

Research Article: Theory/New Concepts | Cognition and Behavior

A Dynamic Connectome Supports the Emergence of Stable Computational Function of Neural Circuits through Reward-Based Learning

David Kappel^{1,2,3}, Robert Legenstein¹, Stefan Habenschuss¹, Michael Hsieh¹ and Wolfgang Maass¹

¹*Institute for Theoretical Computer Science, Graz University of Technology, Graz, Austria*

²*Bernstein Center for Computational Neuroscience, III Physikalisches Institut – Biophysik, Georg-August Universität, Göttingen, Germany*

³*Chair for Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden Dresden, Germany*

DOI: 10.1523/ENEURO.0301-17.2018

Received: 30 August 2017

Revised: 22 March 2018

Accepted: 26 March 2018

Published: 2 April 2018

Author Contributions: DK and RL contributed equally as first author. DK, RL, SH and WM Designed Research. DK and MH Performed Research. DK Analyzed Data. DK, RL and WM Wrote the paper.

Funding: <http://doi.org/10.13039/501100000780>European Commission (EC) #604102 and #720270

Conflict of Interest: Authors report no conflict of interest.

Written under partial support by the Human Brain Project of the European Union #604102 and #720270.

Correspondence should be addressed to David Kappel (david.kappel@phys.uni-goettingen.de)

Cite as: eNeuro 2018; 10.1523/ENEURO.0301-17.2018

Alerts: Sign up at eneuro.org/alerts to receive customized email alerts when the fully formatted version of this article is published.

Accepted manuscripts are peer-reviewed but have not been through the copyediting, formatting, or proofreading process.

Copyright © 2018 Kappel et al.

This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International license, which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

Manuscript Title Page**1. Manuscript Title (50 word maximum)**

A dynamic connectome supports the emergence of stable computational function of neural circuits through reward-based learning

2. Abbreviated Title (50 character maximum)

A dynamic connectome supports stable learning

3. List all Author Names and Affiliations in order as they would appear in the published article

David Kappel^{1,2,3}, Robert Legenstein¹, Stefan Habenschuss¹, Michael Hsieh¹ and Wolfgang Maass¹

¹ Institute for Theoretical Computer Science

Graz University of Technology

Graz, Austria

² Bernstein Center for Computational Neuroscience

III Physikalisches Institut - Biophysik

Georg-August Universität

Göttingen, Germany

³ Chair for Highly-Parallel VLSI-Systems and Neuromorphic Circuits

Technische Universität Dresden

Dresden, Germany

4. Author Contributions

DK and RL contributed equally as first author. DK, RL, SH and WM Designed Research. DK and MH Performed Research. DK Analyzed Data. DK, RL and WM Wrote the paper.

5. Correspondence should be addressed to (include email address)

David Kappel (david.kappel@phys.uni-goettingen.de)

6. Number of Figures

7

7. Number of Tables

1

8. Number of Multimedia

35 0

36 **9. Number of words for Abstract**

37 240

38 **10. Number of words for Significance Statement**

39

40 116

41 **11. Number of words for Introduction**

42 668

43

44 **12. Number of words for Discussion**

45 1656

46 **13. Acknowledgements**

47 Written under partial support by the Human Brain Project of the European Union #604102 and
48 #720270.

49

50 **14. Conflict of Interest**

51

52 Authors report no conflict of interest

53

54 **15. Funding sources**

55 European Commission

A dynamic connectome supports the emergence of stable computational function of neural circuits through reward-based learning

March 23, 2018

Abstract

Synaptic connections between neurons in the brain are dynamic because of continuously ongoing spine dynamics, axonal sprouting, and other processes. In fact, it was recently shown that the spontaneous synapse-autonomous component of spine dynamics is at least as large as the component that depends on the history of pre- and postsynaptic neural activity. These data are inconsistent with common models for network plasticity, and raise the questions how neural circuits can maintain a stable computational function in spite of these continuously ongoing processes, and what functional uses these ongoing processes might have. Here, we present a rigorous theoretical framework for these seemingly stochastic spine dynamics and rewiring processes in the context of reward-based learning tasks. We show that spontaneous synapse-autonomous processes, in combination with reward signals such as dopamine, can explain the capability of networks of neurons in the brain to configure themselves for specific computational tasks, and to compensate automatically for later changes in the network or task. Furthermore we show theoretically and through computer simulations that stable computational performance is compatible with continuously ongoing synapse-autonomous changes. After reaching good computational performance it causes primarily a slow drift of network architecture and dynamics in task-irrelevant dimensions, as observed for neural activity in motor cortex and other areas. On the more abstract level of reinforcement learning the resulting model gives rise to an understanding of reward-driven network plasticity as continuous sampling of network configurations. (240 words)

Keywords

Spine dynamics, rewiring, stochastic synaptic plasticity, synapse-autonomous processes, task-irrelevant dimensions in motor control, reward-modulated STDP, network computation, network plasticity, reinforcement learning, policy gradient, sampling.

Significance statement

Networks of neurons in the brain do not have a fixed connectivity. We address the question how stable computational performance can be achieved by continuously changing neural circuits, and how these networks could even benefit from these changes. We show that the stationary distribution of network configurations provides a level of analysis where these issues can be addressed in a perspicuous manner.

In particular, this theoretical framework allows us to address analytically the questions which rules for reward-gated synaptic rewiring and plasticity would work best in this context, and what impact different levels of activity-independent synaptic processes are likely to have. We demonstrate the viability of this approach through computer simulations and links to experimental data. (116 words)

Introduction

The connectome is dynamic: Networks of neurons in the brain rewire themselves on a time scale of hours to days [Holtmaat et al., 2005, Stettler et al., 2006, Yang et al., 2009, Holtmaat and Svoboda, 2009, Ziv and Ahissar, 2009, Minerbi et al., 2009, Kasai et al., 2010, Loewenstein et al., 2011, Loewenstein et al., 2015, Rumpel and Triesch, 2016, Chambers and Rumpel, 2017, van Ooyen and Butz-Ostendorf, 2017]. This synaptic rewiring manifests in the emergence and vanishing of dendritic spines [Holtmaat and Svoboda, 2009]. Additional structural changes of established synapses are observable as a growth and shrinking of spine heads which take place even in the absence of neural activity [Yasumatsu et al., 2008]. The recent study of [Dvorkin and Ziv, 2016], which includes in Fig. 8 a reanalysis of mouse brain data from [Kasthuri et al., 2015], showed that this spontaneous component is surprisingly large, at least as large as the impact of pre- and postsynaptic neural activity. In addition, Nagaoka and colleagues provide direct evidence in vivo that the baseline turnover of dendritic spines is mediated by activity-independent intrinsic dynamics [Nagaoka et al., 2016]. Furthermore, experimental data also suggest that task-dependent self-configuration of neural circuits is mediated by reward signals in [Yagishita et al., 2014].

Other experimental data show that not only the connectome, but also the dynamics and function of neural circuits is subject to continuously ongoing changes. Continuously ongoing drifts of neural codes were reported in [Ziv et al., 2013, Driscoll et al., 2017]. Further data show that the mapping of inputs to outputs by neural networks that plan and control motor behavior are subject to a random walk on a slow time-scale of minutes to days, that is conjectured to be related to stochastic synaptic rewiring and plasticity [van Beers et al., 2013, Chaisanguanthum et al., 2014, Rokni et al., 2007].

We address two questions that are raised by these data:

- i) How can stable network performance be achieved in spite of the experimentally found continuously ongoing rewiring and activity-independent synaptic plasticity in neural circuits?
- ii) What could be a functional role of these processes?

Similar as [Statman et al., 2014, Loewenstein et al., 2015, Rokni et al., 2007] we model spontaneous synapse-autonomous spine dynamics of each potential synaptic connection i through a stochastic process that modulates a corresponding parameter θ_i . We provide in this article a rigorous mathematical framework for such stochastic spine dynamics and rewiring processes. Our analysis assumes that one can describe the network configuration, i.e., the current state of the dynamic connectome and the strengths of all currently functional synapses, at any time point by a vector $\boldsymbol{\theta}$ that encodes the current values θ_i for all potential synaptic connections i . The stochastic dynamics of this high-dimensional vector $\boldsymbol{\theta}$ defines a Markov chain whose stationary distribution (illustrated in Fig. 1d) provides insight into questions that address the relation between properties of local synaptic processes and the computational function of a neural network.

Based on the well-studied paradigm for reward-based learning in neural networks, we propose the following answer to question i): As long as most of the mass of this stationary distribution lies in regions or low-dimensional manifolds of the parameter space that produce good performance, stable network performance can be assured in spite of continuously ongoing movement of θ [Loewenstein et al., 2015]. Our experimental results suggest that when a computational task has been learnt, most of the subsequent dynamics of θ takes place in task-irrelevant dimensions.

The same model also provides an answer to question ii): Synapse-autonomous stochastic dynamics of the parameter vector θ enables the network not only to find in a high-dimensional regions with good network performance, but also to rewire the network in order to compensate for changes in the task. We analyze how the strength of the stochastic component of synaptic dynamics affects this compensation capability. We arrive at the conclusion that compensation works best for the task considered here if the stochastic component is as large as in experimental data [Dvorkin and Ziv, 2016].

On the more abstract level of reinforcement learning, our theoretical framework for reward-driven network plasticity suggests a new algorithmic paradigm for network learning: policy sampling. Compared with the familiar policy gradient learning [Williams, 1992, Baxter and Bartlett, 2000, Peters and Schaal, 2006] this paradigm is more consistent with experimental data that suggest a continuously ongoing drift of network parameters.

The resulting model for reward-gated network plasticity builds on the approach from [Kappel et al., 2015] for unsupervised learning, that was only applicable to a specific neuron model and a specific STDP-rule. Since the new approach can be applied to arbitrary neuron models, in particular also to large data-based models of neural circuits and systems, it can be used to explore how data-based models for neural circuits and brain areas can attain and maintain a computational function. (668 words)

Results

We first address the design of a suitable theoretical framework for investigating the self-organization of neural circuits for specific computational tasks in the presence of spontaneous synapse-autonomous processes and rewards. There exist well-established models for reward-modulated synaptic plasticity, see e.g. [Frémaux et al., 2010], where reward signals gate common rules for synaptic plasticity, such as STDP. But these rules are lacking two components that we need here:

- an integration of rewiring with plasticity rules that govern the modulation of the strengths of already existing synaptic connections
- a term that reflects the spontaneous synapse-autonomous component of synaptic plasticity and rewiring.

In order to illustrate our approach we consider a neural network scaffold (see Fig. 1a) with a large number of potential synaptic connections between excitatory neurons. Only a subset of these potential connections is assumed to be functional at any point in time.

If one allows rewiring then the concept of a neural network becomes problematic, since the definition of a neural network typically includes its synaptic connections. Hence we refer to the set of neurons of a

network, its set of potential synaptic connections, and its set of definite synaptic connections – such as in our case connections from and to inhibitory neurons (see Fig. 1a) – as a *network scaffold*. A network scaffold \mathcal{N} together with a parameter vector θ that specifies a particular selection of functional synaptic connections out of the set of potential connections and particular synaptic weights for these defines a concrete neural network, to which we also refer as *network configuration*.

For simplicity we assume that only excitatory connections are plastic, but the model can be easily extended to also reflect plasticity of inhibitory synapses. For each potential synaptic connection i , we introduce a parameter θ_i that describes its state both for the case when this potential connection i is currently not functional (this is the case when $\theta_i \leq 0$) and when it is functional (i.e., $\theta_i > 0$). More precisely, θ_i encodes the current strength or weight w_i of this synaptic connection through the formula

$$w_i = \begin{cases} \exp(\theta_i - \theta_0) & \text{if } \theta_i > 0 & (\text{functional synaptic connection}) \\ 0 & \text{if } \theta_i \leq 0 & (\text{non-functional potential connection}) \end{cases}, \quad (1)$$

with a positive offset parameter θ_0 that regulates the initial strength of new functional synaptic connections (we set $\theta_0 = 3$ in our simulations).

The exponential function in Eq. (1) turns out to be useful for relating the dynamics of θ_i to experimental data on the dynamics of synaptic weights. The volume – or image brightness in Ca-imaging – of a dendritic spine is commonly assumed to be proportional to the strength w_i of a synapse [Holtmaat et al., 2005]. The logarithm of this estimate for w_i was shown in Fig. 2i of [Holtmaat et al., 2006] and also in [Yasumatsu et al., 2008, Loewenstein et al., 2011] to exhibit a dynamics similar to that of an Ornstein-Uhlenbeck process, i.e., a random walk in conjunction with a force that draws the random walk back to its initial state. Hence if θ_i is chosen to be proportional to the logarithm of w_i , it is justified to model the spontaneous dynamics of θ_i as an Ornstein-Uhlenbeck process. This is done in our model, as we will explain after Eq. (5) and demonstrate in Fig. 2c. The logarithmic transformation also ensures that additive increments of θ_i yield multiplicative updates of w_i , which have been observed experimentally [Loewenstein et al., 2011].

Altogether our model needs to create a dynamics for θ_i that is not only consistent with experimental data on spontaneous spine dynamics, but is for the case $\theta_i > 0$ also consistent with rules for reward-modulated synaptic plasticity as in [Frémaux et al., 2010]. This suggests to look for plasticity rules of the form

$$d\theta_i = \beta \times (\text{deterministic plasticity rule}) \times dt + \sqrt{2\beta T} d\mathcal{W}_i, \quad (2)$$

where the deterministic plasticity rule could for example be a standard reward-based plasticity rule. We will argue below that it makes sense to include also an activity-independent prior in this deterministic component of rule (2), both for functional reasons and in order to fit data on spontaneous spine dynamics. We will further see that when the activity-independent prior dominates, we obtain the Ornstein-Uhlenbeck process mentioned above. The stochastic term $d\mathcal{W}_i$ in Eq. (2) is an infinitesimal step of a random walk, more precisely for a Wiener process \mathcal{W}_i . A Wiener process is a standard model for Brownian motion in one dimension [Gardiner, 2004]. The term $\sqrt{2\beta T}$ scales the strength of this stochastic component in terms of a “temperature” T and a learning rate β , and is chosen to be of a form that supports analogies to statistical physics. The presence of this stochastic term makes it unrealistic to expect that θ_i converges to a particular

value under the dynamics defined by Eq. (2). In fact, in contrast to many standard differential equations, the stochastic differential equation or SDE (2) does not have a single trajectory of θ_i as solution, but an infinite family of trajectories that result from different random walks.

We propose to focus – instead of the common analysis of the convergence of weights to specific values as invariants – on the most prominent invariant that a stochastic process can offer: the longterm stationary distribution of synaptic connections and weights. The stationary distribution of the vector θ of all synaptic parameters θ_i informs us about the statistics of the infinitely many different solutions of a stochastic differential equation of the form (2). In particular, it informs us about the fraction of time at which particular values of θ will be visited by these solutions (see Methods for details). We show that a large class of reward-based plasticity rules produce in the context of an equation of the form (2) a stationary distribution of θ that can be clearly related to reward expectation for the neural network, and hence to its computational function.

We want to address the question whether reward-based plasticity rules achieve in the context with other terms in Eq. (2) that the resulting stationary distribution of network configurations has most of its mass on highly rewarded network configurations. A key observation is that if the first term on the right-hand-side of (2) can be written for all potential synaptic connections i in the form $\frac{\partial}{\partial \theta_i} \log p^*(\theta)$, where $p^*(\theta)$ is some arbitrary given distribution and $\frac{\partial}{\partial \theta_i}$ denotes the partial derivative with respect to parameter θ_i , then these stochastic processes

$$d\theta_i = \beta \frac{\partial}{\partial \theta_i} \log p^*(\theta) dt + \sqrt{2\beta T} d\mathcal{W}_i \quad (3)$$

give rise to a stationary distribution that is proportional to $p^*(\theta)^{\frac{1}{T}}$. Hence, a rule for reward-based synaptic plasticity that can be written in the form $\frac{\partial}{\partial \theta_i} \log p^*(\theta)$, where $p^*(\theta)$ has most of its mass on highly rewarded network configurations θ , achieves that the network will spend most of its time in highly rewarded network configurations. This will hold even if the network does not converge to or stay in any particular network configuration θ (see Fig. 1d,f for an illustration). Furthermore the role of the temperature T in (3) becomes clearly visible in this result: if T is large the resulting stationary distribution flattens the distribution $p^*(\theta)$, whereas for $0 < T < 1$ the network will remain for larger fractions of the time in those regions of the parameter space where $p^*(\theta)$ achieves its largest values. In fact, if the temperature T converges to 0, the resulting stationary distribution degenerates to one that has all of its mass on the network configuration θ for which $p^*(\theta)$ reaches its global maximum, as in simulated annealing [Kirkpatrick et al., 1983].

