t) \in [0.5, 5]$ and modulation frequency $f_{\text{osc}} = 10$ Hz (see Fig. 8b top). We compared the results in this case to the results when α was kept constant ($\alpha \in \{0.5, 2.5, 5\}$). Fig. 8 Supp. 1 shows sample behavior for the different cases, highlighting the different behavioral regimes (e.g., locking into one solution for $\alpha \equiv 0.5$, see *main text*).

To estimate the probability of the network state encoding a solution, we repeated $N = 100$ simulation runs lasting 20 s each for all four background input setups. For the constant α cases, we report the fraction of network states that encode one of the 3 solutions over all runs. For the oscillatory case, we estimated the fractions of states encoding any interpretation as mean and standard deviation over the 100 runs and plot mean and standard deviation (Fig. 8c). Here, we used a bias input of 40 nS for the 3 biased assemblies.

We then tested the mixing behavior in two ways. First, we estimated the time it took to find all solutions by running the network $N = 100$ times for 20 s for each of the four background input setups. We recorded how long it took for the network state to visit each of the 3 solutions at least once in each of these simulations (Fig. 8d shows mean and standard deviation). If the simulation time was not enough for the network to visit all solutions the runs were discarded (this only occurred for uneven input and $\alpha = 0.5$, where about 1/3 of the runs were discarded). We also estimated the time it took to switch between solutions (i.e., the mode duration) over these simulation runs. Switching times were defined as the difference between the time the network state changed to any solution state and the time the network state next changed to a solution state for a different solution (i.e., difference between solution state onset times, Fig. 8e shows mean and standard deviation). Significance values were calculated using the Wilcoxon rank-sum test.

Fig. 8f shows mean and standard deviation of the fraction of valid states for each interpretation for these $N = 100$ runs. To test whether the interpretations are visited according to the bias input, we set different currents for the three biased assemblies and repeated this analysis (Fig. 8g).

4.10 Relating model features to experimental data

To relate the behavior of the conductance-based model to the experimental data from Jezek et al. (2011), we constructed a circuit consisting of 2 assemblies with recurrent connectivity ($p = 0.1$ for one synapse between each pair of distinct neurons, $w = 2.5$ nS, $E = 0$ mV, synaptic delays randomly chosen from a uniform distribution in $[1, 3]$ ms) and lateral inhibition ($p = 0.5$ for one synapse between each pair of distinct neurons, $w = 5$ nS, $E = -80$ mV, synaptic delay 0.1 ms). One of these assemblies encoded the correct interpretation of the spatial context, neurons in this assembly received a bias input of 10 pA. Background input (oscillating at 8 Hz, i.e., at a medium to high theta frequency) was given to each neuron as in the stimulus disambiguation task. We varied the ratio of inhibitory to excitatory synaptic background input weight in $[0.5, 2]$ as this results in different activity regimes (see above). This leads to $\mathbb{E}[g_{\text{inh}}/g_{\text{exc}}] \in [0.75, 3]$. This model again required current injection, due to the stark influence of the background input (Fig. 9b shows the resulting firing intensity behaviors of stochastic models fitted to LIF neurons at 3 values of the conductance ratio). The injected current required scaling depending on the conductance ratio (current linearly interpolated from $I_{\text{in}} = 40$ nS at $\mathbb{E}[g_{\text{inh}}/g_{\text{exc}}] = 0.75$ to $I_{\text{in}} = 880$ nS at $\mathbb{E}[g_{\text{inh}}/g_{\text{exc}}] = 3$). This model showed flickering behavior similar to the data shown in Jezek et al. (2011) along the entire parameter range, see Fig. 9a.

