

RESEARCH ARTICLE

# Network Plasticity as Bayesian Inference

David Kappel<sup>\*†</sup>, Stefan Habenschuss<sup>‡</sup>, Robert Legenstein, Wolfgang Maass

Institute for Theoretical Computer Science, Graz University of Technology, A-8010 Graz, Austria

<sup>†</sup> These authors contributed equally to this work.

\* [david@igi.tugraz.at](mailto:david@igi.tugraz.at)

## Abstract

General results from statistical learning theory suggest to understand not only brain computations, but also brain plasticity as probabilistic inference. But a model for that has been missing. We propose that inherently stochastic features of synaptic plasticity and spine motility enable cortical networks of neurons to carry out probabilistic inference by sampling from a posterior distribution of network configurations. This model provides a viable alternative to existing models that propose convergence of parameters to maximum likelihood values. It explains how priors on weight distributions and connection probabilities can be merged optimally with learned experience, how cortical networks can generalize learned information so well to novel experiences, and how they can compensate continuously for unforeseen disturbances of the network. The resulting new theory of network plasticity explains from a functional perspective a number of experimental data on stochastic aspects of synaptic plasticity that previously appeared to be quite puzzling.



## OPEN ACCESS

**Citation:** Kappel D, Habenschuss S, Legenstein R, Maass W (2015) Network Plasticity as Bayesian Inference. *PLoS Comput Biol* 11(11): e1004485. doi:10.1371/journal.pcbi.1004485

**Editor:** Jeff Beck, Duke University, UNITED STATES

**Received:** April 8, 2015

**Accepted:** August 3, 2015

**Published:** November 6, 2015

**Copyright:** © 2015 Kappel et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** All relevant data are within the paper and its Supporting Information files.

**Funding:** Written under partial support of the European Union project #604102 The Human Brain Project (HBP) and CHIST-ERA ERA-Net (Project FWF #I753-N23, PNEUMA). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

**Competing Interests:** The authors have declared that no competing interests exist.

## Author Summary

Synaptic connectivity between neurons in the brain and the efficacies (“weights”) of these synaptic connections are thought to encode the long-term memory of an organism. But a closer look at their molecular implementation, as well as imaging experiments over longer periods of time, have shown that synaptic connections are subject to numerous stochastic processes. We propose that this seeming unreliability of synaptic connections is not a bug, but an important feature. It endows networks of neurons with an important experimentally observed but theoretically not understood capability: Automatic compensation for internal and external changes. This perspective of network plasticity requires a new conceptual and mathematical framework, which is provided by this article. Stochasticity of synapses is seen here not as noise of an inherently deterministic system, but as an inherent property, similarly as Brownian motion of particles in a physical system cannot be abstracted away if one wants to understand certain properties of a physical system. In fact, we find that this underlying stochasticity of synaptic connections enables a network of neurons to continuously try out new network configurations while maintaining its functionality.

## Introduction

We reexamine in this article the conceptual and mathematical framework for understanding the organization of plasticity in networks of neurons in the brain. We will focus on synaptic plasticity and network rewiring (spine motility) in this article, but our framework is also applicable to other network plasticity processes. One commonly assumes, that plasticity moves network parameters  $\theta$  (such as synaptic connections between neurons and synaptic weights) to values  $\theta^*$  that are optimal for the current computational function of the network. In learning theory, this view is made precise for example as maximum likelihood learning, where model parameters  $\theta$  are moved to values  $\theta^*$  that maximize the fit of the resulting internal model to the inputs  $\mathbf{x}$  that impinge on the network from its environment (by maximizing the likelihood of these inputs  $\mathbf{x}$ ). The convergence to  $\theta^*$  is often assumed to be facilitated by some external regulation of learning rates, that reduces the learning rate when the network approaches an optimal solution.

This view of network plasticity has been challenged on several grounds. From the theoretical perspective it is problematic because in the absence of an intelligent external controller it is likely to lead to overfitting of the internal model to the inputs  $\mathbf{x}$  it has received, thereby reducing its capability to generalize learned knowledge to new inputs. Furthermore, networks of neurons in the brain are apparently exposed to a multitude of internal and external changes and perturbations, to which they have to respond quickly in order to maintain stable functionality.

Other experimental data point to surprising ongoing fluctuations in dendritic spines and spine volumes, to some extent even in the adult brain [1] and in the absence of synaptic activity [2]. Also a significant portion of axonal side branches and axonal boutons were found to appear and disappear within a week in adult visual cortex, even in the absence of imposed learning and lesions [3]. Furthermore surprising random drifts of tuning curves of neurons in motor cortex were observed [4]. Apart from such continuously ongoing changes in synaptic connections and tuning curves of neurons, massive changes in synaptic connectivity were found to accompany functional reorganization of primary visual cortex after lesions, see e.g. [5].

We therefore propose to view network plasticity as a process that continuously moves high-dimensional network parameters  $\theta$  within some low-dimensional manifold that represents a compromise between overriding structural rules and different ways of fitting the internal model to external inputs  $\mathbf{x}$ . We propose that ongoing stochastic fluctuations (not unlike Brownian motion) continuously drive network parameters  $\theta$  within such low-dimensional manifold. The primary conceptual innovation is the departure from the traditional view of learning as moving parameters to values  $\theta^*$  that represent optimal (or locally optimal) fits to network inputs  $\mathbf{x}$ . We show that our alternative view can be turned into a precise learning model within the framework of probability theory. This new model satisfies theoretical requirements for handling priors such as structural constraints and rules in a principled manner, that have previously already been formulated and explored in the context of artificial neural networks [6, 7], as well as more recent challenges that arise from probabilistic brain models [8]. The low-dimensional manifold of parameters  $\theta$  that becomes the new learning goal in our model can be characterized mathematically as the high probability regions of the posterior distribution  $p^*(\theta|\mathbf{x})$  of network parameters  $\theta$ . This posterior arises as product of a general prior  $p_S(\theta)$  for network parameters (that enforces structural rules) with a term that describes the quality of the current internal model (e.g. in a predictive coding or generative modeling framework: the likelihood  $p_{\mathcal{N}}(\mathbf{x}|\theta)$  of inputs  $\mathbf{x}$  for the current parameter values  $\theta$  of the network  $\mathcal{N}$ ). More precisely, we propose that brain plasticity mechanisms are designed to enable brain networks to sample from this posterior distribution  $p^*(\theta|\mathbf{x})$  through inherent stochastic features of their molecular implementation. In this way synaptic and other plasticity processes are able to carry out probabilistic (or Bayesian) inference through sampling from a posterior distribution

that takes into account both structural rules and fitting to external inputs. Hence this model provides a solution to the challenge of [8] to understand how posterior distributions of weights can be represented and learned by networks of neurons in the brain.

This new model proposes to reexamine rules for synaptic plasticity. Rather than viewing trial-to-trial variability and ongoing fluctuations of synaptic parameters as the result of a sub-optimal implementation of an inherently deterministic plasticity process, it proposes to model experimental data on synaptic plasticity by rules that consist of three terms: the standard (typically deterministic) activity-dependent (e.g., Hebbian or STDP) term that fits the model to external inputs, a second term that enforces structural rules (priors), and a third term that provides the stochastic driving force. This stochastic force enables network parameters to sample from the posterior, i.e., to fluctuate between different possible solutions of the learning task. The stochastic third term can be modeled by a standard formalism (stochastic Wiener process) that had been developed to model Brownian motion. The first two terms can be modeled as drift terms in a stochastic process. A key insight is that one can easily relate details of the resulting more complex rules for the dynamics of network parameters  $\theta$ , which now become stochastic differential equations, to specific features of the resulting posterior distribution  $p^*(\theta|\mathbf{x})$  of parameter vectors  $\theta$  from which the network samples. Thereby, this theory provides a new framework for relating experimentally observed details of local plasticity mechanisms (including their typically stochastic implementation on the molecular scale) to functional consequences of network learning. For example, one gets a theoretically founded framework for relating experimental data on spine motility to experimentally observed network properties, such as sparse connectivity, specific distributions of synaptic weights, and the capability to compensate against perturbations [9].

We demonstrate the resulting new style of modeling network plasticity in three examples. These examples demonstrate how previously mentioned functional demands on network plasticity, such as incorporation of structural rules, automatic avoidance of overfitting, and inherent and immediate compensation for network perturbances, can be accomplished through stochastic local plasticity processes. We focus here on common models for unsupervised learning in networks of neurons: generative models. We first develop the general learning theory for this class of models, and then describe applications to common non-spiking and spiking generative network models. Both structural plasticity (see [10, 11] for reviews) and synaptic plasticity (STDP) are integrated into the resulting theory of network plasticity.

## Results

We present a new theoretical framework for analyzing and understanding local plasticity mechanisms of networks of neurons in the brain as stochastic processes, that generate specific distributions  $p(\theta)$  of network parameters  $\theta$  over which these parameters fluctuate. This framework can be used to analyze and model many types of learning processes. We illustrate it here for the case of unsupervised learning, i.e., learning without a teacher or rewards. Obviously many learning processes in biological organisms are of this nature, especially learning processes in early sensory areas, and in other brain areas that have to provide and maintain on their own an adequate level of functionality, even in the face of internal or external perturbations.

A common framework for modeling unsupervised learning in networks of neurons are generative models, which date back to the 19th century, when Helmholtz proposed that perception could be understood as unconscious inference [12]. Since then the hypothesis of the “generative brain” has been receiving considerable attention, fueling interest in various aspects of the relation between Bayesian inference and the brain [8, 13, 14]. The basic assumption of the “Bayesian brain” theory is that the activity  $\mathbf{z}$  of neuronal networks in the brain can be viewed as an internal

model for hidden variables in the outside world that give rise to sensory experiences  $\mathbf{x}$  (such as the response  $\mathbf{x}$  of auditory sensory neurons to spoken words that are guessed by an internal model  $\mathbf{z}$ ). The internal model  $\mathbf{z}$  is usually assumed to be represented by the activity of neurons in the network, e.g., in terms of the firing rates of neurons, or in terms of spatio-temporal spike patterns. A network  $\mathcal{N}$  of stochastically firing neuron is modeled in this framework by a probability distribution  $p_{\mathcal{N}}(\mathbf{x}, \mathbf{z} | \boldsymbol{\theta})$  that describes the probabilistic relationships between  $N$  input patterns  $\mathbf{x} = \mathbf{x}^1, \dots, \mathbf{x}^N$  and corresponding network responses  $\mathbf{z} = \mathbf{z}^1, \dots, \mathbf{z}^N$ , where  $\boldsymbol{\theta}$  denotes the vector of network parameters that shape this distribution, e.g., via synaptic efficacies and network connectivity. The marginal probability  $p_{\mathcal{N}}(\mathbf{x} | \boldsymbol{\theta}) = \sum_{\mathbf{z}} p_{\mathcal{N}}(\mathbf{x}, \mathbf{z} | \boldsymbol{\theta})$  of the actually occurring inputs  $\mathbf{x} = \mathbf{x}^1, \dots, \mathbf{x}^N$  under the resulting internal model of the neural network  $\mathcal{N}$  with parameters  $\boldsymbol{\theta}$  can then be viewed as a measure for the agreement between this internal model (which carries out “predictive coding” [15]) and its environment (which generates the inputs  $\mathbf{x}$ ).

The goal of network learning is usually described in this probabilistic generative framework as finding parameter values  $\boldsymbol{\theta}^*$  that maximize this agreement, or equivalently the likelihood of the inputs  $\mathbf{x}$  (maximum likelihood learning):

$$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta}} p_{\mathcal{N}}(\mathbf{x} | \boldsymbol{\theta}). \tag{1}$$

Locally optimal parameter solutions are usually determined by gradient ascent on the data likelihood  $p_{\mathcal{N}}(\mathbf{x} | \boldsymbol{\theta})$ .

### Learning a posterior distribution through stochastic synaptic plasticity

In contrast, we assume here that not only a neural network  $\mathcal{N}$ , but also a prior  $p_S(\boldsymbol{\theta})$  for its parameters are given. This prior  $p_S$  can encode both structural constraints (such as sparse connectivity) and structural rules (e.g., a heavy-tailed distribution of synaptic weights). Then the goal of network learning becomes:

$$\text{learn the posterior distribution } p^*(\boldsymbol{\theta} | \mathbf{x}) \text{ defined (up to normalization) by} \tag{2}$$

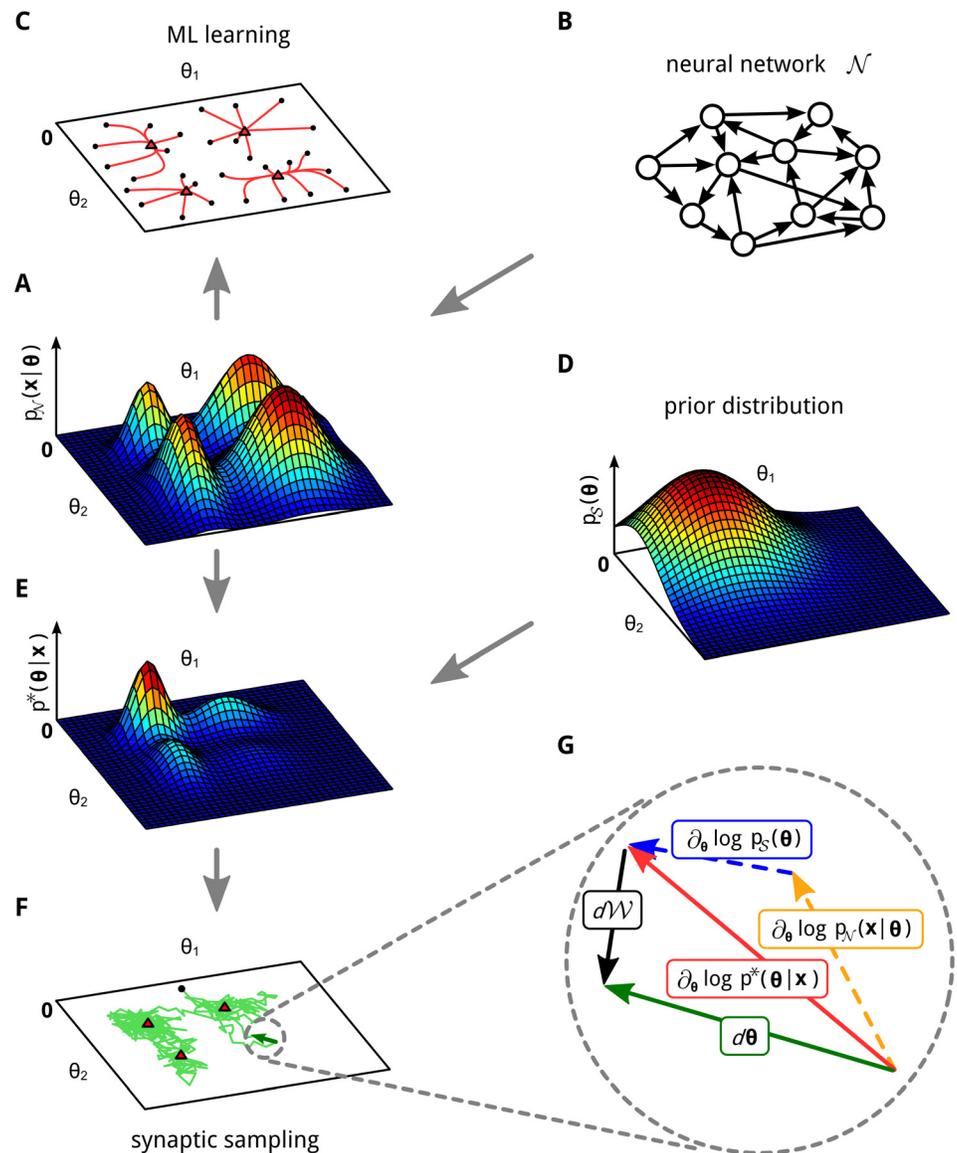
$$p_S(\boldsymbol{\theta}) \cdot p_{\mathcal{N}}(\mathbf{x} | \boldsymbol{\theta}).$$

The patterns  $\mathbf{x} = \mathbf{x}^1, \dots, \mathbf{x}^N$  are assumed here to be regularly reoccurring network inputs.

A key insight (see Fig 1 for an illustration) is that stochastic local plasticity rules for the parameters  $\theta_i$  enable a network to achieve the learning goal Eq (2): The distribution of network parameters  $\boldsymbol{\theta}$  will converge after a while to the posterior distribution  $p^*(\boldsymbol{\theta}) = p^*(\boldsymbol{\theta} | \mathbf{x})$ —and produce samples from it—if each network parameter  $\theta_i$  obeys the dynamics

$$d\theta_i = b \left( \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + \frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(\mathbf{x} | \boldsymbol{\theta}) \right) dt + \sqrt{2b} d\mathcal{W}_i, \tag{3}$$

where the learning rate  $b > 0$  controls the speed of the parameter dynamics. Eq (3) is a stochastic differential equation (see [16]), which differs from commonly considered differential equations through the stochastic term  $d\mathcal{W}_i$  that describes infinitesimal stochastic increments and decrements of a Wiener process  $\mathcal{W}_i$ . A Wiener process is a standard model for Brownian motion in one dimension (more precisely: the limit of a random walk with infinitesimal step size and normally distributed increments  $\mathcal{W}_i^t - \mathcal{W}_i^s \sim \text{NORMAL}(0, t - s)$  between times  $t$  and  $s$ ). Thus in an approximation of Eq (3) for discrete time steps  $\Delta t$  the term  $d\mathcal{W}_i$  can be replaced by Gaussian noise with variance  $\Delta t$  (see Eq (7)). Note that Eq (3) does not have a single solution  $\theta_i(t)$ , but a continuum of stochastic sample paths (see Fig 1F for an example) that each describe one possible time course of the parameter  $\theta_i$ .