We will focus on target distributions $p^*(\theta)$ of the form

$$p^*(\theta) \propto p_S(\theta) \times \mathcal{V}(\theta), \quad (4)$$

where \propto denotes proportionality up to a positive normalizing constant. $p_S(\theta)$ can encode structural priors of the network scaffold \mathcal{N} . For example, it can encode a preference for sparsely connected networks. This happens when $p_S(\theta)$ has most of its mass near $\mathbf{0}$, see Fig. 1c for an illustration. But it could also convey genetically encoded or previously learnt information, such as a preference for having strong synaptic connections between two specific populations of neurons. The term $\mathcal{V}(\theta)$ in Eq. (4) denotes the expected discounted reward associated with a given parameter vector θ (see Fig. 1b). Eq. (3) for the stochastic

dynamics of parameters takes then the form

$$d\theta_i = \beta \left(\frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) + \frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta}) \right) dt + \sqrt{2\beta T} d\mathcal{W}_i. \quad (5)$$

When the term $\frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta})$ vanishes, this equation models spontaneous spine dynamics. We will make sure that this term vanishes for all potential synaptic connections i that are currently not functional, i.e., where $\theta_i \leq 0$. If one chooses a Gaussian distribution as prior $p_S(\boldsymbol{\theta})$, the dynamics of (5) amounts in the case $\frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta}) = 0$ to an Ornstein-Uhlenbeck process. There is currently no generally accepted description of spine dynamics. Ornstein-Uhlenbeck dynamics has previously been proposed as a simple model for experimentally observed spontaneous spine dynamics [Loewenstein et al., 2011, Loewenstein et al., 2015]. Another proposed model uses a combination of multiplicative and additive stochastic dynamics [Statman et al., 2014, Rubinski and Ziv, 2015]. We used in our simulations a Gaussian distribution that prefers small but nonzero weights for the prior $p_S(\boldsymbol{\theta})$. Hence, our model (5) is consistent with previous Ornstein-Uhlenbeck models for spontaneous spine dynamics.

Thus altogether we arrive at a model for the interaction of stochastic spine dynamics with reward where the usually considered deterministic convergence to network configurations $\boldsymbol{\theta}$ that represent local maxima of expected reward $\mathcal{V}(\boldsymbol{\theta})$ (e.g. to the local maxima in Fig. 1b) is replaced by a stochastic model. If the stochastic dynamics of $\boldsymbol{\theta}$ is defined by local stochastic processes of the form (5), as indicated in Fig. 1e, the resulting stochastic model for network plasticity will spend most of its time in network configurations $\boldsymbol{\theta}$ where the posterior $p^*(\boldsymbol{\theta})$, illustrated in Fig. 1d, approximately reaches its maximal value. This provides on the statistical level a guarantee of task performance, in spite of ongoing stochastic dynamics of all the parameters θ_i .

Reward-based rewiring and synaptic plasticity as policy sampling

We assume that all synapses and neurons in the network scaffold \mathcal{N} receive reward signals $r(t)$ at certain times t , corresponding for example to dopamine signals in the brain (see [Collins and Frank, 2016] for a recent discussion of related experimental data). The expected discounted reward $\mathcal{V}(\boldsymbol{\theta})$ that occurs in the second term of Eq. (5) is the expectation of the time integral over all future rewards $r(t)$, while discounting more remote rewards exponentially, see Eq. (6) in Methods. Fig. 1b shows a hypothetical $\mathcal{V}(\boldsymbol{\theta})$ -landscape over two parameters θ_1, θ_2 . The posterior $p^*(\boldsymbol{\theta})$ shown in Fig. 1d is then proportional to the product of $\mathcal{V}(\boldsymbol{\theta})$ (panel b) and the prior (panel c).

The computational behavior of the network configuration, i.e., the mapping of network inputs to network outputs that is encoded by the parameter vector $\boldsymbol{\theta}$, is referred to as a policy in the context of reinforcement learning theory. The parameters $\boldsymbol{\theta}$ (and therefore the policy) are gradually changed through Eq. (5) such that the expected discounted reward $\mathcal{V}(\boldsymbol{\theta})$ is increased: The parameter dynamics follows the gradient of $\log \mathcal{V}(\boldsymbol{\theta})$, i.e., $\frac{d\theta_i}{dt} = \beta \frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta})$, where $\beta > 0$ is a small learning rate. When the parameter dynamics is given solely by the second term in the parenthesis of Eq. (5), $\frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta})$, we recover for the case $\theta_i > 0$ deterministic policy gradient learning [Williams, 1992, Baxter and Bartlett, 2000, Peters and Schaal, 2006].

For a network scaffold \mathcal{N} of spiking neurons, the derivative $\frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta})$ gives rise to synaptic updates

at a synapse i that are essentially given by the product of the current reward signal $r(t)$ and an eligibility trace that depends on pre- or postsynaptic firing times, see *Synaptic dynamics for the reward-based synaptic sampling model* in Methods. Such plasticity rules have previously been proposed by [Seung, 2003, Xie and Seung, 2004, Izhikevich, 2007, Pfister et al., 2006, Florian, 2007, Legenstein et al., 2008, Urbanczik and Senn, 2009]. For non-spiking neural networks, a similar update rule was first introduced by Williams and termed the REINFORCE rule [Williams, 1992].

In contrast to policy gradient, reinforcement learning in the presence of the stochastic last term in Eq. (5) cannot converge to any network configuration. Instead, the dynamics of Eq. (5) produces continuously changing network configurations, with a preference for configurations that both satisfy constraints from the prior $p_S(\theta)$ and provide a large expected reward $\mathcal{V}(\theta)$, see Fig. 1d,f. Hence this type of reinforcement learning samples continuously from a posterior distribution of network configurations. This is rigorously proven in Theorem 1 of Methods. We refer to this reinforcement learning model as *policy sampling*, and to the family of reward-based plasticity rules that are defined by Eq. (5) as *reward-based synaptic sampling*.

Another key difference to previous models for reward-gated synaptic plasticity and policy gradient learning is, apart from the stochastic last term of Eq. (5), that the deterministic first term of Eq. (5) also contains a reward-independent component $\frac{\partial}{\partial \theta_i} \log p_S(\theta)$ that arises from a prior $p_S(\theta)$ for network configurations. In our simulations we consider a simple Gaussian prior $p_S(\theta)$ with mean $\mathbf{0}$ that encodes a preference for sparse connectivity (see Eq. (17)).

It is important that the dynamics of disconnected synapses, i.e., of synapses i with $\theta_i \leq 0$ or equivalently $w_i = 0$, does not depend on pre-/postsynaptic neural activity or reward since non-functional synapses do not have access to such information. This is automatically achieved through our ansatz $\frac{\partial}{\partial \theta_i} \log \mathcal{V}(\theta)$ for the reward-dependent component in Eq. (5), since a simple derivation shows that it entails that the factor w_i appears in front of the term that depends on pre- and postsynaptic activity, see Eq. (15). Instead, the dynamics of θ_i depends for $\theta_i \leq 0$ only on the prior and the stochastic term dW_i . This results in a distribution over waiting times between downwards and upwards crossing of the threshold $\theta_i = 0$ that was found to be similar to the distribution of inter-event times of a Poisson point process, see [Ding and Rangarajan, 2004] for a detailed analysis. This theoretical result suggest a simple approximation of the dynamics of Eq. (5) for currently non-functional synaptic connections, where the process (5) is suspended whenever θ_i becomes negative, and continued with $\theta_i = 0$ after a waiting time that is drawn from an exponential distribution. As in [Deger et al., 2016] this can be realized by letting a non-functional synapse become functional at any discrete time step with some fixed probability (Poisson process). We have compared in Fig. 3c the resulting learning dynamics of the network for this simple approximation with that of the process defined by Eq. (5).

Task-dependent routing of information through the interaction of stochastic spine dynamics with rewards

Experimental evidence about gating of spine dynamics by reward signals in the form of dopamine is available for the synaptic connections from the cortex to the entrance stage of the basal ganglia, the medium spiny neurons (MSNs) in the striatum [Yagishita et al., 2014]. They report that the volumes of their dendritic

spines show significant changes only when pre- and postsynaptic activity is paired with precisely timed delivery of dopamine (see [Yagishita et al., 2014], Fig. 1 E-G, O). More precisely, an STDP pairing protocol followed by dopamine uncaging induced strong LTP in synapses onto MSNs, whereas the same protocol without dopamine uncaging lead only to a minor increase of synaptic efficacies.

MSNs can be viewed as readouts from a large number of cortical areas, that become specialized for particular motor functions, e.g. movements of the hand or leg. We asked whether reward gating of spine dynamics according to the experimental data of [Yagishita et al., 2014] can explain such task dependent specialization of MSNs. More concretely, we asked whether it can achieve that two different distributed activity patterns P_1 , P_2 of upstream neurons in the cortex get routed to two different ensembles of MSNs, and thereby to two different downstream targets T_1 and T_2 of these MSNs (see Fig. 2a,h,i). We assumed that for each upstream activity pattern P_j a particular subset S_j of upstream neurons is most active, $j = 1, 2$. Hence this routing task amounted to routing synaptic input from S_j to those MSNs that project to downstream neuron T_j .

We applied to all potential synaptic connections i from upstream neurons to MSNs a learning rule according to Eq. (5), more precisely, the rule for reward-gated STDP (Eq. (15), Eq. (16) and Eq. (18)) that results from this general framework. The parameters of the model were adapted to qualitatively reproduce the results from Figures 1F,G of [Yagishita et al., 2014] when the same STDP protocol was applied to our model (see Fig. 2d,e). The parameter values are reported in Tab. 1 in Methods. If not stated otherwise, we applied these parameters in all following experiments. In *Role of the prior distribution* in Methods, we further analyze the impact of different prior distributions on task performance and network connectivity.

Our simple model consisted of 20 inhibitory model MSNs with lateral recurrent connections. These received excitatory input from 200 input neurons. The synapses from input neurons to model MSNs were subject to our plasticity rule. Multiple connections were allowed between each pair of input neuron and MSN (see Methods). The MSNs were randomly divided into two assemblies, each projecting exclusively to one of two downstream target areas T_1 and T_2 . Cortical input $\mathbf{x}(t)$ was modeled as Poisson spike trains from the 200 input neurons with instantaneous rates defined by two prototype rate patterns P_1 and P_2 , see Fig. 2h. The task was to learn to activate T_1 -projecting neurons and to silence T_2 -projecting neurons whenever pattern P_1 was presented as cortical input. For pattern P_2 , the activation should be reversed: activate T_2 -projecting neurons and silence those projecting to T_1 . This desired function was defined through a reward signal $r(t)$ that was proportional to the ratio between the mean firing rate of MSNs projecting to the desired target and that of MSNs projecting to the non-desired target area (see Methods).

Fig. 2h shows the firing activity and reward signal of the network during segments of one simulation run. After about 80 minutes of simulated biological time, each group of MSNs had learned to increase its firing rate when the activity pattern P_j associated with its projection target T_j was presented. Fig. 2f shows the average reward throughout learning. After 3 hours of learning about 82% of the maximum reward was acquired on average, and this level was maintained during prolonged learning.

Fig. 2g shows that the parameter vector θ kept moving at almost the same speed even after a high plateau of rewards had been reached. Hence these ongoing parameter changes took place in dimensions that were irrelevant for the reward-level.

Fig. 2i provides snapshots of the underlying “dynamic connectome” [Rumpel and Triesch, 2016] at

different points of time. New synaptic connections that were not present at the preceding snapshot are colored green. One sees that the bulk of the connections maintained a solution of the task to route inputs from S_1 to target area T_1 and inputs from S_2 to target area S_2 . But the identity of these connections, a task-irrelevant dimension, kept changing. In addition the network always maintained some connections to the currently undesired target area, thereby providing the basis for a swift built-up of these connections if these connections would suddenly also become rewarded.

We further examine the exploration along task-irrelevant dimensions in Fig. 2j. Here, the high-dimensional parameter vector over a training experiment of 24 h projected to the first two components of the demixed principal component analysis (dPCA) that best explain the variance of the average reward is shown (see Methods and [Kobak et al., 2016]). The first component (dpc1) explains $> 99.9\%$ of the variance. Movement of the parameter vector mainly takes place along this dimensions during the first 4 hours of learning. After the performance has converged to a high value, exploration continues along other components (dpc2, and higher components) that explain less than 0.1% of the average reward variance.

This simulation experiment showed that reward-gated spine dynamics as analyzed in [Yagishita et al., 2014] is sufficiently powerful from the functional perspective to rewire networks so that each signal is delivered to its intended target.

A model for task-dependent self-configuration of a recurrent network of excitatory and inhibitory spiking neurons

We next asked, whether our simple integrated model for reward-modulated rewiring and synaptic plasticity of neural circuits according to Eq. (5) could also explain the emergence of specific computations in recurrent networks of spiking neurons. As paradigm for a specific computational task we took a simplified version of the task that mice learned to carry out in the experimental setup of [Peters et al., 2014]. There a reward was given whenever a lever was pressed within a given time window indicated by an auditory cue. This task is particular suitable for our context, since spine turnover and changes of network activity were continuously monitored in [Peters et al., 2014] while the animals learned this task.

We adapted the learning task of [Peters et al., 2014] in the following way for our model (see Fig. 3a). The beginning of a trial was indicated through the presentation of a cue input pattern $\mathbf{x}(t)$: a fixed, randomly generated rate pattern for all 200 input neurons that lasted until the task was completed, but at most 10s. As network scaffold \mathcal{N} we took a generic recurrent network of excitatory and inhibitory spiking neurons with connectivity parameters for connections between excitatory and inhibitory neurons according to data from layer 2/3 in mouse cortex [Avermann et al., 2012]. The network consisted of 60 excitatory and 20 inhibitory neurons (see Fig. 3a). Half of the excitatory neurons could potentially receive synaptic connections from the 200 excitatory input neurons. From the remaining 30 neurons we randomly selected one pool D of 10 excitatory neurons to cause downwards movements of the lever, and another pool U of 10 neurons for upwards movements. We refer to the 40 excitatory neurons that were not members of D or U as hidden neurons. All excitatory synaptic connections from the external input (cue) and between the 60 excitatory neurons (including those in the pools D and U) in the network were subjected to reward-based synaptic sampling.

To decode the lever position we filtered the population spikes of D and U with a smoothing kernel. The filtered population spikes of D were then subtracted from those of U to determine the lever position (see Methods for details). When the lever position crossed the threshold +5 after first crossing a lower threshold -5 (black horizontal lines in Fig. 3a,b) within 10 s after cue onset a 400 ms reward window was initiated during which $r(t)$ was set to 1 (red vertical bars in Fig. 3b). Unsuccessful trials were aborted after 10 seconds and no reward was delivered. After each trial a brief holding phase of random length was inserted, during which input neurons were set to a background input rate of 2 Hz.

Thus, the network had to learn without any guidance, except for the reward in response to good performance, to create after the onset of the cue first higher firing in pool D, and then higher firing in pool U. This task was challenging, since the network had no information which neurons belonged to pools D and U. Moreover, the synapses did not “know” whether they connected to hidden neurons, neurons within a pool, hidden neurons and pool-neurons, or input neurons with other neurons. The plasticity of all these different synapses was gated by the same global reward signal. Since the pools D and U were not able to receive direct synaptic connections from the input neurons, the network also had to learn to communicate the presence of the cue pattern via disynaptic connections from the input neurons to these pools.

Network responses before and after learning are shown in Fig. 3b. Initially, the rewarded goal was only reached occasionally, while the turnover of synaptic connections (number of synaptic connections that became functional or became non-functional in a time window of 2 hours) remained very high (see Fig. 3e). After about 3 h, performance improved drastically (Fig. 3c), and simultaneously the turnover of synaptic connections slowed down (Fig. 3e). After learning for 8 hours, the network was able to solve the task in most of the trials, and the average trial duration (movement completion time) had decreased to less than 1 second (851 ± 46 ms, Fig. 3c). Improved performance was accompanied by more stereotyped network activity and lever movement patterns as in the experimental data of [Peters et al., 2014]: compare our Fig. 3d with Fig. 1b and Fig. 2j of [Peters et al., 2014]. In Fig. 3d we show the trial-averaged activity of the 60 excitatory neurons before and after learning for 22 hours. The neurons are sorted in the first two plots of Fig. 3d by the time of maximum activity after movement onset times before learning, and in the 3rd plot resorted according to times of maximum activity after 22 hours of learning (see *Methods*). These plots show that reward-based learning led to a restructuring of the network activity: an assembly of neurons emerged that controlled a sharp upwards movement. Also, less background activity was observed after 22 hours of learning, in particular for neurons with early activity peaks. Lower panels in Fig. 3d show the average lever movement and 10 individual movement traces at the beginning and after 22 hours of learning. Similar as in [Peters et al., 2014] the lever movements became more stereotyped during learning, featuring a sharp upwards movement at cue onset followed by a slower downwards movement in preparation for the next trial.

The synaptic parameter drifts due to stochastic differential equation (5) inherently lead to forgetting. In Fig. 3f we tested this effect by running a very long experiment over 14 simulated days. After 24 h, when the network had learned to reliably solve the task, we stopped the application of the reward but continued the synaptic dynamics. We found that the task could be reliably recalled for more than 5 days. Onset of forgetting was observed after day 6. We wondered whether a simple consolidation mechanism could prevent forgetting in our model. To test this we used the prior distribution $p_S(\theta)$ to stabilize the

synaptic parameters. After 4 simulated days we set the mean of the prior to the current value of the synaptic parameters and reduced the variance, while continuing the synaptic dynamics with the same temperature. A similar mechanism for synaptic consolidation has been recently suggested in [Kirkpatrick et al., 2017]. This mechanism reliably prevents forgetting in our model throughout the simulation time of 14 days. We conclude that the effect of forgetting is quite mild in our model and can be further suppressed by a consolidation mechanism that stabilizes synapses on longer timescales.

