

Weight versus Node Perturbation Learning in Temporally Extended Tasks: Weight Perturbation Often Performs Similarly or Better

Paul Züge¹, Christian Klos¹, and Raoul-Martin Memmesheimer^{1*}

Neural Network Dynamics and Computation, Institute of Genetics, University of Bonn, Germany

 (Received 9 December 2021; revised 9 September 2022; accepted 23 January 2023; published 11 April 2023)

Biological constraints often impose restrictions on plasticity rules such as locality and reward-based rather than supervised learning. Two learning rules that comply with these restrictions are weight (WP) and node (NP) perturbation. NP is often used in learning studies, in particular, as a benchmark; it is considered to be superior to WP and more likely neurobiologically realized, as the number of weights and, therefore, their perturbation dimension typically massively exceed the number of nodes. Here, we show that this conclusion no longer holds when we take two properties into account that are relevant for biological and artificial neural network learning: First, tasks extend in time and/or are trained in batches. This increases the perturbation dimension of NP but not WP. Second, tasks are (comparably) low dimensional, with many weight configurations providing solutions. We analytically delineate regimes where these properties let WP perform as well as or better than NP. Furthermore, we find that the changes in weight space directions that are irrelevant for the task differ qualitatively between WP and NP and that only in WP gathering batches of subtasks in a trial decreases the number of trials required. This may allow one to experimentally distinguish which of the two rules underlies a learning process. Our insights suggest new learning rules which combine for specific task types the advantages of WP and NP. If the inputs are similarly correlated, temporally correlated perturbations improve NP. Using numerical simulations, we generalize the results to networks with various architectures solving biologically relevant and standard network learning tasks. Our findings, together with WP's practicability, suggest WP as a useful benchmark and plausible model for learning in the brain.

DOI: [10.1103/PhysRevX.13.021006](https://doi.org/10.1103/PhysRevX.13.021006)

Subject Areas: Biological Physics,
Computational Physics,
Statistical Physics

I. INTRODUCTION

Different, usually combined, strategies underlie the learning of tasks in humans and other animals [1,2]. Supervised learning allows large, rapid improvements. It is based on observing in which way an action is erroneous and on the ability of the nervous system to use this information for the improvement of neuronal dynamics in a directed manner. This may be implemented by translating an error vector into a vector of suitable synaptic weight updates [3]. Fast learning could be achieved by directly adapting the dynamics [4]. Reward-based learning (reinforcement learning), in contrast, uses only a scalar feedback signal. It is, thus, also applicable if errors are known with little specificity, for example, because there is

sparse, delayed feedback about the cumulative effect of actions, which might tell only whether an action is erroneous but not how the generating neural activity can be improved.

A variety of models for reward-based learning have been developed in the context of theoretical neuroscience and machine learning [5,6]. Two conceptually straightforward implementations of such learning in neural networks are weight perturbation (WP) [7,8] and node perturbation (NP) [9,10]. Their underlying idea is to add perturbations to the weights or to the summed weighted inputs and to correlate them to the change of task performance. If the reward increases due to an attempted perturbation, the weights or the node dynamics are changed in its direction. If the reward decreases, the changes are chosen oppositely. NP and WP are used to model reward-based learning in biological neural networks, due to four properties [4,7,9–14]: (i) They are (with minor modifications) biologically plausible. (ii) They are applicable to a broad variety of networks and tasks. (iii) They are accessible to analytical exploration. (iv) They are optimal in the sense that the average of the generated weight change taken over all noise realizations is

*rm.memmesheimer@uni-bonn.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

along the reward gradient. The schemes' names were originally coined for approaches that directly estimate the individual components of the gradient using single perturbations to each weight or node [15].

WP explores with its random perturbations a space with dimensionality equal to the number of weights. For trials without temporal extent, NP needs only to explore a space with dimensionality equal to the number of nodes. The chain rule, amounting to a simple multiplication with the unweighted input strength, then allows one to translate a desirable change in the summed weighted inputs into a change in a particular weight strength. NP, thus, uses additional information on the structure of the network (namely, the linearity of input summation) to reduce the required exploration.

In linear approximation, the optimal direction of weight changes aligns with the direction of the gradient of the reward. WP and NP seemingly attempt to find this direction by trying out random perturbations. Since the dimension of the space of possible perturbation directions is large, the probability of finding the gradient is small and a lot of exploration is necessary. This impedes WP and NP. The number of weights and, thus, the dimensionality of the perturbation space searched by WP are much larger than the number of nodes. NP is, thus, considered more suitable for reward-based neural network learning [3,9,10,16] and its implementation in biological neural networks [17]. This is supported by quantitative analysis: Ref. [9] considers M linear perceptrons with N random inputs, using a student-teacher task. They find that for WP the optimal convergence rate of the student to the teacher weight matrix is by a factor NM worse than for exact gradient descent (GD). This is consistent with the argument that WP needs to search the NM -dimensional weight space to find the gradient, which is directly computed by GD. Accordingly, NP is worse than gradient descent by the dimensionality M of the node perturbation space.

The prerequisites of the arguments sketched above, however, do not hold in many biological situations. First, tasks in biology often extend in time and have a reward feedback that is temporally distant from the action [1,5,10,14,18]. Second, the effective dimension of neural trajectories and of learning tasks is often comparably low [19–21]. Our article analytically and numerically explores the perturbation-based learning of tasks with these features.

The article is structured as follows. First, we introduce the employed WP and NP learning models. We then derive analytic expressions for the evolution of expected error (negative reward) in linear networks solving temporally extended, low-dimensional reward-based learning tasks. This allows us to identify conditions for which WP outperforms NP as well as the underlying reasons. Furthermore, we delineate distinguishing qualitative characteristics of the weight and error dynamics. Finally, we numerically show that WP is comparably good or outperforms NP in different

biologically relevant and standard network learning applications.

II. RESULTS

A. Learning models and task principles

Our study models the learning of tasks that are temporally extended. Time is split into discrete steps, indexed by $t = 1, \dots, T$, where T is the duration of a trial. During this period, a neural network receives external input and generates output. At the end of a trial, it receives a scalar error feedback E about its performance [8,10,14,17,22]. To quantitatively introduce the learning rules, we consider a neuron i , which may be part of a larger network. It generates in the t th time bin an output firing rate z_{it} , in response to the firing rates r_{jt} of its N presynaptic neurons:

$$z_{it} = g\left(\sum_{j=1}^N w_{ij} r_{jt}\right). \quad (1)$$

Here, w_{ij} is the weight of the synapse (or the total weight of the synapses) from neuron j to neuron i . The generally nonlinear activation function g implements the relation between the total input current and the output firing rate of the neuron [5,23]. We note that the individual synaptic input currents $w_{ij} r_{jt}$ in the model sum up linearly. This is a standard assumption, and it is a requirement for the NP scheme [9,10,15,24]. In the presence of nonlinear dendritic compartments [25–27], each of these could be an independently perturbed node.

We model WP learning by adding in the beginning of a trial a temporally static weight change ξ_{ij}^{WP} to each of the weights w_{ij} [8,22]. The output of the neuron then reads

$$z_{it}^{\text{pert,WP}} = g\left(\sum_{j=1}^N (w_{ij} + \xi_{ij}^{\text{WP}}) r_{jt}^{\text{pert,WP}}\right), \quad (2)$$

where $r_{jt}^{\text{pert,WP}}$ are the input rates, which may have a perturbation due to upstream perturbed weights. ξ_{ij}^{WP} are independent and identically distributed (iid) zero-mean Gaussian white noise perturbations with standard deviation σ_{WP} , $\langle \xi_{ij}^{\text{WP}} \xi_{kl}^{\text{WP}} \rangle = \delta_{ik} \delta_{jl} \sigma_{\text{WP}}^2$, where the angular brackets denote the average over perturbation realizations and δ is the Kronecker delta. The perturbations ξ_{ij}^{WP} change the output, which, in turn, influences the reward received at the end of the trial [Fig. 1(a), left-hand side]. We usually assume that the difference in reward between the perturbed and an unperturbed trial with $\xi_{ij}^{\text{WP}} = 0$ for all i, j is used to estimate the gradient: When the reward increases, for small perturbations this means that the tried perturbation has a positive component along the reward gradient. Consequently, the update is chosen in the direction of that

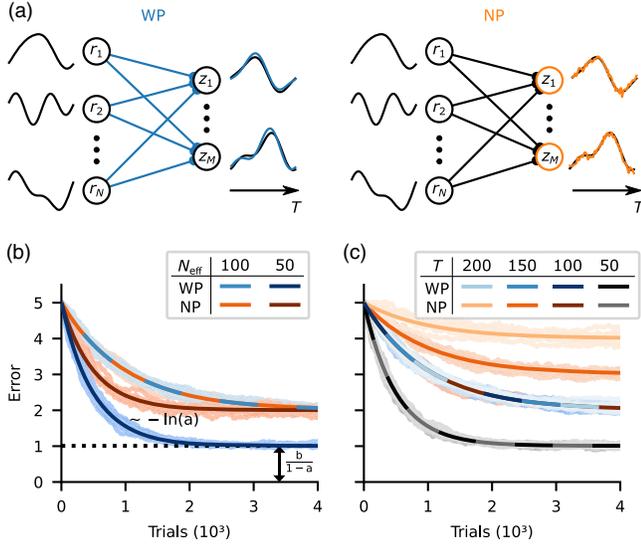


FIG. 1. Learning of temporally extended tasks in linear networks. (a) Schematic setup of WP and NP. The M outputs z_i are weighted sums of the N inputs r_j . Left: WP perturbs the weights at the beginning of a trial; the resulting perturbations of the weighted sums of the inputs and, thus, the outputs reflect the dimensionality and smoothness of the inputs (blue). Right: NP perturbs the weighted sums of the inputs with dynamical noise (orange). (b) WP (blue) works just as well as or better than NP (orange) when learning a single temporally extended input-output mapping. The error decay time decreases for WP and NP likewise with decreasing effective input dimension N_{eff} (light versus dark curves). In contrast, the residual error decreases only for WP. (c) Increased trial duration T does not change the progress of WP learning (blue curves lie on top of each other). In contrast, increasing T hinders NP learning by increasing the residual error (compare the increasingly lighter orange curves for larger T). If T decreases N_{eff} (gray curves), convergence is faster and to a lower residual error in both WP (because of the decrease in N_{eff}) and NP (because of the decrease in N_{eff} and T). (b) shows error curves from simulations (ten runs, shaded) together with analytical curves for the decay of the expected error (solid), for fixed $T, N = 100, M = 10$, and $N_{\text{eff}} \in \{100, 50\}$; $\sigma_{\text{eff}} = 0.04$. Theoretical curves and simulations agree well. For WP and $N_{\text{eff}} = 50$, the decay rate $[-\ln(a)]$ and the residual error (dashed line) are highlighted. (c) shows error curves from simulations and theory similar to (b) for fixed $N = 100$ and $T \in \{200, 150, 100, 50\}$. N_{eff} is set to 100 but cannot be greater than T , such that $T = 50$ forces $N_{\text{eff}} = 50$.

perturbation. When the reward decreases, the update is chosen in the opposite direction. We use the update rule

$$\Delta w_{ij}^{\text{WP}} = -\frac{\eta}{\sigma_{\text{WP}}^2} (E^{\text{pert}} - E) \xi_{ij}^{\text{WP}}, \quad (3)$$

where η is the learning rate, E^{pert} is the error of the perturbed trial, and E is the error of the unperturbed trial. For the delayed non-match-to-sample (DNMS) task, the error of the unperturbed trial is replaced by an average over

the previous errors for biological plausibility. The proportionality of update size and obtained reward implies that, when averaging over the distribution of the perturbations, the weight change

$$\langle \Delta w_{ij}^{\text{WP}} \rangle \approx -\eta \frac{\partial E}{\partial w_{ij}} \quad (4)$$

is parallel to the reward gradient [Eq. (A2)]. Since $\langle \xi_{ij}^{\text{WP}} \rangle = 0$, this holds for any baseline in Eq. (3). The employed choice of baseline E guarantees that for small perturbations the weight change has a positive component in the direction of the reward gradient and, thus, always reduces the error for sufficiently small learning rate η [8]. In fact, it minimizes the update noise, i.e., the fluctuations of updates around the gradient [Eq. (A5)].

WP treats the system as a black box, mapping parameters w onto a scalar error function E . In other words, it uses the information that the weights are fixed parameters during a trial but does not take advantage of specifics of the network structure. This is in contrast to NP, which takes into account some minimal structural knowledge, namely, the linear summation of input currents, but not the constancy of the weights.