**Fig 1. Maximum likelihood (ML) learning vs. synaptic sampling.** **A, B, C:** Illustration of ML learning for two parameters  $\theta = (\theta_1, \theta_2)$  of a neural network  $\mathcal{N}$ . **A:** 3D plot of an example likelihood function. For a fixed set of inputs  $\mathbf{x}$  it assigns a probability density (amplitude on z-axis) to each parameter setting  $\theta$ . **B:** This likelihood function is defined by some underlying neural network  $\mathcal{N}$ . **C:** Multiple trajectories along the gradient of the likelihood function in (A). The parameters are initialized at random initial values (black dots) and then follow the gradient to a local maximum (red triangles). **D:** Example for a prior that prefers small values for  $\theta$ . **E:** The posterior that results as product of the prior (D) and the likelihood (A). **F:** A single trajectory of synaptic sampling from the posterior (E), starting at the black dot. The parameter vector  $\theta$  fluctuates between different solutions, the visited values cluster near local optima (red triangles). **G:** Cartoon illustrating the dynamic forces (plasticity rule Eq (3)) that enable the network to sample from the posterior distribution  $p^*(\theta|\mathbf{x})$  in (E). Magnification of one synaptic sampling step  $d\theta$  of the trajectory in (F) (green). The three forces acting on  $\theta$ : the deterministic drift term (red) is directed to the next local maximum (red triangle), it consists of the first two terms in Eq (3); the stochastic diffusion term  $dW$  (black) has a random direction. See S2 Text for figure details.

doi:10.1371/journal.pcbi.1004485.g001

Rigorous mathematical results based on Fokker-Planck equations (see [Methods](#) and [S1 Text](#) for details) allow us to infer from the stochastic local dynamics of the parameters  $\theta_i$  given by a stochastic differential equation of the form [Eq \(3\)](#) the probability that the parameter vector  $\theta$  can be found after a while in a particular region of the high-dimensional space in which it moves. The key result is that for the case of the stochastic dynamics according to [Eq \(3\)](#) this probability is equal to the posterior  $p^*(\theta|\mathbf{x})$  given by [Eq \(2\)](#). Hence the stochastic dynamics [Eq \(3\)](#) of network parameters  $\theta_i$  enables a network to achieve the learning goal [Eq \(2\)](#): to learn the posterior distribution  $p^*(\theta|\mathbf{x})$ . This posterior distribution is not represented in the network through any explicit neural code, but through its stochastic dynamics, as the unique stationary distribution of a Markov process from which it samples continuously. In particular, if most of the mass of this posterior distribution is concentrated on some low-dimensional manifold, the network parameters  $\theta$  will move most of the time within this low-dimensional manifold. Since this realization of the posterior distribution  $p^*(\theta|\mathbf{x})$  is achieved by sampling from it, we refer to this model defined by [Eq \(3\)](#) (in the case where the parameters  $\theta_i$  represent synaptic parameters) as *synaptic sampling*.

The stochastic term  $d\mathcal{W}_i$  in [Eq \(3\)](#) provides a simple integrative model for a multitude of biological and biochemical stochastic processes that effect the efficacy of a synaptic connection. The mammalian postsynaptic density comprises over 1000 different types of proteins [\[17\]](#). Many of those proteins that effect the amplitude of postsynaptic potentials and synaptic plasticity, for example NMDA receptors, occur in small numbers, and are subject to Brownian motion within the membrane [\[18\]](#). In addition, the turnover of important scaffolding proteins in the postsynaptic density such as PSD-95, which clusters glutamate receptors and is thought to have a substantial impact on synaptic efficacy, is relatively fast, on the time-scale of hours to days, depending on developmental state and environmental condition [\[19\]](#). Also the volume of spines at dendrites, which is assumed to be directly related to synaptic efficacy [\[20, 21\]](#) is reported to fluctuate continuously, even in the absence of synaptic activity [\[2\]](#). Furthermore the stochastically varying internal states of multiple interacting biochemical signaling pathways in the postsynaptic neuron are likely to effect synaptic transmission and plasticity [\[22\]](#).

The contribution of the stochastic term  $d\mathcal{W}_i$  in [Eq \(3\)](#) can be scaled by a temperature parameter  $\sqrt{T}$ , where  $T$  can be any positive number. The resulting stationary distribution of  $\theta$  is proportional to  $p^*(\theta)^{\frac{1}{T}}$ , so that the dynamics of the stochastic process can be described by the energy landscape  $\frac{\log p^*(\theta)}{T}$ . For high values of  $T$  this energy landscape is flattened, i.e., the main modes of  $p^*(\theta)$  become less pronounced. For  $T \rightarrow 0$  the dynamics of  $\theta$  approaches a deterministic process and converges to the next local maximum of  $p^*(\theta)$ . Thus the learning process approximates for low values of  $T$  maximum a posteriori (MAP) inference [\[7\]](#). We propose that this temperature parameter  $T$  is regulated in biological networks of neurons dependent on the developmental state, environment, and behavior of an organism. One can also accommodate a modulation of the dynamics of each individual parameter  $\theta_i$  by a learning rate  $b(\theta_i)$  that depends on its current value (see [Methods](#)).

## Online synaptic sampling

For online learning one assumes that the likelihood  $p_{\mathcal{N}}(\mathbf{x}|\theta) = p_{\mathcal{N}}(\mathbf{x}^1, \dots, \mathbf{x}^N|\theta)$  of the network inputs can be factorized:

$$p_{\mathcal{N}}(\mathbf{x}^1, \dots, \mathbf{x}^N|\theta) = \prod_{n=1}^N p_{\mathcal{N}}(\mathbf{x}^n|\theta), \quad (4)$$

i.e., each network input  $\mathbf{x}^n$  can be explained as being drawn individually from  $p_{\mathcal{N}}(\mathbf{x}^n|\theta)$ , independently from other inputs.

The weight update rule [Eq \(3\)](#) depends on all inputs  $\mathbf{x} = \mathbf{x}^1, \dots, \mathbf{x}^N$ , hence synapses have to keep track of the whole set of all network inputs for the exact dynamics (batch learning). In an online scenario, we assume that only the current network input  $\mathbf{x}^n$  is available for synaptic sampling. One then arrives at the following online-approximation to [Eq \(3\)](#)

$$d\theta_i = b \left( \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + N \frac{\partial}{\partial \theta_i} \log p_N(\mathbf{x}^n | \boldsymbol{\theta}) \right) dt + \sqrt{2b} d\mathcal{W}_i. \quad (5)$$

Note the additional factor  $N$  in the rule. It compensates for the  $N$ -fold summation of the first and last term in [Eq \(5\)](#) when one moves through all  $N$  inputs  $\mathbf{x}^n$ . Although convergence to the correct posterior distribution cannot be guaranteed theoretically for this online rule, we show in [Methods](#) that the rule is a reasonable approximation to the batch-rule [Eq \(3\)](#). Furthermore, all subsequent simulations are based on this online rule, which demonstrates the viability of this approximation.

### Relationship to maximum likelihood learning

Typically, synaptic plasticity in generative network models is modeled as maximum likelihood learning. Time is often discretized into small discrete time steps  $\Delta t$ . For gradient-based approaches the parameter change  $\Delta\theta_i^{ML}$  is then given by the gradient of the log likelihood multiplied with some learning rate  $\eta$ :

$$\Delta\theta_i^{ML} = \eta \frac{\partial}{\partial \theta_i} \log p_N(\mathbf{x}^n | \boldsymbol{\theta}). \quad (6)$$

To compare this maximum likelihood update with synaptic sampling, we consider a version of the parameter dynamics [Eq \(5\)](#) for discrete time (see [Methods](#) for a derivation):

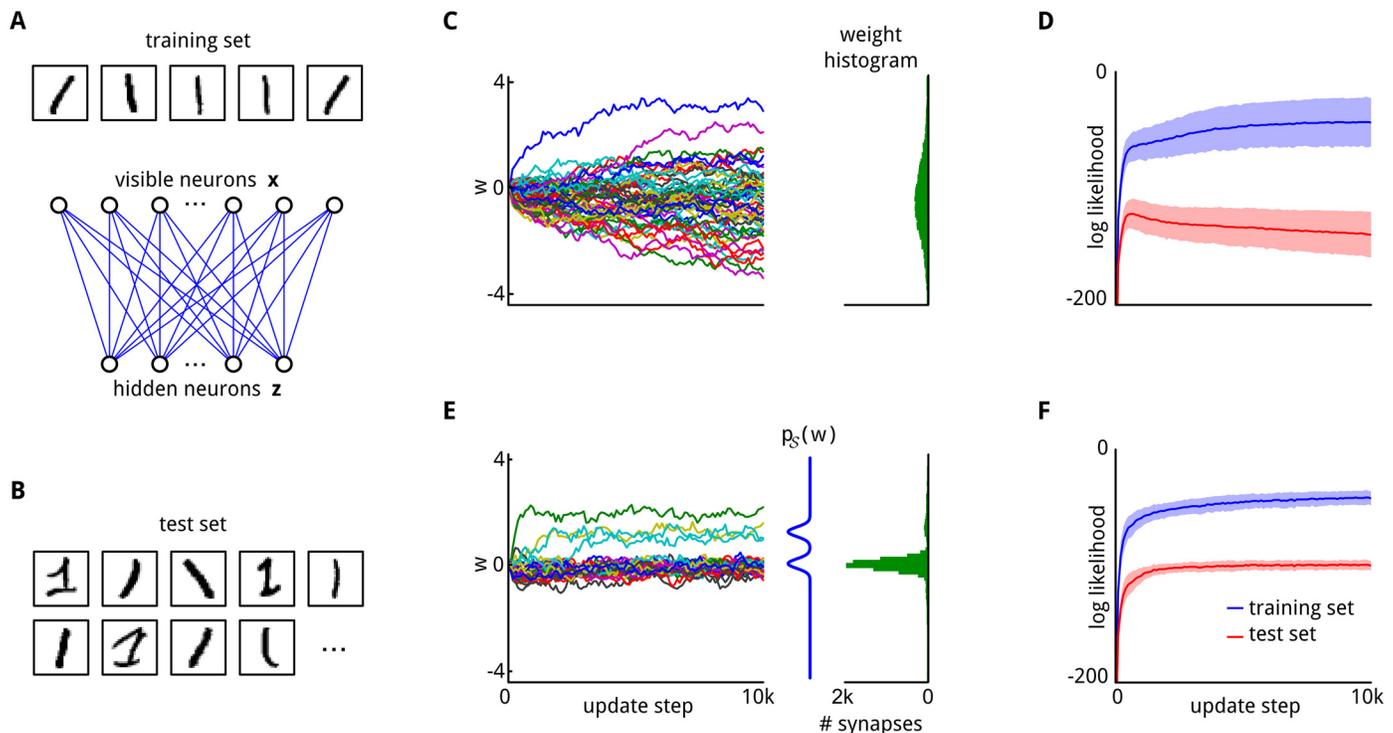
$$\Delta\theta_i = \eta \left( \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + N \frac{\partial}{\partial \theta_i} \log p_N(\mathbf{x}^n | \boldsymbol{\theta}) \right) + \sqrt{2\eta} v_i^t, \quad (7)$$

where the learning rate  $\eta$  is given by  $\eta = b \Delta t$  and  $v_i^t$  denotes Gaussian noise with zero mean and variance 1, drawn independently for each parameter  $\theta_i$  and each update time  $t$ . We see that the maximum likelihood update [Eq \(6\)](#) becomes one term in this online version of synaptic sampling. [Eq \(7\)](#) is a special case of the online Langevin sampler that was introduced in [\[23\]](#).

The first term  $\frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta})$  in [Eq \(7\)](#) arises from the prior  $p_S(\boldsymbol{\theta})$ , and has apparently not been considered in previous rules for synaptic plasticity. An additional novel component is the Gaussian noise term  $v_i^t$  (see also [Fig 1G](#)). It arises because the accumulated impact of the Wiener process  $\mathcal{W}_i$  over a time interval of length  $\Delta t$  is distributed according to a normal distribution with variance  $\Delta t$ . In contrast to traditional maximum likelihood optimization based on additive noise for escaping local optima, this noise term is not scaled down when learning approaches a local optimum. This ongoing noise is essential for enabling the network to sample from the posterior distribution  $p^*(\boldsymbol{\theta})$  via continuously ongoing synaptic plasticity (see [Fig 1F](#)).

### Synaptic sampling improves the generalization capability of a neural network

The previously described theory for learning a posterior distribution over parameters  $\boldsymbol{\theta}$  can be applied to all neural network models  $\mathcal{N}$  where the derivative  $\frac{\partial}{\partial \theta_i} \log p_N(\mathbf{x}^n | \boldsymbol{\theta})$  in [Eq \(5\)](#) can be efficiently estimated. Since this term also has to be estimated for maximum likelihood learning [Eq \(6\)](#), synaptic sampling can basically be applied to all neuron and network models that are amenable to maximum likelihood learning. We illustrate salient new features that result from



**Fig 2. Priors for synaptic weights improve generalization capability.** **A:** The training set, consisting of five samples of a handwritten 1. Below a cartoon illustrating the network architecture of the restricted Boltzmann machine (RBM), composed of a layer of 784 visible neurons  $\mathbf{x}$  and a layer of 9 hidden neurons  $\mathbf{z}$ . **B:** Examples from the test set. It contains many different styles of writing that are not part of the training set. **C:** Evolution of 50 randomly selected synaptic weights throughout learning (on the training set). The weight histogram (right) shows the distribution of synaptic weights at the end of learning. 80 histogram bins were equally spaced between -4 and 4. **D:** Performance of the network in terms of log likelihood on the training set (blue) and on the test set (red) throughout learning. Mean values over 100 trial runs are shown, shaded area indicates std. The performance on the test set initially increases but degrades for prolonged learning. **E:** Evolution of 50 weights for the same network but with a bimodal prior. The prior  $p_S(w)$  is indicated by the blue curve. Most synaptic weights settle in the mode around 0, but a few larger weights also emerge and stabilize in the larger mode. Weight histogram (green) as in (C). **F:** The log likelihood on the test set maintains a constant high value throughout the whole learning session, compare to (D).

doi:10.1371/journal.pcbi.1004485.g002

synaptic sampling (i.e., plasticity rules Eqs (5) or (7)) for some of these models. We begin with the Boltzmann machine [24], one of the oldest generative neural network models. It is currently still extensively investigated in the context of deep learning [25, 26]. We demonstrate in Fig 2D and 2F the improved generalization capability of this model for the learning approach Eq (2) (learning of the posterior), compared with maximum likelihood learning (approach Eq (1)), which had been theoretically predicted by [6] and [7]. But this model for learning the posterior (approach Eq (2)) in Boltzmann machines is now based on local plasticity rules. Note that the Boltzmann machine with synaptic sampling samples simultaneously on two different time scales: In addition to sampling for given parameters  $\theta$  from likely network states in the usual manner, it now samples simultaneously on a slower time scale according to Eq (7) from the posterior of network parameters  $\theta$ .

A Boltzmann machine employs extremely simple non-spiking neuron models with binary outputs. Neuron  $y_i$  outputs 1 with probability  $\sigma(\sum_j w_{ij} y_j + b_i)$ , else 0, where  $\sigma$  is the logistic sigmoid  $\sigma(u) = \frac{1}{1+e^{-u}}$ , with synaptic weights  $w_{ij}$  and bias parameters  $b_i$ . Synaptic connections in a Boltzmann machine are bidirectional, with symmetric weights ( $w_{ij} = w_{ji}$ ). The parameters  $\theta$  for the Boltzmann machine consist of all weights  $w_{ij}$  and biases  $b_i$  in the network. For the special case of a restricted Boltzmann machine (RBM), maximum likelihood learning of these parameters can be done efficiently [27], and therefore RBM's are typically used for deep learning. An

RBM has a layered structure with one layer of visible neurons  $\mathbf{x}$  and a second layer of hidden neurons  $\mathbf{z}$ . Synaptic connections are formed only between neurons on different layers (Fig 2A). The maximum likelihood gradients  $\Delta w_{ij}^{ML} = \frac{\partial}{\partial w_{ij}} \log p_{\mathcal{N}}(\mathbf{x}|\boldsymbol{\theta})$  and  $\Delta b_i^{ML} = \frac{\partial}{\partial b_i} \log p_{\mathcal{N}}(\mathbf{x}|\boldsymbol{\theta})$  can be efficiently approximated for this model, for example

$$\frac{\partial}{\partial w_{ij}} \log p_{\mathcal{N}}(\mathbf{x}^n|\boldsymbol{\theta}) \approx z_i^n x_j^n - \hat{z}_i^n \hat{x}_j^n, \tag{8}$$

where  $x_j^n$  is the output of input neuron  $j$  while input  $\mathbf{x}^n$  is presented, and  $\hat{x}_j^n$  its output during a subsequent phase of spontaneous activity (“reconstruction phase”); analogously for the hidden neuron  $z_j$  (see Methods and S3 Text).