Next we tested whether similar results could be achieved with a simplified version of the stochastic synapse dynamics while a potential synaptic connection i is non-functional, i.e., $\theta_i \leq 0$. Eq. (5) defines for such non-functional synapses an Ornstein-Uhlenbeck process, which yields a heavy-tailed distribution for the waiting time until reappearance (Fig. 3g, left). We tested whether similar learning performance can be achieved if one approximates the distribution by an exponential distribution, for which we chose a mean of 12 h. The small distance between the blue and green curve in Fig. 3c shows that this is in fact the case for the overall computational task that includes a task switch at 24 h that we describe below. Compensation for the task switch was slightly slower when the approximating exponential distribution was used, but the task performance converged to the same result as for the exact rule. This holds in spite of the fact that the approximating exponential distribution is less heavy-tailed (Fig. 3g, right). Altogether these results show that rewiring and synaptic plasticity according to Eq. (5) yields self-organization of a generic recurrent network of spiking neurons so that it can control an arbitrarily chosen motor control task.

Compensation for network perturbations

We wondered whether this model for the task of [Peters et al., 2014] would in addition be able to compensate for a drastic change in the task, an extra challenge that had not been considered in the experiments of [Peters et al., 2014]. To test this we suddenly interchanged the actions that were triggered by the pools D and U at 24 h after learning had started. D now caused upwards and U downwards lever movement.

We found that our model compensated immediately (see the faster movement in the parameter space depicted in Fig. 3h) for this perturbation and reached after about 8 h a similar performance level as before (Fig. 3c). The compensation was accompanied by a substantial increase in the turnover of synaptic connections (Fig. 3e). This observation is similar to findings from experiments that involve learning a new task, see e.g. [Xu et al., 2009]. The turnover rate also remained slightly elevated during the subsequent learning period. Furthermore, a new assembly of neurons emerged that now triggered a sharp onset of activity in the pool D (compare the activity neural traces at $h = 22$ and $h = 46$ in Fig. 3d). Another experimentally observed phenomenon that occurred in our model were drifts of neural codes, which happened also during phases of the experiment without perturbations. Despite these drifts, the task performance stayed constant, similar to experimental data in [Driscoll et al., 2017] (see Fig. 6 in Methods).

In Fig. 3h we further analyzed the profile of synaptic turnover for the different populations of the network scaffold in Fig. 3a. The synaptic parameters were measured immediately before the task change at 24 h and compared to the connectivity after compensation at 48 h for the experiment shown in Fig. 3c (blue). Most synapses (66-75%) were non-functional before and after the task change (stable non-functional). About 20% of the synapses changed their behavior and either became functional or non-functional. Most prominently a large fraction (21.9%) of the synapses from hidden neurons to U became non-functional

while only few (5.9%) new connections were introduced. The connections from hidden to D showed the opposite behavior. This modification of the network connectome reflects the requirement to reliably route information about the presence of the cue pattern encoded in the activity of hidden neurons to the pool D (and not to U) to initiate the lever movement after the task change.

If rewiring was disabled after the task change at 24 h the compensation was significantly delayed and overall performance declined (see black curve in Fig. 3c). Here, we disallowed any turnover of potential synaptic connections such that the connectivity remained the same after 24 h. This result suggests that rewiring is necessary for adapting to the task change. We then asked whether rewiring is also necessary for the initial learning of the task. To answer this question, we performed a simulation where the network connectivity was fixed from the beginning. We found that initial task performance was not significantly worse compared to the setup with rewiring. This indicates that at least for this task, rewiring is necessary for compensating task switches, but not for initial task learning. We expect however that this is not the case for more complex tasks, as indicated by a recent study that used artificial neural networks [Bellec et al., 2017].

A structural difference between stochastic learning models such as policy sampling and learning models that focus on convergence of parameters to a (locally) optimal setting becomes apparent when one tracks the temporal evolution of the network parameters θ over larger periods of time during the previously discussed learning process (Fig. 3i). Although performance no longer improved after 5 hours, both network connectivity and parameters kept changing in task-irrelevant dimensions. For Fig. 3i we randomly selected 5% of the roughly 47000 parameters θ_i and plotted the first 3 principal components of their dynamics. The task change after 24 hours caused the parameter vector θ to migrate to a new region within about 8 hours of continuing learning (see Fig. 7 where the projected parameter dynamics is further analyzed). Again we observe that policy sampling keeps exploring different equally good solutions after the learning process has reached stable performance.

To further investigate the impact of the temperature parameter T on the magnitude of parameter changes, we measured the amplitudes of parameter changes for different values of T . We recorded the synaptic parameters every 20 minutes and measured the average Euclidean distance between successive snapshots of the parameter vectors. We found that a temperature of $T = 0.1$ increased the amplitude of parameter changes by around 150% compared to the case of $T = 0$. A temperature of $T = 0.5$ resulted in an increase of around 400%. Since this increase is induced by additional noise on parameter changes, it can be attributed to enhanced exploration in parameters space.

Role of the prior distribution

Next, we investigated the role of the prior distribution and initial network configuration for the experiment in Fig. 3. Fig. 4 shows the performance and total number of active connections for different parameter settings. As in the previous experiments we used in Fig. 4a,b a Gaussian prior distribution with mean μ and variance σ^2 . The preferred number of active connections changes with the prior, i.e., priors with smaller variance and low mean lead to sparser networks. Convergence to this preferred number can take more than 24 hours depending on the initial connectivity. Different parameter settings can therefore lead to quite different network connectivities at a similar task performance. A too strong prior (e.g., $\mu = -2$,

$\sigma = 0.5$) leads to very sparse networks, thereby preventing learning.

In addition to the Gaussian prior distribution we tested a Laplace prior of the form $p_S(\theta_i) = \frac{1}{2b} \exp(-\frac{|\theta_i|}{b})$, with zero mean and scale parameter $b > 0$ (Fig. 4c). This leads to a constant negative drift term in the parameter dynamics Eq. (5), i.e. $\frac{\partial}{\partial \theta_i} \log p_S(\theta) = -\frac{1}{b}$ for active synaptic connections. A similar mechanism for synaptic weight decay was previously used in [Rokni et al., 2007]. Convergence to sparse connectivities is faster with this prior and good task performance can be reached by networks with less active connections compared to the Gaussian prior. For example, the network with $b = 2$ solved the task in 0.66 seconds on average using roughly 5700 active connections, whereas the best solution for the Gaussian prior was 0.83 seconds on average with typically more than 7500 active connections. Again, for the Laplace prior, parameters that enforced too sparse networks degraded task performance.

We next investigated whether a scaling of the amplitude of the reward signal $r(t)$ while keeping the same prior has an influence on network performance. we introduced a scaling constant c_r that can be used to modulate the amplitude of the reward signal ($c_r = 1$ corresponds to the setting in Fig. 3, see *Methods* for details). We repeated the experiment from Fig. 3 (including the task change after 24 hours) with c_r ranging between 0.1 and 10. For values of c_r smaller than 1 the effect of the second term of the synaptic dynamics (5) is scaled down which results in an overall reduced learning speed and a stronger influence of the prior. Interestingly however, in all cases the network was able to compensate for the task change after 48 hours of simulated biological time (see Fig. 4d, movement completion times of 983 ± 63 ms, 894 ± 41 ms, 820 ± 45 ms, 743 ± 25 ms and 1181 ± 42 ms for $c_r = 0.1, 0.5, 1, 5$, and 10 respectively). In the next section we further investigate the role of the temperature term T that controls the amount of noise in the synaptic dynamics.

Relative contributions of spontaneous and activity-dependent synaptic processes

[Dvorkin and Ziv, 2016] analyzed the correlation of sizes of postsynaptic densities and spine volumes for synapses that shared the same pre- and postsynaptic neuron, called commonly innervated (CI) synapses, and also for synapses that shared in addition the same dendrite (CI_{SD}). Activity-dependent rules for synaptic plasticity, such as Hebbian or STDP rules – on which previous models for network plasticity relied – suggest that the strength of CI and especially CI_{SD} synapses should be highly correlated. But both data from ex-vivo [Kasthuri et al., 2015] and neural circuits in culture [Dvorkin and Ziv, 2016] show that postsynaptic density sizes and spine volumes of CI_{SD} synapses are only weakly correlated, with correlation coefficients between 0.23 and 0.34. Thus even with a conservative estimate that corrects for possible influences of their experimental procedure, more than 50% of the observed synaptic strength appears to result from activity-independent stochastic processes (Fig. 8E of [Dvorkin and Ziv, 2016]). [Bartol Jr et al., 2015] had previously found larger correlations of synaptic strengths of CI_{SD} synapses for a smaller data set (based on 17 CI_{SD} pairs instead of the 72 pairs, 10 triplets, and 2 quadruplets in the ex-vivo data from [Kasthuri et al., 2015]), but the spine volumes differed in these pairs also on average by a factor of around 2.

We asked how such a strong contribution of activity-independent synaptic dynamics affects network learning capabilities, such as the ones that were examined in Fig. 3. We were able to carry out this test because many synaptic connections between neurons that were formed in our model consisted of more than

one synapse. We classified pairs of synapses that had the same pre- and postsynaptic neuron as CI synapses (one could also call them CI_{SD} synapses, since the neuron model did not have different dendrites), and pairs with the same postsynaptic but different presynaptic neurons as non-CI synapses. Example traces of synaptic weights for CI and non-CI synapse pairs of our network model from Fig. 3 are shown in Fig. 5a,b. CI pairs were found to be more strongly correlated than non-CI pairs (Fig. 5c). However also the correlation of CI pairs was quite low, and varied with the temperature parameter T in Eq. (5), see Fig. 5d. The correlation was measured in terms of the Pearson correlation (covariance of synapse pairs normalized between -1 and 1).

Since the correlation of CI pairs in our model depends on the temperature T , we analyzed the model of Fig. 3 for different temperatures (the temperature had been fixed at $T=0.1$ throughout the experiments for Fig. 3). In Fig. 5d the Pearson correlation coefficient for CI synapses is plotted together with the average performance achieved on the task of Fig. 3d-h (24 hours after the task switch) for networks with different temperatures T . The best performing temperature region for the task ($0.01 \leq T \leq 0.5$) roughly coincided with the region of experimentally measured values of Pearson's correlation for CI-synapses. Fig. 5e shows the correlation of 100 CI synapse pairs that emerged from a run with $T = 0.5$. We found a value of $r = 0.239$ in this case. This value is in the order of the lowest experimentally found correlation coefficients in [Dvorkin and Ziv, 2016] (both in culture and ex-vivo, see Fig. 8A-D in [Dvorkin and Ziv, 2016]). The speed of compensation and the overall replenishing of synapses was strongly dependent on the temperature T (see Fig. 5g). For $T = 0$, a complete compensation for the task changes was prevented (performance converged to 2.5 ± 0.2 s during a longer run of 96 hours). The temperature region $0.01 \leq T \leq 0.5$ – that is consistent with experimentally measured Pearson's correlation for CI-synapses – leads to fastest task re-learning, allowing for a compensation within approximately 12 hours of exposure. For $T = 0.15$ we found the best compensation capabilities and the closest match to experimentally measured correlations when the results of [Dvorkin and Ziv, 2016] were corrected for measurement limitations: A correlation coefficient of $r = 0.46 \pm 0.034$ for CI synapses and 0.08 ± 0.015 for non-CI synapse pairs (mean \pm s.e.m. over 5 trials, CI synapses were significantly stronger correlated than non-CI, p-value below 0.005 in all trials. Statistical significance values based on two-tailed Mann-Whitney U test.).

[Dvorkin and Ziv, 2016] further analyzed the ratio of contributions from different processes to the measured synaptic dynamics. They analyzed the contribution of neural activity history dependent processes, which amount for 36% of synapse dynamics in their data, and that of neuron-wide processes that were not specific to presynaptic activity, but specific to the activity of the postsynaptic neuron (8%). Spontaneous synapse-autonomous processes were found to explain 56% of the observed dynamics (see Fig. 8E in [Dvorkin and Ziv, 2016]). The results from our model with $T = 0.15$, that are plotted in Fig. 5f, match these experimentally found values quite well. Altogether we found that the results of [Dvorkin and Ziv, 2016] are best explained by our model for a temperature parameter between $T = 0.5$ (corresponding to their lowest measured correlation coefficient) and $T = 0.15$ (corresponding to their most conservative estimate). This range of parameters coincided with well-functioning learning behavior in our model, which included a test of compensation capability for a change of the task after 24 h (Fig. 5d). Hence our model suggests that a large component of stochastic synapse-autonomous processes, as it occurs in the data, supports efficient network learning and compensation for changes in the task.

Discussion

Recent experimental data ([Nagaoka et al., 2016] and [Dvorkin and Ziv, 2016], where in their Fig. 8 also mouse brain data from [Kasthuri et al., 2015] were reanalyzed) suggest that common models for learning in neural networks of the brain need to be revised, since synapses are subject to powerful processes that do not depend on pre- and postsynaptic neural activity. In addition, experimentally found network rewiring has so far not been integrated into models for reward-gated network plasticity. We have presented a theoretical framework that enables us to investigate and understand reward-based network rewiring and synaptic plasticity in the context of the experimentally found high level of activity-independent fluctuations of synaptic connectivity and synaptic strength. We have shown that the analysis of the stationary distribution of network configurations, in particular the Fokker-Planck equation from theoretical physics, allows us to understand how large numbers of local stochastic processes at different synapses can orchestrate global goal-directed network learning. This approach provides a new normative model for reward-gated network plasticity.

We have shown in Fig. 2 that the resulting model is consistent with experimental data on dopamine-dependent spine dynamics reported in [Yagishita et al., 2014], and that it provides an understanding how these local stochastic processes can produce function-oriented cortical-striatal connectivity. We have shown in Fig. 3 that this model also elucidates reward-based self-organization of generic recurrent neural networks for a given computational task. We chose as benchmark task the production of a specific motor output in response to a cue, like in the experiments of [Peters et al., 2014]. Similarly as reported in [Peters et al., 2014], the network connectivity and dynamics reorganized itself in our model, just driven by stochastic processes and rewards for successful task completion, and reached a high level of performance. Furthermore it maintained this computational function in spite of continuously ongoing further rewiring and network plasticity. A quantitative analysis of the impact of stochasticity on this process has shown in Fig. 5 that the network learns best when the component of synaptic plasticity that does not depend on neural activity is fairly large, as large as reported in the experimental data of [Kasthuri et al., 2015, Dvorkin and Ziv, 2016].

Our approach is based on experimental data for the biological implementation level of network plasticity, i.e., for the lowest level of the Marr hierarchy of models [Marr and Poggio, 1976]. However, we have shown that these experimental data have significant implications for understanding network plasticity on the top level ("what is the functional goal?") and the intermediate algorithmic level ("what is the underlying algorithm?") of the Marr hierarchy. They suggest for the top level that the goal of network plasticity is to evaluate a posterior distribution of network configurations. This posterior integrates functional demands formalized by the expected discounted reward $\mathcal{V}(\theta)$ with a prior $p_S(\theta)$ in a multiplicative manner $p^*(\theta) \propto p_S(\theta) \times \mathcal{V}(\theta)$. Priors can represent structural constraints as well as results of preceding learning experiences and innate programs. Since our model samples from a distribution proportional to $p^*(\theta)^{\frac{1}{T}}$, for $T = 1$, our model suggests to view reward-gated network plasticity as Bayesian inference over network configurations on a slow time scale, see *Probabilistic framework for reward-modulated learning* in *Methods* for details. For a temperature parameter $T \neq 1$, the model samples from a tempered version of the posterior, which generalizes the basic Bayesian approach. This Bayesian perspective also creates a link to previous work

on Bayesian reinforcement learning [Vlassis et al., 2012, Rawlik et al., 2013]. We note however that we do not consider parameter adaptation in our framework to implement full Bayesian learning, as there is no integration over the posterior parameter settings to obtain network outputs (or actions in a reinforcement learning context). Even if one would do that, it would be of little practical use, since the sampling would be much too slow in any but the simplest networks. The experimental data suggest for the intermediate algorithmic level of the Marr hierarchy a strong reliance on stochastic search ("synaptic sampling"). The essence of the resulting model for reward-gated network learning is illustrated in Fig. 1: The traditional view of deterministic gradient ascent (policy gradient) in the landscape (panel b) of reward expectation is first modified through the integration of a prior (panel c), and then through the replacement of gradient ascent by continuously ongoing stochastic sampling (policy sampling) from the posterior distribution of panel d, which is illustrated in panels e and f.