Instead of perturbing the weights directly, NP adds noise to the sum of the inputs:

$$z_{it}^{\text{pert,NP}} = g \left(\sum_{j=1}^N w_{ij} r_{jt}^{\text{pert,NP}} + \xi_{it}^{\text{NP}} \right) \quad (5)$$

[9,10] [see Fig. 1(a), right-hand side]. ξ_{it}^{NP} are iid zero-mean Gaussian white noise perturbations with standard deviation σ_{NP} , $\langle \xi_{it}^{\text{NP}} \xi_{ms}^{\text{NP}} \rangle = \delta_{im} \delta_{ts} \sigma_{\text{NP}}^2$. We note that for temporally extended tasks, in contrast to WP, the noise must be time dependent to explore the space of time-dependent sums of inputs [10]. For temporally constant noise, only the temporal mean of the total input would be varied and, thus, improved.

The NP update rule can be defined as

$$\Delta w_{ij}^{\text{NP}} = -\frac{\eta}{\sigma_{\text{NP}}^2} (E^{\text{pert}} - E) \sum_{t=1}^T \xi_{it}^{\text{NP}} r_{jt} \quad (6)$$

[10], with the eligibility trace $\sum_{t=1}^T \xi_{it}^{\text{NP}} r_{jt}$. As for WP, this yields an average weight update parallel to the reward gradient, which again holds for any baseline of the weight update. The choice of baseline E again minimizes the update noise [Eq. (A6)].

The NP update rule effectively incorporates an error backpropagation step, which reduces to a simple multiplication with r_{jt} due to the linearity of the spatial synaptic input summation. This allows one to perturb only summed inputs instead of individual weights and may be

expected to increase the performance of NP compared to WP [3,10,14,17].

B. Theoretical analysis

We analytically compare WP and NP for temporally extended tasks by training a set of M linear perceptrons with N inputs. The task is to learn the mapping of a single fixed input sequence of duration T to a target output sequence in a reward-based manner. The task choice is motivated first by biological motor tasks that require such a mapping, like the learning of their single song in certain songbirds (see Sec. III). Second, it yields novel insights, as it is the opposite extreme case to having no time dimension and different, random inputs in each trial; this case is treated analytically by Ref. [9] (see the introduction). Third, our findings yield an understanding of the learning performance for more general temporally extended tasks and networks studied later in this article. The analysis shows how learning depends on task dimensions and the structure of the input. Furthermore, it reveals specific disadvantages of WP and NP. Importantly, our theoretical considerations cover very general sequences. In particular, they hold for sequences with and without correlations between subsequent inputs. Furthermore, the sequences can be arbitrarily reordered, also differently in each trial. They may, therefore, also be interpreted as sets or batches of inputs. In Sec. II F, we consider temporally correlated sequences, for which such an interpretation is not useful anymore. In Secs. II G and III, we relax the assumption of exactly repeated input sets.

The perceptrons generate as outputs the product of their $M \times N$ weight matrix w with the inputs

$$z_{it} = \sum_{j=1}^N w_{ij} r_{jt}, \quad (7)$$

where $i = 1, \dots, M$ [Eq. (1) and Fig. 1(a)]. For now, we assume that the target output can be produced with target weights w^* , that is, $z_{it}^* = \sum_{j=1}^N w_{ij}^* r_{jt}$. This condition is alleviated in Sec. II E. The learning goal is to reduce the quadratic deviation of each output from its target, which can be expressed through the weight mismatch $W = w - w^*$ and the input correlation matrix $S_{jk} = (1/T) \sum_{t=1}^T r_{jt} r_{kt}$ [5]:

$$E = \frac{1}{2T} \sum_{i=1}^M \sum_{t=1}^T (z_{it} - z_{it}^*)^2 = \frac{1}{2} \text{tr}[WSW^T]. \quad (8)$$

We note that with this quadratic error function the average weight update [cf. Eq. (4)] follows the gradient exactly, for both WP and NP [Eqs. (A16) and (A17)].

We assume that the inputs are composed of N_{eff} orthogonal latent inputs; all other input components are zero (this is relaxed in Sec. II I). Since there are at most T linearly independent vectors of length T , the effective input

dimension N_{eff} is bounded by $N_{\text{eff}} \leq T$. $T > 1$ thus renders our learning problem nontrivial, by allowing for inputs that are higher dimensional when considering the input-output relation of a single sequence. In biological systems, inputs are low dimensional; N_{eff} is often of the order of 10 (Sec. III), in particular, $N_{\text{eff}} \ll N$. As long as inputs are summed linearly, for clarity we then hypothetically “rotate” the inputs such that only the first N_{eff} inputs are nonzero and equal to the latent ones [Fig. 2(a)]. This allows us to speak about relevant and irrelevant inputs instead of relevant (nonzero) and irrelevant (zero) input space directions. It does not affect the WP or NP learning process, because all perturbations are isotropic and the error function is rotationally invariant [Supplemental Material, Eq. (S1) [28]]. In other words, all results, in particular, the dynamics of the error decay, hold identically for the original networks with nonrotated inputs where all actual inputs may be nonzero. For simplicity in our mathematical analysis, we assume that all latent inputs have the same strength α^2 , i.e., $(1/T) \sum_{i=1}^T r_{it}^2 = \alpha^2$ for the nonzero inputs $i = 1, \dots, N_{\text{eff}}$. A partial treatment of networks with inhomogeneous latent input strength is given in Supplemental Material Sec. IV [28].

C. Error dynamics

To elucidate the learning process and its dependence on the network and task parameters, we analytically derive the evolution of the expected error. This requires the computation of the error signal $E^{\text{pert}} - E$ and weight update after a given perturbation to determine the new error. Subsequent averaging over all perturbations yields the expected error at trial n , $\langle E(n) \rangle$, as a function of $\langle E(n-1) \rangle$, specifically, a linear recurrence relation

$$\langle E(n) \rangle = a \langle E(n-1) \rangle + b \quad (9)$$

(see Appendix B for the detailed derivation). The speed of learning is determined by the convergence factor a , while the per-update error increase b limits the final performance. Learning stops at a finite error when an equilibrium between gradient-related improvement and reward noise-induced deterioration is reached. The recurrence relation is solved by

$$\langle E(n) \rangle = [\langle E(0) \rangle - E_f] a^n + E_f. \quad (10)$$

For $a < 1$, the average error $\langle E(n) \rangle$ converges exponentially at a rate $-\ln(a)$ toward a finite final (residual) error of $E_f = b/(1-a)$, as shown in Fig. 1(b). Usually, in our settings a is sufficiently close to 1 to well approximate the convergence rate by $-\ln(a) \approx 1-a$.

To understand how learning depends on the task parameters, we first consider the speed of learning. The determining convergence factor

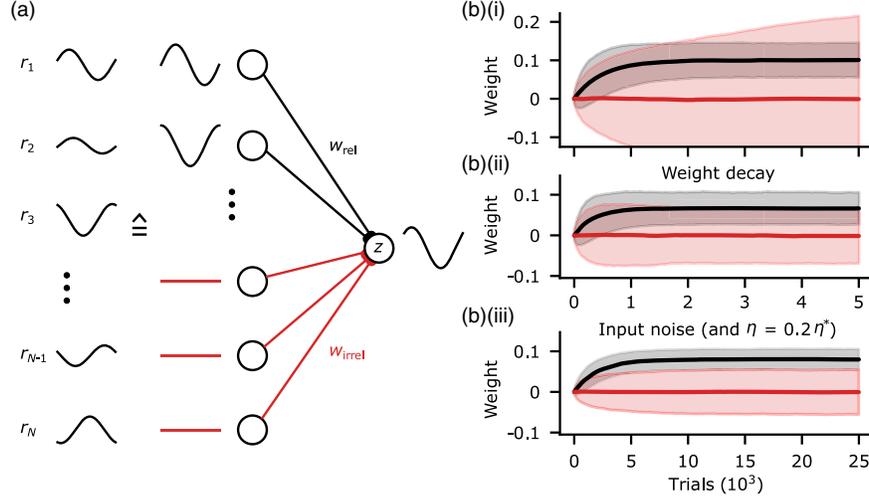


FIG. 2. Hypothetical rotation of inputs and weight diffusion. (a) Because the inputs (left, black) are summed linearly, they can be “rotated” so that for our tasks the first N_{eff} inputs are nonzero and agree with the latent inputs (middle, black). The remaining inputs are then zero (middle, red), and their weights are irrelevant for the output (right, red). (b)(i) In WP with finite perturbation size σ_{WP} , the irrelevant weights diffuse without bounds (red), while the relevant weights converge and fluctuate (black) around the teacher weights. Displayed are the mean (solid) and standard deviation (shaded area) of the weight ensembles. (b)(ii) Weight decay or (b)(iii) input noise confines the diffusion. In (a), $N_{\text{eff}} = 2$, the latent inputs are a sine and a cosine. Parameters in (b)(i) and (b)(ii): $M = 10$, $N = 100$, $N_{\text{eff}} = 50$, $\sigma_{\text{eff}} = 0.04$, teacher weights $w_{\text{rel},i} = 0.1$, and weight decay $\gamma_{\text{WD}} = 0.999$; results are averaged over ten runs. (b)(iii) The same parameters, except $\eta = 0.2\eta^*$ and added iid white input noise with strength $\sigma_{\text{noise}}^2 = 0.5$ (SNR = 2).

$$a = 1 - 2\eta\alpha^2 + \eta^2\alpha^4(MN_{\text{eff}} + 2) \quad (11)$$

[Eq. (B39)] is affected by two opposing effects: On average, updates follow the gradient, thus reducing the error. This is reflected by a reduction of a by $-2\eta\alpha^2(+\eta^2\alpha^4)$, responsible for convergence. However, updates fluctuate, adding a diffusive part to the weight evolution which slows convergence down. Although these fluctuations, having zero mean, do not influence the expected error to linear order, they do so quadratically. Thus, their contribution to a , $\eta^2\alpha^4(MN_{\text{eff}} + 1)$, is quadratic in the learning rate η . It is approximately proportional to MN_{eff} , the number of relevant weights that read out from nonzero inputs: Fluctuations in each of these weights yield the same contribution—the exception being the twice as strong fluctuations along the single gradient-parallel direction, which together with the quadratic effect of the mean update cause the $+2$ in Eq. (11).

The fluctuations originate from a credit assignment problem: Only the perturbation parallel to the error gradient can be credited for causing the linear part of the error signal $E^{\text{pert}} - E$. WP has no way of directly solving the credit assignment problem of identifying this direction. Thus, the perturbations of all MN weights are equally amplified in the constructions of their updates [Eq. (3)] such that all weights fluctuate. This entails fluctuations in the MN_{eff} relevant weights, which influence output and error. NP can at least partially solve the credit assignment problem by using eligibility traces, which are zero for weights that read

out from zero inputs. By projecting each of its (M) T -dimensional output perturbations onto the effectively N_{eff} -dimensional inputs, NP restricts its updates to the MN_{eff} -dimensional subspace of relevant weights. The convergence speed, thus, becomes independent of T as for WP. Interestingly, WP and NP therefore converge at the same speed despite their different numbers of fluctuating weights. The reason is that the fluctuations of the relevant weights are the same for both algorithms.

The balance between the improvement resulting from following the gradient ($\sim\eta$) and the deterioration due to the fluctuations of relevant weights ($\sim\eta^2$) in Eq. (11) is controlled by the learning rate: Small learning rates imply averaging out fluctuations over many updates and, therefore, dominance of gradient following, leading to convergence. For the remainder of the analysis of this setting, both algorithms are compared at their optimal learning rate η^* , which is defined to yield fastest convergence, in other words, to minimize a . This definition is chosen because it is conceptually straightforward, and Eq. (11) directly leads to the simple expressions

$$\eta^* = \frac{1}{(MN_{\text{eff}} + 2)\alpha^2}, \quad a^* = 1 - \frac{1}{MN_{\text{eff}} + 2}. \quad (12)$$

Here, the factor $1/\alpha^2$ in η^* cancels the scaling of the gradient with the input strength and equals the optimal learning rate for GD [Eq. (B5)]. In order to allow for averaging out the update fluctuations, WP and NP learning additionally have to slow down by a factor of

approximately MN_{eff} . Learning diverges for $\eta \rightarrow 2\eta^*$ where $a \rightarrow 1$. Equation (12) shows that WP's convergence rate is worse than GD's by a factor generally smaller than the number of weights. Further, NP's convergence rate is worse by a factor generally larger than the number of nodes. Thus, the number of weights or nodes is insufficient to predict the performance of WP or NP, respectively.

The per-update error increase and the final error, b and E_f , result from finite perturbation sizes. Finite perturbation sizes lead, due to the curved, quadratic error function, to an estimate that is at least slightly incompatible with the linear approximation assumed by the update rules [cf. Eq. (4)]. This is particularly apparent when the output error and, thus, the gradient are (practically) zero: Any finite weight or node perturbation then leads to an increase of the error and, thus, to an opposing weight update instead of no weight modification. This prevents the weights from reaching optimal values and results in a finite final error E_f . The described difference between perturbation-based error estimate and linear approximation is a form of "reward noise." It is nonzero only for finite perturbation size, as reflected by the dependence of b and E_f on σ (which is quadratic due to the quadratic error nonlinearity).