We integrated this maximum likelihood estimate Eq (8) into the synaptic sampling rule Eq (7) in order to test whether a suitable prior  $p_S(\mathbf{w})$  for the weights improves the generalization capability of the network. The network received as input just five samples  $\mathbf{x}^1, \dots, \mathbf{x}^5$  of a handwritten Arabic number 1 from the MNIST dataset (the training set, shown in Fig 2A) that were repeatedly presented. Each pixel of the digit images was represented by one neuron in the visible layer (which consisted of 784 neurons). We selected a second set of 100 samples of the handwritten digit 1 from the MNIST dataset as test set (Fig 2B). These samples include completely different styles of writing that were not present in the training set. After allowing the network to learn the five input samples from Fig 2A for various numbers of update steps (horizontal axis of Fig 2D and 2F), we evaluated the learned internal model of this network  $\mathcal{N}$  for the digit 1 by measuring the average log-likelihood  $\log p_{\mathcal{N}}(\mathbf{x}|\boldsymbol{\theta})$  for the test data. The result is indicated in Fig 2D and 2F for the training samples by the blue curves, and for the new test examples, that were never shown while synaptic plasticity was active, by the red curves.

First, a uniform prior over the synaptic weights was used (Fig 2C), which corresponds to the common maximum likelihood learning paradigm Eq (8). The performance on the test set (shown on vertical axis) initially increases but degrades for prolonged exposure to the training set (length of that prior exposure shown on horizontal axis). This effect is known as overfitting [6, 7]. It can be reduced by choosing a suitable prior  $p_S(\boldsymbol{\theta})$  in the synaptic sampling rule Eq (7). The choice for the prior distribution is best if it matches the statistics of the training samples [6], which has in this case two modes (resulting from black and white pixels). The presence of this prior in the learning rule maintains good generalization capability for test samples even after prolonged exposure to the training set (red curve in Fig 2F).

The improved generalization capability of the network is a result of the prior distribution. It is well known that the prior in Bayesian inference allows to effectively prevent overfitting by making solutions that use fewer or smaller parameters more likely. Similar results would therefore emerge in any other implementation of Bayesian learning in neural networks. A thorough discussion on this topic which is known as *Bayesian regularization* can be found in [6, 7].

As a consequence, the choice of the prior distribution can have a significant impact on the learning result. In S3 Text we compared a set of different priors and demonstrate this effect more systematically. There it can also be seen that if the choice of the prior is bad, the learning performance can even get worse than in the case without a prior.

### Spine motility as synaptic sampling

In the following sections we apply our synaptic sampling framework to networks of spiking neurons and biological models for network plasticity. The number and volume of spines for a synaptic connection is thought to be directly related to its synaptic weight [28]. Experimental studies have provided a wealth of information about the stochastic dynamics of dendritic spines (see e.g. [1, 28–32]). They demonstrate that the volume of a substantial fraction of

dendritic spines varies continuously over time, and that all the time new spines and synaptic connections are formed and existing ones are eliminated. We show that these experimental data on spine motility can be understood as special cases of synaptic sampling. The synaptic sampling framework is however very general, and many different models for spine motility can be derived from it as special cases. We demonstrate this here for one simple model, induced by the following assumptions:

1. We restrict ourselves to plasticity of excitatory synapses, although the framework is general enough to apply to inhibitory synapses as well.
2. In accordance with experimental studies [28], we require that spine sizes have a multiplicative dynamics, i.e., that the amount of change within some given time window is proportional to the current size of the spine.
3. We assume here for simplicity that a synaptic connection between two neurons is realized by a single spine and that there is a single parameter  $\theta_i$  for each potential synaptic connection  $i$ .

The last requirement can be met by encoding the state of the synapse in an abstract form, that represents synaptic connectivity and synaptic plasticity in a single parameter  $\theta_i$ . We define that negative values of  $\theta_i$  represent a current disconnection and positive values represent a functional synaptic connection. The distance of the current value of  $\theta_i$  from zero indicates how likely it is that the synapse will soon reconnect (for negative values) or withdraw (for positive values), see Fig 3A. In addition the synaptic parameter  $\theta_i$  encodes for positive values the synaptic efficacy  $w_i$ , i.e., the resulting EPSP amplitudes, by a simple mapping  $w_i = f(\theta_i)$ .

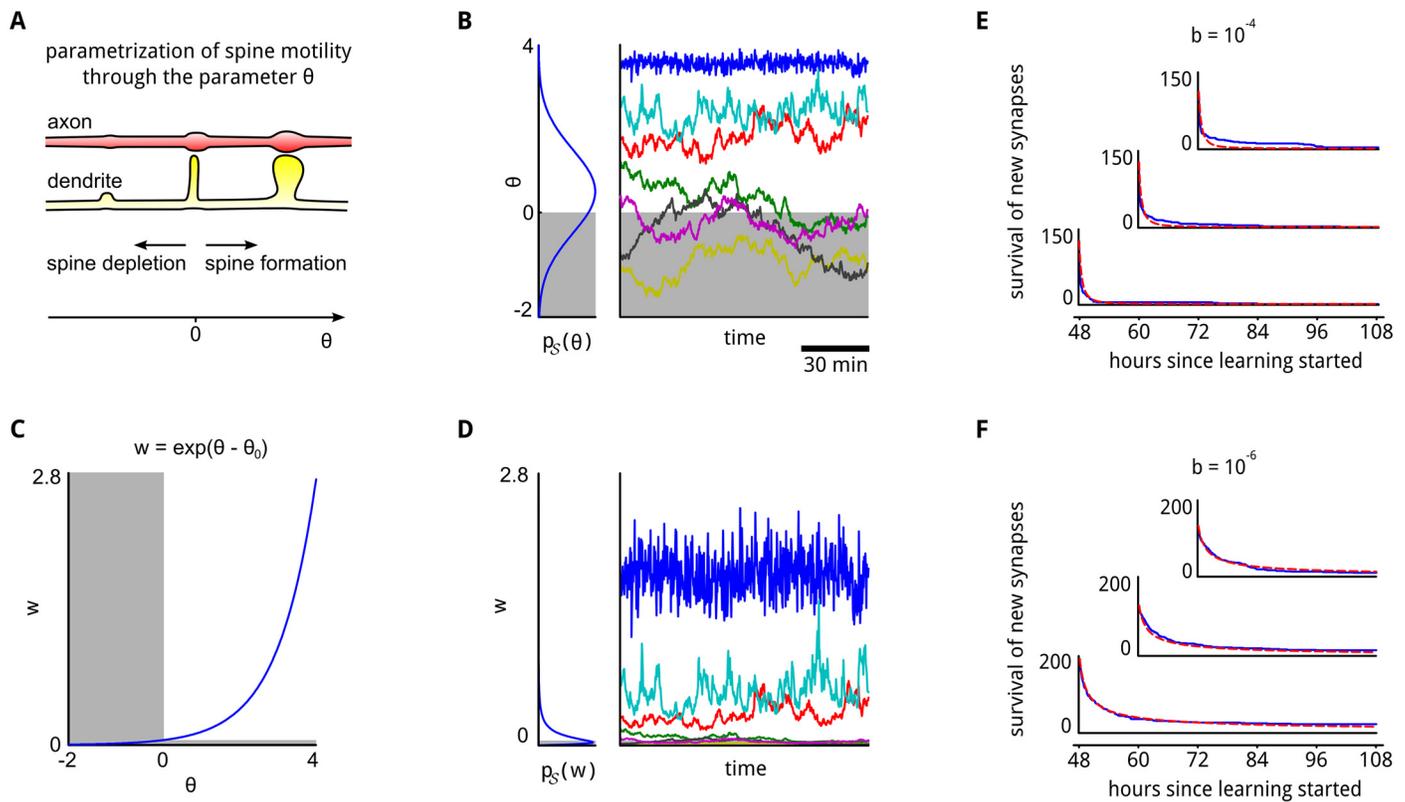
A large class of mapping functions  $f$  is supported by our theory (see S4 Text for details). The second assumption which requires multiplicative synaptic dynamics supports an exponential function  $f$  in our model, in accordance with previous models of spine motility [28]. Thus, we assume in the following that the efficacy  $w_i$  of synapse  $i$  is given by

$$w_i = \exp(\theta_i - \theta_0), \tag{9}$$

see Fig 3C. Note that for a large enough offset  $\theta_0$ , negative parameter values  $\theta_i$  (which model a non-functional synaptic connection) are automatically mapped onto a tiny region close to zero in the  $w$ -space, so that retracted spines have essentially zero synaptic efficacy. The general rule for online synaptic sampling Eq (5) for the exponential mapping Eq (9) can be written as (see S4 Text)

$$d\theta_i = b \left( \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + N w_i \frac{\partial}{\partial w_i} \log p_N(\mathbf{x}^n | \mathbf{w}) \right) dt + \sqrt{2b} d\mathcal{W}_i. \tag{10}$$

In Eq (10) the multiplicative synaptic dynamics becomes explicit. The gradient  $\frac{\partial}{\partial w_i} \log p_N(\mathbf{x}^n | \mathbf{w})$ , i.e., the activity-dependent contribution to synaptic plasticity, is weighted by  $w_i$ . Hence, for negative values of  $\theta_i$  (non-functional synaptic connection), the activities of the pre- and post-synaptic neurons have negligible impact on the dynamics of the synapse. Assuming a large enough  $\theta_0$ , retracted synapses therefore evolve solely according to the prior  $p_S(\boldsymbol{\theta})$  and the random fluctuations  $d\mathcal{W}_i$ . For large values of  $\theta_i$  the opposite is the case. The influence of the prior  $\frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta})$  and the Wiener process  $d\mathcal{W}_i$  become negligible, and the dynamics is dominated by the activity-dependent likelihood term. Large synapses can therefore become quite stable if the presynaptic activity is strong and reliable (see Fig 3B). Through the use of parameters  $\boldsymbol{\theta}$  which determine both synaptic connectivity and synaptic efficacies, the synaptic sampling framework provides a unified model for structural and synaptic plasticity. The prior



**Fig 3. Integration of spine motility into the synaptic sampling model.** **A:** Illustration of the parametrization of spine motility. Values  $\theta > 0$  indicate a functional synaptic connection. **B:** A Gaussian prior  $p_S(\theta)$ , and a few stochastic sample trajectories of  $\theta$  according to the synaptic sampling rule Eq (10). Negative values of  $\theta$  (gray area) are interpreted as non-functional connections. Some stable synaptic connections emerge (traces in the upper half), whereas other synaptic connections come and go (traces in lower half). All traces, as well as survival statistics shown in (E, F), are taken from the network simulation described in detail in the next section and S5 Text. **C:** The exponential function maps synapse parameters  $\theta$  to synaptic efficacies  $w$ . Negative values of  $\theta$ , corresponding to (retracted) spines are mapped to a tiny region close to zero in the  $w$ -space. **D:** The Gaussian prior in the  $\theta$ -space translates to a log-normal distribution in the  $w$ -space. The traces from (B) are shown in the right panel transformed into the  $w$ -space. Only persistent synaptic connections contribute substantial synaptic efficacies. **E, F:** The emergent survival statistics of newly formed synaptic connections, (i.e., formed during the preceding 12 hours) evaluated at three different start times throughout learning (blue traces, axes are aligned with start times of the analyses). The survival statistics exhibit in our synaptic sampling model a power-law behavior (red curves, see S5 Text). The time-scale (and exponent of the power-law) depends on the learning rate  $b$  in Eq (10), and can assume any value in our quite general model (shown is  $b = 10^{-4}$  in (E) and  $b = 10^{-6}$  in (F)).

doi:10.1371/journal.pcbi.1004485.g003

distribution can have significant impact on the spine motility, encouraging for example sparser or denser synaptic connectivity. If the activity-dependent second term in Eq (10), that tries to maximize the likelihood, is small (e.g., because  $\theta_i$  is small or parameters are near a mode of the likelihood) then Eq (10) implements an Ornstein Uhlenbeck process. This prediction of our model is consistent with a previous analysis which showed that an Ornstein Uhlenbeck process is a viable model for synaptic spine motility [28].

The weight dynamics that emerges through the stochastic process Eq (10) is illustrated in the right panel of Fig 3D. A Gaussian parameter prior  $p_S(\theta_i)$  results in a log-normal prior  $p_S(w_i)$  in a corresponding stochastic differential equation for synaptic efficacies  $w_i$ ; (see S4 Text for details).

The last term (noise term) in our synaptic sampling rule Eq (10) predicts that eliminated connections spontaneously regrow at irregular intervals. In this way they can test whether they can contribute to explaining the input. If they cannot contribute, they disappear again. The resulting power-law behavior of the survival of newly formed synaptic connections (Fig 3E and 3F) matches corresponding new experimental data [32] and is qualitatively similar to earlier

experimental results which revealed a quick decay of transient dendritic spines [30, 31, 33]. Functional consequences of this structural plasticity are explored in the following sections.

## Fast adaptation of synaptic connections and weights to a changing input statistics

We will explore in this and the next section implications of the synaptic sampling rule Eq (10) for network plasticity in simple generative spike-based neural network models.

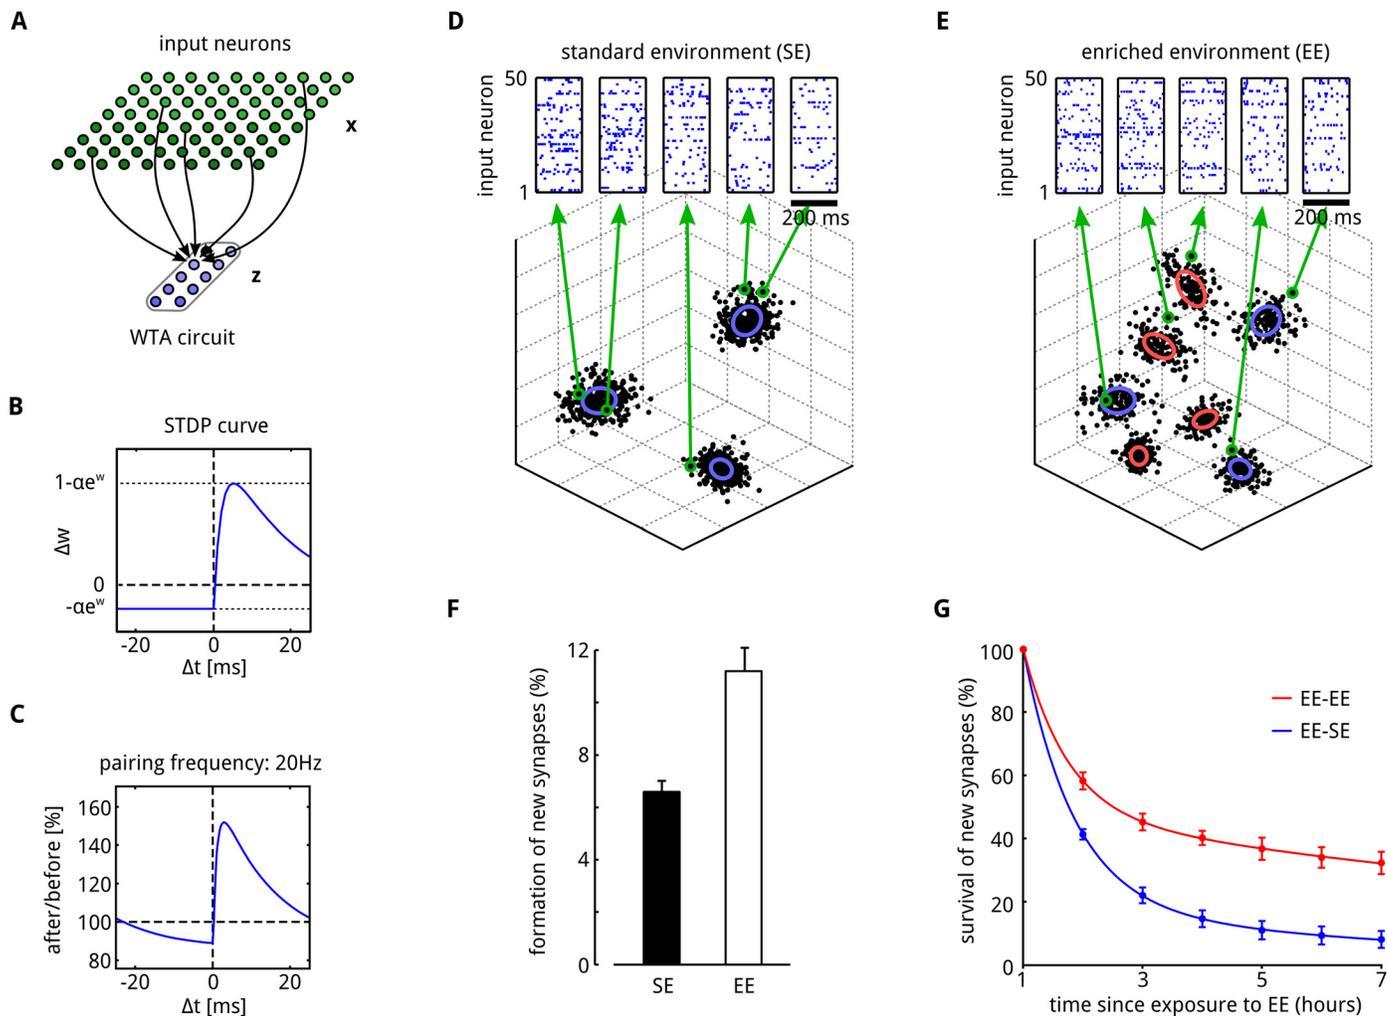
The main types of spike-based generative neural network models that have been proposed are [34–37]. We focus here on the type of models introduced by [36–38], since these models allow an easy estimation of the likelihood gradient (the second term in Eq (10)) and can relate this likelihood term to STDP. Since these spike-based neural network models have non-symmetric synaptic connections (that model chemical synapses between pyramidal cells in the cortex), they do not allow to regenerate inputs  $x$  from internal responses  $z$  by running the network backwards (like in a Boltzmann machine). Rather they are *implicit* generative models, where synaptic weights from inputs to hidden neurons are interpreted as implicit models for presynaptic activity, given that the postsynaptic neuron fires.