Fig. 9c shows the firing rate of the neurons in the model over the phase of the background input (mean values over $N = 100$ runs lasting $T = 50$ s). At every conductance ratio, we defined the start of the cycle according to the minimum firing rate of the network (as in Jezek et al., 2011). We calculated mixed network states (defined as states in which both assemblies were more than 20% active, similar to Jezek et al., 2011) over the phase (Fig. 9d shows mean and standard deviation over the 100 runs for 3 conductance ratios). Fig. 9e shows how the mean probability of mixed states within the first and second half of the cycle change over the parameter range. We also fitted stochastic models to data at each conductance ratio and found that the response functions intersect at $p = 0.5$ around $\mathbb{E}[g_{\text{inh}}/g_{\text{exc}}] = 1.8$.

4.11 Simulation details

The simulations of sampling experiments with current-based neurons were performed with sbs (Breitwieser et al., 2020) version 1.8.2 with slight modifications. This framework was executed with PyNN (Davison et al., 2009) version 0.9.1 and NEST (Peyser et al., 2017) version 2.14.0 with a time resolution of $\Delta t = 0.1$ ms.

The simulations using conductance-based models were performed using Brian2 (Stimberg et al., 2019) version 2.4.2 with a time resolution of $\Delta t = 0.05$ ms.

Acknowledgments

The authors gratefully acknowledge funding from the state of Baden-Württemberg through bwHPC, from the German Research Foundation (DFG) through grant no INST 39/963-1 FUGG (bwForCluster NEMO), from the European Union under grant agreements 604102, 720270, 785907, 945539 (HBP), from the Helmholtz Association Initiative and Networking Fund under project number SO-092 (Advanced Computing Architectures, ACA), from the Austrian Science Fund (FWF): I 3251-N33, and from the Manfred Stärk Foundation. We thank Wolfgang Maass for helpful discussions.

References

- A. R. Adamantidis, C. Gutierrez Herrera, and T. C. Gent. Oscillating circuitries in the sleeping brain. *Nature Reviews Neuroscience*, 20(12):746–762, 2019. doi: 10.1038/s41583-019-0223-4.
- L. Aitchison and M. Lengyel. The Hamiltonian brain: efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLOS Computational Biology*, 12(12):e1005186, 2016. doi: 10.1371/journal.pcbi.1005186.
- M. Al-Shedivat, R. Naous, G. Cauwenberghs, and K. N. Salama. Memristors empower spiking neurons with stochasticity. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 5(2):242–253, 2015. doi: 10.1109/jetcas.2015.2435512.
- M. Bartos, I. Vida, M. Frotscher, J. R. Geiger, and P. Jonas. Rapid signaling at inhibitory synapses in a dentate gyrus interneuron network. *Journal of Neuroscience*, 21(8):2687–2698, 2001. doi: 10.1523/JNEUROSCI.21-08-02687.2001.
- E. Başar, C. Başar-Eroglu, S. Karakaş, and M. Schürmann. Gamma, alpha, delta, and theta oscillations govern cognitive processes. *International Journal of Psychophysiology*, 39(2-3):241–248, 2001. doi: 10.1016/s0167-8760(00)00145-8.
- P. Berkes, G. Orbán, M. Lengyel, and J. Fiser. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013):83–87, 2011. doi: 10.1126/science.1195870.
- J. Binas, G. Indiveri, and M. Pfeiffer. Spiking analog vlsi neuron assemblies as constraint satisfaction problem solvers. In *2016 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 2094–2097. IEEE, IEEE, 2016. doi: 10.1109/ISCAS.2016.7538992.
- A. Borji. Pros and cons of gan evaluation measures. *Computer Vision and Image Understanding*, 179:41–65, 2019. doi: 10.1016/j.cviu.2018.10.009.
- O. Breitwieser, A. Baumbach, A. Korcsak-Gorzo, J. Klähn, M. Brixner, and M. Petrovici. sbs: Spike-based sampling (v1.8.2), 2020.
- O. Breuleux, Y. Bengio, and P. Vincent. Unlearning for better mixing. Technical report, Universite de Montreal/DIRO, 2010.
- N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8(3):183–208, 2000. doi: 10.1023/A:1008925309027.
- L. Buesing, J. Bill, B. Nessler, and W. Maass. Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLOS Computational Biology*, 7(11):e1002211, 2011. doi: 10.1371/journal.pcbi.1002211.
- T. J. Buschman and E. K. Miller. Shifting the spotlight of attention: evidence for discrete computations in cognition. *Frontiers in human neuroscience*, 4:194, 2010. doi: 10.3389/fnhum.2010.00194.
- G. Buzsáki. *Rhythms of the Brain*. Oxford University Press, USA, 2006.
- G. Buzsáki and A. Draguhn. Neuronal oscillations in cortical networks. *Science*, 304(5679):1926–1929, 2004. doi: 10.1126/science.1099745.