For a fair comparison of WP and NP, we choose σ_{WP} and σ_{NP} such that they lead to the same effective perturbation strength σ_{eff}^2 , as measured by the total induced output variance. This leads to $\sigma_{\text{NP}}^2 = \sigma_{\text{eff}}^2$ and $\sigma_{\text{WP}}^2 = 1/(\alpha^2 N_{\text{eff}}) \cdot \sigma_{\text{eff}}^2$ [Eq. (A22)]. Evaluated at the optimal learning rate η^* , the leading-order term of the final error is

$$E_f^{\text{WP}} = \frac{b_{\text{WP}}^*(\eta^*)}{1 - a^*} \approx \frac{1}{8} \sigma_{\text{eff}}^2 \cdot M^2 N_{\text{eff}}, \quad (13)$$

$$E_f^{\text{NP}} = \frac{b_{\text{NP}}^*(\eta^*)}{1 - a^*} \approx \frac{1}{8} \sigma_{\text{eff}}^2 \cdot M^2 T. \quad (14)$$

Importantly, the final error of WP is here generally smaller, by a factor $N_{\text{eff}}/T \leq 1$. To understand this, we focus for both WP and NP on the output perturbations that they generate. By perturbing the weights, WP induces output perturbations that are linear combinations of the inputs. These are confined to the effectively MN_{eff} -dimensional subspace in which also the (realizable part of the) output error gradient $(z - z^*)/T$ lies. NP, on the other hand, creates an entirely random MT -dimensional perturbation vector [Fig. 1(a)]. Only the projection of this vector onto the output gradient is useful for learning. This projection is smaller for NP's random vector, since the vector has effectively a larger dimensionality than the output perturbation vector of WP, at the same length. NP compensates this deficit by amplifying the smaller gradient projection more strongly. It, thus, achieves the same mean update and convergence speed as WP. However, it also more strongly amplifies the reward noise that comes with larger perturbation sizes, which results in a larger final error. The scaling

of E_f with $M^2 N_{\text{eff}}$ or $M^2 T$ reflects the effective output perturbation dimensions, MN_{eff} or MT of WP or NP, and additionally the general scaling of errors with M [Eq. (8), Supplemental Material Sec. I, and Eqs. (S55)–(S57) [28]].

Taken together, we observe that here WP learning works just as well as or better than NP. Both algorithms have the same speed of convergence, but the final error E_f of WP is smaller than or equal compared to NP. The rate of convergence decreases with increasing M and N_{eff} . Longer trial durations T harm NP by linearly increasing E_f . Larger effective input dimensionality N_{eff} similarly harms WP. This result differs from the observation in Ref. [9] that WP converges much (N -times) slower than NP. The reason is that our networks learn a single temporally extended input-output relation, while those in Ref. [9] learn the weights of a teacher network, by trials with random input of duration $T = 1$. We explore the relation between the results in detail in Sec. II H.

D. Weight diffusion

When the input has less than maximal dimensionality, $N_{\text{eff}} < N$, only certain combinations of weights read out nonzero components of the input. This becomes particularly clear for the considered rotated inputs: If WP adds a perturbation to a weight mediating zero input, to an irrelevant weight, the output and the error remain unchanged. This missing feedback leads to an unbounded diffusionlike random walk of irrelevant weights. For unrotated inputs, the weight strength diffuses in irrelevant weight space directions. We see below (Sec. II G) that the weight diffusion harms performance when learning multiple input-output patterns.

We find that, for WP in the limit of infinitesimally small perturbations $\sigma_{\text{WP}} \rightarrow 0$, all weights initially change and then converge (Supplemental Material, Fig. S1 [28]). This is because the learning-induced drift of relevant and the diffusion of irrelevant weights both stop when the error converges to zero: The error E is quadratic, such that for infinitesimally small perturbations $E^{\text{pert}} = E$ at its minimum. In contrast, for finite perturbations a residual error remains and weights continue to change. In particular, irrelevant weights continue to diffuse [Fig. 2(b)(i)]. The quantitative details of the weight diffusion process can be analytically understood [Supplemental Material Sec. II, Eqs. (S61) and (S64) [28]]. Standard mechanisms such as an exponential weight decay [29,30] confine their growth. Simultaneously, they bias the relevant weights toward zero and therewith increase the residual error [Fig. 2(b)(ii)]. Also, input noise confines the weight spread, by adding error feedback to irrelevant weights. Simultaneously, the noise increases the final task performance error and enforces a lower learning rate; see Fig. 2(b)(iii), Sec. III, and Figs. 5(a) and 5(b).

NP does not generate weight diffusion in noise-free networks: The rotated inputs make it obvious that in NP the eligibility trace [Eq. (6)] selects only the weights from

relevant inputs to be updated, since for irrelevant inputs we have $r_{jt} = 0$ for all t such that $\Delta w_{ij}^{\text{NP}} = 0$. Input noise renders irrelevant inputs and their weight updates nonzero, such that irrelevant weights diffuse also for NP (Sec. II I).

Differences in weight spread and updates between WP and NP suggest experimental measurements to distinguish which one of them underlies learning of a certain task: As long as the noise is weak compared to the signal such that the task can be satisfied with high precision, the spread of irrelevant weights is with NP much smaller than with WP [Fig. 5(b)]. Furthermore, a large variance in the weight updates that is independent of presynaptic activity together with a resulting weight spread that is largest for weight directions that read out weak latent inputs point to WP [Fig. 5(b) and Supplemental Material Sec. IV, Eq. (S126) [28]]. Consistent with this, prominent random walklike weight changes that are unrelated to neuronal activity and task learning are common in biological neural networks [31]. Weight updates whose variance scales with input strength but whose final spread is independent of it [Fig. 5(b) and Supplemental Material Sec. IV, Eq. (S128) [28]] point to NP.

E. Unrealizable targets

In the previous sections, we assume that the target outputs z_{it}^* could be exactly realized by setting the perceptron weights w equal to some target weights w^* , $z_{it}^* = \sum_{j=1}^N w_{ij}^* r_{jt}$. In general, however, the target outputs may contain components d that cannot be generated by the network, which is limited to producing linear combinations of the inputs. Unrealizable components are orthogonal to all inputs when interpreted as T -dimensional vectors, $\sum_{t=1}^T d_{it} r_{jt} = 0 \forall i, j$. The target may be written as a sum of realizable and orthogonal unrealizable parts, $z_{it}^* = \sum_{j=1}^N w_{ij}^* r_{jt} + d_{it}$. An illustration of such a target is given in Fig. 3(a). In practice, unrealizable targets occur, for

example, in machine learning classification tasks, see Sec. II N.

WP induces output perturbations $\delta z_{it} = \sum_{j=1}^N \xi_{ij}^{\text{WP}} r_{jt}$, which are linear combinations of the inputs. The components of $z^{\text{pert}} - z^*$ that are orthogonal to all inputs, d , thus always remain unchanged, irrespective of the current student weights and applied perturbations. This leads to the same constant additive contribution $E_{\text{opt}} = 1/(2T)\text{tr}[dd^T]$ to the perturbed and unperturbed errors E^{pert} and E [Eq. (8)]. It cancels in the weight update rule [Eq. (3)] such that WP learning is unchanged and Eq. (10) still holds when shifting its final error to $E_{f,\text{unr}}^{\text{WP}} = E_f^{\text{WP}} + E_{\text{opt}}$ [Eq. (S119) [28]]. E_{opt} marks the minimum error that necessarily remains even with $w = w^*$, due to the unrealizable components.

In contrast, NP perturbs the outputs with white noise. This noise generally has a nonzero component along d , which affects E_{pert} . Since such a component cannot be realized through an update of the weights, the resulting change of the error is noninstructive and represents reward noise that adds noise to the updates. Consequently, while the convergence factor a remains unchanged, the final error of NP increases more strongly than for WP [Fig. 3(b)]. At the optimal learning rate, the increase in final error due to unrealizable target components is twice that of WP: $E_{f,\text{unr}}^{\text{NP}} \approx E_f^{\text{NP}} + 2E_{\text{opt}}$ [Supplemental Material Sec. I, Eq. (S58) [28]]. For $E_{\text{opt}} > E_f^{\text{NP}}$, the coupling of node perturbations to unrealizable target components becomes NP's main contribution to the part of the final error that exceeds the unavoidable E_{opt} .

F. Input and perturbation correlations

Our results hold for very general sequences. In particular, correlations in the input may be present or absent without affecting learning. Furthermore, the T input-output relations can be temporally permuted. These invariances follow

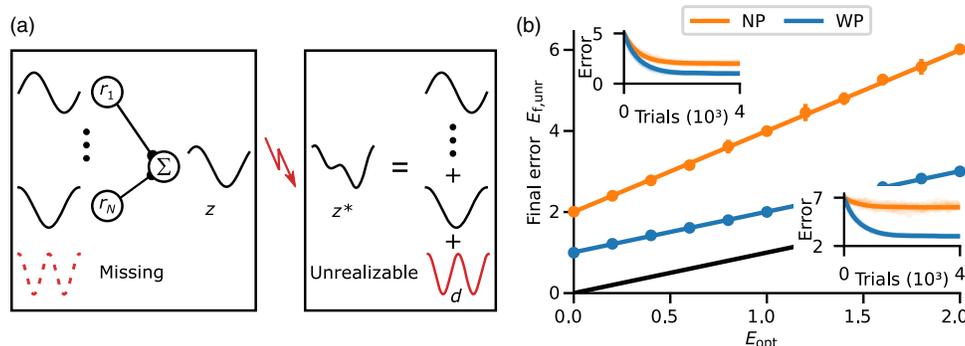


FIG. 3. Unrealizable target components harm NP learning. (a) General targets may contain a component that is perpendicular to any input and, thus, unrealizable (red). (b) Final error after convergence as a function of the error E_{opt} that necessarily remains, since the target is unrealizable. The final error of WP (blue) is shifted only by E_{opt} ; that of NP (orange) increases twice as fast, by approximately $2E_{\text{opt}}$. Data points: mean and standard deviation (averaged over ten simulated runs) of the final error. Curves: theoretical predictions. Black: E_{opt} . Insets: error dynamics for $E_{\text{opt}} = 0$ (left) and $E_{\text{opt}} = 2$ (right). Parameters: $M = 10$, $N, T = 100$, $N_{\text{eff}} = 50$, and $\sigma_{\text{eff}} = 0.04$.

straightforwardly from the weight update equations (Supplemental Material Sec. V [28]). We note that overly strong input correlations reduce the effective temporal dimension of the input such that its effective dimension cannot be kept up [cf. Fig. 1(c)], which affects learning.

WP generates output perturbations that are automatically adapted to the inputs (Sec. II E). To adapt NP to tasks with temporal input correlations, we modify it to time-correlated NP, NP_c. We generate the correlated perturbations by temporal low-pass filtering of white noise, with filtering time constant $\tau_{\text{corr}}^{\text{pert}}$. Other possibilities are to compose them from low-frequency Fourier modes or to use perturbations that are piecewise constant. As the correlation of the perturbation decays during $\tau_{\text{corr}}^{\text{pert}}$, for a reasonable representation of the perturbation we need to specify it at $T_{\text{eff}}^{\text{pert}} = T/(\tau_{\text{corr}}^{\text{pert}} + 1)$ time points with temporal distance $\tau_{\text{corr}}^{\text{pert}} + 1$. For vanishingly short correlations, $\tau_{\text{corr}}^{\text{pert}} = 0$, $T_{\text{eff}}^{\text{pert}} = T$ and NP_c equals NP; a large filtering time constant generates perturbations that vary slowly and have small effective temporal dimension $T_{\text{eff}}^{\text{pert}}$.