We focus in this section on a simple model for an ubiquitous cortical microcircuit motif: an ensemble of pyramidal cells with lateral inhibition, often referred to as Winner-Take-All (WTA) circuit. It has been proposed that this microcircuit motif provides for computational analysis an important bridge between single neurons and larger brain systems [39]. We employ a simple form of divisive normalization (as proposed by [39]; see [Methods](#)) to model lateral inhibition, thereby arriving at a theoretically tractable version of this microcircuit motif that allows us to compute the maximum likelihood term (second term in Eq (10)) in the synaptic sampling rule. We assumed Gaussian prior distributions  $p_S(\theta_i)$ , with mean  $\mu$  and variance  $\sigma^2$  over the synaptic parameters  $\theta_i$  (as in [Fig 3B](#)). Then the synaptic sampling rule Eq (10) yields for this model

$$d\theta_i = b \left( \frac{1}{\sigma^2} (\mu - \theta_i) + N w_i S(t) (x_i(t) - \alpha e^{w_i}) \right) dt + \sqrt{2b} dW_i, \quad (11)$$

where  $S(t)$  denotes the spike train of the postsynaptic neuron and  $x_i(t)$  denotes the weight-normalized value of the sum of EPSPs from presynaptic neuron  $i$  at time  $t$  (i.e., the summed EPSPs that would arise for weight  $w_i = 1$ ; see [Methods](#) for details).  $\alpha$  is a parameter that scales the impact of synaptic plasticity depending on the current synaptic efficacy. The resulting activity-dependent component  $S(t)(x_i(t) - \alpha e^{w_i})$  of the likelihood term is a simplified version of the standard STDP learning rule ([Fig 4B and 4C](#)), like in [36, 40]. Synaptic plasticity (STDP) for connections from input neurons to pyramidal cells in the WTA circuit can be understood from the generative aspect as fitting a mixture of Poisson (or other exponential family) distributions to high-dimensional spike inputs [36, 37]. The factor  $w_i = \exp(\theta_i - \theta_0)$  had been discussed in [36], because it is compatible with the underlying generative model, but provides in addition a better fit to the experimental data of [41]. We examine in this section emergent properties of network plasticity in this simple spike-based neural network under the synaptic sampling rule Eq (11).

It is well documented that cortical dendritic spines are transient and that spine turnover is enhanced by novel experience and training [33, 42, 43]. For example, enhanced spine formation as a consequence of sensory enrichment was found in mouse somatosensory cortex [33]. In this study the animals were exposed to a new sensory environment by adding additional objects to their home cage. This sensory enrichment resulted in a rapid increase in the formation of new spines. If the exposure to the enriched environment was only brief, the newly formed spines quickly decayed.



**Fig 4. Adaptation of synaptic connections to changing input statistics through synaptic sampling.** **A:** Illustration of the network architecture. A WTA circuit consisting of ten neurons **z** receives afferent stimuli from input neurons **x** (few connections shown for a single neuron in **z**). **B:** The STDP learning curve that arises from the likelihood term in Eq (11). **C:** Measured STDP curve that results from a related STDP rule for a moderate pairing frequency of 20 Hz, as in [41]. (Figure adapted from [36]). **D, E:** Each sensory experience was modeled by 200 ms long spiking activity of 1000 input neurons, that covered some 3D data space with Gaussian tuning curves (the results do not depend on the finite dimension of the data space, we chose 3 dimension for easier visualization). Insets show the firing activity of randomly chosen 50 of the 1000 input neurons for the sample data points marked by green circles. Objects in the environment were represented by Gaussian clusters (ellipses) in this finite dimensional data space. **F:** During learning phase 1 (3 hours) only samples from SE were presented to the network, in phase 2 (which lasted 1 hour) samples from EE. Shortly after the transition from SE to EE the number of newly formed synaptic connections significantly increases (compare to Fig. 1h in [33]). **G:** Comparison of the survival of synapses for a network with persistent exposure to EE (EE-EE condition) and a network that was returned to SE (EE-SE condition). Newly formed synaptic connections are transient and quickly decay after formation. A significantly larger fraction of synapses persists if the network continuously receives EE inputs (compare to Fig. 2c in [33]). The dots show means of measurements taken every 30 minutes, the lines represent two-term exponential fits ( $r^2 = 1$ ). The results in (F, G) show means over 5 trial runs. Error bars indicate STD.

doi:10.1371/journal.pcbi.1004485.g004

We wondered whether these experimentally observed effects also emerge in our synaptic sampling model. As in [33] we exposed the network to different sensory environments to study these effects. Sensory experiences typically involve several processing steps and interactions between multiple brain systems, and precise knowledge about their cortical representation is still missing. Therefore we used here a simple symbolic representation of the sensory environment. We represented each sensory experience by a point in some finite dimensional space which is covered by the tuning curves of a large number of input neurons. Their spike output was then communicated to the WTA circuit in the form of 200 ms-long spike patterns of the 1000 input

neurons (see Fig 4D and 4E and Methods for details). Independently drawn sensory experiences were presented sequentially and synaptic sampling according to Eq (11) was applied continuously to all synapses from the 1000 input neurons to the ten neurons in the WTA circuit.

Each environment was represented as a mixture of Gaussians (clusters) of points in the finite-dimensional sensory space. Each cluster could represent for example different sensory experiences with some object in the environment. Consequently we modelled an enriched environment (EE) simply by adding a few new clusters to the standard environment (SE). In phase 1 the network was exposed to an environment with 3 clusters (standard environment (SE), see Fig 4D). After 3 hours the network input was enriched by adding 4 additional clusters (enriched environment (EE), see Fig 4E). We found that exposure to EE significantly increased the rate of new synapse formation as in the experimental result of [33] (Fig 4F).

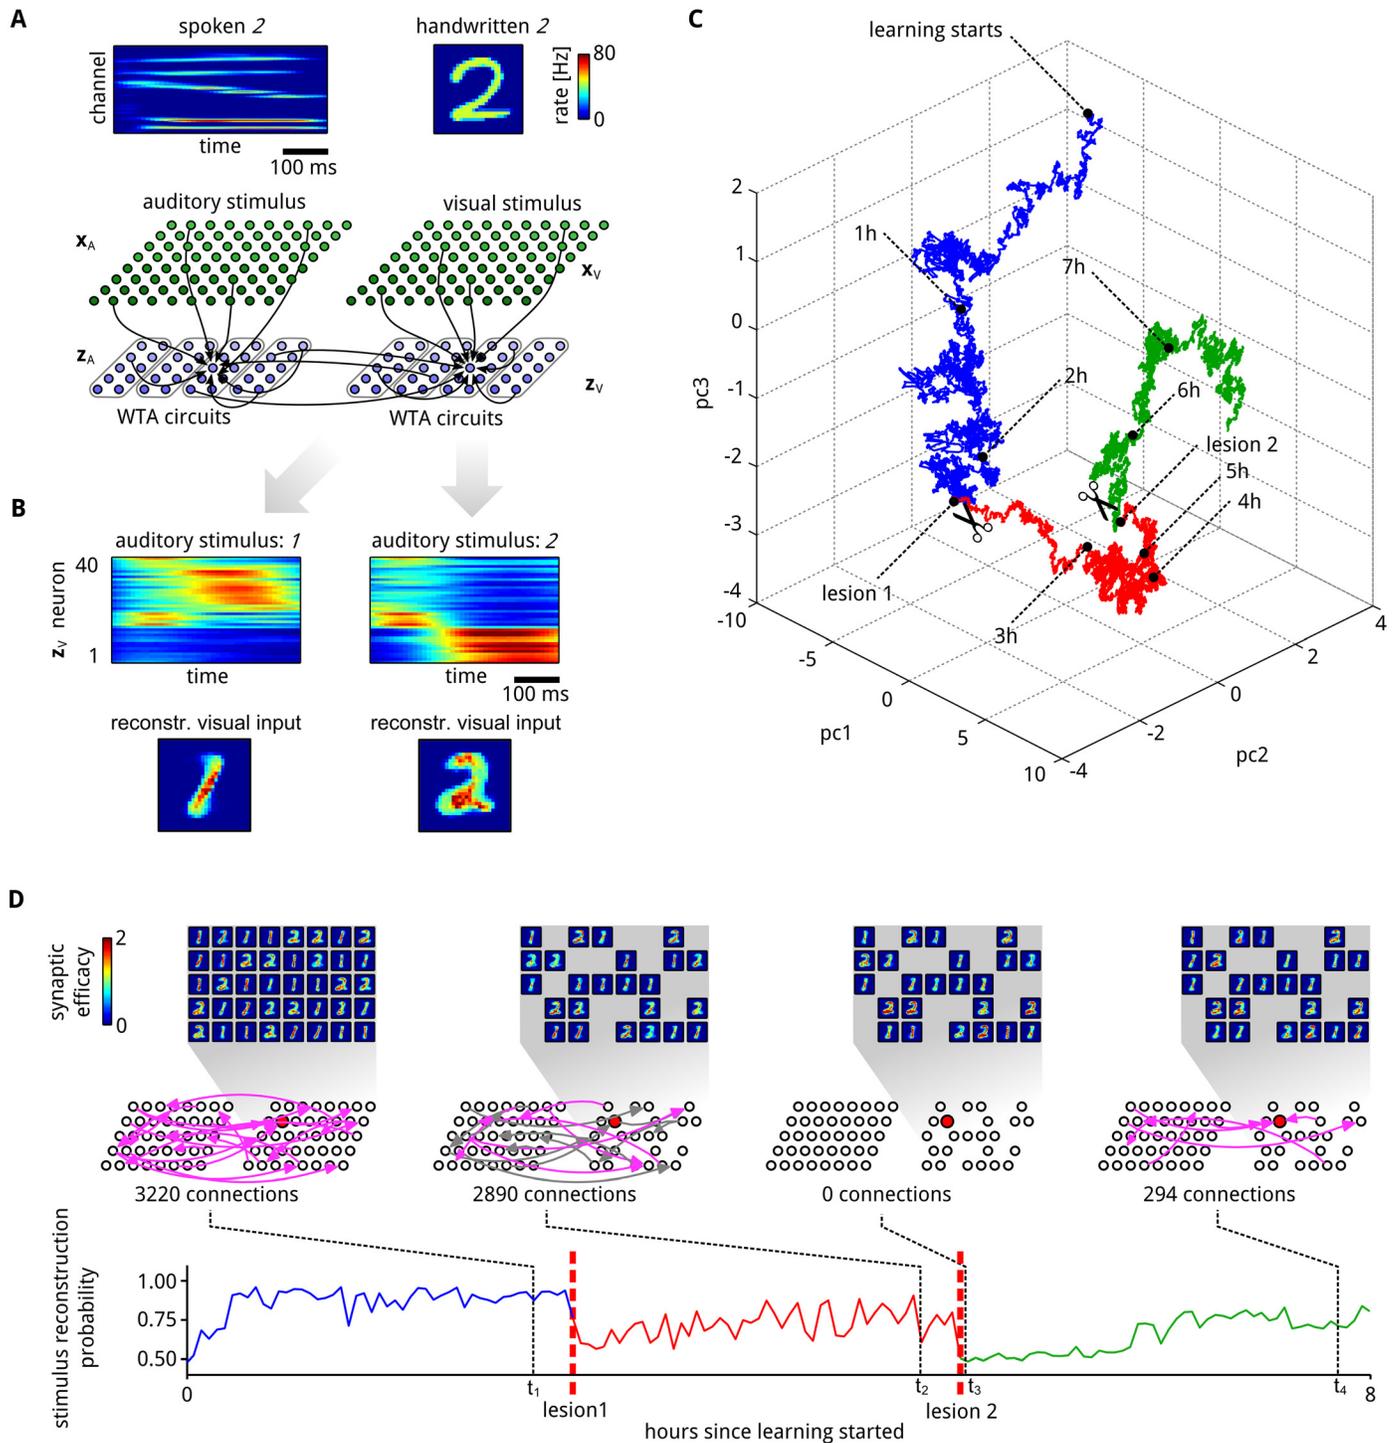
Most of the newly formed synapses decayed within a few hours after return to the standard environment (EE-SE situation, see Fig 4G). In this case only about about 8% become stable. A fraction of about 30% becomes stable when the enriched environment was maintained (EE-EE situation). These results qualitatively reproduce the findings from mouse barrel cortex (compare Figures 1h and 2c in [33]). Note that we used here relatively large update rates  $b$  to keep simulation times in a feasible range, which results in spine dynamics on the time scale of hours instead of days as in biological synapses [33].

### Inherent network compensation capability through synaptic sampling

Numerous experimental data show that the same function of a neural circuit is achieved in different individuals with drastically different parameters, and also that a single organism can compensate for disturbances by moving to a new parameter vector [9, 44–47]. These results suggest that there exists some low-dimensional submanifold of values for the high-dimensional parameter vector  $\theta$  of a biological neural network that all provide stable network function (degeneracy). We propose that the previously discussed posterior distribution of network parameters  $\theta$  provides a mathematical model for such low-dimensional submanifold. Furthermore we propose that the underlying continuous stochastic fluctuation  $dW$  provides a driving force that automatically moves network parameters (with high probability) to a functionally more attractive regime when the current solution performs worse because of perturbations, such as lesions of neurons or network connections. This compensation capability is not an add-on to the synaptic sampling model, but an inherent feature of its organization.

We demonstrate this inherent compensation capability in Fig 5 for a generative spiking neural network with synaptic parameters  $\theta$  that regulate simultaneously structural plasticity and synaptic plasticity (dynamics of weights) as in Figs 3 and 4. The prior  $p_S(\theta)$  for these parameters is here the same as in the preceding section (see Fig 4G on the left). But in contrast to the previous section we consider here a network that allows us to study the self-organization of connections *between* hidden neurons. The network consists of eight WTA-circuits, but in contrast to Fig 4 we allow here arbitrary excitatory synaptic connections between neurons within the same or different ones of these WTA circuits. This network models multi-modal sensory integration and association in a simplified manner. Two populations of “auditory” and “visual” input neurons  $\mathbf{x}_A$  and  $\mathbf{x}_V$  project onto corresponding populations  $\mathbf{z}_A$  and  $\mathbf{z}_V$  of hidden neurons (each consisting of one half of the WTA circuits, see lower panel of Fig 5A). Only a fraction of the potential synaptic connections became functional (see Fig. S2A in S6 Text) through the synaptic sampling rule Eq (11) that integrates structural and synaptic plasticity. Synaptic weights and connections were not forced to be symmetric or bidirectional.