- C. Capone, E. Pastorelli, B. Golosio, and P. S. Paolucci. Sleep-like slow oscillations improve visual classification through synaptic homeostasis and memory association in a thalamo-cortical model. *Scientific reports*, 9(1):1–11, 2019. doi: 10.1038/s41598-019-45525-0.
- R. J. Compton, D. Gearinger, and H. Wild. The wandering mind oscillates: EEG alpha power is enhanced during moments of mind-wandering. *Cognitive, Affective and Behavioral Neuroscience*, 19(5):1184–1191, 2019. doi: 10.3758/s13415-019-00745-9.
- A. Davison, D. Brüderle, J. Eppler, J. Kremkow, E. Muller, D. Pecevski, L. Perrinet, and P. Yger. Pynn: a common interface for neuronal network simulators. *Frontiers in Neuroinformatics*, 2:11, 2009. doi: 10.3389/neuro.11.011.2008.
- A. Destexhe. Self-sustained asynchronous irregular states and up-down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons. *Journal of Computational Neuroscience*, 27(3):493–506, 2009. doi: 10.1007/s10827-009-0164-4.
- A. Destexhe, M. Rudolph, and D. Paré. The high-conductance state of neocortical neurons in vivo. *Nature Reviews Neuroscience*, 4(9):739–751, 2003. doi: 10.1038/nrn1198.
- S. Diekelmann and J. Born. Slow-wave sleep takes the leading role in memory reorganization. *Nature Reviews Neuroscience*, 11(3):218–218, 2010. doi: 10.1038/nrn2762-c2.
- D. Dold, I. Bytschok, A. F. Kungl, A. Baumbach, O. Breitwieser, W. Senn, J. Schemmel, K. Meier, and M. A. Petrovici. Stochasticity from function—why the Bayesian brain may need no noise. *Neural Networks*, 119:200–213, 2019. doi: 10.1016/j.neunet.2019.08.002.
- R. Echeveste, L. Aitchison, G. Hennequin, and M. Lengyel. Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference. *Nature Neuroscience*, 23(9):1138–1149, 2020. doi: 10.1038/s41593-020-0671-1.
- A. S. Ecker, P. Berens, G. A. Keliris, M. Bethge, N. K. Logothetis, and A. S. Tolias. Decorrelated neuronal firing in cortical microcircuits. *Science*, 327(5965):584–587, 2010. doi: 10.1126/science.1179867.
- J. Fiser, P. Berkes, G. Orbán, and M. Lengyel. Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3):119–130, 2010. doi: 10.1016/j.tics.2010.01.003.
- G. A. Fonseca Guerra and S. B. Furber. Using stochastic spiking neural networks on spinnaker to solve constraint satisfaction problems. *Frontiers in Neuroscience*, 11:714, 2017. doi: 10.3389/fnins.2017.00714.
- I. Grothe, S. D. Neitzel, S. Mandon, and A. K. Kreiter. Switching neuronal inputs by differential modulations of gamma-band phase-coherence. *Journal of Neuroscience*, 32(46):16172–16180, 2012. doi: 10.1523/jneurosci.0890-12.2012.
- S. Habenschuss, Z. Jonke, and W. Maass. Stochastic computations in cortical microcircuit models. *PLOS Computational Biology*, 9(11), 2013. doi: 10.1371/journal.pcbi.1003311.
- B. Haider, A. Duque, A. R. Hasenstaub, and D. A. McCormick. Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. *Journal of Neuroscience*, 26(17):4535–4545, 2006. doi: 10.1523/jneurosci.5297-05.2006.
- M. Häusser and A. Roth. Estimating the time course of the excitatory synaptic conductance in neocortical pyramidal cells using a novel voltage jump method. *Journal of Neuroscience*, 17(20):7606–7625, 1997. doi: 10.1523/jneurosci.17-20-07606.1997.
- F. Huang and Y. LeCun. The norb dataset, 2004. URL <https://cs.nyu.edu/~ylclab/data/norb-v1.0-small/>.
- E. M. Izhikevich. Polychronization: computation with spikes. *Neural Computation*, 18(2):245–282, 2006. doi: 10.1162/089976606775093882.
- K. Jezek, E. J. Henriksen, A. Treves, E. I. Moser, and M.-B. Moser. Theta-paced flickering between place-cell maps in the hippocampus. *Nature*, 478(7368):246–249, 2011. doi: 10.1002/hipo.22743.
- R. Jolivet, A. Rauch, H.-R. Lüscher, and W. Gerstner. Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *Journal of Computational Neuroscience*, 21(1):35–49, 2006. doi: 10.1007/s10827-006-7074-5.
- Z. Jonke, S. Habenschuss, and W. Maass. Solving constraint satisfaction problems with networks of spiking neurons. *Frontiers in neuroscience*, 10:118, 2016. doi: 10.3389/fnins.2016.00118.