If the inputs have similar temporal correlations, the filtering of the perturbation concentrates its perturbative

power on the realizable output subspace, since this is spanned by the inputs. This reduces the update noise due to the quadratic reward noise from finite perturbation sizes [cf. Eq. (S112) [28]], because perturbations that are more aligned with the output gradient require less amplification to yield a sizable expected update along the weight gradient. Furthermore, it reduces the linear reward noise resulting from coupling to unrealizable target components [cf. Eqs. (S108) and (S80) [28]]. Both noise reductions lower the final error. We find that NP_c robustly improves upon NP over a range of filtering time constants similar to that of the inputs (Fig. 4). If the perturbations become too smooth, however, they also lose power in the realizable output subspace, which slows the learning of realizable target components with higher frequencies down. This is because the suppressed modes with their comparably small amplitude contribute little to the projection of the perturbations onto the T -dimensional output gradients (despite being prominent in the latter), which results in a small contribution to the error signal [Eq. (S3) [28]]. This contribution, in turn, determines the magnitude of the mean weight update used to match an output mode to the target. Consequently, the update part used to match the high-frequency modes is small, and the matching takes a

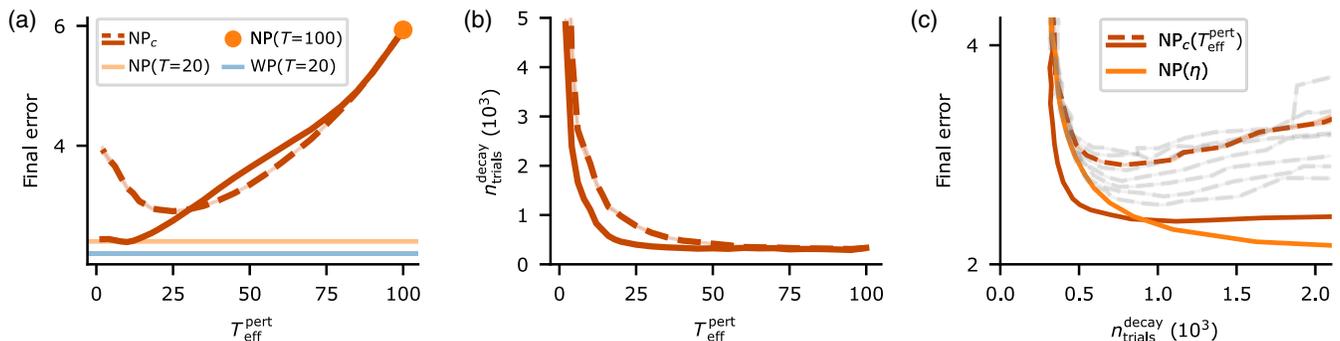


FIG. 4. Temporally correlated perturbations improve NP if the input has similar correlations. (a) Final error of NP_c versus the effective time dimension $T_{\text{eff}}^{\text{pert}}$ of its perturbations; smaller $T_{\text{eff}}^{\text{pert}}$ means smoother perturbations. The inputs are constructed with $T_{\text{eff}}^{\text{input}} = 20$ (red continuous line) or $T_{\text{eff}}^{\text{input}} = 100$ (uncorrelated, red dashed line). The final error of NP_c decreases compared to that of NP (orange dot) for $T_{\text{eff}}^{\text{pert}} < T = 100$. For correlated inputs and $T_{\text{eff}}^{\text{pert}} \leq T_{\text{eff}}^{\text{input}} = 20$, NP_c's final error reaches that of NP in a reduced task with $T = 20$ (orange line; WP, blue line); for uncorrelated inputs and small $T_{\text{eff}}^{\text{pert}}$, it increases again. This increase is not caused by lack of convergence. (b) Convergence time, measured as the number of trials until 95% of the final error correction is reached, versus $T_{\text{eff}}^{\text{pert}}$. Learning gets considerably slower if the correlation times of the perturbation are longer than those of the input, i.e., $T_{\text{eff}}^{\text{pert}} < T_{\text{eff}}^{\text{input}}$. The same curve markings as in (a). (c) Simultaneous plot of error and convergence time of NP_c [red line, data as in (a) and (b) for correlated (solid line) or uncorrelated input (dashed line); thus, curves are obtained by varying $T_{\text{eff}}^{\text{pert}}$; faint gray line, curves for uncorrelated input with different learning rates η] and NP (orange, curves obtained by varying the learning rate η). For correlated inputs and similarly correlated perturbations, NP_c yields a true improvement over NP: It has simultaneously a smaller error and smaller convergence time. For considerably longer correlation times, the final error saturates, but the convergence time increases (region with $n_{\text{trials}}^{\text{decay}} > 1000$). NP with reduced learning rate η there yields a smaller final error at the same convergence time. For uncorrelated input, NP_c does not yield a true improvement. In particular, the smaller final error (a) can be achieved at smaller convergence times by NP with reduced η , even if NP_c is allowed to additionally adapt its learning rate (gray curves). Parameters: $N_{\text{eff}} = M = 10$, $N = T = 100$, $\sigma_{\text{eff}} = 0.04$, $E_{\text{opt}} = 2$, and 20000 trials. (a),(b) Mean and standard error of the mean (SEM, partly occluded). (c) Mean values and SEM of final error (red and orange lines). In gray curves, η is modified relative to the red curve by a factor ranging from 0.5 (lowest curve) to 1.2 (highest curve) in steps of 0.1.

long time. With optimal perturbation correlation time $\tau_{\text{corr}}^{\text{pert}} \approx \tau_{\text{corr}}^{\text{input}}$ (Sec. IV), NPc approaches the performance (as measured by final error and convergence speed) of NP on a substitute task with $T_{\text{eff}}^{\text{input}}$ bins; see Figs. 4(a) and 4(b). This indicates that our analytical considerations for NP transfer to those of NPc with optimally correlated noise if we take into account that the correlations effectively reduce the trial duration to $T_{\text{eff}}^{\text{input}}$. In particular, in Fig. 4, optimal NPc performs worse than WP due to the low dimensionality and the remaining temporal extension of the task as well as the unrealizable target components.

G. Multiple subtasks

In general learning tasks, inputs and targets may vary from trial to trial. To obtain an intuition for how this affects the speed of WP and NP learning, we here consider a simplified case: The goal is to solve a task with an overall effective input dimension of $N_{\text{eff}}^{\text{task}}$. The task has the same properties as the tasks before where each trial was identical. In particular, it has $N_{\text{eff}}^{\text{task}}$ orthogonal latent inputs of strength α^2 , and the inputs are rotated such that only the first $N_{\text{eff}}^{\text{task}}$ inputs are nonzero. The task is, however, not presented as a whole, but in pieces: In each trial, a random subset of $N_{\text{eff}}^{\text{trial}}$ out of the first $N_{\text{eff}}^{\text{task}}$ inputs are active to train the network. The error in an individual trial then depends only on its $MN_{\text{eff}}^{\text{trial}}$ trial-relevant weights, while the performance on the full task depends on the $MN_{\text{eff}}^{\text{task}}$ task-relevant weights. The ratio $N_{\text{eff}}^{\text{task}}/N_{\text{eff}}^{\text{trial}} = P$ marks the number of trials needed to gather information on all task-relevant weights.

NP updates only the weights relevant in a trial (Sec. II D). Also, for tasks consisting of multiple subtasks, it can thus operate at the learning rate that is optimal for a trial, $\eta_{\text{NP}}^* = 1/[(MN_{\text{eff}}^{\text{trial}} + 2)\alpha^2]$ [cf. Eq. (12)]. Because an update improves only $MN_{\text{eff}}^{\text{trial}}$ of the $MN_{\text{eff}}^{\text{task}}$ task-relevant weights, the convergence rate $-\ln a \approx 1 - a$ of the expected error, averaged over the input distribution, is smaller by a factor of $1/P$ than for a single input pattern [Supplemental Material Sec. III, Eq. (S79) [28]]:

$$a_{\text{NP}}^* = 1 - \frac{1}{P MN_{\text{eff}}^{\text{trial}} + 2}. \quad (15)$$

WP, on the other hand, updates all weights such that the weights that are irrelevant for the trial are changed randomly (Sec. II D). This worsens the performance for the inputs of other trials. Because there are now $MN_{\text{eff}}^{\text{task}}$ task-relevant weights whose fluctuations hinder learning, WP has an optimal learning rate of only $\eta_{\text{WP}}^* = 1/[(MN_{\text{eff}}^{\text{task}} + 2)\alpha^2]$. As for NP, each trial's progress is only on $1/P$ of the task-relevant weights, such that the optimal convergence factor for WP on the full task is [Supplemental Material Sec. III, Eq. (S78) [28]]

$$a_{\text{WP}}^* = 1 - \frac{1}{P MN_{\text{eff}}^{\text{task}} + 2}. \quad (16)$$

The convergence of WP is, thus, slower than that of NP by roughly $1/P$, the ratio of $N_{\text{eff}}^{\text{task}}$ and $N_{\text{eff}}^{\text{trial}}$ [Fig. 6(c)].

Our results have concrete implications for learning of multiple actions such as sequences of movements [32]. They can be learned by splitting them into subsets, which are called (mini)batches in machine learning. If we assume for simplicity that individual data points are pairwise orthogonal and have no time dimension, each batch corresponds in our terminology to a subtask, the number of batches to P , the dimensionality of the input data to $N_{\text{eff}}^{\text{task}}$, and the batch size N_{batch} to $N_{\text{eff}}^{\text{trial}}$. For $MN_{\text{eff}}^{\text{trial}} \gg 2$, Eqs. (15) and (16) thus imply that the convergence rate of NP is independent of the batch size while that of WP is proportional to the batch size and reaches NP's convergence rate for full batch learning (Supplemental Material, Fig. S5 [28]). The same holds for the optimal learning rates as α^2 scales inversely with the batch size [Supplemental Material Sec. III, Eqs. (S90), (S92), and (S93) [28]].

H. Comparison with Ref. [9]

Reference [9] investigates how a student network learns the responses of a teacher network to arbitrary input with GD, NP, and WP, using patterns without temporal extent. In contrast to our tasks with typically $N_{\text{eff}} < N$ or $N_{\text{eff}}^{\text{task}} < N$, successful learning in Ref. [9] requires one to match all weights of the teacher network. In other words, the student network is trained at its capacity limit, where (only) one weight configuration fulfills the task. It learns from random input patterns and the teacher's responses to them. This is a special case of the setup introduced in Sec. II G, where (i) the task dimension equals the input dimension, $N_{\text{eff}}^{\text{task}} = N$ (since the employed random input patterns lie in arbitrary directions of input space), and (ii) there is no temporal extent of the tasks, $T = 1$. The latter implies that a single input pattern has effective dimension $N_{\text{eff}}^{\text{trial}} = 1$: All N inputs in a single pattern are linearly dependent, since they are scalar, temporal vectors of length 1 (we could rotate the input pattern such that it has only one nonzero entry).

A comparison of our results for the convergence speed in the described special case with those of Ref. [9] reveals that they agree approximately when straightforwardly setting $N_{\text{eff}}^{\text{trial}} = 1$, $N_{\text{eff}}^{\text{task}} = N$, and $P = N_{\text{eff}}^{\text{task}}/N_{\text{eff}}^{\text{trial}} = N$ in Eqs. (15) and (16):

| Algorithm | Our results | Results from Ref. [9] |
|-----------|--|--|
| WP | $a_{\text{WP}}^* = 1 - \frac{1}{N MN + 2}$ | $a_{\text{WP}}^* = 1 - \frac{1}{N+2} \frac{1}{MN+2}$ |
| NP | $a_{\text{NP}}^* = 1 - \frac{1}{N M+1+2}$ | $a_{\text{NP}}^* = 1 - \frac{1}{N+2} \frac{1}{M+2}$ |

The reason for the remaining difference is that the individual inputs in Ref. [9] are drawn from a Gaussian distribution, without subsequent normalization to the same strength like in our scheme. We obtain $P \rightarrow N + 2$ and, thus, perfect

agreement if we adapt our setting such that the summed input strength $\sum_{j=1}^N \alpha_j^2$ fluctuates as it does in Ref. [9].

I. Input noise

Inputs in biological neural networks are noisy. To investigate the impact of input noise on WP and NP, we add white noise to all input neurons. Noise in the relevant inputs causes a finite error that remains even for optimal weights. Moreover, the irrelevant weights are not completely irrelevant anymore: They mediate noise (instead of zero) input, have an optimal value of zero, and lead to significant output error if they become too large. Since the input noise is different in the perturbed and unperturbed trials, it becomes a source of additional reward noise. We, thus, expect that larger input noise requires stronger weight or node perturbations to ensure that the beneficial, error gradient-related part of the reward signal is not dominated by reward noise (cf. also Ref. [33]). Furthermore, an increase in overall noise and additional weights that become more and more important should require the integration of more trials to extract gradient information. We, therefore, expect a reduction of the optimal learning rate with increasing input noise strength. Our numerical simulations confirm these points; see Fig. 5(a) (increase in task error and optimal error) and Supplemental Material Fig. S6 [28] (estimation of optimal learning rates and perturbation sizes).