As in the previous demonstrations we do not use external rewards or teacher-inputs for guiding network plasticity. Rather, we allow the model to discover on its own regularities in its



**Fig 5. Inherent compensation for network perturbations.** **A:** A spike-based generative neural network (illustrated at the bottom) received simultaneously spoken and handwritten representations of the same digit (and for tests only spoken digits, see (B)). Stimulus examples for spoken and written digit 2 are shown at the top. These inputs are presented to the network through corresponding firing rates of “auditory” ( $x_A$ ) and “visual” ( $x_V$ ) input neurons. Two populations  $z_A$  and  $z_V$  of 40 neurons, each consisting of four WTA circuits like in Fig 4, receive exclusively auditory or visual inputs. In addition, arbitrary lateral excitatory connections between these “hidden” neurons are allowed. **B:** Assemblies of hidden neurons emerge that encode the presented digit (1 or 2). Top panel shows PETH of all neurons from  $z_V$  for stimulus 1 (left) and 2 (right) after learning, when only an auditory stimulus is presented. Neurons are sorted by the time of their highest average firing. Although only auditory stimuli are presented, it is possible to reconstruct an internally generated “guessed” visual stimulus that represents the same digit (bottom). **C:** First three PCA components of the temporal evolution of a subset  $\theta'$  of network parameters  $\theta$  (see text). Two major lesions were applied to the network. In the first lesion (transition to red) all neurons that significantly encode stimulus 2 were removed from

the population  $\mathbf{z}_V$ . In the second lesion (transition to green) all currently existing synaptic connections between neuron in  $\mathbf{z}_A$  and  $\mathbf{z}_V$  were removed, and not allowed to regrow. After each lesion the network parameters  $\theta$  migrate to a new manifold. **D**: The generative reconstruction performance of the “visual” neurons  $\mathbf{z}_V$  for the test case when only an auditory stimulus is presented was tracked throughout the whole learning session, including lesions 1 and 2 (bottom panel). After each lesion the performance strongly degrades, but reliably recovers. Insets show at the top the synaptic weights of neurons in  $\mathbf{z}_V$  at 4 time points  $t_1, \dots, t_4$ , projected back into the input space like in Fig 4E. Network diagrams in the middle show ongoing network rewiring for synaptic connections between the hidden neurons  $\mathbf{z}_A$  and  $\mathbf{z}_V$ . Each arrow indicates a functional connection between two neurons. To keep the figure uncluttered only subsets of synapses are shown (1% randomly drawn from the total set of possible lateral connections). Connections at time  $t_2$  that were already functional at time  $t_1$  are plotted in gray. The neuron whose parameter vector  $\theta$  is tracked in (C) is highlighted in red. The text under the network diagrams shows the total number of functional connections between hidden neurons at the time point.

doi:10.1371/journal.pcbi.1004485.g005

network inputs. The “auditory” hidden neurons  $\mathbf{z}_A$  on the left in Fig 5A received temporal spike patterns from the auditory input neurons  $\mathbf{x}_A$  that were generated from spoken utterings of the digit 1 and 2 (which lasted between 320 ms and 520 ms). Simultaneously we presented to the “visual” hidden neurons  $\mathbf{z}_V$  on the right for the same time period a (symbolic) visual representation of the same digit (randomly drawn from the MNIST database like in Fig. 2).

The emergent associations between the two populations  $\mathbf{z}_A$  and  $\mathbf{z}_V$  of hidden neurons were tested by presenting auditory input only and observing the activity of the “visual” hidden neurons  $\mathbf{z}_V$ . Fig 5B shows the emergent activity of the neurons  $\mathbf{z}_V$  when only the auditory stimulus was presented (visual input neurons  $\mathbf{x}_V$  remained silent). The generative aspect of the network can be demonstrated by reconstructing for this case the visual stimulus from the activity of the “visual” hidden neurons  $\mathbf{z}_V$ . Fig 5B shows reconstructed visual stimuli from a single run where only the auditory stimuli  $\mathbf{x}_A$  for digits 1 (left) and 2 (right) were presented to the network. Digit images were reconstructed by multiplying the synaptic efficacies of synapses from neurons in  $\mathbf{x}_V$  to neurons in  $\mathbf{z}_V$  (which did not receive any input from  $\mathbf{x}_V$  in this experiment) with the instantaneous firing rates of the corresponding  $\mathbf{z}_V$ -neurons.

Interestingly we found that synaptic sampling significantly outperforms the pure deterministic STDP updates introduced in [38], which do not impose a prior distribution over synaptic parameters. The structural prior that favors solutions with only a small number of large synaptic weights seems to be beneficial for this task as it allows to learn few but pronounced associations between the neurons (see S6 Text).

In order to investigate the inherent compensation capability of synaptic sampling, we applied two lesions to the network within a learning session of 8 hours. In the first lesion all neurons (16 out of 40) that became tuned for digit 2 in the preceding learning (see Fig 5D and S6 Text) were removed. The lesion significantly impaired the performance of the network in stimulus reconstruction, but it was able to recover from the lesion after about one hour of continuing network plasticity according to Eq (11) (Fig 5D). The reconstruction performance of the network was measured here continuously through the capability of a linear readout neuron from the visual ensemble to classify the current auditory stimulus (1 or 2).

In the second lesion all synaptic connections between hidden neurons that were present after recovery from the first lesion were removed and not allowed to regrow (2936 synapses in total). After about two hours of continuing synaptic sampling 294 new synaptic connections between hidden neurons emerged. These made it again possible to infer the auditory stimulus from the activity of the remaining 24 hidden neurons in the population  $\mathbf{z}_V$  (in the absence of any input from the population  $\mathbf{x}_V$ ), at about 75% of the performance level before the second lesion (see bottom panel of Fig 5D).

In order to illustrate the ongoing network reconfiguration we track in Fig 5C the temporal evolution of a subset  $\theta$  of network parameters (35 parameters  $\theta_i$  associated with the potential synaptic connections of the neuron marked in red in the middle of Fig 5D from or to other hidden neurons, excluding those that were removed at lesion 2 and not allowed to regrow). The

first three PCA components of this 35-dimensional parameter vector are shown. The vector  $\theta$  fluctuates first within one region of the parameter space while probing different solutions to the learning problem, e.g., high probability regions of the posterior distribution (blue trace). Each lesion induced a fast switch to a different region (red, green), accompanied by a recovery of the visual stimulus reconstruction performance (see Fig 5D).

The random fluctuations were found to be an integral part of the fast recovery from lesions. In S6 Text we analyzed the impact of the diffusion term in Eq (11) on the learning speed. We found that it acts as a temperature parameter that allows to scale the speed of exploration in the parameter space (see also the Methods for a detailed derivation).

Altogether this experiment showed that continuously ongoing synaptic sampling maintains stable network function also in a more complex network architecture. Another consequence of synaptic sampling was that the neural codes (assembly sequences) that emerged for the two digit classes within the hidden neurons  $z_A$  and  $z_V$  (see Fig. S2B in S6 Text) drifted over larger periods of time (also in the absence of lesions), similarly as observed for place cells in [48] and for tuning curves of motor cortex neurons in [4].

## Discussion

We have shown that stochasticity may provide an important function for network plasticity. It enables networks to sample parameters from some low-dimensional manifold in a high-dimensional parameter space that represents attractive combinations of structural constraints and rules (such as sparse connectivity and heavy-tailed distributions of synaptic weights) and a good fit to empirical evidence (e.g., sensory inputs). We have developed a normative model for this new learning perspective, where the traditional gold standard of maximum likelihood optimization is replaced by theoretically optimal sampling from a posterior distribution of parameter settings, where regions of high probability provide a theoretically optimal model for the low-dimensional manifold from which parameter settings should be sampled. The postulate that networks should learn such posterior distributions of parameters, rather than maximum likelihood values, had been proposed already for quite some while for artificial neural networks [6, 7], since such organization of learning promises better generalization capability to new examples. The open problem how such posterior distributions could be learned by networks of neurons in the brain, in a way that is consistent with experimental data, has been highlighted in [8] as a key challenge for computational neuroscience. We have presented here such a model, whose primary innovation is to view experimentally found trial-to-trial variability and ongoing fluctuations of parameters such as spine volumes no longer as a nuisance, but as a functionally important component of the organization of network learning, since it enables sampling from a distribution of network configurations. The mathematical framework that we have presented provides a normative model for evaluating such empirically found stochastic dynamics of network parameters, and for relating specific properties of this “noise” to functional aspects of network learning.

Reports of trial-to-trial variability and ongoing fluctuations of parameters related to synaptic weights are ubiquitous in experimental studies of synaptic plasticity and its molecular implementation, from fluctuations of proteins such as PSD-95 [19] in the postsynaptic density that are thought to be related to synaptic strength, over intrinsic fluctuations in spine volumes and synaptic connections [1–3, 5, 28, 31, 32], to surprising shifts of neural codes on a larger time scale [4, 48]. These fluctuations may have numerous causes, from noise in the external environment over noise and fluctuations of internal states in sensory neurons and brain networks, to noise in the pre- and postsynaptic molecular machinery that implements changes in synaptic efficacies on various time scales [18]. One might even hypothesize, that it would be

very hard for this molecular machinery to implement synaptic weights that remain constant in the absence of learning, and deterministic rules for synaptic plasticity, because the half-life of many key proteins that are involved is relatively short, and receptors and other membrane-bound proteins are subject to Brownian motion. In this context the finding that neural codes shift over time [4, 48] appears to be less surprising. In fact, our model predicts (see [S6 Text](#)) that also stereotypical assembly sequences that emerge in our model through learning, similarly as in the experimental data of [49], are subject to such shifts on a larger time scale. However it should be pointed out that our model is agnostic with regard to the time scale on which these changes occur, since this time scale can be defined arbitrarily through the parameter  $b$  (learning rate) in [Eq \(3\)](#).

The model that we have presented makes no assumptions about the actual sources of noise. It only assumes that salient network parameters are subject to stochastic processes, that are qualitatively similar to those which have been studied and modeled in the context of Brownian motion of particles as random walk on the microscale. One can scale the influence of these stochastic forces in the model by a parameter  $T$  that regulates the “temperature” of the stochastic dynamics of network parameters  $\theta$ . This parameter  $T$  regulates the tradeoff between trying out different regions (or modes) of the posterior distribution of  $\theta$  (exploration), and staying for longer time periods in a high probability region of the posterior (exploitation). We conjecture that this parameter  $T$  varies in the brain between different brain regions, and possibly also between different types of synaptic connections within a cortical column. For example, spine turnover is increased for large values of  $T$ , and network parameters  $\theta$  can move faster to a new peak in the posterior distribution, thereby supporting faster learning (and faster forgetting). Since spine turnover is reported to be higher in the hippocampus than in the cortex [50], such higher value of  $T$  could for example be more adequate for modeling network plasticity in the hippocampus. This model would then also support the hypothesis of [50] that memories are more transient in the hippocampus. In addition  $T$  is likely to be regulated on a larger time scale by developmental processes, and on a shorter time scale by neuromodulators and attentional control. The view that synaptic plasticity is stochastic had already been explored through simulation studies in [4, 51]. Artificial neural networks were trained in [51] through supervised learning with high learning rates and high amounts of noise both on neuron outputs and synaptic weight changes. The authors explored the influence of various combinations of noise levels and learning rates on the success of learning, which can be understood as varying the temperature parameters  $T$  in the synaptic sampling framework. In order to measure this parameter  $T$  experimentally in a direct manner, one would have to apply repeatedly the same plasticity induction protocol to the same synapse, with a complete reset of the internal state of the synapse between trials, and measure the resulting trial-to-trial variability of changes of its synaptic efficacy. Since such complete reset of a synaptic state appears to be impossible at present, one can only try to approximate it by the variability that can be measured between different instances of the same type of synaptic connection.

We have shown that the Fokker-Planck equation, a standard tool in physics for analyzing the temporal evolution of the spatial probability density function for particles under Brownian motion, can be used to create bridges between details of local stochastic plasticity processes on the microscale and the probability distribution of the vector  $\theta$  of all parameters on the network level. This theoretical result provides the basis for the new theory of network plasticity that we are proposing. In particular, this link allows us to derive rules for synaptic plasticity which enable the network to learn, and represent in a stochastic manner, a desirable posterior distribution of network parameters; in other words: to approximate Bayesian inference.

We find that resulting rules for synaptic plasticity contain the familiar term for maximum likelihood learning. But another new term, apart from the Brownian-motion-like stochastic term, is the term  $\frac{\partial}{\partial \theta_i} \log p_S(\theta_i)$  that results from a prior distributions  $p_S(\theta_i)$ , which could actually be different for each biological parameter  $\theta_i$  and enforce structural requirements and preferences of networks of neurons in the brain. Some systematic dependencies of changes in synaptic weights (for the same pairing of pre- and postsynaptic activity) on their current values had already been reported in [41, 52–54]. These can be modeled as impact of priors. Other potential functional benefits of priors (on emergent selectivity of neurons) have recently been demonstrated in [55] for a restricted Boltzmann machine. An interesting open question is whether the non-local learning rules of their model can be approximated through biologically more realistic local plasticity rules, e.g. through synaptic sampling. We have also demonstrated in Figs 3 and 4 that suitable priors can model experimental data from [32] and [33] on the survival statistics of dendritic spines. The transient behavior of synaptic turnover in our model fits a two-term exponential function, the long-term (stationary) behavior is well described by a power-law. Both findings are in accordance with experimental data.

The results reported in [56] suggest that learned neural representations integrate experience with a priori beliefs about the sensory environment. The model presented here could be used to further investigate this hypothesis. Also the Fokker-Planck formalism was previously applied to describe the dynamics of dendritic spines in hippocampus [57]. The methods described there to integrate experimental data into computational models could be combined with the synaptic sampling framework to further improve the fit to biology.

Finally, we have demonstrated in Figs 4 and 5 that suitable priors for network parameters  $\theta_i$  that model spine volumes endow a neural network with the capability to respond to changes in the input distribution and network perturbations with a network rewiring that maintains or restores the network function, while simultaneously observing structural constraints such as sparse connectivity.

Our model underlines the importance of further experimental investigation of priors for network parameters. How are they implemented on a molecular level? What role does gene regulation have in their implementation? How does the history of a synapse affect its prior? In particular, can consolidation of a synaptic weight  $\theta_i$  be modeled in an adequate manner as a modification of its prior? This would be attractive from a functional perspective, because according to our model it both allows long-term storage of learned information and flexible network responses to subsequent perturbations.

Besides the use of parameter priors, dropout [58] and dropconnect [59] can be used to avoid overfitting in artificial neural networks. In particular, dropconnect, which drops randomly chosen synaptic connections during training, is reminiscent of stochastic synaptic release in biological neuronal networks. In synaptic sampling, synaptic parameters are assumed to be stochastic, however, this stochastic dynamics evolves on a much slower time scale than stochastic release, which was not modeled in our simulations. An interesting open question is whether synaptic sampling combined with stochastic synaptic release would further improve generalization capabilities of spiking neural networks in a similar manner as dropconnect for artificial neural networks.

We have focused in the examples for our model on the plasticity of synaptic weights and synaptic connections. But the synaptic sampling framework can also be used for studying the plasticity of other synaptic parameters, e.g., parameters that control the short term dynamics of synapses, i.e., their individual mixture of short term facilitation and depression. The corresponding parameters  $U$ ,  $D$ ,  $F$  of the models from [60, 61] are known to depend in a systematic manner on the type of pre- and postsynaptic neuron [62], indicative of a corresponding prior.

However also a substantial variability within the same type of synaptic connections, had been found [62]. Hence it would be interesting to investigate functional properties and experimentally testable consequences of stochastic plasticity rules of type Eq (5) for  $U, D, F$ , and to compare the results with those of previously considered deterministic plasticity rules for  $U, D, F$  (see e.g., [63]).

Early theoretical work on activity-dependent formation and elimination of synapses has been used to model ocular dominance in the visual cortex [64, 65]. Theoretical models for structural plasticity have also shown that simple plasticity models combined with mechanisms for rewiring are able to model cortical reorganization after lesions [66, 67]. In [68] a model was presented that combines structural plasticity and STDP. This model was able to reproduce the existence of transient and persistent spines in the cortex. A recently introduced probabilistic model of structural plasticity was also able to reproduce the statistics of the number of synaptic connections between pairs of neurons in the cortex [69]. Furthermore a simple model of structural synaptic plasticity has been introduced that was able to explain cognitive phenomena such as graded amnesia and catastrophic forgetting [70]. In contrast to these previous studies, the goal of the current work was to establish a model of structural plasticity that follows from a first functional principle, that is, sampling from the posterior distribution over parameters.

We have demonstrated that this framework provides a new and principled way of modeling structural plasticity [10, 11]. The challenge to find a biologically plausible way of modeling structural plasticity as Bayesian inference has been highlighted by [8]. In addition, the proposed framework does not treat rewiring and synaptic plasticity separately, but provides a unified theory for both phenomena, that can be directly related to functional aspects of the network via the resulting posterior distribution. We have shown in Figs 3 and 4 that this rule produces a population of persistent synapses that remain stable over long periods of time, and another population of transient synaptic connections which disappear and reappear randomly, thereby supporting automatic adaptation of the network structure to changes in the distribution of external inputs (Fig 4) and network perturbation (Fig 5).

On a more general level we propose that a framework for network plasticity where network parameters are sampled continuously from a posterior distribution will automatically be less brittle than previously considered maximum likelihood learning frameworks. The latter require some intelligent supervisor who recognizes that the solution given by the current parameter vector is no longer useful, and induces the network to resume plasticity. In contrast, plasticity processes remain active all the time in our sampling-based framework. Hence network compensation for external or internal perturbations is automatic and inherent in the organization of network plasticity.

The need to rethink observed parameter values and plasticity processes in biological networks of neurons in a way which takes into account their astounding variability and compensation capabilities has been emphasized by Eve Marder (see e.g. [9, 47, 71]) and others. This article has introduced a new conceptual and mathematical framework for network plasticity that promises to provide a foundation for such rethinking of network plasticity.