- M. Katkov, S. Romani, and M. Tsodyks. Memory Retrieval from First Principles. *Neuron*, 94(5):1027–1032, 2017. doi: 10.1016/j.neuron.2017.03.048.
- S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science (New York, N.Y.)*, 220(4598): 671–80, 1983.
- W. Klimesch. Alpha-band oscillations, attention, and controlled access to stored information. *Trends in Cognitive Sciences*, 16(12):606–617, 2012. doi: 10.1016/j.tics.2012.10.007.
- J. G. Klinzing, N. Niethard, and J. Born. Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience*, 22(10):1598–1610, 2019. doi: 10.1038/s41593-019-0467-3.
- C. Koch, M. Rapp, and I. Segev. A brief history of time (constants). *Cerebral cortex*, 6(2):93–101, 1996. doi: 10.1093/cercor/6.2.93.
- A. F. Kungl, S. Schmitt, J. Klähn, P. Müller, A. Baumbach, D. Dold, A. Kugele, E. Müller, C. Koke, M. Kleider, et al. Accelerated physical emulation of Bayesian inference in spiking neural networks. *Frontiers in Neuroscience*, 13:1201, 2019. doi: 10.3389/fnins.2019.01201.
- M. E. Larkum, W. Senn, and H. Lüscher. Top-down Dendritic Input Increases the Gain of Layer 5 Pyramidal Neurons. *Cerebral Cortex*, 14(10):1059–1070, 2004. doi: 10.1093/cercor/bhh065.
- Y. LeCun and C. Cortes. MNIST handwritten digit database, 2010. URL <http://yann.lecun.com/exdb/mnist/>.
- L. Leng, R. Martel, O. Breitwieser, I. Bytschok, W. Senn, J. Schemmel, K. Meier, and M. A. Petrovici. Spiking neurons with short-term synaptic plasticity form superior generative networks. *Scientific Reports*, 8(1):10651, 2018. doi: 10.1038/s41598-018-28999-2.
- M. Lundqvist, A. M. Bastos, and E. K. Miller. Preservation and changes in oscillatory dynamics across the cortical hierarchy. *Journal of Cognitive Neuroscience*, 32(10):2024–2035, 2020. doi: 10.1162/jocn_a-01600.
- C. Lustenberger, M. R. Boyle, A. A. Foulser, J. M. Mellin, and F. Fröhlich. Functional role of frontal alpha oscillations in creativity. *Cortex*, 67:74–82, 2015. doi: 10.1016/j.cortex.2015.03.012.
- E. Marinari and G. Parisi. Simulated tempering: A new monte carlo scheme. *Europhysics Letters (EPL)*, 19(6):451–458, 1992. doi: 10.1209/0295-5075/19/6/002.
- A. Marzecová, A. Schettino, A. Widmann, I. SanMiguel, S. A. Kotz, and E. Schröger. Attentional gain is modulated by probabilistic feature expectations in a spatial cueing task: ERP evidence. *Scientific Reports*, 8(1):1–14, 2018. doi: 10.1038/s41598-017-18347-1.
- D. A. McCormick, B. W. Connors, J. W. Lighthall, and D. A. Prince. Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *Journal of neurophysiology*, 54(4):782–806, 1985. doi: 10.1152/jn.1985.54.4.782.
- M. Megias, Z. Emri, T. Freund, and A. Gulyas. Total number and distribution of inhibitory and excitatory synapses on hippocampal CA1 pyramidal cells. *Neuroscience*, 102(3):527–540, 2001. doi: 10.1016/s0306-4522(00)00496-6.
- S. Mensi, R. Naud, and W. Gerstner. From stochastic nonlinear integrate-and-fire to generalized linear models. In *Advances in Neural Information Processing Systems*, pages 1377–1385, 2011.
- L. Muller and A. Destexhe. Propagating waves in thalamus, cortex and the thalamocortical system: experiments and models. *Journal of Physiology-Paris*, 106(5-6):222–238, 2012. doi: 10.1016/j.jphysparis.2012.06.005.
- R. M. Neal. Sampling from multimodal distributions using tempered transitions. *Statistics and computing*, 6(4):353–366, 1996. doi: 10.1007/bf00143556.
- J. O’Keefe and M. L. Recce. Phase relationship between hippocampal place units and the eeg theta rhythm. *Hippocampus*, 3(3):317–330, 1993. doi: 10.1002/hipo.450030307.
- G. Orbán, P. Berkes, J. Fiser, and M. Lengyel. Neural variability and sampling-based probabilistic representations in the visual cortex. *Neuron*, 92(2):530–543, 2016. doi: 10.1016/j.neuron.2016.09.038.
- D. Pecevski, L. Buesing, and W. Maass. Probabilistic inference in general graphical models through sampling in stochastic networks of spiking neurons. *PLOS Computational Biology*, 7(12):e1002294, 2011. doi: 10.1371/journal.pcbi.1002294.