We find that in the presence of input noise the irrelevant weights diffuse in WP and NP [Figs. 2(b)(iii) and 5(b)]. They settle at a finite spread, which contributes to the error [Fig. 5(c)]. The diffusion stops because WP and NP update irrelevant weights on average toward zero, due to their generation of errors. For WP, the final spread increases with decreasing noise strength and reaches infinity for zero noise. [Note that Fig. 5(b) does not show the final, stationary spread.] For NP, the diffusion of irrelevant weights is caused by their noise-induced updates and is in contrast to the noise-free case. Their final spread is independent of the noise strength and discontinuously drops to zero at zero noise. To explain this, we identify (weak) noise input with (weak) deterministic input and apply our findings for noise-free networks with inhomogeneous input strength distribution: For WP, each weight contributes equally to the final error, and the final spread scales like one over the square root of the input strength [Supplemental Material Sec. IV, Eq. (S126) [28]]. For NP, weights related to small inputs contribute only little to the error, and the final spread is independent of input strength [Supplemental Material Sec. IV, Eq. (S128) [28]].

The simulations [Fig. 5(a)] and our analytical understanding also show that for finite learning time or when introducing a weight-limiting mechanism there is no discontinuity in the error when increasing the input noise

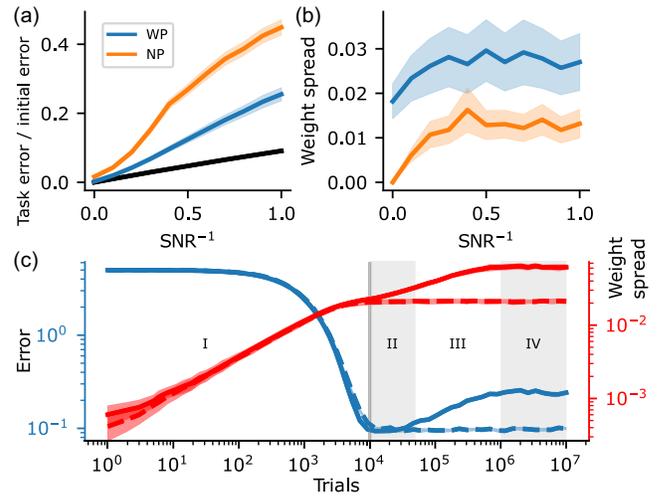


FIG. 5. Influence of input noise on the task error and on the spread of irrelevant weights. (a) Task error (fraction of the initial error) of WP (blue) and NP (orange) at 10 000 trials (curves, means; shading, standard deviations) as a function of inverse signal-to-noise ratio (SNR^{-1}). The noise in the relevant inputs renders the optimal error (black) nonzero. The plot covers SNRs ranging from infinity down to 1. (b) Spread (standard deviation) of irrelevant weights (mean and standard deviation) at 10 000 trials. (c) Evolution of error (blue, mean and SEM) and spread of irrelevant weights (red, mean and SEM) in WP proceeds for weak input noise ($\text{SNR}^{-1} = 0.1$) in four phases (solid curve, phases indicated by numerals). Appropriate weight decay stops the dynamics in phase II and induces long-term results as in the noise-free case [dashed curve; cf. also Fig. 2(b)(ii)]. (a) and (b) use the learning rates that minimize error after 10000 trials [gray vertical line in (c)] for a given noise level (Supplemental Material, Fig. S6 [28]). This implies that WP’s error in (a) is the one in phase II. Parameters: $M = N_{\text{eff}} = 10$, $N = T = 100$, $\alpha^2 = N/N_{\text{eff}} = 10$, $\sigma_{\text{eff}} = 0.04$, $E_{\text{opt}} = 0$, and $\gamma_{\text{WD}} = 0.99998$.

from zero to some finite value. The previous error analysis, therefore, stays valid as the limit of weak input noise. For WP, this is because limiting the learning time or the weights limits the final spread of irrelevant weights. This happens in such a way that sufficiently small noise has only a negligible effect on the output (Supplemental Material Sec. IV [28]). For NP and weak input noise, the final spread of irrelevant weights is approximately equal to that of the relevant weights and, thus, also limited.

Finally, we observe that WP learning proceeds for weak input noise in four phases [Fig. 5(c)]. In the first phase, the relevant weights are learned, such that the error decreases. Since the noise is small, the error due to its presence in relevant inputs is small. In the second phase, the error remains approximately constant, at a low level. In the third phase, the irrelevant weights, which have been diffusing all the time, become so large that they amplify the input noise sufficiently to influence the output despite the small noise amplitude. The error therefore increases. In the fourth and final phase, this error becomes so large that WP counteracts the further diffusion of weights. The error

therefore saturates, at a high level. The four phases can be clearly temporally separated. [Note the logarithmic axis scale chosen in Fig. 5(c).] We note that we observe divergence of the error for sufficiently large irrelevant weights if the learning rate is too large or the perturbations are too weak (Supplemental Material, Fig. S6 [28]). If learning stops in the second phase, the contributions from irrelevant weights can be neglected. The same holds if a limiting mechanism stops the diffusion of irrelevant weights at the level that they reach in the second phase, while only mildly affecting the relevant weights, because they converge at a shorter timescale [Fig. 5(c), dashed line: network with weight decay; cf. also Fig. 2(b)(ii)]. Interestingly, WP can then reach a lower (final) error than NP and is less affected by input noise; cf. Fig. 5(a).

J. Conclusions from the theoretical analysis and new learning rules

Our theoretical analysis reveals a simple reason for the differences between WP and NP: WP produces better perturbations, while NP better solves the credit assignment problem. Output perturbations caused by WP lie, in contrast to NP, always in the realizable output subspace and do not interfere with unrealizable target components. On the other hand, NP updates only (trial-)relevant weights, while WP updates all weights such that the (trial-)irrelevant weights change randomly. Small input noise does not change the overall picture. When single trials capture only a small part of the full task, WP learning slows down. Training in batches reduces the disadvantage.

Based on these insights, we introduce two novel learning rules, WPO and hybrid perturbation (HP) [Figs. 6(a) and 6(b)]. WPO adds a simple modification to WP: not to update currently irrelevant weights, i.e., weights whose inputs are zero (or close to it). This solves part of WP's credit assignment problem, as changing the weights

does not improve the output, and it avoids diffusion of irrelevant weights. The improvement is especially large when inputs are sparse such that many inputs are (close to) zero [Figs. 6(a) and 6(c)], which might be frequently the case in biological neural networks [34–36]. HP aims to combine the advantages of WP and NP by generating the output perturbations like WP, through perturbing the weights, and generating updates like NP, using its eligibility trace. The learning rule performs well when all latent inputs have (approximately) the same strength α^2 [Figs. 6(b) and 6(c)].

WPO and HP perform for the tasks used in the theoretical analysis section as well as WP and NP or better than both [Fig. 6(c)]. WPO is, however, benefited by the assumption of rotated inputs (in contrast to WP and NP), as it renders the input maximally sparse. Furthermore, the latent inputs have equal strengths, benefiting HP. We observe only slight improvements of WPO over WP for the reservoir computing and MNIST (Modified National Institute of Standards and Technology database) task, due to the lack of coding sparseness in our networks. HP performs much worse than WP and NP in the reservoir computing and similar to NP in the MNIST task. We explain this by the relevance of weak inputs [Supplemental Material Sec. VI, Eq. (S147) [28]]. Adding appropriately equalizing preprocessing layers may mitigate HP's problems. Furthermore, weak inputs may be irrelevant for biological learning. Since HP generates perturbations of the same class as the inputs and suppresses the learning of weights related to small inputs, we expect it to also work well for correlated input (as in Fig. 4) and in the presence of input noise.

K. Simulated learning experiments

In the following, we apply WP and NP to more general networks and temporally extended tasks with nonlinearities. We cover reservoir computing for dynamical

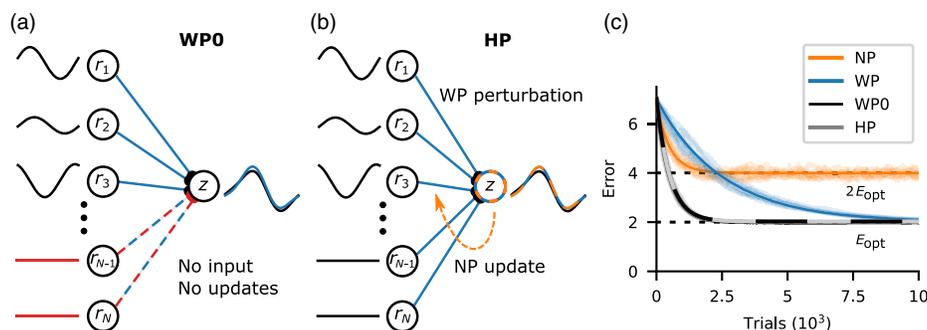


FIG. 6. New learning rules and learning of tasks consisting of multiple subtasks. (a) WPO does not update weights that mediate zero input, avoiding their diffusion. (b) Hybrid perturbation (HP): NP scheme with output perturbations induced by WP. (c) WP converges approximately $P = 5$ (number of subtasks) times slower than NP, but in the presence of unrealizable target components [or for finite σ_{eff} and $N_{\text{eff}}^{\text{trial}} < T$; Supplemental Material Sec. III, Eqs. (S87)–(S89) [28]] to a lower final error. For the used maximally sparse and equally strong inputs, WPO and HP combine the higher convergence rate of NP with the low final error of WP. Error curves (solid, theoretical predictions; shaded, ten exemplary runs) are for $M = 10$, $N, T = 100$, $N_{\text{eff}}^{\text{task}} = 50$, $N_{\text{eff}}^{\text{trial}} = 10$, negligible σ_{eff} , and $E_{\text{opt}} = 2$.

pattern generation, learning of recurrent weights in a delayed non-match-to-sample task, and a temporally extended, reward-based learning version of MNIST. The results confirm and extend our findings for analytically tractable tasks: They often show similar or superior performance of WP in temporally extended tasks relevant for biology and machine learning.

L. Reservoir computing-based drawing task

In reservoir computing schemes, an often low-dimensional input is given to a recurrent nonlinear network. The network effectively acts as a nonlinear filter bench: It expands the input and its recent past by applying a set of nonlinear functions to them. Each unit outputs one such function, which depends on the interactions within the recurrent network. Like a “computational reservoir,” the network thereby provides in its current state the results of manifold nonlinear computations on the current and past inputs. A desired result can be extracted by training a simple, often linear readout of the reservoir neurons’ activities. Reservoir computing schemes are widely used as models of neurobiological computations [37–41], since learning in them is simpler and seems more easily achievable with biological machinery than learning of full recurrent and multilayer networks. Furthermore, the schemes explain the presence of inhomogeneity and apparent randomness in neuron properties and connectivity in biological neural networks as helpful for enriching the computational reservoir. Here, we find that, when learning temporally extended output patterns with a reservoir computing scheme, WP can learn as well as or better than NP and NPc.

We consider a recurrently connected reservoir of $N = 500$ rate neurons driven by five external inputs of length $T = 500$. Inspired by the behaviorally relevant task of reproducing a movement from memory—here, drawing a figure—the task is to generate the x and y coordinates of a

butterfly trajectory [42,43] at the $M = 2$ outputs by training a linear readout [Fig. 7(a)]. The trajectory is nontrivial in that it is not realizable from the external inputs. In fact, it requires reading out from many reservoir modes [Fig. 7(b), dashed gray line].

Formally, the task is similar to the setting discussed above, with the difference that there is a wide distribution of different, nonzero input strengths α_μ^2 . The evolution of expected error is then best described by splitting the error $E = \sum_{\mu=1}^N E^\mu$ into different error components, each of which is associated with the weights that read out from a latent input $r_{\mu t}$ [Supplemental Material Sec. IV, Eq. (S103) [28]]. In WP and NP, the evolution of the error components follows a matrix exponential where different components decay at different rates and interfere with each other. Components that decay relatively quickly may be the main source of improvements in the beginning of training, whereas more slowly decaying components dominate the error toward the end. This effect can be seen in the approximately piecewise linear error decay in the logarithmic plot in Fig. 7(c).