This model explains a number of experimental data that had not been addressed by previous models. Continuously ongoing stochastic sampling of network configurations suggests that synaptic connectivity does not converge to a fixed point solution but rather undergoes permanent modifications (Fig. 3h,i). This behavior is compatible with reports of continuously ongoing spine dynamics and axonal sprouting even in the adult brain [Holtmaat and Svoboda, 2009, Yasumatsu et al., 2008, Stettler et al., 2006, Yamahachi et al., 2009, Loewenstein et al., 2011, Holtmaat et al., 2005, Loewenstein et al., 2015]. Recently proposed models to maintain stable network function in the presence of highly volatile spine dynamics suggest that subsets of connections are selectively stabilized to support network function [Berry and Nedivi, 2017, Mongillo et al., 2017]. Our result shows that high task performance can be reached in spiking neural networks in the presence of high volatility of all synapses. Still our model can be extended with a process that selectively stabilizes synapses on longer timescales as demonstrated in Fig. 3f. In addition, our model predicts that not only synaptic spine dynamics but also *changes* of synaptic efficacies show a large stochastic component on all timescales.

The continuously ongoing parameter changes induce continuously ongoing changes in the assembly sequences that accompany and control a motor response (see Fig. 3d). These changes do not impair the performance of the network, but rather enable the network to explore different but equally good solutions when exposed for many hours to the same task (see Fig. 3i). Such continuously ongoing drifts of neural codes in functionally less relevant dimensions have already been observed experimentally in some brain areas [Ziv et al., 2013, Driscoll et al., 2017]. Our model also suggests that the same computational function is realized by the same neural circuit in different individuals with drastically different parameters, a feature which has already been addressed in [Tang et al., 2010, Grashow et al., 2010, Marder, 2011, Prinz et al., 2004]. In fact, this *degeneracy* of neural circuits is thought to be an important property of biological neural networks [Marder, 2011, Prinz et al., 2004, Marder and Goaillard, 2006]. Our model networks automatically compensate for disturbances by moving their continuously ongoing sampling of network configurations to a new region of the parameter space, as illustrated by the response to the disturbance marked by * in Fig. 3i.

Our theoretical framework is consistent with experimental data that showed drifts of neural representations in motor learning [Rokni et al., 2007]. In that article, a stochastic plasticity model was proposed that is structurally similar to our model. It was shown in computer simulations that a simple feed forward rate-based neural network is able to retain stable functionality despite of such stochastic parameter changes.

The authors hypothesized that this is the case because network parameters move on a sub-manifold in parameter space with constant performance. Our theoretical framework provides a mathematical justification for their hypothesis in general, but also refines these statements. It shows that the network samples network configurations (including the rewiring of connections that was not considered in [Rokni et al., 2007]) from a well-defined distribution. The manifold that is visited during the learning process is given by the high-probability regions of this distribution, but in principle, also sub-optimal regions could be visited. Such sub-optimal regions are however highly unlikely if the parameter space is overcomplete, i.e., if large volumes of the parameter space lead to good performance. Hence, in comparison with [Rokni et al., 2007], this work provides the following features: (a) It provides a quantitative mathematical framework for the qualitative descriptions in [Rokni et al., 2007] that allows a rigorous understanding of the plasticity processes, (b) it includes synaptic rewiring, reproducing experimental data on this topic and providing a hypothesis on its computational role, and (c), it is able to tackle the case of recurrent spiking neural networks as compared to feed forward rate models.

We have shown in Fig. 3f that despite these permanent parameter drifts, the task performance in our model remains stable for many simulated days if reward delivery is stopped. At the same time, the model is also able to continuously adapt to changes in the task Fig. 3c-e. These results suggest that our model keeps a quite good balance between stability and plasticity [Abraham and Robins, 2005], which has already been shown to be one important functional aspect of network rewiring [Fauth et al., 2015]. Furthermore, we have shown in Fig. 3f, that the structural priors over synaptic parameters can be utilized to stabilize synaptic parameters similar to previous models of synaptic consolidation [Fusi et al., 2005, Kirkpatrick et al., 2017]. In addition, more complex prior distributions over multiple synapses could be utilized to model homeostatic processes and clustering of synapses. The latter has been suggested as a mechanism to tackle the stability-plasticity dilemma [Fares and Stepanyants, 2009].

In conclusion the mathematical framework presented in this article provides a principled way of understanding the complex interplay of deterministic and stochastic processes that underlie the implementation of goal-directed learning in neural circuits of the brain. It also offers a solution to the problem how reliable network computations can be achieved with a “dynamic connectome” [Rumpel and Triesch, 2016]. We have argued that the stationary distribution of the high-dimensional parameter vector θ that results from large numbers of local stochastic processes at the synapses provides a time-invariant perspective of salient properties of a network. Standard reward-gated plasticity rules can achieve that this stationary distribution has most of its mass on regions in the parameter space that provide good network performance. The stochastic component of synaptic dynamics can flatten or sharpen the resulting stationary distribution, depending on whether the scaling parameter T (“temperature”) of the stochastic component is larger or smaller than 1. A functional benefit of this stochastic component is that the network keeps exploring its parameter space even after a well-performing region has been found, providing one mechanism to tackle the exploration-exploitation dilemma (see Fig. 2j). This enables the network to migrate quickly and automatically to a better performing region when the network or task changes. We found in the case of the motor learning task of Fig. 3 that a temperature T around 0.15, which lies in the same range as related experimental data (see Fig. 5d), suffices to provide this functionally important compensation capability. The same mathematical framework can also be applied to artificial neural networks, leading to a novel

brain-inspired learning algorithm that uses rewiring to train deep networks under the constraint of very sparse connectivity [Bellec et al., 2017]. (1657 words)

Methods

Probabilistic framework for reward-modulated learning

The classical goal of reinforcement learning is to maximize the expected future discounted reward $\mathcal{V}(\boldsymbol{\theta})$ given by

$$\mathcal{V}(\boldsymbol{\theta}) = \left\langle \int_0^\infty e^{-\frac{\tau}{\tau_e}} r(\tau) d\tau \right\rangle_{p(\mathbf{r}|\boldsymbol{\theta})} . \quad (6)$$

In Eq. (6) we integrate over all future rewards $r(\tau)$, while discounting more remote rewards exponentially with a discount rate τ_e , which for simplicity was set equal to 1 s in this paper. We find (see Eq. (15)) that this time constant τ_e is immediately related to the experimentally studied time window or eligibility trace for the influence of dopamine on synaptic plasticity [Yagishita et al., 2014]. This property is true in general for reward-based learning rules that make use of eligibility traces and is not unique to our model. The expectation in Eq. (6) is taken with respect to the distribution $p(\mathbf{r}|\boldsymbol{\theta})$ over sequences $\mathbf{r} = \{r(\tau), \tau \geq 0\}$ of future rewards that result from the given set of synaptic parameters $\boldsymbol{\theta}$. The stochasticity of the reward sequence \mathbf{r} arises from stochastic network inputs, stochastic network responses, and stochastic reward delivery. The resulting distribution $p(\mathbf{r}|\boldsymbol{\theta})$ of reward sequences \mathbf{r} for the given parameters $\boldsymbol{\theta}$ can also include influences of network initial conditions by assuming some distribution over these initial conditions. Network initial conditions include for example initial values of neuron membrane voltages and refractory states of neurons. The role of initial conditions on network learning is further discussed below when we consider the online learning scenario in *Reward-modulated synaptic plasticity approximates gradient ascent on the expected discounted reward*.

There exists a close relationship between reinforcement learning and Bayesian inference [Vlassis et al., 2012, Rawlik et al., 2013, Botvinick and Toussaint, 2012]. To make this relationship apparent, we define our model for reward-gated network plasticity by introducing a binary random variable v_b that represents the currently expected future discounted reward in a probabilistic manner. The likelihood $p_{\mathcal{N}}(v_b = 1 | \boldsymbol{\theta})$ is determined in this theoretical framework by the expected future discounted reward Eq. (6) that is achieved by a network with parameter set $\boldsymbol{\theta}$ (see e.g., [Rawlik et al., 2013]):

$$p_{\mathcal{N}}(v_b = 1 | \boldsymbol{\theta}) \equiv \frac{1}{\mathcal{Z}_{\mathcal{V}}} \mathcal{V}(\boldsymbol{\theta}) , \quad (7)$$

where $\mathcal{Z}_{\mathcal{V}}$ denotes a constant, that assures that Eq. (7) is a correctly normalized probability distribution. Thus reward-based network optimization can be formalized as maximizing the likelihood $p_{\mathcal{N}}(v_b = 1 | \boldsymbol{\theta})$ with respect to the network configuration $\boldsymbol{\theta}$. Structural constraints can be integrated into a stochastic model for network plasticity through a prior $p_S(\boldsymbol{\theta})$ over network configurations. Hence reward-gated network optimization amounts from a theoretical perspective to learning of the posterior distribution $p^*(\boldsymbol{\theta}|v_b = 1)$, which by Bayes' rule is defined (up to normalization) by $p_S(\boldsymbol{\theta}) \cdot p_{\mathcal{N}}(v_b = 1 | \boldsymbol{\theta})$. Therefore, the learning goal can be formalized in a compact form as evaluating the posterior distribution $p^*(\boldsymbol{\theta}|v_b = 1)$ of network

parameters θ under the constraint that the abstract learning goal $v_b = 1$ is achieved.

More generally, one is often interested in a tempered version of the posterior

$$p_T^*(\theta) \equiv \frac{1}{\mathcal{Z}} p^*(\theta|v_b = 1)^{\frac{1}{T}}, \quad (8)$$

where \mathcal{Z} is a suitable normalization constant and $T > 0$ is the temperature parameter that controls the “sharpness” of $p_T^*(\theta)$. For $T = 1$, $p_T^*(\theta)$ is given by the original posterior, $T < 1$ emphasizes parameter values with high probability in the posterior, while $T > 1$ leads to parameter distributions $p_T^*(\theta)$ which are more uniformly distributed than the posterior.

Analysis of policy sampling

Here we prove that the stochastic parameter dynamics Eq. (5) samples from the tempered posterior distribution $p_T^*(\theta)$ given in Eq. (8). In *Results* we suppressed time-dependencies in order to simplify notation. We reiterate Eq. (3) with explicit time-dependencies of parameters:

$$d\theta_i(t) = \beta \left. \frac{\partial}{\partial \theta_i} \log p^*(\theta|v_b = 1) \right|_{\theta(t)} dt + \sqrt{2\beta T} d\mathcal{W}_i, \quad (9)$$

where the notation $\left. \frac{\partial}{\partial \theta_i} f(\theta) \right|_{\theta(t)}$ denotes the derivative of $f(\theta)$ with respect to θ_i evaluated at the current parameter values $\theta(t)$. By Bayes’ rule, the derivative of the log posterior is the sum of the derivatives of the log prior and the log likelihood:

$$\frac{\partial}{\partial \theta_i} \log p^*(\theta|v_b = 1) = \frac{\partial}{\partial \theta_i} \log p_S(\theta) + \frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(v_b = 1|\theta) = \frac{\partial}{\partial \theta_i} \log p_S(\theta) + \frac{\partial}{\partial \theta_i} \log \mathcal{V}(\theta),$$

which allows us to rewrite Eq. (9) as

$$d\theta_i(t) = \beta \left(\left. \frac{\partial}{\partial \theta_i} \log p_S(\theta) \right|_{\theta(t)} + \left. \frac{\partial}{\partial \theta_i} \log \mathcal{V}(\theta) \right|_{\theta(t)} \right) dt + \sqrt{2\beta T} d\mathcal{W}_i, \quad (10)$$

which is identical to the form Eq. (5), where the contributions of $p_S(\theta)$ and $\mathcal{V}(\theta)$ are given explicitly.

The fundamental property of the synaptic sampling dynamics Eq. (9) is formalized in Theorem 1 and proven below. Before we state the theorem, we briefly discuss its statement in simple terms. Consider some initial parameter setting $\theta(0)$. Over time, the parameters change according to the dynamics (9). Since the dynamics include a noise term, the exact value of the parameters $\theta(t)$ at some time $t > 0$ cannot be determined. However, it is possible to describe the exact distribution of parameters for each time t . We denote this distribution by $p_{\text{FP}}(\theta, t)$, where the “FP” subscript stands for “Fokker-Planck” since the evolution of this distribution is described by the Fokker-Planck equation (11) given below. Note that we make the dependence of this distribution on time explicit in this notation. It can be shown that for the dynamics (11), $p_{\text{FP}}(\theta, t)$ converges to a well-defined and unique *stationary distribution* in the limit of large t . Of practical relevance is the so-called burn-in time after which the distribution of parameters is very close to the stationary distribution. Note that the parameters will continue to change. Nevertheless, at any time

t after the burn in, we can expect the parameter vector $\boldsymbol{\theta}(t)$ to be situated at a particular value with the probability (density) given by the stationary distribution, see Fig. 1d,f. Any distribution that is *invariant* under the parameter dynamics is a stationary distribution. Here, invariance means: when one starts with an invariant distribution over parameters in the Fokker-Planck equation, the dynamics are such that this distribution will be kept forever (we will use this below in the proof of Theorem 1). Theorem 1 states that the parameter dynamics leaves $p_T^*(\boldsymbol{\theta})$ given in Eq. (8) invariant, i.e., it is a stationary distribution of the network parameters. Note that in general, the stationary distribution may not be uniquely defined. That is, it could happen that for two different initial parameter values, the network reaches two different stationary distributions. Theorem 1 further states that for the synaptic sampling dynamics, the stationary distribution is unique, i.e., the distribution $p_T^*(\boldsymbol{\theta})$ is reached from *any* initial parameter setting when the conditions of the theorem apply. We now state Theorem 1 formally. To simplify notation we drop in the following the explicit time dependence of the synaptic parameters $\boldsymbol{\theta}$.

Theorem 1. *Let $p^*(\boldsymbol{\theta} | v_b = 1)$ be a strictly positive, continuous probability distribution over parameters $\boldsymbol{\theta}$, twice continuously differentiable with respect to $\boldsymbol{\theta}$, and let $\beta > 0$. Then the set of stochastic differential equations Eq. (9) leaves the distribution $p_T^*(\boldsymbol{\theta})$ (8) invariant. Furthermore, $p_T^*(\boldsymbol{\theta})$ is the unique stationary distribution of the sampling dynamics.*

Proof. The proof is analogous to the one provided in [Kappel et al., 2015]. The stochastic differential equation Eq. (9) translates into a Fokker-Planck equation [Gardiner, 2004] that describes the evolution of the distribution over parameters $\boldsymbol{\theta}$

$$\frac{\partial}{\partial t} p_{\text{FP}}(\boldsymbol{\theta}, t) = \sum_i -\frac{\partial}{\partial \theta_i} \left(\beta \frac{\partial}{\partial \theta_i} \log p^*(\boldsymbol{\theta} | v_b = 1) \right) p_{\text{FP}}(\boldsymbol{\theta}, t) + \frac{\partial^2}{\partial \theta_i^2} (\beta T p_{\text{FP}}(\boldsymbol{\theta}, t)), \quad (11)$$

where $p_{\text{FP}}(\boldsymbol{\theta}, t)$ denotes the distribution over network parameters at time t . To show that $p_T^*(\boldsymbol{\theta})$ leaves the distribution invariant, we have to show that $\frac{\partial}{\partial t} p_{\text{FP}}(\boldsymbol{\theta}, t) = 0$ (i.e., $p_{\text{FP}}(\boldsymbol{\theta}, t)$ does not change) if we set $p_{\text{FP}}(\boldsymbol{\theta}, t)$ to $p_T^*(\boldsymbol{\theta})$ on the right hand side of Eq. (11). Plugging in the presumed stationary distribution $p_T^*(\boldsymbol{\theta})$ for $p_{\text{FP}}(\boldsymbol{\theta}, t)$ on the right hand side of Eq. (11), one obtains

$$\begin{aligned} \frac{\partial}{\partial t} p_{\text{FP}}(\boldsymbol{\theta}, t) &= \sum_i -\frac{\partial}{\partial \theta_i} \left(\beta \frac{\partial}{\partial \theta_i} \log p^*(\boldsymbol{\theta} | v_b = 1) p_T^*(\boldsymbol{\theta}) \right) + \frac{\partial^2}{\partial \theta_i^2} (\beta T p_T^*(\boldsymbol{\theta})) \\ &= \sum_i -\frac{\partial}{\partial \theta_i} \left(\beta p_T^*(\boldsymbol{\theta}) \frac{\partial}{\partial \theta_i} \log p^*(\boldsymbol{\theta} | v_b = 1) \right) + \frac{\partial}{\partial \theta_i} \left(\beta T \frac{\partial}{\partial \theta_i} p_T^*(\boldsymbol{\theta}) \right) \\ &= \sum_i -\frac{\partial}{\partial \theta_i} \left(\beta p_T^*(\boldsymbol{\theta}) \frac{\partial}{\partial \theta_i} \log p^*(\boldsymbol{\theta} | v_b = 1) \right) + \frac{\partial}{\partial \theta_i} \left(\beta T p_T^*(\boldsymbol{\theta}) \frac{\partial}{\partial \theta_i} \log p_T^*(\boldsymbol{\theta}) \right), \end{aligned}$$

which by inserting $p_T^*(\boldsymbol{\theta}) = \frac{1}{\mathcal{Z}} p^*(\boldsymbol{\theta} | v_b = 1)^{\frac{1}{T}}$, with normalizing constant \mathcal{Z} , becomes

$$\begin{aligned} \frac{\partial}{\partial t} p_{\text{FP}}(\boldsymbol{\theta}, t) &= \frac{1}{\mathcal{Z}} \sum_i -\frac{\partial}{\partial \theta_i} \left(\beta p^*(\boldsymbol{\theta}) \frac{\partial}{\partial \theta_i} \log p^*(\boldsymbol{\theta} | v_b = 1) \right) + \frac{\partial}{\partial \theta_i} \left(\beta T p^*(\boldsymbol{\theta}) \frac{1}{T} \frac{\partial}{\partial \theta_i} \log p^*(\boldsymbol{\theta} | v_b = 1) \right) \\ &= \sum_i 0 = 0. \end{aligned}$$

This proves that $p_T^*(\theta)$ is a stationary distribution of the parameter sampling dynamics Eq. (9). Since β is strictly positive, this stationary distribution is also unique (see Section 3.7.2 in [Gardiner, 2004]).