## Methods

### Details to *Learning a posterior distribution through stochastic synaptic plasticity*

Here we prove that  $p^*(\theta) = p(\theta | \mathbf{x})$  is the unique stationary distribution of the parameter dynamics Eq (3) that operate on the network parameters  $\theta = (\theta_1, \dots, \theta_M)$ . Convergence to this stationary distribution then follows for strictly positive  $p^*(\theta)$ . In fact, we prove here a more

general result for parameter dynamics given by

$$d\theta_i = \left( b(\theta_i) \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + b(\theta_i) \frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(\mathbf{x}|\boldsymbol{\theta}) + T b'(\theta_i) \right) dt + \sqrt{2Tb(\theta_i)} d\mathcal{W}_i \tag{12}$$

for  $i = 1, \dots, M$  and  $b'(\theta_i) := \frac{\partial}{\partial \theta_i} b(\theta_i)$ . This dynamics includes a temperature parameter  $T$  and a sampling-speed factor  $b(\theta_i)$  that can in general depend on the current value of the parameter  $\theta_i$ . The temperature parameter  $T$  can be used to scale the diffusion term (i.e., the noise). The sampling-speed factor controls the speed of sampling, i.e., how fast the parameter space is explored. It can be made dependent on the individual parameter value without changing the stationary distribution. For example, the sampling speed of a synaptic weight can be slowed down if it reaches very high or very low values. Note that the dynamics Eq (3) is a special case of the dynamics Eq (12) with unit temperature  $T = 1$  and constant sampling speed  $b(\theta_i) \equiv b$ . We show that the stochastic dynamics Eq (12) leaves the distribution

$$p^*(\boldsymbol{\theta}) \equiv \frac{1}{\mathcal{Z}} q^*(\boldsymbol{\theta}) \tag{13}$$

invariant, where  $\mathcal{Z}$  is a normalizing constant  $\mathcal{Z} = \int q^*(\boldsymbol{\theta}) d\boldsymbol{\theta}$  and

$$q^*(\boldsymbol{\theta}) = p(\boldsymbol{\theta} | \mathbf{x})^\dagger. \tag{14}$$

Note that the stationary distribution  $p^*(\boldsymbol{\theta})$  is shaped by the temperature parameter  $T$ , in the sense that  $p^*(\boldsymbol{\theta})$  is a flattened version of the posterior for high temperature. The result is formalized in the following theorem, which is proven in detail in S1 Text:

**Theorem 1.** *Let  $p(\mathbf{x}, \boldsymbol{\theta})$  be a strictly positive, continuous probability distribution over continuous or discrete states  $\mathbf{x} = \mathbf{x}^1, \dots, \mathbf{x}^N$  and continuous parameters  $\boldsymbol{\theta} = (\theta_1, \dots, \theta_M)$ , twice continuously differentiable with respect to  $\boldsymbol{\theta}$ . Let  $b(\boldsymbol{\theta})$  be a strictly positive, twice continuously differentiable function. Then the set of stochastic differential Eq (12) leaves the distribution  $p^*(\boldsymbol{\theta})$  invariant. Furthermore,  $p^*(\boldsymbol{\theta})$  is the unique stationary distribution of the sampling dynamics.*

**Online approximation.** We show here that the rule Eq (5) is a reasonable approximation to the batch-rule Eq (3). According to the dynamics Eq (12), synaptic plasticity rules that implement synaptic sampling have to compute the log likelihood derivative  $\frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(\mathbf{x}|\boldsymbol{\theta})$ . We assume that every  $\tau_x$  time units a different input  $\mathbf{x}^n$  is presented to the network. For simplicity, assume that  $\mathbf{x}^1, \dots, \mathbf{x}^N$  are visited in a fixed regular order. Under the assumption that input patterns are drawn independently, the likelihood of the generative model factorizes

$$p_{\mathcal{N}}(\mathbf{x}, |\boldsymbol{\theta}) = \prod_{n=1}^N p_{\mathcal{N}}(\mathbf{x}^n | \boldsymbol{\theta}). \tag{15}$$

The derivative of the log likelihood is then given by

$$\frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(\mathbf{x}|\boldsymbol{\theta}) = \sum_{n=1}^N \frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(\mathbf{x}^n | \boldsymbol{\theta}). \tag{16}$$

Using Eq (16) in the dynamics Eq (12), one obtains

$$d\theta_i = \left( b(\theta_i) \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + b(\theta_i) \sum_{n=1}^N \frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(\mathbf{x}^n | \boldsymbol{\theta}) + T b'(\theta_i) \right) dt + \sqrt{2Tb(\theta_i)} d\mathcal{W}_i. \tag{17}$$

Hence, the parameter dynamics depends at any time on all network inputs and network responses.

This “batch” dynamics does not map readily onto a network implementation because the weight update requires at any time knowledge of all inputs  $\mathbf{x}^n$ . We provide here an online approximation for small sampling speeds. To obtain an online learning rule, we consider the parameter dynamics

$$d\theta_i = \left( b(\theta_i) \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + Nb(\theta_i) \frac{\partial}{\partial \theta_i} \log p_N(\mathbf{x}^n | \boldsymbol{\theta}) + T b'(\theta_i) \right) dt + \sqrt{2Tb(\theta_i)} d\mathcal{W}_i. \tag{18}$$

As in the batch learning setting, we assume that each input  $\mathbf{x}^n$  is presented for a time interval of  $\tau_x$ . Integrating the parameter changes Eq (18) over one full presentation of the data  $\mathbf{x}$ , i.e., starting from  $t = 0$  with some initial parameter values  $\boldsymbol{\theta}(0)$  up to time  $t = N\tau_x$ , we obtain for slow sampling speeds ( $N\tau_x b(\theta_i) \ll 1$ )

$$\theta_i(N\tau_x) - \theta_i(0) \approx N\tau_x \left( b(\theta_i) \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + b(\theta_i) \sum_{n=1}^N \frac{\partial}{\partial \theta_i} \log p_N(\mathbf{x}^n | \boldsymbol{\theta}) + T b'(\theta_i) \right) + \sqrt{2Tb(\theta_i)} (\mathcal{W}_i^{N\tau_x} - \mathcal{W}_i^0).$$

This is also what one obtains when integrating Eq (17) for  $N\tau_x$  time units (for slow  $b(\theta_i)$ ). Hence, for slow enough  $b(\theta_i)$ , Eq (18) is a good approximation of optimal weight sampling. The update rule Eq (5) follows from Eq (18) for  $T = 1$  and  $b(\theta_i) \equiv b$ .

**Discrete time approximation.** Here we provide the derivation for the approximate discrete time learning rule Eq (7). For a discrete time parameter update at time  $t$  with discrete time step  $\Delta t$  during which  $\mathbf{x}^n$  is presented, a corresponding rule can be obtained by short integration of the continuous time rule Eq (18) over the time interval from  $t$  to  $t + \Delta t$ :

$$\begin{aligned} \Delta\theta_i &= \Delta t \left( b(\theta_i) \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + Nb(\theta_i) \frac{\partial}{\partial \theta_i} \log p_N(\mathbf{x}^n | \boldsymbol{\theta}) + T b'(\theta_i) \right) \\ &\quad + \sqrt{2Tb(\theta_i)} (\mathcal{W}_i^{t+\Delta t} - \mathcal{W}_i^t) \\ &= \Delta t \left( b(\theta_i) \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + Nb(\theta_i) \frac{\partial}{\partial \theta_i} \log p_N(\mathbf{x}^n | \boldsymbol{\theta}) + T b'(\theta_i) \right) \\ &\quad + \sqrt{2T\Delta t b(\theta_i)} v_i^t, \end{aligned} \tag{19}$$

where  $v_i^t$  denotes Gaussian noise  $v_i^t \sim \text{NORMAL}(0, 1)$ . The update rule Eq (7) is obtained by choosing a constant  $b(\theta) \equiv b$ ,  $T = 1$ , and defining  $\eta = \Delta t b$ .

**Synaptic sampling with hidden states.** When there is a direct relationship between network parameters  $\boldsymbol{\theta}$  and the distribution over input patterns  $\mathbf{x}^n$ , the parameter dynamics can directly be derived from the derivative of the data log likelihood and the derivative of the parameter prior. Typically however, generative models for brain computation assume that the network response  $\mathbf{z}^n$  to input pattern  $\mathbf{x}^n$  represents in some manner the value of hidden variables that explain the current input pattern. In the presence of hidden variables, maximum likelihood learning cannot be applied directly, since the state of the hidden variables is not known from the observed data. The expectation maximization algorithm [7] can be used to overcome this problem. We adopt this approach here. In the online setting, when pattern  $\mathbf{x}^n$  is applied to the network, it responds with network state  $\mathbf{z}^n$  according to  $p_N(\mathbf{z} | \mathbf{x}^n, \boldsymbol{\theta})$ , where the current network parameters are used in this inference process. The parameters are updated in

parallel according to the dynamics

$$d\theta_i = \left( b(\theta_i) \frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + Nb(\theta_i) \frac{\partial}{\partial \theta_i} \log p_N(\mathbf{x}^n, \mathbf{z}^n | \boldsymbol{\theta}) + T b'(\theta_i) \right) dt + \sqrt{2Tb(\theta_i)} dW_i. \tag{20}$$

Note that in comparison with the dynamics Eq (18), the likelihood term now also contains the current network response  $\mathbf{z}^n$ . It can be shown that this dynamics leaves the stationary distribution

$$p^*(\boldsymbol{\theta}) \equiv \frac{1}{\mathcal{Z}} p(\boldsymbol{\theta} | \mathbf{x}, \mathbf{z})^\dagger, \tag{21}$$

invariant, where  $\mathcal{Z}$  is again a normalizing constant (the dynamics Eq (20) is again the online-approximation). Hence, in this setup, the network samples concurrently from circuit states (given  $\boldsymbol{\theta}$ ) and network parameters (given the network state  $\mathbf{z}^n$ ), which can be seen as a sampling-based version of online expectation maximization.

### Details to Improving the generalization capability of a neural network through synaptic sampling

For learning the distribution over different writings of digit 1 with different priors in Fig 2, a restricted Boltzmann machine (RBM) with 748 visible and 9 hidden neurons was used. A detailed definition of the RBM model and additional details to the simulations are given in S3 Text.

**Network inputs.** Handwritten digit images were taken from the MNIST dataset [72]. In MNIST, each instance of a handwritten digit is represented by a 784-dimensional vector  $\mathbf{x}^n$ . Each entry is given by the gray-scale value of a pixel in the  $28 \times 28$  pixel image of the handwritten digit. The pixel values were scaled to the interval [0, 1]. In the RBM, each pixel was represented by a single visible neuron. When an input was presented to the network, the output of a visible neuron was set to 1 with probability as given by the scaled gray-scale value of the corresponding pixel.

**Learning procedure.** In each parameter update step the contrastive divergence algorithm of [27] was used to estimate the likelihood gradients. Therefore, each update step consisted of a “wake” phase, a “reconstruction” phase, and the update of the parameters. The “wake” samples were generated by setting the outputs of the visible neurons to the values of a randomly chosen digit  $\mathbf{x}^n$  from the training set and drawing the outputs  $\mathbf{z}^n$  of all hidden layer neurons for the given visible output. The “reconstruction” activities  $\hat{\mathbf{x}}_j^n$  and  $\hat{\mathbf{z}}_i^n$  were generated by starting from this state of the hidden neurons and then drawing outputs of all visible neurons. After that, the hidden neurons were again updated and so on. In this way we performed five cycles of alternating visible and hidden neuron updates. The outputs of the network neurons after the fifth cycle were taken as the resulting “reconstruction” samples  $\hat{\mathbf{x}}_j^n$  and  $\hat{\mathbf{z}}_i^n$  and used for the parameter updates Eqs (22)–(24) given below. This update of parameters concluded one update step.

Log likelihood derivatives for the biases  $b_i^{\text{hid}}$  of hidden neurons are approximated in the contrastive divergence algorithm [27] as  $\frac{\partial}{\partial b_i^{\text{hid}}} \log p_N(\mathbf{x}^n, \mathbf{z}^n | \boldsymbol{\theta}) \approx z_i^n - \hat{z}_i^n$  (the derivatives for visible biases  $b_j^{\text{vis}}$  are analogous). Using Eq (7), the synaptic sampling update rules for the biases are thus given by

$$\Delta b_i^{\text{hid}} = \eta N(z_i^n - \hat{z}_i^n) + \sqrt{2\eta} v_i^t, \tag{22}$$

$$\Delta b_j^{\text{vis}} = \eta N(x_j^n - \hat{x}_j^n) + \sqrt{2\eta} v_j^t. \tag{23}$$

Note that the parameter prior does not show up in these equations since no priors were used for the biases in our experiments. Contrastive divergence approximates the log likelihood derivatives for the weights  $w_{ij}$  as  $\frac{\partial}{\partial w_{ij}} \log p_{\mathcal{N}}(\mathbf{x}^n, \mathbf{z}^n | \boldsymbol{\theta}) \approx z_i^n x_j^n - \hat{z}_i^n \hat{x}_j^n$ . This leads to the synaptic sampling rule

$$\Delta w_{ij} = \eta \left( \frac{\partial}{\partial w_{ij}} \log p_{\mathcal{S}}(\mathbf{w}) + N \left( z_i^n x_j^n - \hat{z}_i^n \hat{x}_j^n \right) \right) + \sqrt{2\eta} v_{ij}^t. \tag{24}$$

In the simulations, we used this rule with  $\eta = 10^{-4}$  and  $N = 100$ . Learning started from random initial parameters drawn from a Gaussian distribution with standard deviation 0.25 and means at 0 and -1 for weights  $w_{ij}$  and biases  $(b_i^{\text{hid}}, b_j^{\text{vis}})$ , respectively.

To compare learning with and without parameter priors, we performed simulations with an uninformative (i.e., uniform) prior on weights (Fig 2C and 2D), which was implemented by setting  $\frac{\partial}{\partial w_{ij}} \log p_{\mathcal{S}}(\mathbf{w})$  to zero. In simulations with a parameter prior (Fig 2E and 2F), we used a local prior for each weight in order to obtain local plasticity rules. In other words, the prior  $p_{\mathcal{S}}(\mathbf{w})$  was assumed to factorize into priors for individual weights  $p_{\mathcal{S}}(\mathbf{w}) = \prod_{i,j} p_{\mathcal{S}}(w_{ij})$ . For each individual weight prior, we used a bimodal distribution implemented by a mixture of two Gaussians

$$p_{\mathcal{S}}(w_{ij}) = 0.5 \text{NORMAL}(w_{ij} | \mu_1, \sigma_1) + 0.5 \text{NORMAL}(w_{ij} | \mu_2, \sigma_2), \tag{25}$$

with means  $\mu_1 = 1.0$ ,  $\mu_2 = 0.0$ , and standard deviations  $\sigma_1 = \sigma_2 = 0.15$ .

### Details to Fast adaptation to changing input statistics

**Spike-based Winner-Take-All network model.** Network neurons were modeled as stochastic spike response neurons with a firing rate that depends exponentially on the membrane voltage [73, 74]. The membrane potential  $u_k(t)$  of neuron  $k$  is given by

$$u_k(t) = \sum_i w_{ki} x_i(t) + \beta_k(t), \tag{26}$$

where  $x_i(t)$  denotes the (unweighted) input from input neuron  $i$ ,  $w_{ki}$  denotes the efficacy of the synapse from input neuron  $i$ , and  $\beta_k(t)$  denotes a homeostatic adaptation current (see below). The input  $x_i(t)$  models the influence of additive excitatory postsynaptic potentials (EPSPs) on the membrane potential of the neuron. Let  $t_i^{(1)}, t_i^{(2)}, \dots$  denote the spike times of input neuron  $i$ . Then,  $x_i(t)$  is given by

$$x_i(t) = \sum_f \epsilon(t - t_i^{(f)}), \tag{27}$$

where  $\epsilon$  is the response kernel for synaptic input, i.e., the shape of the EPSP, that had a double-exponential form in our simulations:

$$\epsilon(s) = \Theta(s) \left( e^{-\frac{s}{\tau_f}} - e^{-\frac{s}{\tau_r}} \right), \tag{28}$$

with the rise-time constant  $\tau_r = 2$  ms, the fall-time constant  $\tau_f = 20$  ms.  $\Theta(\cdot)$  denotes the Heaviside step function. The instantaneous firing rate  $\rho_k(t)$  of network neuron  $k$  depends exponentially on the membrane potential and is subject to divisive lateral inhibition  $I_{\text{lat}}(t)$  (described below):

$$\rho_k(t) = \frac{\rho_{\text{net}}}{I_{\text{lat}}(t)} \exp(u_k(t)), \tag{29}$$

where  $\rho_{\text{net}} = 100$  Hz scales the firing rate of neurons. Such exponential relationship between the membrane potential and the firing rate has been proposed as a good approximation to the firing properties of cortical pyramidal neurons [73]. Spike trains were then drawn from independent Poisson processes with instantaneous rate  $\rho_k(t)$  for each neuron. We denote the resulting spike train of the  $k^{\text{th}}$  neuron by  $S_k(t)$ .

**Homeostatic adaptation current.** Each output spike caused a slow depressing current, giving rise to the adaptation current  $\beta_k(t)$ . This implements a slow homeostatic mechanism that regulates the output rate of individual neurons (see [75] for details). It was implemented as

$$\beta_k(t) = \gamma \sum_f \kappa(t - t_k^{(f)}), \tag{30}$$

where  $t_k^{(f)}$  denotes the  $f$ -th spike of neuron  $k$  and  $\kappa$  is an adaptation kernel that was modeled as a double exponential (Eq (28)) with time constants  $\tau_r = 12$  s and  $\tau_f = 30$  s. The scaling parameter  $\gamma$  was set to  $\gamma = -8$ .