Figure 7(c) compares the performance of WP, NP, and NPc in the drawing task. Perturbation size is finite, $\sigma_{\text{eff}} = 5 \times 10^{-3}$. WP converges faster initially, which may be typical for tasks with distributed input strengths [Supplemental Material Sec. IV, Eq. (S117) [28]]. It also achieves a lower final error. This is compatible with the observation that the effective dimension of the reservoir dynamics, as measured by the participation ratio ($\text{PR} \approx 5$), is much smaller than the temporal extent of the task: The resulting smaller effective perturbation dimension of WP (MPR versus MT for NP versus $\text{MT}_{\text{eff}}^{\text{pert}}$ for NPc) yields an advantage for WP [Figs. 1(b) and 4(a)]. NPc reaches a lower final error than NP. We optimize its $T_{\text{eff}}^{\text{pert}}$ using a parameter scan (Supplemental Material, Fig. S8 [28]). From our simulations, we cannot completely exclude that

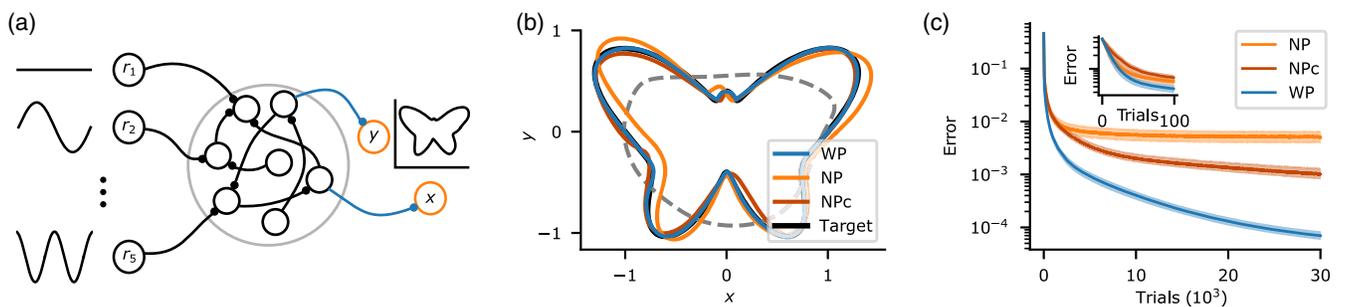


FIG. 7. WP outperforms NP on a reservoir computing-based drawing task. (a) Schematic of the recurrent, fixed reservoir receiving five external inputs. Only readout weights are learned. (b) Target (black) and final outputs of WP (blue), NP (orange), and NPc (red). A least squares fit (gray, dashed line) using only the first five principle components of the reservoir dynamics demonstrates that the task critically depends on reading out further, weaker dynamical components. (c) Error curves on a logarithmic scale. WP reaches a lower final error than NP and NPc, with NPc improving on NP; cf. also (b). Inset: early error evolution. There is a considerable improvement already during the first 50 trials. The curves show median (solid) and interquartile range between first and third quartile (shaded) over 1000 runs of the same network with different noise configurations.

the improvement is (in part) due to an effective decrease in learning rate, resulting from a longer correlation time in the perturbations than in some relevant inputs. For the biologically less relevant case of infinitesimal perturbation sizes, the performances of WP and NP are similar (Supplemental Material, Fig. S7 [28]) (compatible with Fig. 1 with $b = 0$). Toward larger trial numbers, the convergence of WP becomes slower: WP has difficulties with adjusting weights mediating weak inputs, since the impact of their perturbation on the error is small; the same effect underlies the weight diffusion in Fig. 2. Simulations indicate that the convergence is slower only by a constant factor of the order of 1 and that the optimal learning rate can be well estimated from the participation ratio (Supplemental Material, Fig. S9 [28]).

M. Delayed non-match-to-sample task

To ensure analytical tractability and for simplicity, so far we made a few biologically implausible assumptions. Specifically, only connection weights to linear units were trained, each trial consisted of a perturbed and an unperturbed run, and mostly the exact same input was used in each trial. In the following, we show that our findings generalize to settings without these assumptions. For this, we consider the learning of a DNMS task (temporal XOR) by nonlinear recurrent networks. DNMS tasks and closely related variants are widely used in both experiment [44] and theory [14,45], where they serve as simple working memory-reliant, not linearly separable decision-making tasks. We use the same setting as Ref. [14], which shows that a new variant of NP is able to solve the DNMS task. In particular, the setting is not adjusted to WP. The network consists of 200 nonlinear rate neurons receiving input from two external units u_1 and u_2 . One of the network neurons, whose rate we denote with z , serves as its output [Fig. 8(a)]. In each trial, the network receives two input pulses, where each pulse is a 200-ms-long period with either u_1 or u_2 set to 1, and subsequently has to output 1 for 200 ms if different inputs are presented and -1 if the same inputs are presented [Fig. 8(b)]. There is a 200-ms-long delay period after each input pulse.

We train all recurrent weights using the usual update rules [Eqs. (3) and (6)] but replace the error of the unperturbed trial by an exponential average of the errors of the previous trials [12–14]. Hence, each trial now consists only of a perturbed and not additionally an unperturbed run. We first assume that the exact perturbations ξ are accessible for the weight update, which seems biologically plausible for WP (cf. Sec. III) but less so for NP (cf. Sec. III and Ref. [14]). Therefore, we also compare WP and NP to the biologically plausible version of NP proposed by Ref. [14], which avoids this assumption: In the weight update rule, it approximates the exact node perturbations ξ^{NP} with a nonlinearly modulated difference

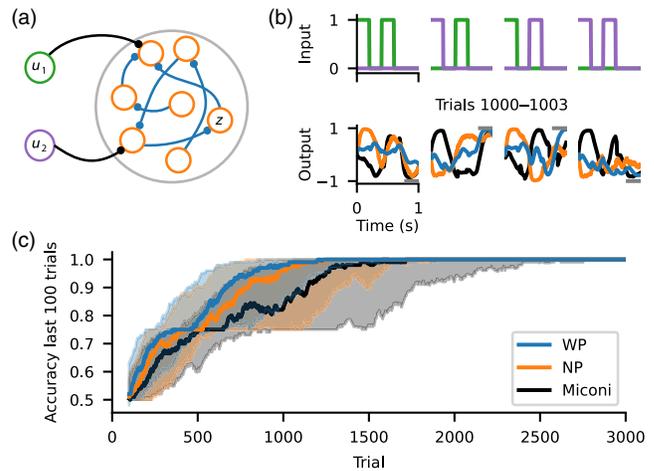


FIG. 8. WP performs as well as NP on a DNMS task. (a) Schematic of the recurrent network with inputs u_1 and u_2 and output z . All network weights are learned; i.e., for WP, all network weights (blue) are perturbed, and for NP, all network nodes (orange) are perturbed. (b) Inputs and outputs during example trials. Top row: inputs u_1 (green) and u_2 (purple) for the four different trial types. Bottom row: outputs for WP (blue), NP (orange), and the version of NP proposed by Ref. [14] (black) for trials 1000–1003 for the inputs shown above. Gray bars show target outputs. (c) Accuracy during training. WP (blue) performs similarly well as NP (orange) and the version of NP used by Ref. [14] (black). There is a noticeable transient slowdown at an accuracy of 75%, which corresponds to the successful learning of three out of the four different trial types. Solid lines show the median, and shaded areas represent the interquartile range between the first and third quartile using 100 network instances.

between the momentary input to a neuron and its short term temporal average (see Sec. IV for more details).

Figure 8(c) shows the performance of the three update rules in terms of their accuracy over the last 100 trials, where a trial is considered successful if the mean absolute difference between z and the target output is smaller than 1. We find that all update rules learn the task comparably well and reach perfect accuracy within at most 2000 trials when considering the median of network instances. Thus, our previous findings that WP can perform as well as or better than NP in simplified settings extend to the considered biologically plausible setup. That means WP can perform well for nonlinear neuron models, recurrent connectivity, and when the error of the unperturbed network is not available. Furthermore, the results indicate that approximating the perturbation as in Ref. [14] only mildly impacts the performance of NP for the considered task. Finally, we apply NPC to the task, which does not yield an improvement over NP (Supplemental Material, Fig. S12 [28]). Together with the similar performance of WP and NP, this indicates that the temporal dimension of the perturbation has little effect on task performance, perhaps because the period in which the target value needs to be assumed is rather short and the output is otherwise unconstrained.

N. MNIST

Finally, we apply WP and NP to MNIST classification. We use batches of images to train the networks. Each time step thereby corresponds to the presentation of one image, and the networks receive error feedback only at the end of a batch. This allows us to test how well WP and NP work on a more complicated, temporally extended task and in networks with a multilayer structure. In addition, it allows us to study how our analytical results for the learning of multiple input patterns (Sec. II G) extend to real-world tasks.

We use a two-layer feed-forward network with ten output neurons and 100 neurons in the hidden layer [Fig. 9(a)]. It learns via the rules [Eqs. (3) and (6)], where T equals the batch size N_{batch} . Hence, the perturbation is different for each image in the case of NP, while it is the same for WP. We test WP and NP for batch sizes of $N_{\text{batch}} \in \{1, 10, 100, 1000\}$. For each batch size, we determine the best-performing learning rates η and perturbation strengths σ_{WP}^2 and σ_{NP}^2 via grid searches. The perturbation strength has, however, little impact on performance, indicating that the final error is not restricted by reward noise due to finite size perturbations [Eqs. (13) and (14)].

We find that for WP the performance improves drastically with increasing batch size [Fig. 9(b)]. The final test accuracy is only about 69% for a batch size of 1 but reaches 92% for $N_{\text{batch}} = 1000$. Simultaneously, the optimal learning rate increases strongly, by a factor of approximately 50 [Supplemental Material, Fig. S10(c) and Table S2 [28]]. For comparison, the stochastic gradient descent (SGD) rule, which implements supervised not reward-based learning, reaches accuracies of 95%–98% for the considered batch sizes. In contrast, the learning curves of NP appear to be entirely independent of

the batch size [Fig. 9(b)]; the final test accuracy is always about 86% and the optimal learning rate is constant as well. We also apply NPc to the task. The inputs are temporally uncorrelated, because the elements of the batches are drawn randomly. Based on our previous observations for uncorrelated input [Fig. 4(c)], we therefore expect that NPc performs similar to or worse than NP. The numerical experiments confirm this: Performance deteriorates with increasing perturbation correlation time; the effect is more pronounced with larger batch size (Supplemental Material Fig. S12 [28]). In conclusion, larger batch sizes, as commonly used in machine learning, favor WP, while smaller batch sizes favor NP (and NPc).

An improvement of WP with batch size and NP's independence of it are in agreement with our theoretical analysis Sec. II G. However, from this analysis we also expected that WP's learning rate can reach at most that of NP for large batch size. NP's slower convergence suggests that it is more susceptible to deviations of the network architecture from linear, single-layer networks. Indeed, when using single-layer networks, NP's performance improves, while the opposite holds for WP and SGD (Supplemental Material Fig. S11 [28]). In a single-layer linear network with realizable targets, NP performs better than WP even for large batch sizes (Supplemental Material Fig. S11 [28]), consistent with our analytical findings that training with different subtasks (here, different batches) harms WP (Sec. II G) while the absence of unrealizable targets benefits NP (Sec. II E).

The results are particularly remarkable when naively comparing the number of perturbed nodes and weights: For the network considered here, there are only 110 output and hidden nodes but 79 510 weights (including biases). Nevertheless, WP can clearly outperform NP. Also, a

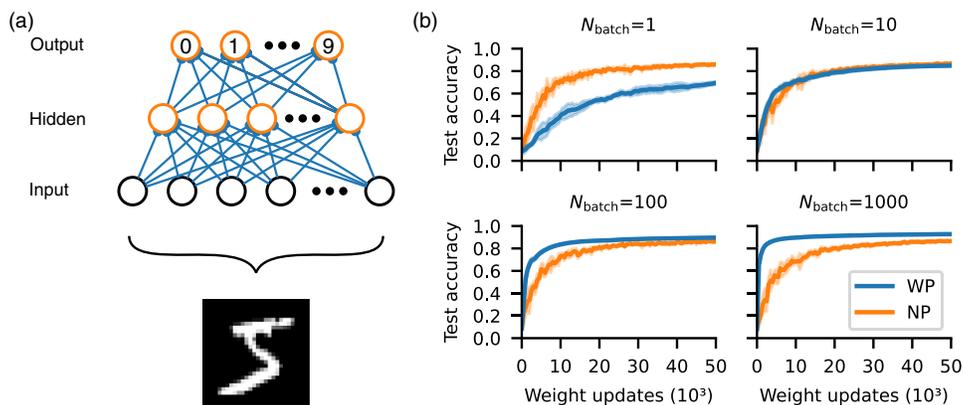


FIG. 9. WP can outperform NP on MNIST. (a) Schematic of the used fully connected, two-layer network. All network weights are learned; i.e., for WP all network weights (blue) are perturbed, and for NP all network nodes (orange) are perturbed. (b) Test accuracy as a function of the number of weight updates for WP (blue) and NP (orange) for different batch sizes. NP does not profit from increasing the batch size and always reaches a final accuracy of approximately 86%. WP improves considerably with increasing batch sizes and reaches a final accuracy of approximately 92% for $N_{\text{batch}} = 1000$. Solid lines show the mean, and shaded areas show the standard deviation using five network instances.

comparison of the actual perturbation dimensions cannot explain the better performance of WP in, e.g., Fig. 9(b) lower left (WP pert. dim., 79 510; NP pert. dim., $110 \times T = 11\,000$).