- M. A. Petrovici. *Form versus function: theory and models for neuronal substrates*. Springer International Publishing, 1 edition, 2016. doi: 10.1007/978-3-319-39552-4.
- M. A. Petrovici, J. Bill, I. Bytschok, J. Schemmel, and K. Meier. Stochastic inference with spiking neurons in the high-conductance state. *Physical Review E*, 94(4):042312, 2016. doi: 10.1103/physreve.94.042312.
- J. D. Pettigrew. Searching for the Switch: Neural Bases for Perceptual Rivalry Alternations. *Brain and Mind*, 2:85–118, 2001. doi: 10.1023/A:1017929617197.
- A. Peyser, A. Sinha, S. B. Vennemo, T. Ippen, J. Jordan, S. Graber, A. Morrison, G. Trensche, T. Fardet, H. Mørk, J. Hahne, J. Schuecker, M. Schmidt, S. Kunkel, D. Dahmen, J. M. Eppler, S. Diaz, D. Terhorst, R. Deepu, P. Weidel, I. Kitayama, S. Mahmoudian, D. Kappel, M. Schulze, S. Appukuttan, T. Schumann, H. C. Tunç, J. Mitchell, M. Hoff, E. Müller, M. M. Carvalho, B. Zajzon, and H. E. Plesser. *Nest 2.14.0*, 2017.
- D. Probst, M. A. Petrovici, I. Bytschok, J. Bill, D. Pecevski, J. Schemmel, and K. Meier. Probabilistic inference in discrete spaces can be implemented into networks of lif neurons. *Frontiers in Computational Neuroscience*, 9:13, 2015. doi: 10.3389/fncom.2015.00013.
- J. H. Reynolds and L. Chelazzi. Attentional modulation of visual processing. *Annual Review of Neuroscience*, 27:611–647, 2004. doi: 10.1146/annurev.neuro.26.041002.131039.
- K. Roy, A. Jaiswal, and P. Panda. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575(7784):607–617, 2019. doi: 10.1038/s41586-019-1677-2.
- R. Salakhutdinov. Learning in markov random fields using tempered transitions. In *Proceedings of the 22nd International Conference on Neural Information Processing Systems, NIPS'09*, page 1598–1606, Red Hook, NY, USA, 2009. Curran Associates Inc. ISBN 9781615679119.
- R. Salakhutdinov. Learning deep boltzmann machines using adaptive mcmc. In *ICML 2010 - Proceedings, 27th International Conference on Machine Learning*, pages 943–950, 2010.
- C. Savin, P. Dayan, and M. Lengyel. Optimal recall from bounded metaplastic synapses: predicting functional adaptations in hippocampal area ca3. *PLoS Computational Biology*, 10(2):e1003489, 2014. doi: 10.1371/journal.pcbi.1003489.