**Lateral inhibition.** Divisive inhibition [39] between the  $K$  neurons in the WTA network was implemented in an idealized form [36]

$$I_{\text{lat}}(t) = \sum_{l=1}^K \exp(u_l(t)). \tag{31}$$

This form of lateral inhibition, that assumes an idealized access to neuronal membrane potentials, has been shown to implement a well-defined generative network model [36], see below.

**Synaptic sampling in spike-based Winner-Take-All networks as stochastic STDP.** It has been shown in [37] that the WTA-network defined above implicitly defines a generative model that is a mixture of Poissonian distributions. In this generative model, inputs  $\mathbf{x}^n$  are assumed to be generated in dependence on the value of a hidden multinomial random variable  $h^n$  that can take on  $K$  possible values  $1, \dots, K$ . Each neuron  $k$  in the WTA circuit corresponds to one value  $k$  of this hidden variable. In the generative model, for a given value of  $h^n = k$ , the value of an input  $x_i^n$  is then distributed according to a Poisson distribution with a mean that is determined by the synaptic weight  $w_{ki}$  from input neuron  $i$  to network neuron  $k$ :

$$p_{\mathcal{N}}(x_i^n | h^n = k, \mathbf{w}) = \text{POISSON}(x_i^n | \alpha e^{w_{ki}}), \tag{32}$$

with a scaling parameter  $\alpha > 0$ . In other words, the synaptic weight  $w_{ki}$  encodes (in log-space) the firing rate of input neuron  $i$ , given that the hidden cause is  $k$ . For a given hidden cause, inputs are assumed to be independent, hence one obtains the probability of an input vector for a given hidden cause as

$$p_{\mathcal{N}}(\mathbf{x}^n | h^n = k, \mathbf{w}) = \prod_i \text{POISSON}(x_i^n | \alpha e^{w_{ki}}). \tag{33}$$

The network implements inference in this generative model, i.e., for a given input  $\mathbf{x}^n$ , the firing rate of network neuron  $z_k$  is proportional to the posterior probability  $p(h^n = k | \mathbf{x}^n, \mathbf{w})$  of the corresponding hidden cause. An online maximum likelihood learning rule for this generative model was derived in [37]. It changes synaptic weights according to

$$\frac{\partial}{\partial w_{ki}} \log p_{\mathcal{N}}(\mathbf{x}^n | \mathbf{w}) \approx S_k(t) (x_i(t) - \alpha e^{w_{ki}}), \tag{34}$$

where  $S_k(t)$  denotes the spike train of the postsynaptic neuron and  $x_i(t)$  denotes the weight-normalized value of the sum of EPSPs from presynaptic neuron  $i$  at time  $t$  (i.e., the summed EPSPs that would arise for weight  $w_{ki} = 1$ ). To define the synaptic sampling learning rule

completely, we also need to define the parameter prior. In our experiments, we used a simple Gaussian prior on each parameter  $p_S(\theta) = \prod_{k,i} \text{NORMAL}(\theta_{ki}|\mu, \sigma^2)$  with  $\mu = 0.5$  and  $\sigma = 1$ . The derivative of the log-prior is given by

$$\frac{\partial}{\partial \theta_{ki}} \log p_S(\theta) = \frac{1}{\sigma^2} (\mu - \theta_{ki}). \tag{35}$$

Inserting Eqs (34) and (35) into the general form Eq (10), we find that the synaptic sampling rule is given by

$$d\theta_{ki} = b \left( \frac{1}{\sigma^2} (\mu - \theta_{ki}) + N w_{ki} S_k(t) (x_i(t) - \alpha e^{w_{ki}}) \right) dt + \sqrt{2b} dW_{ki}, \tag{36}$$

which corresponds to rule Eq (11) with double indices  $ki$  replaced by single parameter indexing  $i$  to simplify notation.

**Simulation details for spiking network simulations.** Computer simulations of spiking neural networks (Figs 3, 4 and 5) were based on adapted event-based simulation software from [38]. In all spiking neural network simulations, synaptic weights were updated according to the rule Eq (11) with parameters  $N = 100$ ,  $\alpha = e^{-2}$ , and  $b = 10^{-4}$ , except for panel Fig 3F where  $b = 10^{-6}$  was used as a control. In the simulations, we directly implemented the time-continuous evolution of the network parameters in an event-based update scheme. Before learning, initial synaptic parameters were independently drawn from the prior distribution  $p_S(\theta)$ .

For the mapping Eq (9) from synaptic parameters  $\theta_{ki}$  to synaptic efficacies  $w_{ki}$ , we used as offset  $\theta_0 = 3$ . This results in synaptic weights that shrink to small values ( $< 0.05$ ) when synaptic parameters are below zero. In the simulation, we clipped the synaptic weights to zero for negative synaptic parameters  $\theta$  to account for retracted synapses. More precisely, the actual weights  $\hat{w}_{ki}$  used for the computation of the membrane potential Eq (26) were given by  $\hat{w}_{ki} = \max\{0, w_{ki} - \exp(-\theta_0)\}$ . To avoid numerical problems, we clipped the synaptic parameters at -5 and the maximum amplitude of instantaneous parameter changes to  $5b$ .

**Network inputs.** The spatiotemporal spike patterns in Fig 4 are realizations of Poisson spike trains, each representing a certain point in the 3-dimensional sensory environment (a unit cube). Each of the 1000 input neurons was assigned to a Gaussian tuning curve with  $\sigma = 0.3$ . Tuning curve centers were independently and equally scattered over the unit cube. For each sensory experience the firing rate of an individual input neuron was given by the support of sensory experience under the neuron’s tuning curve (normalized between 0 Hz and 80 Hz). In addition an offset of 5 Hz background noise was added. The patterns had a duration of 200 ms. During that time the firing rates of input neurons were kept fixed and independent Poisson spike trains were drawn.

The two environments (SE and EE) in Fig 4 were realized by Gaussian mixture models. The means of the Gaussians were randomly placed around the center of the unit cube (each component was independently drawn from  $\text{NORMAL}(0.5, 0.2)$ ). The covariance matrices of the Gaussian cluster centers were randomly given by  $0.04\mathbb{I} + 0.01\xi$ , where  $\mathbb{I}$  is the 3-dimensional identity matrix and  $\xi$  is a matrix of randomly drawn values from  $\text{NORMAL}(0, 1)$ . Sensory experiences were generated by randomly selecting one Gaussian cluster (with equal probability) and then drawing a sample position from the corresponding multivariate Gaussian.

**Learning schedule and data analysis.** The network was first exposed to samples from the standard environment (SE, Fig 4D) for 3 hours (54000 input sample presentations). In the second learning phase input samples from the enriched environment (EE, Fig 4E) were given for 1 hour (18000 samples). In the third phase samples from either SE (EE-SE condition) or EE

(EE-EE condition) were presented for additional 5 hours (90000 samples, the two cases are compared in Fig 4G).

Formation rates of synaptic connections shown in Fig 4F represent the number of spines that were formed during a time window of  $\Delta t = 30$  minutes, i.e. the number of synaptic connections that were not present ( $\theta_i \leq 0$ ) at time  $t - \Delta t$  but at time  $t$ . The SE condition in Fig 4F was evaluated at the end of learning phase 1, the EE condition was evaluated at the beginning of EE exposure.

For the survival plots in Fig 4G the newly formed synaptic connections at the end of the EE condition were taken into account (see above). Networks from the EE-EE or EE-SE condition were compared. The presence of synaptic connections ( $\theta_i > 0$ ) was evaluated in intervals of 30 minutes. The plot in Fig 4E and 4F show mean values and standard deviations over 5 individual trial runs.

### Details to *Inherent compensation capabilities of networks with synaptic sampling*

Here we provide details to the network model and spiking inputs for the recurrent WTA circuits (Fig 5). Additional details to the data analysis and performance evaluation are provided in S6 Text.

**Network model.** In Fig 5 two recurrently connected ensembles, each consisting of four WTA circuits, were used. The parameters of neuron and synapse dynamics were as described in the previous section. All synapses, lateral and feedforward, were subject to the same learning rule Eq (11). Lateral connections within and between the WTA Circuit neurons were unconstrained (allowing potentially all-to-all connectivity). Connections from input neurons were constraint as shown in Fig 5. The lateral synapses were treated in the same way as synapses from input neurons but had a synaptic delay of 5 ms.

**Network inputs.** Handwritten digit images for Fig 5 were taken from the MNIST dataset [72]. Each pixel was represented by a single afferent neuron. Gray scale values were scaled to 0–50 Hz Poisson input rate and 1 Hz input noise was added on top. These Poisson rates were kept fixed for each example input digit for the duration of the input presentation.

The spoken digit presentations in Fig 5 were given by reconstructed cochleagrams of speech samples of isolated spoken digits from the TI 46 dataset (also used in [40, 76]). Each of the 77 channels of the cochleagrams was represented by 10 afferent neurons, giving a total of 770. Cochleagrams were normalized between 0 Hz and 80 Hz and used to draw individual Poisson spike trains for each afferent neuron. In addition 1 Hz Poisson noise was added on top. We used 10 different utterances of digits 1 and 2 of a single speaker. We selected 7 utterances for training and 3 for testing. For training, one randomly selected utterance from the training set was presented together with a randomly chosen instance of the corresponding handwritten digit. The spike patterns for the written digits (see above) had the same duration as the spoken digits. Each digit presentation was padded with 25 ms, 1 Hz Poisson noise before and after the digit pattern.

For test trials in which only the auditory stimulus was presented, the activity of the visual input neurons was set to 1 Hz throughout the whole pattern presentation. The learning rate  $b$  was set to zero during these trials. The PETH plots were computed over 100 trial responses of the network to the same stimulus class (e.g. presentation of digit 1). Spike patterns for input stimuli were randomly drawn in each trial for the given rates. Spike trains were then filtered with a Gaussian filter with  $\sigma = 50$  ms and summed in a time-discrete matrix with 10 ms bin length. Maximum firing times were assigned to the time bin with the highest PETH amplitude for each neuron.

## Supporting Information

### S1 Text. Proof of Theorem 1.

(PDF)

### S2 Text. Supporting information to [Fig 1](#).

(PDF)

### S3 Text. Supporting information to *Synaptic sampling improves the generalization capability of a neural network*.

(PDF)

### S4 Text. Supporting information to *Spine motility as synaptic sampling*.

(PDF)

### S5 Text. Supporting information to [Fig 3](#).

(PDF)

### S6 Text. Supporting information to *Inherent network compensation capability through synaptic sampling*.

(PDF)

### S1 Code. Source Code.

(TAR.GZ)

## Acknowledgments

We would like to thank Seth Grant, Christopher Harvey, Jason MacLean and Simon Rumpel for helpful comments.

## Author Contributions

Conceived and designed the experiments: DK RL WM SH. Performed the experiments: DK. Analyzed the data: DK. Wrote the paper: WM RL DK SH. Developed the theory: SH RL DK.

## References

- Holtmaat A, Svoboda K. Experience-dependent structural synaptic plasticity in the mammalian brain. *Nature Reviews Neuroscience*. 2009; 10(9):647–658. doi: [10.1038/nrn2699](#) PMID: [19693029](#)
- Yasumatsu N, Matsuzaki M, Miyazaki T, Noguchi J, Kasai H. Principles of long-term dynamics of dendritic spines. *The Journal of Neuroscience*. 2008; 28(50):13592–13608. doi: [10.1523/JNEUROSCI.0603-08.2008](#) PMID: [19074033](#)
- Stettler DD, Yamahachi H, Li W, Denk W, Gilbert CD. Axons and synaptic boutons are highly dynamic in adult visual cortex. *Neuron*. 2006; 49:877–887. doi: [10.1016/j.neuron.2006.02.018](#) PMID: [16543135](#)
- Rokni U, Richardson AG, Bizzi E, Seung HS. Motor learning with unstable neural representations. *Neuron*. 2007; 54(4):653–666. doi: [10.1016/j.neuron.2007.04.030](#) PMID: [17521576](#)
- Yamahachi H, Marik SA, McManus JNJ, Denk W, Gilbert CD. Rapid axonal sprouting and pruning accompany functional reorganization in primary visual cortex. *Neuron*. 2009; 64(5):719–729. doi: [10.1016/j.neuron.2009.11.026](#) PMID: [20005827](#)
- MacKay DJ. Bayesian interpolation. *Neural Computation*. 1992; 4(3):415–447. doi: [10.1162/neco.1992.4.3.415](#)
- Bishop CM. *Pattern Recognition and Machine Learning*. New York: Springer; 2006.
- Pouget A, Beck JM, Ma WJ, Latham PE. Probabilistic brains: knowns and unknowns. *Nature Neuroscience*. 2013; 16(9):1170–1178. doi: [10.1038/nn.3495](#) PMID: [23955561](#)
- Marder E. Variability, compensation and modulation in neurons and circuits. *Proceedings of the National Academy of Sciences*. 2011; 108(3):15542–15548. doi: [10.1073/pnas.1010674108](#)

10. May A. Experience-dependent structural plasticity in the adult human brain. *Trends in Cognitive Sciences*. 2011; 15(10):475–482. doi: [10.1016/j.tics.2011.08.002](https://doi.org/10.1016/j.tics.2011.08.002) PMID: [21906988](https://pubmed.ncbi.nlm.nih.gov/21906988/)
11. Caroni P, Donato F, Muller D. Structural plasticity upon learning: regulation and functions. *Nature Reviews Neuroscience*. 2012; 13(7):478–490. doi: [10.1038/nrn3258](https://doi.org/10.1038/nrn3258) PMID: [22714019](https://pubmed.ncbi.nlm.nih.gov/22714019/)
12. Hatfield G. Perception as Unconscious Inference. In: *Perception and the Physical World: Psychological and Philosophical Issues in Perception*. John Wiley & Sons, Ltd; 2002. p. 115–143.
13. Rao RPN, Olshausen BA, Lewicki MS. *Probabilistic Models of the Brain*. MIT Press; 2002.
14. Doya K, Ishii S, Pouget A, Rao RPN. *Bayesian Brain: Probabilistic Approaches to Neural Coding*. MIT Press; 2007.
15. Winkler I, Denham S, Mill R, Böhm TM, Bendixen A. Multistability in auditory stream segregation: a predictive coding view. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2012; 367(1591):1001–1012. doi: [10.1098/rstb.2011.0359](https://doi.org/10.1098/rstb.2011.0359)
16. Gardiner CW. *Handbook of Stochastic Methods*. 3rd ed. Springer; 2004.
17. Coba MP, Pocklington AJ, Collins MO, Kopanitsa MV, Uren RT, Swamy S, et al. Neurotransmitters Drive Combinatorial Multistate Postsynaptic Density Networks. *Science Signaling*. 2009; 2(68):1–11. doi: [10.1126/scisignal.2000102](https://doi.org/10.1126/scisignal.2000102)
18. Ribault C, Sekimoto K, Triller A. From the stochasticity of molecular processes to the variability of synaptic transmission. *Nature Reviews Neuroscience*. 2011; 12(7):375–387. doi: [10.1038/nrn3025](https://doi.org/10.1038/nrn3025) PMID: [21685931](https://pubmed.ncbi.nlm.nih.gov/21685931/)
19. Gray NW, Weimer RM, Bureau I, Svoboda K. Rapid redistribution of synaptic PSD-95 in the neocortex in vivo. *PLoS Biology*. 2006; 4(11):e370. doi: [10.1371/journal.pbio.0040370](https://doi.org/10.1371/journal.pbio.0040370) PMID: [17090216](https://pubmed.ncbi.nlm.nih.gov/17090216/)
20. Engert F, Bonhoeffer T. Dendritic spine changes associated with hippocampal long-term synaptic plasticity. *Nature*. 1999; 399(6731):66–70. doi: [10.1038/19978](https://doi.org/10.1038/19978) PMID: [10331391](https://pubmed.ncbi.nlm.nih.gov/10331391/)
21. Ho VM, Lee JA, Martin KC. The cell biology of synaptic plasticity. *Science*. 2011; 334(6056):623–628. doi: [10.1126/science.1209236](https://doi.org/10.1126/science.1209236) PMID: [22053042](https://pubmed.ncbi.nlm.nih.gov/22053042/)
22. Bhalla US, Iyengar R. Emergent properties of networks of biological signaling pathways. *Science*. 1999; 283(5400):381–387. doi: [10.1126/science.283.5400.381](https://doi.org/10.1126/science.283.5400.381) PMID: [9888852](https://pubmed.ncbi.nlm.nih.gov/9888852/)
23. Welling M, Teh YW. Bayesian learning via stochastic gradient Langevin dynamics. In: *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*; 2011. p. 681–688.
24. Ackley DH, Hinton GE, Sejnowski TJ. A Learning Algorithm for Boltzmann Machines. *Cognitive Science*. 1985; 9:147–169. doi: [10.1207/s15516709cog0901\\_7](https://doi.org/10.1207/s15516709cog0901_7)
25. Hinton G, Osindero S, Teh YW. A fast learning algorithm for deep belief nets. *Neural Computation*. 2006; 18(7):1527–1554. doi: [10.1162/neco.2006.18.7.1527](https://doi.org/10.1162/neco.2006.18.7.1527) PMID: [16764513](https://pubmed.ncbi.nlm.nih.gov/16764513/)
26. Salakhutdinov R, Hinton G. An Efficient Learning Procedure for Deep Boltzmann Machines. *Neural Computation*. 2012; 24:1967–2006. doi: [10.1162/NECO\\_a\\_00311](https://doi.org/10.1162/NECO_a_00311) PMID: [22509963](https://pubmed.ncbi.nlm.nih.gov/22509963/)
27. Hinton GE. Training products of experts by minimizing contrastive divergence. *Neural Computation*. 2002; 14(8):1771–1800. doi: [10.1162/089976602760128018](https://doi.org/10.1162/089976602760128018) PMID: [12180402](https://pubmed.ncbi.nlm.nih.gov/12180402/)
28. Löwenstein Y, Kuras A, Rumpel S. Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex in vivo. *The Journal of Neuroscience*. 2011; 31(26):9481–9488. doi: [10.1523/JNEUROSCI.6130-10.2011](https://doi.org/10.1523/JNEUROSCI.6130-10.2011)
29. Trachtenberg JT, Chen BE, Knott GW, Feng G, Sanes JR, Welker E, et al. Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature*. 2002; 420(6917):788–794. doi: [10.1038/nature01273](https://doi.org/10.1038/nature01273) PMID: [12490942](https://pubmed.ncbi.nlm.nih.gov/12490942/)
30. Zuo Y, Lin A, Chang P, Gan WB. Development of long-term dendritic spine stability in diverse regions of cerebral cortex. *Neuron*. 2005; 46(2):181–189. doi: [10.1016/j.neuron.2005.04.001](https://doi.org/10.1016/j.neuron.2005.04.001) PMID: [15848798](https://pubmed.ncbi.nlm.nih.gov/15848798/)
31. Holtmaat AJ, Trachtenberg JT, Wilbrecht L, Shepherd GM, Zhang X, Knott GW, et al. Transient and Persistent Dendritic Spines in the Neocortex In Vivo. *Neuron*. 2005; 45:279–291. doi: [10.1016/j.neuron.2005.01.003](https://doi.org/10.1016/j.neuron.2005.01.003) PMID: [15664179](https://pubmed.ncbi.nlm.nih.gov/15664179/)
32. Löwenstein Y, Yanover U, Rumpel S. Predicting the dynamics of network connectivity in the neocortex. *The Journal of Neuroscience*. 2015; 35(36):12535–1. doi: [10.1523/JNEUROSCI.2917-14.2015](https://doi.org/10.1523/JNEUROSCI.2917-14.2015)
33. Yang G, Pan F, Gan WB. Stably maintained dendritic spines are associated with lifelong memories. *Nature*. 2009; 462(7275):920–924. doi: [10.1038/nature08577](https://doi.org/10.1038/nature08577) PMID: [19946265](https://pubmed.ncbi.nlm.nih.gov/19946265/)
34. Brea J, Senn W, Pfister JP. Matching Recall and Storage in Sequence Learning with Spiking Neural Networks. *The Journal of Neuroscience*. 2013; 33(23):9565–9575. doi: [10.1523/JNEUROSCI.4098-12.2013](https://doi.org/10.1523/JNEUROSCI.4098-12.2013) PMID: [23739954](https://pubmed.ncbi.nlm.nih.gov/23739954/)