III. DISCUSSION

Our results show that WP performs better than NP for tasks where long trials capture most of the task’s content. This might seem paradoxical, as NP incorporates more structural knowledge on the network, namely, the linear summation of inputs. However, WP accounts for the fact that the weights in a neural network are (approximately) static. Furthermore, by perturbing the weights, it implicitly accounts for low input dimensionality and generates only realizable output changes. Therefore, it generates better tentative perturbations. This leads to less noise in the reward signal and better performance (smaller final error and sometimes faster convergence) in the tasks where WP is superior to NP.

Our theoretical analysis shows that the lower noise in WP first results from an effective perturbation dimension that is lower than NP’s if the temporal extent of a task is larger than its input dimensionality, $T > N_{\text{eff}}$. Second, factors such as the attempt of NP to realize unrealizable targets contribute. Temporally extended tasks with durations on the order of seconds and low dimensionality occur frequently in biology, for example, in motor learning and working memory tasks. In line with perturbation-based learning, biological movements are endowed with noise, which helps their learning and refinement [46]. The associated neuronal dynamics in the brain are confined to a low-dimensional space, a property shared by many types of biological and artificial neural network activity [47–49]. The dynamics for simple movements as investigated in typical experiments are embedded in spaces of dimension of the order of 10 [19]. This indicates low effective input dimensionality N_{eff} at the different processing stages. The effective muscle activation dimensionality is similarly low [20,50]. Neurons under *in vivo* conditions can faithfully follow input fluctuations on a timescale of 10 ms [51], and significant changes in neuronal trajectories happen on a timescale of 100 ms [19,21,52]. For the learning of a movement of duration 1 s, this suggests an effective temporal dimension of about 10–100 similar to the expected input dimension. This implies that WP as well as NP and NPc are promising candidates for the learning of simple movements. Our results indicate that WP is superior if the movements are longer lasting or lower dimensional.

We explicitly study the learning of movement generation (drawing task) and of a working memory task (DNMS). The numerical simulations show that WP performs similarly well or better compared to NP. In a task generally investigating the learning of complicated nonlinear, temporally extended input-output tasks (MNIST), WP outperforms NP as soon as the tasks have sufficient temporal extent.

As another concrete application, consider the learning of the single song in certain birds. A single, stereotypical input sequence in a “conductor area” (HVC) may drive the circuit [35,53]. The effective input dimension N_{eff} is, thus, at most as large as the temporal dimension T of the task. Based on recent experiments, Ref. [53] proposes that the output of the tutor and experimenter area (LMAN) is modified by reinforcement learning via NP, such that it guides the motor area (RA) to learn the right dynamics. Our analytical results predict that WP is as well or better suited to achieve this task, since $N_{\text{eff}} \leq T$. Earlier work suggests that WP [22] or NP [17] may directly mediate the learning of the connections from HVC to RA. Because of HVC’s very sparse activity, WP0 is highly suitable for such learning. Reward-based learning of mappings between conductor sequences and downstream neural networks may also be important for different kinds of precisely timed motor activity [54,55] and for sequential memory [56,57].

Biological neural networks are inherently noisy. We find that WP and NP induce two types of weight update noise: credit assignment and reward noise. We understand their impact analytically. Additional feedback or output noise implies additional reward noise with like effects. We, thus, additionally study only the impact of input noise on the learning of linear networks. Our simulations show that the convergence time and the final error increase with the input noise strength. The increase is smaller in WP than in NP. We further find that our results with zero noise are recovered in the limit of small noise compared to the strength of the relevant latent inputs, if the learning time is finite. The same holds for WP also if the weights are appropriately limited, for example, due to weight decay. The fact that animals can learn to perform tasks with high precision, i.e., with small final error, indicates that the case of small noise may be the biologically relevant one. The results of WP and NP learning with noisy inputs can be understood analytically from our findings on inhomogeneous input distributions (Supplemental Material Sec. IV [28]). Also, the considered DNMS and MNIST tasks contain input noise: the DNMS task because of the randomly chosen initial conditions and the MNIST task because of the naturally noise-afflicted input images. We conclude that the finding of similar or better performance of WP compared to NP in temporally extended, low-dimensional tasks may readily apply to biologically plausible, noisy networks.

WP and NP have biologically plausible implementations. NP requires that the plastic synapses can keep track of their input and the somatic perturbations (which may arrive from a tutor or experimenter neuron). Biologically plausible mechanisms for this have been proposed for tasks with both immediate reward [12,13] and reward at a temporally extended trial’s end [14]. Their underlying idea is to assume that the perturbation fluctuates more quickly than the other input. The present fluctuation can then be approximately isolated by subtracting a short-term

temporal average of the past overall input from the present one [12,13]. This difference replaces the injected perturbation in the eligibility trace. For tasks with late reward, the eligibility trace needs to integrate a nonlinearly modulated version of the described difference [14]. This prevents the cancellation of a perturbation's effect by the subsequent change in the average that it evokes, because the peak in the original perturbation is sharper and higher than the one in the average. We use this learning model of Ref. [14] in Fig. 8. The biological implementation of WP may be even simpler. A neural network needs to generate labile random weight changes and keep track of them. They should be approximately constant during a task and enhanced, deleted, or reversed by a subsequent reward signal. Experiments on timescales from minutes to days find spontaneous changes in the synaptic weights, which have similar strength as changes due to activity-dependent plasticity [31]. Such changes might generate the perturbations required for our WP scheme. Previous work suggests also synaptic unreliability to provide the perturbations for WP [58]. This fits into our scheme of static weight perturbations if neurons spike once during a trial or if they burst once and the synaptic transmission is restricted to a single time bin. Another source of the required randomness may be fluctuations of activity-dependent plasticity, while the deterministic baseline acts as a useful prior. If the baseline is unrelated to the task, it is with high probability orthogonal to task-relevant directions (due to the high-dimensional weight space) and not harm learning, similar to the weight diffusion in WP. In this way, the fluctuations of activity-dependent plasticity, rather than their deterministic part, may be the source of learning.

Modulation of weight changes by reward is observed in various experiments [59,60]. As an example, the potentiation of synapses is enhanced or reversed depending on the presence or absence of a temporally close dopamine reward signal [61]. Also, other factors play a role; potentiation can, for example, be reversed within a “grace period” of tens of minutes by a change of environment [62]. The consolidation and amplification of changes may be dependent on plasticity-related proteins, which are upregulated by reward and for which the synapses compete (synaptic tagging hypothesis) [60,63]. *A posteriori* modifications of tentative synaptic weight changes are also assumed in the reinforcement learning scheme of reward-modulated Hebbian plasticity [64,65], which is closely related to WP.

WP applies a random perturbation vector to the weights, measures the error change to obtain a reinforcement signal, and applies as weight update the perturbation modulated by the reinforcement signal. The improvement, therefore, follows on average the weight gradient. A related approach is to randomly perturb and accept the perturbation if it leads to a better

performance [22]. This simple instance of an evolutionary strategy [66,67] is also applicable if there is no gradient. Our results for WP suggest that this and related evolutionary learning strategies might benefit from tasks that are low dimensional, as reported previously [68], and not be harmed by their temporal extension. The sketched simple evolutionary learning may in the brain generate structural improvements: Experiments show that synaptic turnover occurs in the presence but also spontaneously in the absence of neuronal activity [69–72]. This may reflect the random elimination and creation of synapses and their consolidation by activity-dependent plasticity [73–78]. The basis of consolidation is that mainly weak synapses are removed such that strengthening through Hebbian learning causes the long-term presence of a synapse. Furthermore, Hebbian learning counteracts spontaneous synaptic weight changes, which could otherwise weaken useful synapses and ultimately lead to their removal. Network models show that restructuring with subsequent selective consolidation can recruit sparse available connectivity for task learning, prevent catastrophic forgetting, and may explain the benefits of dividing learning into several temporally distinct phases [73,75,77]. Furthermore, it may explain the experimental observation that there are commonly multiple synapses between connected neurons [74,76,78]. The signal for the strengthening of a tentatively established synapse may be interpreted as a reinforcement signal for its presence. This signal is generated if the pre- and postsynaptic neurons are coactive, due to external stimulation, or recall mediated by other, already strengthened synapses. In contrast to WP, the reinforcement signal is, therefore, specific to a neuronal connection (consisting of the possible multiple direct synapses from a presynaptic to a postsynaptic neuron), which simplifies learning. In particular, any useful new synapse is consolidated (unless the coactivity of the pair of connected neurons stops due to other changes). In contrast, in WP, tentatively applied useful weight changes can be easily reverted due to harmful ones in other parts of the weight perturbation vector. In a related model, Ref. [79], consolidatory strengthening is implemented with NP instead of Hebbian learning. This allows one to learn tasks based on a global reward signal. The synaptic weight fluctuations in the model induce random changes in task-irrelevant directions similar to the weight diffusion that we observe in WP. Our results suggest to directly exploit the synaptic weight fluctuations for consolidatory synaptic strengthening by using WP when learning low-dimensional and temporally extended reward-based tasks.

WP is proposed in several variants. They differ in (i) the task setup, for example, instantaneous [7,9] or temporally extended tasks [8,22,58], (ii) the implementing network, for example, rate [22] or spike-based [58,80] networks, (iii) the perturbation scheme, where all weights

are simultaneously perturbed [7,8,58] or only one weight at a time [15], (iv) the computation of the weight update, by correlating reward and perturbation [7–9,11,58] or direct estimation of the gradient components (for the single weight perturbation scheme) [9,15], (v) the estimation of the success of the perturbed network, which may involve a comparison of the obtained reward to an unperturbed baseline [8,9] or a running average [22,58] or it may consider the reward only [7,58], and (vi) the weight update, which may be proportional to the success of the network [7–9,11,58] or independent of its size as long as there is an improvement [22]. A similar diversity of NP variants exists [9–14,16,17,24,79,81,82].

The tasks considered in our article are temporally extended. The reward is provided at the end of the trial but influenced by earlier output states. This is consistent with many tasks in biology [1,5,14,22] and with the learning schemes by Refs. [8,10,14,22]. We choose a WP rule that is biologically plausible, as it involves simultaneous perturbations to all weights and correlates reward and weight change. The success measure compares the obtained reward to the reward of an unperturbed network in order to reduce the update noise [8,9]. In particular, this avoids unfavorable perturbations being associated with positive reward feedback. Finally, the weight update is proportional to the measured success in order to ensure that it occurs on average parallel to the reward gradient. The choices are identical to those by Refs. [8,9] for temporally not extended tasks. Specifically, the results in Ref. [9] appear as a special case of our results for multiple input patterns; if the task dimension is maximal, single trials have no temporal extent, and the inputs have fluctuating amplitude (see Sec. II H).

We choose the NP scheme such that it matches the WP scheme. It is a discrete-time version of the NP scheme proposed by Ref. [10] and an extension of the scheme by Ref. [9] to temporally extended tasks. In biologically plausible implementations of WP and NP, the reward should be compared to an intrinsically generated prediction, such as an average of previous rewards [12–14] or the reward of another perturbed trial [82]. In the delayed non-match-to-sample task, we thus replace our standard unperturbed baseline by such an average. This also allows a direct comparison with the NP scheme by Ref. [14]. In Sec. II I, the perturbed and unperturbed trials have different input noise, such that E is no longer the exact unperturbed counterpart of E^{pert} .

To exploit correlations in the inputs with a node perturbation learning rule, we introduce NPc, which is identical to standard NP apart from using temporally correlated, smoothed node perturbations. We find that the temporal correlations are usually beneficial if also the inputs are (similarly) correlated. This is in contrast to Ref. [13], which observes a detrimental effect already of short correlations for a node perturbation variant that

relies on high-frequency perturbations. Other previous studies inject white noise perturbations only [10,12,16,17]. Our simulations with linear networks indicate that the perturbation correlation time of NPc should optimally match that of the inputs; NPc then performs similar to NP in a task with reduced temporal extension.

NP is studied in various concrete neurobiological settings. Previous work uses feedforward networks with NP to model the learning of coordinate transforms in the visual system [83], birdsong [17,53], and motor output [12,84]. Reference [13] shows that reservoir computers with NP trained, fed back readouts can learn periodic inputs, routing, and working memory tasks. Reference [14] uses a fully plastic recurrent network for the learning of a delayed non-match-to-sample, a selective integration, and a motor control task. Finally, NP is often employed for reference and comparison [85–91]. WP is considered less in studies of neurobiological learning. It is implemented in early feedforward network models of birdsong [92] and binary output task learning [58,80]. Furthermore, it is occasionally used for comparison [86,88]. Very recently, Ref. [4] has shown that recurrent neural networks can be pretrained with WP and the reservoir computing scheme to thereafter learn with static weights to generate fixed point activity.

The results of our present article using feedforward, reservoir computing and fully plastic recurrent networks suggest that for many tasks WP is at least as suitable as NP, while the implementation may be even simpler. This indicates that WP is a useful benchmark and a similarly plausible model for learning in the brain as NP. Experimentally measurable features of the learning and weight dynamics may allow one to distinguish the learning rules in biological neural networks.