The unique stationary distribution of Eq. (11) is given by $p_T^*(\theta) = \frac{1}{Z} p^*(\theta|v_b = 1)^{\frac{1}{T}}$, i.e. $p_T^*(\theta)$ is the only solution for which $\frac{\partial}{\partial t} p_{\text{FP}}(\theta, t)$ becomes 0, which completes the proof. \square

The uniqueness of the stationary distribution follows because each parameter setting can be reached from any other parameter setting with non-zero probability (ergodicity). The stochastic process can therefore not get trapped in cycles or absorbed into a subregion of the parameter space. The time spent in a certain region of the parameter space is therefore directly proportional to the probability of that parameter region under the posterior distribution. The proof requires that the posterior distribution is smooth and differentiable with respect to the synaptic parameters. This is not true in general for a spiking neural networks. In our simulations we used a stochastic neuron model (defined in the next section). As the reward landscape in our case is defined by the *expected* discounted reward (see below), a probabilistic network tends to smoothen this landscape and therefore the posterior distribution.

Network model

Plasticity rules for this general framework were derived based on a specific spiking neural network model, which we describe in the following. All reported computer simulations were performed with this network model. We considered a general network scaffold \mathcal{N} of K neurons with potentially asymmetric recurrent connections. Neurons are indexed in an arbitrary order by integers between 1 and K . We denote the output spike train of a neuron k by $z_k(t)$. It is defined as the sum of Dirac delta pulses positioned at the spike times $t_k^{(1)}, t_k^{(2)}, \dots$, i.e., $z_k(t) = \sum_l \delta(t - t_k^{(l)})$. Potential synaptic connections are also indexed in an arbitrary order by integers between 1 and K_{syn} , where K_{syn} denotes the number of potential synaptic connections in the network. We denote by PRE_i and POST_i the index of the pre- and postsynaptic neuron of synapse i , respectively, which unambiguously specifies the connectivity in the network. Further, we define SYN_k to be the index set of synapses that project to neuron k . Note that this indexing scheme allows us to include multiple (potential) synaptic connections between a given pair of neurons. We included this experimentally observed feature of biological neuronal networks in all our simulations. We denote by $w_i(t)$ the synaptic efficacy of the i -th synapse in the network at time t .

Network neurons were modeled by a standard stochastic variant of the spike response model [Gerstner et al., 2014]. In this model, the membrane potential of a neuron k at time t is given by

$$u_k(t) = \sum_{i \in \text{SYN}_k} y_{\text{PRE}_i}(t) w_i(t) + \vartheta_k(t), \quad (12)$$

where $\vartheta_k(t)$ denotes the slowly changing bias potential of neuron k , and $y_{\text{PRE}_i}(t)$ denotes the trace of the (unweighted) postsynaptic potentials (PSPs) that neuron PRE_i leaves in its postsynaptic synapses at time t . More precisely, it is defined as $y_{\text{PRE}_i}(t) = z_{\text{PRE}_i}(t) * \epsilon(t)$ given by spike trains filtered with a PSP kernel of the form $\epsilon(t) = \Theta(t) \frac{\tau_r}{\tau_m - \tau_r} \left(e^{-\frac{t}{\tau_m}} - e^{-\frac{t}{\tau_r}} \right)$, with time constants $\tau_m = 20$ ms and $\tau_r = 2$ ms, if not stated otherwise. Here $*$ denotes convolution and $\Theta(\cdot)$ is the Heaviside step function, i.e. $\Theta(x) = 1$ for $x \geq 0$ and 0 otherwise.

The synaptic weights $w_i(t)$ in Eq. (12) were determined by the synaptic parameters $\theta_i(t)$ through the mapping Eq. (1) for $\theta_i(t) > 0$. Synaptic connections with $\theta_i(t) \leq 0$ were interpreted as not functional (disconnected) and $w_i(t)$ was therefore set to 0 in that case.

The bias potential $\vartheta_k(t)$ in Eq. (12) implements a slow adaptation mechanism of the intrinsic excitability, which ensures that the output rate of each neuron stays near the firing threshold and the neuron maintains responsiveness [Desai et al., 1999, Fan et al., 2005]. We used a simple adaptation mechanism which was updated according to

$$\tau_{\vartheta} \frac{d\vartheta_k(t)}{dt} = \nu_0 - z_k(t), \quad (13)$$

where $\tau_{\vartheta} = 50$ s is the time constant of the adaptation mechanism and $\nu_0 = 5$ Hz is the desired output rate of the neuron. In our simulations, the bias potential $\vartheta_k(t)$ was initialized at -3 and then followed the dynamics given in Eq. (13). We found that this regularization significantly increased the performance and learning speed of our network model. In [Remme and Wadman, 2012] a similar mechanism was proposed to balance activity in networks of excitatory and inhibitory neurons. The regularization used here can be seen as a simplified version of this mechanism that regulates the mean firing rate of each excitatory neuron using a simple linear control loop and thereby stabilizes the output behavior of the network.

We used a simple refractory mechanism for our neuron model. The firing rate, or intensity, of neuron k at time t is defined by the function $f_k(t) = f(u_k(t), \rho_k(t))$, where $\rho_k(t)$ denotes a refractory variable that measures the time elapsed since the last spike of neuron k . We used an exponential dependence between membrane potential and firing rate, such that the instantaneous firing rate of the neuron k at time t can be written as

$$f_k(t) = f(u_k, \rho_k) = \exp(u_k) \Theta(\rho_k - t_{\text{ref}}). \quad (14)$$

Furthermore, we denote by $f_{\text{post}_i}(t)$ the firing rate of the neuron postsynaptic to synapse i . If not stated otherwise we set the refractory time t_{ref} to 5 ms. In addition, a subset of neurons was clamped to some given firing rates (input neurons), such that $f_k(t)$ of these input neurons was given by an arbitrary function. We denote the spike train from these neurons by $\mathbf{x}(t)$, the network input.

Synaptic dynamics for the reward-based synaptic sampling model

Here, we provide additional details on how the synaptic parameter dynamics Eq. (5) was computed. We will first provide an intuitive interpretation of the equations and then provide a detailed derivation in the next section. The second term $\frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta})$ of Eq. (5) denotes the gradient of the expected future discounted reward Eq. (6). In general, optimizing this function has to account for the case where rewards are provided after some delay period. It is well known that this *distal reward problem* can be solved using plasticity mechanisms that make use of eligibility traces in the synapses that are triggered by near coincident spike patterns, but their consolidation into the synaptic weights is delayed and gated by the reward signal $r(t)$ [Sutton and Barto, 1998, Izhikevich, 2007]. The theoretically optimal shape for these eligibility traces can be derived using the reinforcement learning theory and depends on the choice of network model. For the spiking neural network model described above, the gradient $\frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta})$ can be estimated through a

plasticity mechanism that uses an eligibility trace $e_i(t)$ in each synapse i which gets updated according to

$$\frac{de_i(t)}{dt} = -\frac{1}{\tau_e}e_i(t) + w_i(t)y_{\text{PRE}_i}(t)(z_{\text{POST}_i}(t) - f_{\text{POST}_i}(t)), \quad (15)$$

where $\tau_e = 1$ s is the time constant of the eligibility trace. Recall that PRE_i denotes the index of the presynaptic neuron and POST_i the index of the postsynaptic neuron for synapse i . In Eq. (15) $z_{\text{POST}_i}(t)$ denotes the postsynaptic spike train, $f_{\text{POST}_i}(t)$ denotes the instantaneous firing rate (Eq. (14)) of the postsynaptic neuron and $w_i(t)y_{\text{PRE}_i}(t)$ denotes the postsynaptic potential under synapse i .

The last term of Eq. (15) shares salient properties with standard STDP learning rules, since plasticity is enabled by the presynaptic term $y_{\text{PRE}_i}(t)$ and gated by the postsynaptic term $(z_{\text{POST}_i}(t) - f_{\text{POST}_i}(t))$ (see [Pfister et al., 2006]). The latter term also regularizes the plasticity mechanism such that synapses stop growing if the firing probability $f_{\text{POST}_i}(t)$ of the postsynaptic neuron is already close to one.

The eligibility trace Eq. (15) is multiplied by the reward $r(t)$ and integrated in each synapse i using a second dynamic variable

$$\frac{dg_i(t)}{dt} = -\frac{1}{\tau_g}g_i(t) + \left(\frac{r(t)}{\hat{r}(t)} + \alpha\right)e_i(t), \quad (16)$$

where $\hat{r}(t)$ is a low-pass filtered version of $r(t)$ (described below). The variable $g_i(t)$ combines the eligibility trace and the reward, and averages over the time scale τ_g . α is a constant offset on the reward signal. This parameter can be set to an arbitrary value without changing the stationary dynamics of the model (see next section). In our simulations, this offset α was chosen slightly above 0 ($\alpha = 0.02$) such that small parameter changes were also present without any reward, as observed in [Yagishita et al., 2014]. Furthermore, α does not have to be chosen constant. E.g. this term can be used to incorporate predictions about the reward outcome by setting α to the negative of output of a *critic* network that learns to predict future reward. This approach has been previously studied in [Frémaux et al., 2013] to model experimental data of [Schultz et al., 1997, Schultz, 2002]. In the experiment Fig. 4d we included in Eq. (16) the scaling constant c_r to modulate the reward term $\left(\frac{r(t)}{\hat{r}(t)} + \alpha\right)$.

In the next section we show that $g_i(t)$ approximates the gradient of the expected future reward with respect to the synaptic parameter. In our simulations we found that incorporating the low-pass filtered eligibility traces (Eq. (16)) into the synaptic parameters works significantly better than using the eligibility traces directly for weight updates, although the latter approach was taken in a number of previous studies (see e.g. [Pfister et al., 2006, Legenstein et al., 2008, Urbanczik and Senn, 2009]). Eq. (16) essentially combines the eligibility trace with the reward and smoothens the resulting trace with a low-pass filter with time constant τ_g . This time constant has been chosen to be in the order spontaneous decay of disinhibited CaMKII in the synapse which is closely related to spine enlargement in the dopamine-gated STDP protocol of [Yagishita et al., 2014] (c.f. their Fig. 3F and Fig. 4C).

$\hat{r}(t)$ in Eq. (16) is a low-pass filtered version of $r(t)$ that scales the synaptic updates. It was implemented through $\tau_a \frac{d\hat{r}(t)}{dt} = -\hat{r}(t) + r(t)$, with $\tau_a = 50$ s. The value of τ_a has been chosen to equal τ_g based on theoretical considerations (see *Online learning*). This scaling of the reward signal has the following effect. If the current reward $r(t)$ exceeds the average reward $\hat{r}(t)$, the effect of the neuromodulatory signal $r(t)$ will be greater than 1. On the other hand, if the current reward is below average synaptic updates will be

weighted by a term significantly lower than 1. Therefore, parameter updates are preferred for which the current reward signal exceeds the average.

Similar plasticity rules with eligibility traces in spiking neural networks have previously been proposed by several authors [Seung, 2003, Xie and Seung, 2004, Izhikevich, 2007, Pfister et al., 2006, Florian, 2007, Legenstein et al., 2008, Urbanczik and Senn, 2009, Frémaux et al., 2010, Frémaux et al., 2013]. In [Frémaux et al., 2013] also a method to estimate the neural firing rate $f_{\text{post}_i}(t)$ from back-propagating action potentials in the synapses has been proposed. The main difference to these previous approaches is that the activity-dependent last term in Eq. (15) is scaled by the current synaptic weight $w_i(t)$. This weight-dependence of the update equations induces multiplicative synaptic dynamics and is a consequence of the exponential mapping Eq. (1) (see derivation in the next section). This is an important property for a network model that includes rewiring. Note, that for retracted synapses ($w_i(t) = 0$), both $e_i(t)$ and $g_i(t)$ decay to zero (within few minutes in our simulations). Therefore, we find that the dynamics of retracted synapses is only driven by the first (prior) and last (random fluctuations) term of Eq. (5) and are independent from the network activity. Thus, retracted synapses spontaneously reappear also in the absence of reward after a random amount of time.

The first term in Eq. (5) is the gradient of the prior distribution. We used a prior distribution that pulls the synaptic parameters towards $\theta_i(t) = 0$ such that unused synapses tend to disappear and new synapses are permanently formed. If not stated otherwise we used independent Gaussian priors for the synaptic parameters

$$p_S(\boldsymbol{\theta}) = \prod_i p_S(\theta_i(t)) , \quad \text{with} \quad p_S(\theta_i(t)) = \frac{1}{\sigma \sqrt{2\pi}} \exp\left(-\frac{(\theta_i(t) - \mu)^2}{2\sigma^2}\right) ,$$

where σ is the standard deviation of the prior distribution. Using this, we find that the contribution of the prior to the online parameter update equation is given by

$$\frac{\partial}{\partial \theta_i} \log p_S(\boldsymbol{\theta}) = \frac{1}{\sigma^2} (\mu - \theta_i(t)) . \quad (17)$$

Finally by plugging Eq. (17) and (16) into Eq. (5) the synaptic parameter changes at time t are given by

$$d\theta_i(t) = \beta \left(\frac{1}{\sigma^2} (\mu - \theta_i(t)) + g_i(t) \right) dt + \sqrt{2\beta T} d\mathcal{W}_i . \quad (18)$$

We tuned the parameters of the prior distribution by hand to achieve good results on the task presented in Fig. 3 (see Fig. 4 for a comparison of different prior distributions). These parameters were given by $\sigma = 2$ and $\mu = 0$ and were used throughout all experiments if not stated otherwise. By inspecting Eq. (18) it becomes immediately clear that the parameter dynamics follow an Ornstein-Uhlenbeck process if the activity-dependent second term is inactive (in the absence of reward), i.e. if $g_i(t) = 0$. In this case the dynamics are given by the deterministic drift towards the mean value μ and the stochastic diffusion fueled by the Wiener process \mathcal{W}_i . The temperature T and the standard deviation σ scale the contribution of these two forces.

Reward-modulated synaptic plasticity approximates gradient ascent on the expected discounted reward

We first consider a theoretical setup where the network is operated in arbitrarily long episodes such that in each episode a reward sequence \mathbf{r} is encountered. The reward sequence \mathbf{r} can be any discrete or real-valued function that is positive and bounded. The episodic scenario is useful to derive exact batch parameter update rules, from which we will then deduce online learning rules. Due to stochastic network inputs, stochastic network responses, and stochastic reward delivery, the reward sequence \mathbf{r} is stochastic.

The classical goal of reinforcement learning is to maximize the function $\mathcal{V}(\boldsymbol{\theta})$ of discounted expected rewards Eq. (6). Policy gradient algorithms perform gradient ascent on $\mathcal{V}(\boldsymbol{\theta})$ by changing each parameter θ_i in the direction of the gradient $\partial \log \mathcal{V}(\boldsymbol{\theta}) / \partial \theta_i$. Here, we show that the parameter dynamics Eq. (15), (16) approximate this gradient, i.e., $g_i(t) \approx \partial \log \mathcal{V}(\boldsymbol{\theta}) / \partial \theta_i$.