35. Boerlin M, Machens CK, Denève S. Predictive coding of dynamical variables in balanced spiking networks. *PLoS Computational Biology*. 2013; 9(11):e1003258. doi: [10.1371/journal.pcbi.1003258](https://doi.org/10.1371/journal.pcbi.1003258) PMID: [24244113](https://pubmed.ncbi.nlm.nih.gov/24244113/)
36. Nessler B, Pfeiffer M, Buesing L, Maass W. Bayesian Computation Emerges in Generic Cortical Microcircuits through Spike-Timing-Dependent Plasticity. *PLoS Computational Biology*. 2013; 9(4):e1003037. doi: [10.1371/journal.pcbi.1003037](https://doi.org/10.1371/journal.pcbi.1003037) PMID: [23633941](https://pubmed.ncbi.nlm.nih.gov/23633941/)
37. Habenschuss S, Puhf H, Maass W. Emergence of optimal decoding of population codes through STDP. *Neural Computation*. 2013; 25:1–37. doi: [10.1162/NECO\\_a\\_00446](https://doi.org/10.1162/NECO_a_00446)
38. Kappel D, Nessler B, Maass W. STDP Installs in Winner-Take-All Circuits an Online Approximation to Hidden Markov Model Learning. *PLoS Computational Biology*. 2014; 10(3):e1003511. doi: [10.1371/journal.pcbi.1003511](https://doi.org/10.1371/journal.pcbi.1003511) PMID: [24675787](https://pubmed.ncbi.nlm.nih.gov/24675787/)
39. Carandini M. From circuits to behavior: a bridge too far? *Nature Neuroscience*. 2012; 15(4):507–509. PMID: [22449960](https://pubmed.ncbi.nlm.nih.gov/22449960/)
40. Klampfl S, Maass W. Emergence of dynamic memory traces in cortical microcircuit models through STDP. *The Journal of Neuroscience*. 2013; 33(28):11515–11529. doi: [10.1523/JNEUROSCI.5044-12.2013](https://doi.org/10.1523/JNEUROSCI.5044-12.2013) PMID: [23843522](https://pubmed.ncbi.nlm.nih.gov/23843522/)
41. Sjöström PJ, Turrigiano GG, Nelson SB. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*. 2001; 32(6):1149–1164. doi: [10.1016/S0896-6273\(01\)00542-6](https://doi.org/10.1016/S0896-6273(01)00542-6) PMID: [11754844](https://pubmed.ncbi.nlm.nih.gov/11754844/)
42. Hofer SB, Mrcsic-Flogel TD, Bonhoeffer T, Hübener M. Experience leaves a lasting structural trace in cortical circuits. *Nature*. 2009; 457(7227):313–317. doi: [10.1038/nature07487](https://doi.org/10.1038/nature07487) PMID: [19005470](https://pubmed.ncbi.nlm.nih.gov/19005470/)
43. Kuhlman SJ, O'Connor DH, Fox K, Svoboda K. Structural plasticity within the barrel cortex during initial phases of whisker-dependent learning. *The Journal of Neuroscience*. 2014; 34(17):6078–6083. doi: [10.1523/JNEUROSCI.4919-12.2014](https://doi.org/10.1523/JNEUROSCI.4919-12.2014) PMID: [24760867](https://pubmed.ncbi.nlm.nih.gov/24760867/)
44. Tang LS, Goeritz ML, Caplan JS, Taylor AL, Fisek M, Marder E. Precise temperature compensation of phase in a rhythmic motor pattern. *PLoS Biology*. 2010; 8(8):e1000469. doi: [10.1371/journal.pbio.1000469](https://doi.org/10.1371/journal.pbio.1000469) PMID: [20824168](https://pubmed.ncbi.nlm.nih.gov/20824168/)
45. Grashow R, Brookings T, Marder E. Compensation for variable intrinsic neuronal excitability by circuit-synaptic interactions. *The Journal of Neuroscience*. 2010; 20(27):9145–9156. doi: [10.1523/JNEUROSCI.0980-10.2010](https://doi.org/10.1523/JNEUROSCI.0980-10.2010)
46. Marder E, Taylor AL. Multiple models to capture the variability in biological neurons and networks. *Nature Neuroscience*. 2011; 14(2):133–138. doi: [10.1038/nn.2735](https://doi.org/10.1038/nn.2735) PMID: [21270780](https://pubmed.ncbi.nlm.nih.gov/21270780/)
47. Prinz AA, Bucher D, Marder E. Similar network activity from disparate circuit parameters. *Nature Neuroscience*. 2004; 7(12):1345–1352. doi: [10.1038/nn1352](https://doi.org/10.1038/nn1352) PMID: [15558066](https://pubmed.ncbi.nlm.nih.gov/15558066/)
48. Ziv Y, Burns LD, Cocker ED, Hamel EO, Ghosh KK, Kitch L, et al. Long-term dynamics of CA1 hippocampal place codes. *Nature Neuroscience*. 2013; 16(3):264–266. doi: [10.1038/nn.3329](https://doi.org/10.1038/nn.3329) PMID: [23396101](https://pubmed.ncbi.nlm.nih.gov/23396101/)
49. Harvey CD, Coen P, Tank DW. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*. 2012; 484:62–68. doi: [10.1038/nature10918](https://doi.org/10.1038/nature10918) PMID: [22419153](https://pubmed.ncbi.nlm.nih.gov/22419153/)
50. Attardo A, Fitzgerald JE, Schnitzer MJ. Impermanence of dendritic spines in live adult CA1 hippocampus. *Nature*. 2015; 523:592–596. doi: [10.1038/nature14467](https://doi.org/10.1038/nature14467) PMID: [26098371](https://pubmed.ncbi.nlm.nih.gov/26098371/)
51. Ajemian R, Dâ??Ausilio A, Moorman H, Bizzi E. A theory for how sensorimotor skills are learned and retained in noisy and nonstationary neural circuits. *Proceedings of the National Academy of Sciences*. 2013; 110(52):5078–E5087. doi: [10.1073/pnas.1320116110](https://doi.org/10.1073/pnas.1320116110)
52. Liao D, Jones A, Malinow R. Direct measurement of quantal changes underlying long-term potentiation in CA1 hippocampus. *Neuron*. 1992; 9(6):1089–1097. doi: [10.1016/0896-6273\(92\)90068-O](https://doi.org/10.1016/0896-6273(92)90068-O) PMID: [1334418](https://pubmed.ncbi.nlm.nih.gov/1334418/)
53. Bi GQ, Poo MM. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neuroscience*. 1998; 18(24):10464–10472. PMID: [9852584](https://pubmed.ncbi.nlm.nih.gov/9852584/)
54. Montgomery JM, Pavlidis P, Madison DV. Pair recordings reveal all-silent synaptic connections and the postsynaptic expression of long-term potentiation. *Neuron*. 2001; 29(3):691–701. doi: [10.1016/S0896-6273\(01\)00244-6](https://doi.org/10.1016/S0896-6273(01)00244-6) PMID: [11301028](https://pubmed.ncbi.nlm.nih.gov/11301028/)
55. Xiong H, Szedmak S, Rodríguez-Sánchez A, Piater J. Towards sparsity and selectivity: Bayesian learning of restricted Boltzmann machine for early visual features. In: *Artificial Neural Networks and Machine Learning ICANN 2014*, Springer International Publishing. 2014;p. 419–426.
56. Fiser J, Berkes P, Orbán G, Lengyel M. Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*. 2010; 14(3):119–130. doi: [10.1016/j.tics.2010.01.003](https://doi.org/10.1016/j.tics.2010.01.003) PMID: [20153683](https://pubmed.ncbi.nlm.nih.gov/20153683/)

57. O'Donnell C, Nolan MF, van Rossum MC. Dendritic spine dynamics regulate the long-term stability of synaptic plasticity. *The Journal of Neuroscience*. 2011; 31(45):16142–16156. doi: [10.1523/JNEUROSCI.2520-11.2011](https://doi.org/10.1523/JNEUROSCI.2520-11.2011) PMID: [22072667](https://pubmed.ncbi.nlm.nih.gov/22072667/)
58. Hinton GE, Srivastava N, Krizhevsky A, Sutskever I, Salakhutdinov RR. Improving neural networks by preventing co-adaptation of feature detectors. arXiv preprint. 2012;arXiv:12070580.
59. Wan L, Zeiler M, Zhang S, Cun YL, Fergus R. Regularization of neural networks using dropconnect. In: *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*; 2013. p. 1058–1066.
60. Varela JA, Sen K, Gibson J, Fost J, Abbott LF, Nelson SB. A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *The Journal of Neuroscience*. 1997; 17:220–224.
61. Markram H, Wang Y, Tsodyks M. Differential signaling via the same axon of neocortical pyramidal neurons. *Proceedings of the National Academy of Sciences*. 1998; 95:5323–5328. doi: [10.1073/pnas.95.9.5323](https://doi.org/10.1073/pnas.95.9.5323)
62. Markram H, Toledo-Rodriguez M, Wang Y, Gupta A, Silberberg G, Wu C. Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*. 2004; 5(10):793–807. doi: [10.1038/nrn1519](https://doi.org/10.1038/nrn1519) PMID: [15378039](https://pubmed.ncbi.nlm.nih.gov/15378039/)
63. Natschlaeger T, Maass W, Zador A. Efficient Temporal Processing with Biologically Realistic Dynamic Synapses. *Network: Computation in Neural Systems*. 2001; 12:75–87. doi: [10.1080/net.12.1.75.87](https://doi.org/10.1080/net.12.1.75.87)
64. Elliott T, Shadbolt NR. Competition for neurotrophic factors: ocular dominance columns. *The Journal of Neuroscience*. 1998; 18(15):5850–5858. PMID: [9671672](https://pubmed.ncbi.nlm.nih.gov/9671672/)
65. Elliott T, Shadbolt N. Competition for neurotrophic factors: mathematical analysis. *Neural Computation*. 1998; 10(8):1939–1981. doi: [10.1162/089976698300016927](https://doi.org/10.1162/089976698300016927) PMID: [9804667](https://pubmed.ncbi.nlm.nih.gov/9804667/)
66. Butz M, van Ooyen A. A simple rule for dendritic spine and axonal bouton formation can account for cortical reorganization after focal retinal lesions. *PLoS Computational Biology*. 2013; 9:e1003259. doi: [10.1371/journal.pcbi.1003259](https://doi.org/10.1371/journal.pcbi.1003259) PMID: [24130472](https://pubmed.ncbi.nlm.nih.gov/24130472/)
67. Butz M, Steenbuck ID, van Ooyen A. Homeostatic structural plasticity can account for topology changes following deafferentation and focal stroke. *Frontiers in Neuroanatomy*. 2014; 8:115. doi: [10.3389/fnana.2014.00115](https://doi.org/10.3389/fnana.2014.00115) PMID: [25360087](https://pubmed.ncbi.nlm.nih.gov/25360087/)
68. Deger M, Helias M, Rotter S, Diesmann M. Spike-timing dependence of structural plasticity explains cooperative synapse formation in the neocortex. *PLoS Computational Biology*. 2012; 8(9):e1002689. doi: [10.1371/journal.pcbi.1002689](https://doi.org/10.1371/journal.pcbi.1002689) PMID: [23028287](https://pubmed.ncbi.nlm.nih.gov/23028287/)
69. Fauth M, Wörgötter F, Tetzlaff C. The formation of multi-synaptic connections by the interaction of synaptic and structural plasticity and their functional consequences. *PLoS Computational Biology*. 2015; 11(1):e1004031. doi: [10.1371/journal.pcbi.1004031](https://doi.org/10.1371/journal.pcbi.1004031) PMID: [25590330](https://pubmed.ncbi.nlm.nih.gov/25590330/)
70. Knoblauch A, Körner E, Körner U, Sommer FT. Structural Synaptic Plasticity Has High Memory Capacity and Can Explain Graded Amnesia, Catastrophic Forgetting, and the Spacing Effect. *PLoS One*. 2014; 9(5):e96485. doi: [10.1371/journal.pone.0096485](https://doi.org/10.1371/journal.pone.0096485) PMID: [24858841](https://pubmed.ncbi.nlm.nih.gov/24858841/)
71. Marder E, Goaillard JM. Variability, compensation and homeostasis in neuron and network function. *Nature Reviews Neuroscience*. 2006; 7(7):563–574. doi: [10.1038/nrn1949](https://doi.org/10.1038/nrn1949) PMID: [16791145](https://pubmed.ncbi.nlm.nih.gov/16791145/)
72. LeCun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*. 1998; 86(11):2278–2324. doi: [10.1109/5.726791](https://doi.org/10.1109/5.726791)
73. Jolivet R, Rauch A, Lüscher HR, Gerstner W. Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *Journal of Computational Neuroscience*. 2006; 21(1):35–49. doi: [10.1007/s10827-006-7074-5](https://doi.org/10.1007/s10827-006-7074-5) PMID: [16633938](https://pubmed.ncbi.nlm.nih.gov/16633938/)
74. Mensi S, Naud R, Gerstner W. From stochastic nonlinear integrate-and-fire to generalized linear models. In: *Advances in Neural Information Processing Systems*. vol. 24; 2011. p. 1377–1385.
75. Habenschuss S, Bill J, Nessler B. Homeostatic plasticity in Bayesian spiking networks as Expectation Maximization with posterior constraints. In: *Advances in Neural Information Processing Systems*. vol. 25; 2012. p. 782–790.
76. Hopfield JJ, Brody CD. What is a moment? Transient synchrony as a collective mechanism for spatio-temporal integration. *Proceedings of the National Academy of Sciences*. 2001; 98(3):1282–1287. doi: [10.1073/pnas.98.3.1282](https://doi.org/10.1073/pnas.98.3.1282)