- A. Sengupta, P. Panda, P. Wijesinghe, Y. Kim, and K. Roy. Magnetic tunnel junction mimics stochastic cortical spiking neurons. *Scientific Reports*, 6(1):1–8, 2016. doi: 10.1038/srep30039.
- M. N. Shadlen and W. T. Newsome. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience*, 18(10):3870–3896, 1998. doi: 10.1523/jneurosci.18-10-03870.1998.
- Y. Shu, A. Hasenstaub, and D. A. McCormick. Turning on and off recurrent balanced cortical activity. *Nature*, 423(6937):288–293, 2003. doi: 10.1038/nature01616.
- V. S. Sohal and M. E. Hasselmo. Changes in GABAB modulation during a theta cycle may be analogous to the fall of temperature during annealing. *Neural Computation*, 10(4):869–882, 1998a. doi: 10.1162/089976698300017539.
- V. S. Sohal and M. E. Hasselmo. GABAB modulation improves sequence disambiguation in computational models of hippocampal region CA3. *Hippocampus*, 8(2):171–193, 1998b. doi: 10.1002/(sici)1098-1063(1998)8:2<171::aid-hipo9>3.0.co;2-o.
- L. Stålesen Ramfjord, E. Hertenstein, K. Fehér, C. Mikutta, C. L. Schneider, C. Nissen, and J. G. Maier. Local sleep and wakefulness—the concept and its potential for the understanding and treatment of insomnia disorder. *Somnologie*, 24(2):116–120, 2020. doi: 10.1007/s11818-020-00245-w.
- F. Stella, P. Baracska, J. O’Neill, and J. Csicsvari. Hippocampal Reactivation of Random Trajectories Resembling Brownian Diffusion. *Neuron*, 102(2):450–461.e7, 2019. doi: 10.1016/j.neuron.2019.01.052.
- M. Steriade, I. Timofeev, N. Dürmüller, and F. Grenier. Dynamic properties of corticothalamic neurons and local cortical interneurons generating fast rhythmic (30–40 hz) spike bursts. *Journal of Neurophysiology*, 79(1):483–490, 1998. doi: 10.1152/jn.1998.79.1.483.
- M. Stimberg, R. Brette, and D. F. Goodman. Brian 2, an intuitive and efficient neural simulator. *eLife*, 8:e47314, 2019. doi: 10.7554/eLife.47314.

- X.-J. Wang. Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological Reviews*, 90(3): 1195–1268, 2010. doi: 10.1152/physrev.00035.2008.
- Y. Zerlaut, S. Chemla, F. Chavane, and A. Destexhe. Modeling mesoscopic cortical dynamics using a mean-field model of conductance-based networks of adaptive exponential integrate-and-fire neurons. *Journal of Computational Neuroscience*, 44(1):45–61, 2018. doi: 10.1007/s10827-017-0668-2.