It is natural to assume that the reward signal $r(\tau)$ only depends indirectly on the parameters $\boldsymbol{\theta}$, through the history of network spikes $z_k(\tau)$ up to time τ , which we write as $\mathbf{z}(\tau) = \{z_k(s) \mid 0 \leq s < \tau, 1 \leq k \leq K\}$, i.e., $p_{\mathcal{N}}(r(t), \mathbf{z}(t) \mid \boldsymbol{\theta}) = p(r(t) \mid \mathbf{z}(t)) p_{\mathcal{N}}(\mathbf{z}(t) \mid \boldsymbol{\theta})$. We can first expand the expectation $\langle \cdot \rangle_{p(\mathbf{r} \mid \boldsymbol{\theta})}$ in Eq. (6) to be taken over the joint distribution $p(\mathbf{r}, \mathbf{z} \mid \boldsymbol{\theta})$ over reward sequences \mathbf{r} and network trajectories \mathbf{z} . The derivative

$$\frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta}) = \frac{1}{\mathcal{V}(\boldsymbol{\theta})} \frac{\partial}{\partial \theta_i} \mathcal{V}(\boldsymbol{\theta}) = \frac{1}{\mathcal{V}(\boldsymbol{\theta})} \frac{\partial}{\partial \theta_i} \left\langle \int_0^\infty e^{-\frac{\tau}{\tau_e}} r(\tau) d\tau \right\rangle_{p(\mathbf{r}, \mathbf{z} \mid \boldsymbol{\theta})} \quad (19)$$

can be evaluated using the well-known identity $\frac{\partial}{\partial x} \langle f(a) \rangle_{p(a \mid x)} = \langle f(a) \frac{\partial}{\partial x} \log p(a \mid x) \rangle_{p(a \mid x)}$:

$$\begin{aligned} \frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta}) &= \frac{1}{\mathcal{V}(\boldsymbol{\theta})} \left\langle \int_0^\infty e^{-\frac{\tau}{\tau_e}} r(\tau) \frac{\partial}{\partial \theta_i} \log p(r(\tau), \mathbf{z}(\tau) \mid \boldsymbol{\theta}) d\tau \right\rangle_{p(\mathbf{r}, \mathbf{z} \mid \boldsymbol{\theta})} \\ &= \frac{1}{\mathcal{V}(\boldsymbol{\theta})} \left\langle \int_0^\infty e^{-\frac{\tau}{\tau_e}} r(\tau) \frac{\partial}{\partial \theta_i} (\log p(r(\tau) \mid \mathbf{z}(\tau)) + \log p_{\mathcal{N}}(\mathbf{z}(\tau) \mid \boldsymbol{\theta})) d\tau \right\rangle_{p(\mathbf{r}, \mathbf{z} \mid \boldsymbol{\theta})} \\ &= \left\langle \int_0^\infty e^{-\frac{\tau}{\tau_e}} \frac{r(\tau)}{\mathcal{V}(\boldsymbol{\theta})} \frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(\mathbf{z}(\tau) \mid \boldsymbol{\theta}) d\tau \right\rangle_{p(\mathbf{r}, \mathbf{z} \mid \boldsymbol{\theta})}. \end{aligned} \quad (20)$$

Here, $p_{\mathcal{N}}(\mathbf{z}(\tau) \mid \boldsymbol{\theta})$ is the probability of observing the spike train $\mathbf{z}(\tau)$ in the time interval 0 to τ . For the definition of the network \mathcal{N} given above, the gradient $\frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(\mathbf{z}(\tau) \mid \boldsymbol{\theta})$ of this distribution can be directly evaluated. Using Eq. (12) and (1) we get [Pfister et al., 2006]

$$\begin{aligned} \frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(\mathbf{z}(\tau) \mid \boldsymbol{\theta}) &= \frac{\partial w_i}{\partial \theta_i} \frac{\partial}{\partial w_i} \int_0^\tau z_{\text{POST}_i}(s) \log(f_{\text{POST}_i}(s)) - f_{\text{POST}_i}(s) ds \\ &\approx \int_0^\tau w_i y_{\text{PRE}_i}(s) (z_{\text{POST}_i}(s) - f_{\text{POST}_i}(s)) ds, \end{aligned} \quad (21)$$

where we have used that by construction only the rate function $f_{\text{POST}_i}(s)$ depends on the parameter θ_i . If one discretizes time and assumes that rewards and parameter updates are only realized at the end of each episode, the REINFORCE rule is recovered [Williams, 1992].

In Eq. (21) we used the approximation $\frac{\partial w_i}{\partial \theta_i} \approx w_i$. This expression ignores the discontinuity of Eq. (1)

at $\theta_i = 0$, where the function is not differentiable. In practice we found that this approximation is quite accurate if θ_0 is large enough such that $\exp(\theta_i - \theta_0)$ is close to zero (which is the case for $\theta_0 = 3$ in our simulation). In control experiments we also used a smooth function $w_i = \exp(\theta_i - \theta_0)$ (without the jump at $\theta_i = 0$) for which Eq. (21) is exact, and found that this yields results that are not significantly different from the ones that use the mapping Eq. (1).

Online learning

Eq. (20) defines a batch learning rule with an average taken over learning episodes where in each episode network responses and rewards are drawn according to the distribution $p(\mathbf{r}, \mathbf{z}|\boldsymbol{\theta})$. In a biological setting, there are typically no clear episodes but rather a continuous stream of network inputs and rewards and parameter updates are performed continuously (i.e., learning is online). The analysis of online policy gradient learning is far more complicated than the batch scenario, and typically only approximate results can be obtained that however perform well in practice, see e.g., [Seung, 2003, Xie and Seung, 2004] for discussions.

In order to arrive at an online learning rule for this scenario, we consider an estimator of Eq. (20) that approximates its value at each time $t > \tau_g$ based on the recent network activity and rewards during time $[t - \tau_g, t]$ for some suitable $\tau_g > 0$. We denote the estimator at time t by $G_i(t)$ where we want $G_i(t) \approx \frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta})$ for all $t > \tau_g$. To arrive at such an estimator, we approximate the average over episodes in Eq. (20) by an average over time where each time point is treated as the start of an episode. The average is taken over a long sequence of network activity that starts at time t and ends at time $t + \tau_g$. Here, one systematic difference to the batch setup is that one cannot guarantee a time-invariant distribution over initial network conditions as we did there since those will depend on the current network parameter setting. However, under the assumption that the influence of initial conditions (such as initial membrane potentials and refractory states) decays quickly compared to the time scale of the environmental dynamics, it is reasonable to assume that the induced error is negligible. We thus rewrite Eq. (20) in the form (we use the abbreviation $PSP_i(s) = w_i(s) y_{PRE_i}(s)$).

$$\frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta}) \approx G_i(t) = \frac{1}{\tau_g} \int_t^{t+\tau_g} \int_{\zeta}^{t+\tau_g} e^{-\frac{\tau-\zeta}{\tau_e}} \frac{r(\tau)}{\mathcal{V}(\boldsymbol{\theta})} \int_{\zeta}^{\tau} PSP_i(s) (z_{POST_i}(s) - f_{POST_i}(s)) ds d\tau d\zeta ,$$

where τ_g is the length of the sequence of network activity over which the empirical expectation is taken. Finally, we can combine the second and third integral into a single one, rearrange terms and substitute s and τ so that integrals run into the past rather than the future, to obtain

$$G_i(t) \approx \frac{1}{\tau_g} \int_{t-\tau_g}^t \frac{r(\tau)}{\mathcal{V}(\boldsymbol{\theta})} \int_0^{\tau} e^{-\frac{s}{\tau_e}} PSP_i(\tau - s) (z_{POST_i}(\tau - s) - f_{POST_i}(\tau - s)) ds d\tau , \quad (22)$$

We now discuss the relationship between $G_i(t)$ and Eq. (15), (16) to show that the latter equations approximate $G_i(t)$. Solving Eq. (15) with zero initial condition $e_i(0) = 0$ yields

$$e_i(t) = \int_0^t e^{-\frac{s}{\tau_e}} PSP_i(t - s) (z_{POST_i}(t - s) - f_{POST_i}(t - s)) ds . \quad (23)$$

This corresponds to the inner integral in Eq. (22) and we can write

$$G_i(t) \approx \frac{1}{\tau_g} \int_{t-\tau_g}^t \frac{r(\tau)}{\mathcal{V}(\boldsymbol{\theta})} e_i(\tau) d\tau = \left\langle \frac{r(t)}{\mathcal{V}(\boldsymbol{\theta})} e_i(t) \right\rangle_{\tau_g} \approx \left\langle \frac{r(t)}{\hat{r}(t)} e_i(t) \right\rangle_{\tau_g}, \quad (24)$$

where $\langle \cdot \rangle_{\tau_g}$ denotes the temporal average from $t - \tau_g$ to t and $\hat{r}(t)$ estimates the expected discounted reward through a slow temporal average.

Finally, we observe that any constant α can be added to $r(\tau)/\mathcal{V}(\boldsymbol{\theta})$ in Eq. (20) since

$$\left\langle \int_0^\infty e^{-\frac{\tau}{\tau_e}} \alpha \frac{\partial}{\partial \theta_i} \log p_{\mathcal{N}}(\mathbf{z}(\tau) | \boldsymbol{\theta}) d\tau \right\rangle_{p(\mathbf{r}, \mathbf{z} | \boldsymbol{\theta})} = 0 \quad (25)$$

for any constant α (cf. [Williams, 1992, Urbanczik and Senn, 2009]).

Hence, we have $G_i(t) \approx \left\langle \left(\frac{r(t)}{\hat{r}(t)} + \alpha \right) e_i(t) \right\rangle_{\tau_g}$. Eq. (16) implements this in the form of a running average and hence $g_i(t) \approx G_i(t) \approx \frac{\partial}{\partial \theta_i} \log \mathcal{V}(\boldsymbol{\theta})$ for $t > \tau_g$. Note that this result assumes that the parameters $\boldsymbol{\theta}$ change slowly on the time-scale of τ_g and τ_g has to be chosen significantly longer than the time constant of the eligibility trace τ_e such that the estimator works reliably, so we require $\tau_e < \tau_g < \frac{1}{\beta}$. The time constant τ_a to estimate the average reward $\mathcal{V}(\boldsymbol{\theta})$ through $\tau_a \frac{d\hat{r}(t)}{dt} = -\hat{r}(t) + r(t)$ should be on the same order as the time constant τ_g for estimating the gradient. We selected both to be 50 s in our simulations. Simulations using the batch model outlined above and the online learning model showed qualitatively the same behavior for the parameters used in our experiments (data not shown).

Simulation details

Simulations were performed with NEST [Gewaltig and Diesmann, 2007] using an in-house implementation of the synaptic sampling model; additional tests were run in Matlab R2011b (Mathworks). The code/software described in the paper is freely available online at [Kappel et al., 2017]. The differential equations of the neuron and synapse models were approximated using the Euler method, with fixed time steps $\Delta t = 1$ ms. All network variables were updated based on this time grid, except for the synaptic parameters $\theta_i(t)$ according to Eq. (18) which were updated only every 100 ms to reduce the computation time. Control experiments with $\Delta t = 0.1$ ms, and 1 ms update steps for all synaptic parameters showed no significant differences. If not stated otherwise synaptic parameters were initially drawn from a Gaussian distribution with $\mu = -0.5$ and $\sigma = 0.5$ and the temperature was set to $T = 0.1$. The learning rate for the synaptic dynamics was chosen to be $\beta = 10^{-5}$ and synaptic delays were 1 ms. Synaptic parameter changes were clipped at $\pm 4 \times 10^{-4}$ and synaptic parameters θ_i were not allowed to exceed the interval $[-2, 5]$ for the sake of numerical stability.

Details to: Task-dependent routing of synaptic connections through the interaction of stochastic spine dynamics with rewards

The number of potential excitatory synaptic connections between each pair of input and MSN neurons was initially drawn from a Binomial distribution ($p = 0.5$, $n = 10$). The connections then followed the reward-based synaptic sampling dynamics Eq. (5) as described above. Lateral inhibitory connections were fixed

and thus not subject to learning. These connections between MSN neurons were drawn from a Bernoulli distribution with $p = 0.5$ and synaptic weights were drawn from a Gaussian distribution with $\mu = -1$ and $\sigma = 0.2$, truncated at zero. Two subsets of ten neurons were connected to either one of the targets T_1 or T_2 .

To generate the input patterns we adapted the method from [Kappel et al., 2015]. The inputs were representations of a simple symbolic environment, realized by Poisson spike trains that encoded sensory experiences P_1 or P_2 . The 200 input neurons were assigned to Gaussian tuning curves ($\sigma = 0.2$) with centers independently and equally scattered over the unit cube. The sensory experiences P_1 and P_2 were represented by two different, randomly selected points in this 3-dimensional space. The stimulus positions were overlaid with small-amplitude jitter ($\sigma = 0.05$). For each sensory experience the firing rate of an individual input neuron was given by the support of the sensory experience under the input neuron’s tuning curve (maximum firing rate was 60 Hz). An additional offset of 2 Hz background noise was added. The lengths of the spike patterns were uniformly drawn from the interval [750 ms, 1500 ms]. The spike patterns were alternated with time windows (durations uniformly drawn from the interval [1000 ms, 2000 ms]) during which only background noise of 2 Hz was presented.

The network was rewarded if the assembly associated to the current sensory experience fired stronger than the other assembly. More precisely, we used a sliding window of 500 ms length to estimate the current output rate of the neural assemblies. Let $\hat{\nu}_1(t)$ and $\hat{\nu}_2(t)$ denote the estimated output rates of neural pools projecting to T_1 and T_2 , respectively, at time t and let $I(t)$ be a function that indicates the identity of the input pattern at time t , i.e. $I(t) = 1$ if pattern P_1 was present and $I(t) = -1$ if pattern P_2 was present. If $I(t)(\hat{\nu}_1(t) - \hat{\nu}_2(t)) < 0$ the reward was set to $r(t) = 0$. Otherwise the reward signal was given by $r(t) = S(\frac{1}{5}(I(t)\hat{\nu}_1(t) - I(t)\hat{\nu}_2(t) - \nu_0))$, where $\nu_0 = 25$ Hz is a soft firing threshold and $S(\cdot)$ denotes the logistic sigmoid function. The reward was recomputed every 10 ms. During the presentation of the background patterns no reward was delivered.

In Fig. 2d,e we tested our reward-gated synaptic plasticity mechanism with the reward-modulated STDP pairing protocol reported in [Yagishita et al., 2014]. We applied the STDP protocol to 50 synapses and reported mean and s.e.m values of synaptic weight changes in Fig. 2d,e. Briefly, we presented 15 pre/post pairings; one per 10 seconds. In each pre/post pairing 10 presynaptic spikes were presented at a rate of 10 Hz. Each presynaptic spike was followed ($\Delta t = 10$ ms) by a brief postsynaptic burst of 3 spikes (100 Hz). The total duration of one pairing was thus 1 s indicated by the gray shaded rectangle in Fig. 2e. During the pairings the membrane potential was set to $u(t) = -2.4$ and Eq. (14),(15), (16) and (18) solved for each synapse. Reward was delivered here in the form of a rectangular-shaped wave of constant amplitude 1 and duration 300 ms to mimic puff application of dopamine. Rewards were delivered for each pre/post pairing and reward delays were relative to the onset of the STDP pairings. The time constants τ_e and τ_g , the reward offset α and the temperature T of the synapse model were chosen to qualitatively match the results of Fig. 1 and Fig. 4 of [Yagishita et al., 2014] (see Tab. 1). The value of τ_a for the estimation of the average reward has been chosen to equal τ_g based on theoretical considerations (see *Online learning*). We found that the parameters of the prior had relatively small effect on the synaptic dynamics on timescales of one hour.

Synaptic parameter changes in Fig. 2g were measured by taking snapshots of the synaptic parameter

vectors every 4 minutes. Parameter changes were measured in terms of the Euclidean norm of the difference between two successively recorded vectors. The values were then normalized by the maximum value of the whole experiment and averages over 5 trials were reported.

To generate the dPCA projection of the synaptic parameters in Fig. 2j, we adopted the methods of [Kobak et al., 2016]. We randomly selected a subset of 500 synaptic parameters to compute the projection. We sorted the parameter vectors by the average reward achieved over a time window of 10 minutes and binned them into 10 equally spaced bins. The dPCA algorithm was then applied on this dataset to yield the projection matrix and estimated fractions of explained reward variance. The projection matrix was then applied to the whole trajectory of network parameters and the first 2 components were plotted. The trajectory was projected onto the estimated expected reward surface based on the binned parameter vectors.

Details to: A model for task-dependent self-configuration of a recurrent network of excitatory and inhibitory spiking neurons

Neuron and synapse parameters were as reported above, except for the inhibitory neurons for which we used faster dynamics with a refractory time $t_{\text{ref}} = 2$ ms and time constants $\tau_m = 10$ ms and $\tau_r = 1$ ms for the PSP kernel. The network connectivity between excitatory and inhibitory neurons was as suggested in [Avermann et al., 2012]. Excitatory (pools D, U and hidden) and inhibitory neurons were randomly connected with connection probabilities given in Table 2 in [Avermann et al., 2012]. Connections include lateral inhibition between excitatory and inhibitory neurons. The connectivity to and from inhibitory neurons was kept fixed throughout the simulation (not subject to synaptic plasticity or rewiring). The connection probability from excitatory to inhibitory neurons was given by 0.575. The synaptic weights were drawn from a Gaussian distribution (truncated at zero) with $\mu = 0.5$ and $\sigma = 0.1$. Inhibitory neurons were connected to their targets with probability 0.6 (to excitatory neurons) and 0.55 (to inhibitory neurons) and the synaptic weights were drawn from a truncated normal distribution with $\mu = -1$ and $\sigma = 0.2$. The number of potential excitatory synaptic connections between each pair of excitatory neurons was drawn from a Binomial distribution ($p = 0.5$, $n = 10$). These connections were subject to the reward-based synaptic sampling and rewiring described above. In the resulting network scaffold around 49% of connections consisted of multiple synapses.

To infer the lever position from the network activity, we weighted spikes from the neuron pool D with -1 and spikes from U with $+1$, summed them and then filtered them with a long PSP kernel with $\tau_r = 50$ ms (rise) and $\tau_m = 500$ ms (decay). The cue input pattern was realized by the same method that was used to generate the patterns P_1 and P_2 outlined above. If a trial was completed successfully the reward signal $r(t)$ was set to 1 for 400 ms and was 0 otherwise. After each trial a short holding phase was inserted during which the input neurons were set to 2 Hz background noise. The lengths of these holding phases were uniformly drawn from the interval [1 s, 2 s]. At the time points marked by *, the reward policy was changed by switching the decoding functions of the neural pools D and U and by randomly re-generating the input cue pattern.

To identify the movement onset times in Fig. 3d we adapted the method from [Peters et al., 2014]. Lever movements were recorded at a sampling rate of 5 ms. Lever velocities were estimated by taking the

difference between subsequent time steps and filtering them with a moving average filter of 5 time steps length. A Hilbert transform was applied to compute the envelope of the lever velocities. The movement onset time for each trial was then defined as the time point where the estimated lever velocity exceeded a threshold of 1.5 in the upward movement direction. If this value was never reached throughout the whole trial the time point of maximum velocity was used (most cases at learning onset).

The trial-averaged activity traces in Fig. 3d were generated by filtering the spiking activity of the network with a Gaussian kernel with $\sigma = 75$ ms. The activity traces were aligned with the movement onset times (indicated by black arrows in Fig. 3d) and averaged across 100 trials. The resulting activity traces were then normalized by the neuron's mean activity over all trials and values below the mean were clipped. The resulting activity traces were normalized to the unit interval.

Turnover statistics of synaptic connections in Fig. 3e were measured as follows. The synaptic parameters were recorded in intervals of 2 hours. The number of synapses that appeared (crossed the threshold of $\theta_i(t) = 0$ from below) or disappeared (crossed $\theta_i(t) = 0$ from above) between two measurements were counted and the total number was reported as turnover rate.

For the consolidation mechanism in Fig. 3f we used a modified version of the algorithm where we introduced for each synaptic parameter θ_i an independent mean μ_i for the prior distribution $p_S(\boldsymbol{\theta})$. After 4 simulated days we set μ_i to the current value of θ_i for each synaptic parameter and the standard deviation σ was set to 0.05. Simulation of the synaptic parameter dynamics was then continued for 10 subsequent days.

For the approximation of simulating retracted potential synaptic connections in Fig. 3c,g we paused evaluation of the SDE (5) for $\theta_i \leq 0$. Instead, synaptic parameters of retracted connections were randomly set to values above zero after random waiting times drawn from an Exponential distribution with a mean of 12 hours. When a connection became functional at time t we set $\theta_i(t) = 10^{-5}$ and reset the eligibility trace $e_i(t)$ and gradient estimator $g_i(t)$ to zero and then continued the synaptic dynamics according to (5). Histograms in Fig. 3f were computed over bins of 2 hours width.

In Fig. 6 we further analyzed the trial-averaged activity at three different time points (18 h, 19 h and 20 h) where the performance was stable (see Fig. 3c). Drifts of neural codes on fast time scales could also be observed during this phase of the experiment.

Since the 3D illustration of the PCA projection in Fig. 3i is ambiguous, the corresponding 2D projections are shown in Fig. 7. The projection to the first 2 components (pc1 and pc2) show the migration of synaptic parameters to a new region after the task change. The first 3 principal components explain 82% of the total variance in the parameter dynamics.

Details to: Compensation for network perturbations

The black curve in Fig. 3c shows the learning curve of a network for which rewiring was disabled after the task change at 24 h. Here, synaptic parameters were not allowed to cross the threshold at $\theta_i = 0$ and thus could not change sign after 24 h. Apart from this modification the synaptic dynamics evolved according to Eq. (18) as above with $T = 0.1$.

For the analysis of synaptic turnover in Fig. 3g we recorded the synaptic parameters at $t_1 = 24$ h and $t_2 = 48$ h. We then classified each potential synaptic connection i into one of four classes, stable non-

functional: $(\theta_i(t_1) \leq 0) \wedge (\theta_i(t_2) \leq 0)$, transient decaying: $(\theta_i(t_1) > 0) \wedge (\theta_i(t_2) \leq 0)$, transient emerging: $(\theta_i(t_1) \leq 0) \wedge (\theta_i(t_2) > 0)$ and stable functional: $(\theta_i(t_1) > 0) \wedge (\theta_i(t_2) > 0)$.

In Fig. 3h we randomly selected 5% of the synaptic parameters θ_i and recorded their traces over a learning experiment of 48 hours (1 sample per minute). The principal component analysis (PCA) was then computed over these traces, treating the parameter vectors at each time point as one data sample. The high-dimensional trace was then projected to the first three principal components in Fig. 3h, and colored according to the average movement completion time that was acquired by the network at the corresponding time points.

Details to: Relative contribution of spontaneous and activity-dependent processes to synaptic plasticity

Synaptic weights in Fig. 5a,b were recorded in intervals of 10 minutes. We selected all pairs of synapses with common pre- and postsynaptic neurons as CI synapses and synapse pairs with the same post- but not the same presynaptic neuron as non-CI synapses. In Fig. 5d-f we took a snapshot of the synaptic weights after 48 hours of learning and computed the Pearson correlation of all CI and non-CI pairs for random subsets of around 5000 pairs. Data for 100 randomly chosen CI synapse pairs are plotted of Fig. 5e.

In Fig. 5f we analyzed the contribution of activity-dependent and spontaneous processes in our model. [Dvorkin and Ziv, 2016] reported that a certain degree of the stochasticity in their results could be attributed to their experimental setup. The maximum detectable correlation coefficient was limited to $0.76 - 0.78$, due to the variability of light fluorescence intensities which were used to estimate the sizes of postsynaptic densities. Since in our computer simulations we could directly read out values of the synaptic parameters we were not required to correct our results for noise sources in the experimental procedure (see p. 16ff and equations on p. 18 of [Dvorkin and Ziv, 2016]). This is also reflected in our data by the fact that we got a correlation coefficient that was close to 1.0 in the case $T = 0$ (see Fig. 5d). Following the procedure of [Dvorkin and Ziv, 2016] we estimated in our model the contributions of activity history dependent and spontaneous synapse-autonomous processes as in Fig. 8E of [Dvorkin and Ziv, 2016]. Using the assumption of zero measurement error and thus a theoretically achievable maximum correlation coefficient of $r = 1.0$. The Pearson correlation of CI-synapses was given by 0.46 ± 0.034 and that of non-CI synapses by 0.08 ± 0.015 . Therefore, we estimated the fraction of contributions of specific activity histories to synaptic changes (for $T = 0.15$) as $0.46 - 0.08 = 0.38$ and of spontaneous synapse-autonomous processes as $1.0 - 0.46 = 0.54$ [Dvorkin and Ziv, 2016]. The remaining 8% (measured correlation between non-CI synapses) resulted from processes that were not specific to presynaptic input, but specific to the activity of the postsynaptic neuron (neuron-wide processes).

Figure legends

1 Illustration of the theoretical framework.

(a) A neural network scaffold \mathcal{N} of excitatory (blue triangles) and inhibitory (purple circles) neurons. Potential synaptic connections (broken blue arrows) of only two excitatory neurons are shown to keep the figure uncluttered. Synaptic connections (black connections) from and to inhibitory neurons are assumed to be fixed for simplicity. (b) A reward landscape for two parameters $\theta = \{\theta_1, \theta_2\}$ with several local optima. Z-amplitude and color indicate the expected reward $\mathcal{V}(\theta)$ for given parameters θ (X-Y plane). (c) Example prior that prefers small values for θ_1 and θ_2 . (d) The posterior distribution $p^*(\theta)$ that results as product of the prior from panel c and the expected discounted reward of panel b. (e) Illustration of the dynamic forces (plasticity rule Eq. (5)) that act on θ in each sampling step $d\theta$ (black) while sampling from the posterior distribution. The deterministic term (red), which consists of the first two terms (prior and reward expectation) in Eq. (5), is directed to the next local maximum of the posterior. The stochastic term $d\mathcal{W}$ (green) of Eq. (5) has a random direction. (f) A single trajectory of policy sampling from the posterior distribution of panel d under Eq. (5), starting at the black dot. The parameter vector θ fluctuates between different solutions, and moves primarily along the task-irrelevant dimension θ_2 .

2 Reward-based routing of input patterns.

(a) Illustration of the network scaffold. A population of 20 model MSNs (blue) receives input from 200 excitatory input neurons (green) that model cortical neurons. Potential synaptic connections between these 2 populations of neurons were subject to reward-based synaptic sampling. In addition, fixed lateral connections provided recurrent inhibitory input to the MSNs. The MSNs were divided into two groups, each projecting exclusively to one of two target areas T_1 and T_2 . Reward was delivered whenever the network managed to route an input pattern P_i primarily to that group of MSNs that projected to target area T_i . (b) Illustration of the model for spine dynamics. Five potential synaptic connections at different states are shown. Synaptic spines are represented by circular volumes with diameters proportional to $\sqrt[3]{w_i}$ for functional connections, assuming a linear correlation between spine-head volume and synaptic efficacy w_i [Matsuzaki et al., 2001]. (c) Dynamics of weights w_i in log-scale for 10 potential synaptic connections i when the activity-dependent term $\frac{\partial}{\partial \theta_i} \log \mathcal{V}(\theta) dt$ in Eq. (5) is set equal to zero. Consistent with experimental data (see e.g. Fig. 2i of [Holtmaat et al., 2006]) the dynamics is in this case consistent with an Ornstein-Uhlenbeck process in the logarithmic scale. Weight values are plotted relative to the initial value at time 0. (d, e) Dynamics of a model synapse when a reward-modulated STDP pairing protocol as in [Yagishita et al., 2014] was applied. (d) Reward delivery after repeated firing of the presynaptic neuron before the postsynaptic neuron resulted in a strong weight increase (left). This effect was reduced without reward (right), and prevented completely if no presynaptic stimulus was applied. Values in d and e represent percentage of weight changes relative the pairing onset time (dashed line, means and s.e.m. over 50 synapses). Compare with Fig. 1F,G in [Yagishita et al., 2014]. (e) Dependence of resulting changes in synaptic weights in our model as a function of the delay of reward delivery. Gray shaded rectangle indicates the time window of STDP pairing application. Reward delays denote time between pairing and

reward onset. Compare to Figure 1O in [Yagishita et al., 2014]. **(f)** The average reward achieved by the network increased quickly during learning according to Eq. (5) (mean over 5 independent trial runs; shaded area indicates s.e.m.). **(g)** Synaptic parameters kept changing throughout the experiment in **f**. The magnitude of the change of the synaptic parameter vector θ is shown (mean \pm s.e.m. as in **f**; Euclidean norm, normalized to the maximum value). The parameter change peaks at the onset of learning, but remains high (larger than 80% of the maximum value) even when stable performance has been reached. **(h)** Spiking activity of the network during learning. Activities of 20 randomly selected input neurons and all MSNs are shown. 3 salient input neurons (belonging to pools S_1 or S_2 in **i**) are highlighted. Most neurons have learnt to fire at a higher rate for the input pattern P_j that corresponds to the target area T_j to which they are projecting. Bottom: reward delivered to the network. **(i)** Dynamics of network rewiring throughout learning. Snapshots of network configurations for the times t indicated below the plots are shown. Gray lines indicate active connections between neurons; connections that were not present at the preceding snapshot are highlighted in green. All output neurons and two subsets of input neurons that fire strongly in pattern P_1 or P_2 are shown (pools S_1 and S_2 , 20 neurons each). Numbers denote total counts of functional connections between pools. The connectivity was initially dense and then rapidly restructured and became sparser. Rewiring took place all the time throughout learning. **(j)** Analysis of random exploration in task-irrelevant dimensions of the parameter space. Projection of the parameter vector θ to the two dPCA components that best explain the variance of the average reward. dpc1 explains more than 99.9% of the reward variance (dpc2 and higher dimensions less than 0.1%). A single trajectory of the high-dimensional synaptic parameter vector over 24 hours of learning projected onto dpc1 and dpc2 is shown. Amplitude on the y-axis denotes the estimated average reward (in fractions of the total maximum achievable reward). After converging to a region of high reward (movement mainly along dpc1) network continues to explore task-irrelevant dimensions (movement mainly along dpc2).

3 Reward-based self-configuration and compensation capability of a recurrent neural network.

(a) Network scaffold and task schematic. Symbol convention as in Fig. 1a. A recurrent network scaffold of excitatory and inhibitory neurons (large blue circle); a subset of excitatory neurons received input from afferent excitatory neurons (indicated by green shading). From the remaining excitatory neurons, two pools D and U were randomly selected to control lever movement (blue shaded areas). Bottom inset: stereotypical movement that had to be generated to receive a reward. **(b)** Spiking activity of the network at learning onset and after 22 hours of learning. Activities of random subsets of neurons from all populations are shown (hidden: excitatory neurons of the recurrent network, which are not in pool D or U). Bottom: lever position inferred from the neural activity in pools D and U. Rewards are indicated by red bars. Gray shaded areas indicate cue presentation. **(c)** Task performance quantified by the average time from cue presentation onset to movement completion. The network was able to solve this task in less than 1 seconds on average after about 8 hours of learning. A task change was introduced at time 24 h (asterisk; function of D and U switched), which was quickly compensated by the network. Using a simplified version of the learning rule, where the re-introduction of non-functional potential connections was approximated

using exponentially distributed waiting times (green), yielded similar results (see also panel e). If the connectome was kept fixed after the task change at 24 h performance was significantly worse (black). (d) Trial-averaged network activity (top) and lever movements (bottom). Activity traces are aligned to movement onsets (arrows). Y-axis of trial-averaged activity plots are sorted by the time of highest firing rate within the movement at various times during learning: sorting of the first and second plot is based on the activity at $t = 0$ h, third and fourth by that at $t = 22$ h, fifth is resorted by the activity at $t = 46$ h. Network activity is clearly restructured through learning with particularly stereotypical assemblies for sharp upward movements. Bottom: average lever movement (black) and 10 individual movements (gray). (e) Turnover of synaptic connections for the experiment shown in d. Y-axis is clipped at 3000. Turnover rate during the first two hours was around 12.000 synapses ($\sim 25\%$) and then decreased rapidly. Another increase in spine turnover rate can be observed after the task change at time 24 h. (f) Effect of forgetting due to parameter diffusion over 14 simulated days. Application of reward was stopped after 24 hours when the network had learned to reliably solve the task. Parameters subsequently continue to evolve according to the SDE (5). Onset of forgetting can be observed after day 6. A simple consolidation mechanism triggered after 4 days reliably prevents forgetting. (g) Histograms of time intervals between disappearance and reappearance of synapses (waiting times) for the exact (upper plot) and approximate (lower plot) learning rule. (h) Relative fraction of potential synaptic connections that were stably non-functional, transiently decaying, transiently emerging or stably function during the re-learning phase for the experiment shown in d. (i) PCA of a random subset of the parameters θ_i . The plot suggests continuing dynamics in task-irrelevant dimensions after the learning goal has been reached (indicated by red color). When the function of the neuron pools U and D was switched after 24 h, the synaptic parameters migrated to a new region. All plots show means over 5 independent runs (error bars: s.e.m.).

4 Impact of the prior distribution and reward amplitude on the synaptic dynamics.

Task performance and total number of active synaptic connections throughout learning for different prior distributions and distribution of initial synaptic parameters. Synaptic parameters were initially drawn from a Gaussian distribution with mean μ_{init} and $\sigma = 0.5$. (a) Comparison of the task performance and number of active synapses for the parameter set used in Fig. 3 and (b) Gaussian prior distribution with different parameters. (c) In addition a Laplace prior with different parameters was tested. The prior distribution and the initial synaptic parameters had a marked effect on the task performance and overall network connectivity. (d) Impact of the reward amplitude on the synaptic dynamics. Task performance is here measured for different values of c_r to scale the amplitude of the reward signal. Dashed lines denote the task switch as in Fig. 3.

5 Contribution of spontaneous and neural activity-dependent processes to synaptic dynamics

(a,b) Evolution of synaptic weights w_i plotted against time for a pair of CI synapses in a, and non-CI synapses in b, for temperature $T = 0.5$. (c) Pearson's correlation coefficient computed between synaptic weights of CI and non-CI synapses of a network with $T = 0.5$ after 48 h of learning as in Fig. 3c,d. CI

synapses were only weakly, but significantly stronger correlated than non-CI synapses. **(d)** Impact of T on correlation of CI synapses (X-axis) and learning performance (Y-axis). Each dot represents averaged data for one particular temperature value, indicated by the color. Values for T were 1.0, 0.75, 0.5, 0.35, 0.2, 0.15, 0.1, 0.01, 0.001, 0.0. These values are proportional to the small vertical bars above the color bar. The performance (measured in movement completion time) is measured after 48 hours for the learning experiment as in Fig. 3c,d, where the network changed completely after 24 h. Good performance was achieved for a range of temperature values between 0.01 and 0.5. Too low (< 0.01) or too high (> 0.5) values impaired learning. Means + s.e.m. over 5 independent trials are shown. **(e)** Synaptic weights of 100 pairs of CI synapses that emerged from a run with $T = 0.5$. Pearson's correlation is 0.239, comparable to the experimental data in Fig. 8A-D of [Dvorkin and Ziv, 2016]. **(f)** Estimated contributions of activity history dependent (green), spontaneous synapse-autonomous (blue) and neuron-wide (gray) processes to the synaptic dynamics for a run with $T = 0.15$. The resulting fractions are very similar to those in the experimental data, see Fig. 8E of [Dvorkin and Ziv, 2016]. **(g)** Evolution of learning performance and total number of active synaptic connections for different temperatures as in **d**. Compensation for task perturbation was significantly faster with higher temperatures. Temperatures larger than 0.5 prevented compensation. Overall number of synapses was decreasing for temperatures $T < 0.1$ and increasing for $T \geq 0.1$.

6 Drifts of neural codes while performance remained constant.

Trial-averaged network activity as in Fig. 3d evaluated at three different times selected from a time window where the network performance was stable (see Fig. 3c). Each column shows the same trial-averaged activity plot but subject to different sorting. Rows correspond to one sorting criterion based on one evaluation time.

7 2D projections of the PCA analysis in Fig. 3i.

The 3D projection as in Fig. 3i (top right) and the corresponding 2D projections are shown.

Tables

1 Parameters of the synapse model Eq. (15), (16) and (18).

Parameter values were found by fitting the experimental data of [Yagishita et al., 2014]. If not stated otherwise, these values were used in all experiments.

<i>symbol</i>	<i>value</i>	<i>description</i>
T	0.1	temperature
τ_e	1 s	time constant of eligibility trace
τ_g	50 s	time constant of gradient estimator
τ_a	50 s	time constant to estimate the average reward
α	0.02	offset to reward signals
β	10^{-5}	learning rate
μ	0	mean of prior
σ	2	std of prior

References

- Abraham and Robins, 2005. Abraham, W. C. and Robins, A. (2005). Memory retention—the synaptic stability versus plasticity dilemma. *Trends in neurosciences*, 28(2):73–78.
- Avermann et al., 2012. Avermann, M., Tamm, C., Mateo, C., Gerstner, W., and Petersen, C. (2012). Microcircuits of excitatory and inhibitory neurons in layer 2/3 of mouse barrel cortex. *Journal of Neurophysiology*, 107(11):3116–3134.
- Bartol Jr et al., 2015. Bartol Jr, T. M., Bromer, C., Kinney, J., Chirillo, M. A., Bourne, J. N., Harris, K. M., and Sejnowski, T. J. (2015). Nanoconnectomic upper bound on the variability of synaptic plasticity. *Elife*, 4:e10778.
- Baxter and Bartlett, 2000. Baxter, J. and Bartlett, P. L. (2000). Direct gradient-based reinforcement learning. In *The 2000 IEEE International Symposium on Circuits and Systems*, volume 3, pages 271–274. IEEE.
- Bellec et al., 2017. Bellec, G., Kappel, D., Maass, W., and Legenstein, R. (2017). Deep rewiring: Training very sparse deep networks. *arXiv preprint arXiv:1711.05136*.
- Berry and Nedivi, 2017. Berry, K. P. and Nedivi, E. (2017). Spine dynamics: Are they all the same? *Neuron*, 96(1):43–55.
- Botvinick and Toussaint, 2012. Botvinick, M. and Toussaint, M. (2012). Planning as inference. *Trends in Cognitive Sciences*, 16(10):485–488.
- Chaisanguanthum et al., 2014. Chaisanguanthum, K. S., Shen, H. H., and Sabes, P. N. (2014). Motor variability arises from a slow random walk in neural state. *Journal of Neuroscience*, 34(36):12071–12080.
- Chambers and Rumpel, 2017. Chambers, A. R. and Rumpel, S. (2017). A stable brain from unstable components: Emerging concepts, implications for neural computation. *Neuroscience*, 357:172–184.
- Collins and Frank, 2016. Collins, A. G. and Frank, M. J. (2016). Surprise! dopamine signals mix action, value and error. *Nature Neuroscience*, 19(1):3.
- Deger et al., 2016. Deger, M., Seholzer, A., and Gerstner, W. (2016). Multi-contact synapses for stable networks: a spike-timing dependent model of dendritic spine plasticity and turnover. *arXiv preprint arXiv:1609.05730*.
- Desai et al., 1999. Desai, N. S., Rutherford, L. C., and Turrigiano, G. G. (1999). Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nature Neuroscience*, 2(6):515–520.
- Ding and Rangarajan, 2004. Ding, M. and Rangarajan, G. (2004). First passage time problem: A fokker-planck approach. In Wille, L., editor, *New Directions in Statistical Physics*, pages 31–46. Springer.

- Driscoll et al., 2017. Driscoll, L. N., Pettit, N. L., Minderer, M., Chettih, S. N., and Harvey, C. D. (2017). Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell*, 170:986–999.
- Dvorkin and Ziv, 2016. Dvorkin, R. and Ziv, N. E. (2016). Relative contributions of specific activity histories and spontaneous processes to size remodeling of glutamatergic synapses. *PLoS Biology*, 14(10):e1002572.
- Fan et al., 2005. Fan, Y., Fricker, D., Brager, D. H., Chen, X., Lu, H.-C., Chitwood, R. A., and Johnston, D. (2005). Activity-dependent decrease of excitability in rat hippocampal neurons through increases in I_h . *Nature Neuroscience*, 8(11):1542–1551.
- Fares and Stepanyants, 2009. Fares, T. and Stepanyants, A. (2009). Cooperative synapse formation in the neocortex. *Proceedings of the National Academy of Sciences*, 106(38):16463–16468.
- Fauth et al., 2015. Fauth, M., Wörgötter, F., and Tetzlaff, C. (2015). The formation of multi-synaptic connections by the interaction of synaptic and structural plasticity and their functional consequences. *PLoS computational biology*, 11(1):e1004031.
- Florian, 2007. Florian, R. V. (2007). Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. *Neural Computation*, 19(6):1468–1502.
- Frémaux et al., 2010. Frémaux, N., Sprekeler, H., and Gerstner, W. (2010). Functional requirements for reward-modulated spike-timing-dependent plasticity. *Journal of Neuroscience*, 30(40):13326–13337.
- Frémaux et al., 2013. Frémaux, N., Sprekeler, H., and Gerstner, W. (2013). Reinforcement learning using a continuous time actor-critic framework with spiking neurons. *PLoS Comput Biol*, 9(4):e1003024.
- Fusi et al., 2005. Fusi, S., Drew, P. J., and Abbott, L. (2005). Cascade models of synaptically stored memories. *Neuron*, 45(4):599–611.
- Gardiner, 2004. Gardiner, C. (2004). *Handbook of Stochastic Methods*. 3rd ed. Springer.
- Gerstner et al., 2014. Gerstner, W., Kistler, W. M., Naud, R., and Paninski, L. (2014). *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press.
- Gewaltig and Diesmann, 2007. Gewaltig, M.-O. and Diesmann, M. (2007). NEST (NEural Simulation Tool). *Scholarpedia*, 2(4):1430.
- Grashow et al., 2010. Grashow, R., Brookings, T., and Marder, E. (2010). Compensation for variable intrinsic neuronal excitability by circuit-synaptic interactions. *The Journal of Neuroscience*, 20(27):9145–9156.
- Holtmaat and Svoboda, 2009. Holtmaat, A. and Svoboda, K. (2009). Experience-dependent structural synaptic plasticity in the mammalian brain. *Nature Reviews Neuroscience*, 10(9):647–658.
- Holtmaat et al., 2006. Holtmaat, A., Wilbrecht, L., Knott, G. W., Welker, E., and Svoboda, K. (2006). Experience-dependent and cell-type-specific spine growth in the neocortex. *Nature*, 441(7096):979–983.

- Holtmaat et al., 2005. Holtmaat, A. J., Trachtenberg, J. T., Wilbrecht, L., Shepherd, G. M., Zhang, X., Knott, G. W., and Svoboda, K. (2005). Transient and persistent dendritic spines in the neocortex in vivo. *Neuron*, 45:279–291.
- Izhikevich, 2007. Izhikevich, E. M. (2007). Solving the distal reward problem through linkage of stdp and dopamine signaling. *Cerebral cortex*, 17(10):2443–2452.
- Kappel et al., 2015. Kappel, D., Habenschuss, S., Legenstein, R., and Maass, W. (2015). Network plasticity as Bayesian inference. *PLoS Computational Biology*, 11(11):e1004485.
- Kappel et al., 2017. Kappel, D., Hoff, M., Subramoney, A. (2017). IGITUGraz/spore-nest-module: SPORE version 2.14.0 (Version v2.14.0). *Zenodo*. <http://doi.org/10.5281/zenodo.1043486>
- Kasai et al., 2010. Kasai, H., Fukuda, M., Watanabe, S., Hayashi-Takagi, A., and Noguchi, J. (2010). Structural dynamics of dendritic spines in memory and cognition. *Trends in neurosciences*, 33(3):121–129.
- Kasthuri et al., 2015. Kasthuri, N., Hayworth, K. J., Berger, D. R., Schalek, R. L., Conchello, J. A., Knowles-Barley, D. L., Vazquez-Reina, A., Kaynig, V., Jones, T. R., Roberts, M., Morgan, J. L., Tapia, J. C., Seung, H. S., Roncal, W. G., Vogelstein, J. T., Burns, R., Sussman, D. L., Priebe, C. E., Pfister, H., and Lichtman, J. (2015). Saturated reconstruction of a volume of neocortex. *Cell*, 3:648–661.
- Kirkpatrick et al., 2017. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, page 201611835.
- Kirkpatrick et al., 1983. Kirkpatrick, S., Vecchi, M., et al. (1983). Optimization by simulated annealing. *Science*, 220(4598):671–680.
- Kobak et al., 2016. Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C. E., Kepecs, A., Mainen, Z. F., Romo, R., Qi, X.-L., Uchida, N., and Machens, C. K. (2016). Demixed principal component analysis of neural population data. *Elife*, 5:e10989.
- Legenstein et al., 2008. Legenstein, R., Pecevski, D., and Maass, W. (2008). A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. *PLoS Computational Biology*, 4(10):e1000180.
- Loewenstein et al., 2011. Loewenstein, Y., Kuras, A., and Rumpel, S. (2011). Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex in vivo. *The Journal of Neuroscience*, 31(26):9481–9488.
- Loewenstein et al., 2015. Loewenstein, Y., Yanover, U., and Rumpel, S. (2015). Predicting the dynamics of network connectivity in the neocortex. *The Journal of Neuroscience*, 35(36):12535–12544.
- Marder, 2011. Marder, E. (2011). Variability, compensation, and modulation in neurons and circuits. *Proceedings of the National Academy of Sciences*, 108(Supplement 3):15542–15548.

- Marder and Goaillard, 2006. Marder, E. and Goaillard, J.-M. (2006). Variability, compensation and homeostasis in neuron and network function. *Nature Reviews Neuroscience*, 7(7):563–574.
- Marr and Poggio, 1976. Marr, D. and Poggio, T. (1976). From understanding computation to understanding neural circuitry. Technical report, Massachusetts Institute of Technology, Cambridge, MA, USA.
- Matsuzaki et al., 2001. Matsuzaki, M., Ellis-Davies, G. C., Nemoto, T., Miyashita, Y., Iino, M., and Kasai, H. (2001). Dendritic spine geometry is critical for ampa receptor expression in hippocampal cal pyramidal neurons. *Nature Neuroscience*, 4(11):1086–1092.
- Minerbi et al., 2009. Minerbi, A., Kahana, R., Goldfeld, L., Kaufman, M., Marom, S., and Ziv, N. E. (2009). Long-term relationships between synaptic tenacity, synaptic remodeling, and network activity. *PLoS Biology*, 7(6):e1000136.
- Mongillo et al., 2017. Mongillo, G., Rumpel, S., and Loewenstein, Y. (2017). Intrinsic volatility of synaptic connections a challenge to the synaptic trace theory of memory. *Current opinion in neurobiology*, 46:7–13.
- Nagaoka et al., 2016. Nagaoka, A., Takehara, H., Hayashi-Takagi, A., Noguchi, J., Ishii, K., Shirai, F., Yagishita, S., Akagi, T., Ichiki, T., and Kasai, H. (2016). Abnormal intrinsic dynamics of dendritic spines in a fragile x syndrome mouse model in vivo. *Scientific Reports*, 6.
- Peters et al., 2014. Peters, A. J., Chen, S. X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature*, 510(7504):263–267.
- Peters and Schaal, 2006. Peters, J. and Schaal, S. (2006). Policy gradient methods for robotics. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2219–2225. IEEE.
- Pfister et al., 2006. Pfister, J.-P., Toyoizumi, T., Barber, D., and Gerstner, W. (2006). Optimal spike-timing-dependent plasticity for precise action potential firing in supervised learning. *Neural Computation*, 18(6):1318–1348.
- Prinz et al., 2004. Prinz, A. A., Bucher, D., and Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nature Neuroscience*, 7(12):1345–1352.
- Rawlik et al., 2013. Rawlik, K., Toussaint, M., and Vijayakumar, S. (2013). On stochastic optimal control and reinforcement learning by approximate inference. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 3052–3056. AAAI Press.
- Remme and Wadman, 2012. Remme, M. W. and Wadman, W. J. (2012). Homeostatic scaling of excitability in recurrent neural networks. *PLoS Computational Biology*, 8(5):e1002494.
- Rokni et al., 2007. Rokni, U., Richardson, A. G., Bizzi, E., and Seung, H. S. (2007). Motor learning with unstable neural representations. *Neuron*, 54(4):653–666.

- Rubinski and Ziv, 2015. Rubinski, A. and Ziv, N. E. (2015). Remodeling and tenacity of inhibitory synapses: relationships with network activity and neighboring excitatory synapses. *PLoS computational biology*, 11(11):e1004632.
- Rumpel and Triesch, 2016. Rumpel, S. and Triesch, J. (2016). The dynamic connectome. *e-Neuroforum*, 7(3):48–53.
- Schultz, 2002. Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, 36(2):241–263.
- Schultz et al., 1997. Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306):1593–1599.
- Seung, 2003. Seung, H. S. (2003). Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron*, 40(6):1063–1073.
- Statman et al., 2014. Statman, A., Kaufman, M., Minerbi, A., Ziv, N. E., and Brenner, N. (2014). Synaptic size dynamics as an effectively stochastic process. *PLoS Computational Biology*, 10(10):e1003846.
- Stettler et al., 2006. Stettler, D. D., Yamahachi, H., Li, W., Denk, W., and Gilbert, C. D. (2006). Axons and synaptic boutons are highly dynamic in adult visual cortex. *Neuron*, 49:877–887.
- Sutton and Barto, 1998. Sutton, R. S. and Barto, A. G. (1998). *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.
- Tang et al., 2010. Tang, L. S., Goeritz, M. L., Caplan, J. S., Taylor, A. L., Fisek, M., and Marder, E. (2010). Precise temperature compensation of phase in a rhythmic motor pattern. *PLoS Biology*, 8(8):e1000469.
- Urbanczik and Senn, 2009. Urbanczik, R. and Senn, W. (2009). Reinforcement learning in populations of spiking neurons. *Nature Neuroscience*, 12(3):250–252.
- van Beers et al., 2013. van Beers, R. J., Brenner, E., and Smeets, J. B. (2013). Random walk of motor planning in task-irrelevant dimensions. *Journal of neurophysiology*, 109(4):969–977.
- van Ooyen and Butz-Ostendorf, 2017. van Ooyen, A. and Butz-Ostendorf, Markus, e. (2017). *The Rewiring Brain*. Academic Press.
- Vlassis et al., 2012. Vlassis, N., Ghavamzadeh, M., Mannor, S., and Poupart, P. (2012). Bayesian reinforcement learning. In *Reinforcement Learning*, pages 359–386. Springer.
- Williams, 1992. Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Xie and Seung, 2004. Xie, X. and Seung, H. S. (2004). Learning in neural networks by reinforcement of irregular spiking. *Physical Review E*, 69(4):041909.

- Xu et al., 2009. Xu, T., Yu, X., Perlik, A. J., Tobin, W. F., Zweig, J. A., Tennant, K., Jones, T., and Zuo, Y. (2009). Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature*, 462(7275):915–919.
- Yagishita et al., 2014. Yagishita, S., Hayashi-Takagi, A., Ellis-Davies, G. C., Urakubo, H., Ishii, S., and Kasai, H. (2014). A critical time window for dopamine actions on the structural plasticity of dendritic spines. *Science*, 345(6204):1616–1620.
- Yamahachi et al., 2009. Yamahachi, H., Marik, S. A., McManus, J. N. J., Denk, W., and Gilbert, C. D. (2009). Rapid axonal sprouting and pruning accompany functional reorganization in primary visual cortex. *Neuron*, 64(5):719–729.
- Yang et al., 2009. Yang, G., Pan, F., and Gan, W.-B. (2009). Stably maintained dendritic spines are associated with lifelong memories. *Nature*, 462(7275):920–924.
- Yasumatsu et al., 2008. Yasumatsu, N., Matsuzaki, M., Miyazaki, T., Noguchi, J., and Kasai, H. (2008). Principles of long-term dynamics of dendritic spines. *The Journal of Neuroscience*, 28(50):13592–13608.
- Ziv and Ahissar, 2009. Ziv, N. E. and Ahissar, E. (2009). Neuroscience: New tricks and old spines. *Nature*, 462(7275):859–861.
- Ziv et al., 2013. Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L., Gama, A. E., and Schnitzer, M. J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nature Neuroscience*, 16(3):264–266.