IV. MATERIALS AND METHODS

A. Analytical error dynamics

To analytically compute the dynamics of the expected error, we consider an arbitrary perturbation ξ . This determines the error change $E^{\text{pert}} - E$ and the resulting weight update Δw via Eqs. (3) and (6). Δw , in turn, determines the new weights and via Eq. (8) the error $E(n+1)$ after the update. $E(n+1)$ is, thus, a function of ξ , the weight mismatch $W(n)$ before the update, and the input correlations S :

$$E(n+1) = \frac{1}{2} \text{tr}\{[W(n) + \Delta w(\xi)]S[W(n) + \Delta w(\xi)]^T\}. \quad (17)$$

Averaging over perturbations and using Isserlis' theorem yields an equation for the expected error $\langle E(n+1) \rangle$. When assuming that all latent inputs have the same strength, $\langle E(n+1) \rangle$ becomes a function of the error $\langle E(n) \rangle$ before the update and the system parameters, leading to Eq. (9). The detailed derivation is given in Appendix B.

B. Numerical simulations accompanying the theoretical analysis

In the numerical experiments in Sec. II B, the N_{eff} nonzero inputs are orthonormal functions, superpositions of sines, scaled by $\alpha^2 = N/N_{\text{eff}}$ to keep the total input strength $\alpha^2 N_{\text{eff}}$ for different N_{eff} constant. Targets z_{it} are obtained by linearly combining these functions using teacher weights $w_{ij}^* = 0.1$ and adding as an unrealizable component a further, appropriately scaled, orthonormal function. Learning rates are η^* .

C. Input and perturbation correlations

NPc is applicable to general network dynamics (cf. Fig. 7 and Figs. S7, S8, and S12 [28]). In Fig. 4, we apply it to linear networks with correlated inputs that are constructed similar to the perturbations, by low-pass filtering N_{eff} white noise traces (filtering time constant $\tau_{\text{corr}}^{\text{input}}$ and effective temporal dimension $T_{\text{eff}}^{\text{input}}$). Subsequently, we orthonormalize them, which somewhat modifies the correlation times. The realizable components of a target are linear combinations of these correlated inputs weighted by $w_{ij}^* = 0.1$. We assume that there is an additional unrealizable component, which, for simplicity, contains all modes orthogonal to the inputs with equal strength, such that $E_{\text{opt}} = 2$.

For $T_{\text{eff}}^{\text{pert}} \geq T_{\text{eff}}^{\text{input}}$, NPc operates at $\eta_{\text{NPc}} = \eta_{\text{NP}}^*$, the optimal learning rate for NP. For $T_{\text{eff}}^{\text{pert}} < N_{\text{eff}}$, we use the optimal learning rate of NP for a task with reduced $N_{\text{eff}} = T_{\text{eff}}^{\text{pert}}$, i.e., $\eta_{\text{NPc}} \approx (N_{\text{eff}}/T_{\text{eff}}^{\text{pert}}) \cdot \eta_{\text{NP}}^*$. Our intuition is that perturbations with temporal dimension $T_{\text{eff}}^{\text{pert}} < N_{\text{eff}}$ can only improve an $MT_{\text{eff}}^{\text{pert}}$ -dimensional subspace of the weights. For NPc, $T_{\text{eff}}^{\text{pert}} < N_{\text{eff}}$ therefore reduces the learnable number of weights from MN_{eff} to effectively $MT_{\text{eff}}^{\text{pert}}$ independent ones. This is like in Fig. 1(c), gray curves, where T restricts N_{eff} . As effectively fewer weights are learned, a higher learning rate can be chosen. We perform additional simulations with an unadjusted (smaller) learning rate that confirms our choice, as convergence otherwise becomes much slower.

For a given $T_{\text{eff}}^{\text{pert}}$, simulations of 20 000 trials are repeated 1000 times (randomly generated inputs change between runs but not within a run). The final error of a run is computed by averaging over the last 500 trials (in which the error is approximately constant; Supplemental Material Fig. S4 [28]) to determine the mean over runs and its SEM. To determine $n_{\text{trials}}^{\text{decay}}$, we use ten samples of 100 runs each. For each sample, we compute the mean error over runs and additionally smooth it with a centered temporal running average of window size 20. $n_{\text{trials}}^{\text{decay}}$ is then the trial for which the described average drops for the first time below $E_{f,\text{unr}} + 0.05 \cdot [E(0) - E_{f,\text{unr}}]$. Figure 4(b) reports the mean and standard error of the mean of $n_{\text{trials}}^{\text{decay}}$ over all samples.

Figure 4(c) repeats the same analysis for NP with η varied from $0.05 \cdot \eta_{\text{NP}}^*$ to η_{NP}^* . Gray curves in Fig. 4(c) (NPc with η adjusted by a factor of 0.5, 0.6, ..., 1.2) use 100 repetitions.

D. Input noise

We extend the basic theory task, in which a single mapping from an effectively N_{eff} -dimensional input signal r_{jt} onto target outputs $z_{it}^* = \sum_{j=1}^N w_{ij}^* r_{jt}$ is learned, by adding independent white noise to each input at each time step:

$$r_{jt}^{\text{noisy}} = r_{jt} + \chi_{jt}, \quad \langle \chi_{jt} \chi_{ks} \rangle = \sigma_{\text{noise}}^2 \delta_{jk} \delta_{ts}. \quad (18)$$

As the added white noise has a rotationally symmetric distribution in the space of input neurons, we can still without loss of generality rotate the input space such that each of the first N_{eff} input neurons carries a signal component and additional noise, while the remaining inputs are purely noisy. We note that, because the noise in the task-relevant inputs is amplified by the weights, their optimal values are closer to zero than those of the noise-free task.

SNR is defined as the ratio of the total (summed) power in the input signal to that in the noise. Averages are taken over 100 repetitions [Figs. 5(a)–5(c)], the last 1000 of 100 000 trials [Figs. 5(a) and 5(b)], and over irrelevant weights [Fig. 5(c)].

E. Reservoir computing task

The $N = 500$ rate neurons of the fully connected recurrent reservoir network evolve according to

$$x_{jt} = \gamma x_{j,t-1} + (1 - \gamma) \left(\sum_{k=1}^N w_{jk}^{\text{rec}} r_{k,t-1} + \sum_{q=1}^{N_{\text{inputs}}} w_{jq}^{\text{in}} r_{qt}^{\text{in}} \right). \quad (19)$$

The rate of neuron k is $r_{kt} = \tanh(x_{kt})$. Their decay time constant is $\tau = 10$ time steps, i.e., $\gamma = e^{-1/\tau}$. Recurrent weights w^{rec} are drawn from a centered normal distribution; the weight matrix is thereafter normalized to ensure that the real part of its largest eigenvalue is $g_{\text{rec}} = 1$. Input weights w^{in} are drawn from $w_{jk}^{\text{in}} \sim \mathcal{N}(0, 1/N_{\text{in}})$. Creating various instances of such random networks shows that performance and participation ratio $\text{PR} = (\sum_{\mu=1}^N \alpha_{\mu}^2)^2 / \sum_{\mu=1}^N \alpha_{\mu}^4$ are rather independent of the instance. The participation ratio gives an estimate of the dimensionality of the reservoir dynamics [19,93]. Generally, we observe $\text{PR} \approx 5$; for example, in the network in Fig. 7, $\text{PR} \approx 5.3$. The $N_{\text{in}} = 5$ inputs to the reservoir are orthogonal to each other, $r_{1t}^{\text{in}} = 1$, $r_{2t}^{\text{in}} = \sqrt{2} \sin(\omega t)$, $r_{3t}^{\text{in}} = \sqrt{2} \cos(\omega t)$, $r_{4t}^{\text{in}} = \sqrt{2} \sin(2\omega t)$, $r_{5t}^{\text{in}} = \sqrt{2} \cos(2\omega t)$, $\omega = 2\pi/T$, and $T = 500$ time steps. The trained linear readout produces $M = 2$ outputs $z_{it} = \sum_{j=1}^N w_{ij} r_{jt}$. Their target z_{it}^* is, up to scaling, the same

as in Ref. [42]: $z_{1t}^* = \text{radius}_t \cos(\omega t)$, $z_{2t}^* = \text{radius}_t \sin(\omega t)$, with $\text{radius}_t = 0.1 \cdot [9 - \sin(\omega t) + 2 \sin(3\omega t) + 2 \sin(5\omega t) - \sin(7\omega t) + 3 \cos(2\omega t) - 2 \cos(4\omega t)]$. Already 100 time steps before the task starts, the reservoir is initialized and given external input. By the time the task begins, network activity is enslaved by the external input and settles down to a periodic orbit. Technically, we record the reservoir activity traces r_{jt} once for the entire training of w , because they are the same in each trial. The value of the participation ratio motivates us to construct an optimal readout reading out the largest five principal components via the least squares fit $w_{ij}^{\text{LS}} = \sum_{k=1}^N \sum_{t=1}^T z_{it}^* r_{kt} S_{kj}^{\text{pinv}}$ [Fig. 7(b), dashed gray line]. Here, S^{pinv} is the pseudoinverse of the reduced correlation matrix of the reservoir that is obtained by setting all eigenvalues of S except the largest five to zero. Including six principal components does not qualitatively change the result.

From the theoretical analysis Eq. (12), we obtain an estimate $\eta^* = 1/[(\text{MPR} + 2)\overline{\alpha^2}]$ for the optimal learning rate, by setting $N_{\text{eff}} \rightarrow \text{PR}$ and $\alpha^2 \rightarrow \overline{\alpha^2} = (1/\text{PR}) \sum_{\mu=1}^N \alpha_{\mu}^2$. $\overline{\alpha^2}$ is the strength of each latent input when we assume that the total input strength is generated by PR equally strong ones. We verify by a grid search that this estimated value yields for both WP and NP close to optimal performance, as measured by the error after 10 000 trials with infinitesimally small σ_{eff}^2 , showing that it indeed maximizes convergence speed. We, therefore, choose it as the learning rate for our task with infinitesimal (Supplemental Material Fig. S7 [28]) and also with finite perturbation size (Fig. 7), since the theoretical analysis yields independence of η^* from σ_{eff}^2 [Eq. (12)]. For NPC, we use the same learning rate as for NP and WP and determine the optimal $T_{\text{eff}}^{\text{pert,opt}} = 18$ by minimizing the error after 30000 trials (Supplemental Material Fig. S8 [28]). Simulations with finite perturbations use $\sigma_{\text{eff}} = 5 \times 10^{-3}$. A scan over σ_{eff} confirms that the final error depends quadratically on it, as predicted by the theory.

F. Delayed non-match-to-sample task

The fully connected recurrent network has $N = 200$ rate neurons. The dynamics of neuron i , $i = 4, \dots, N$, are governed by

$$\tau \dot{x}_i = -x_i(t) + \sum_{j=1}^N w_{ij}^{\text{rec}} r_j(t) + \sum_{q=1}^2 w_{iq}^{\text{in}} u_q(t), \quad (20)$$

with time constant $\tau = 30$ ms. The constant activations $x_1(t) = x_2(t) = 1$ and $x_3(t) = -1$ provide biases [14]. The rate of each neuron i , $i = 1, \dots, N$, is given by $r_i(t) = \tanh[x_i(t)]$. $z(t) = r_4(t)$ is the network output. We use the forward Euler-method with step size $dt = 1$ ms to simulate the dynamics and draw the initial activations from a uniform distribution, $x_i(0) \sim \mathcal{U}(-0.1, 0.1)$ for $i = 4, \dots, N$. Recurrent weights are drawn from a Gaussian distribution,

$w_{ij}^{\text{rec}} \sim \mathcal{N}(0, g^2/N)$, with $g = 1.5$. Input weights are drawn from a uniform distribution, $w_{iq}^{\text{in}} \sim \mathcal{U}(-1, 1)$.

All recurrent weights w_{ij}^{rec} are trained. The error function of WP and NP is the mean squared difference between the output z and the target within the last 200 ms of each trial. For each of the different trial types k , $k = 1, \dots, 4$, we use an exponential average of the previous errors $E^{\text{pert}}(n_k)$ for this trial type (n_k indexes the trials of type k) as the error baseline:

$$E_k(n_k) = E_k(n_k - 1) + \frac{1}{\tau_E} [E^{\text{pert}}(n_k) - E_k(n_k - 1)], \quad (21)$$

where $\tau_E = 4$. To get the best-performing learning parameters, we perform a grid search, which yields $\eta^{\text{WP}} = 1 \times 10^{-5}$, $\sigma^{\text{WP}} = 4.64 \times 10^{-3}$, $\eta^{\text{NP}} = 1 \times 10^{-5}$, and $\sigma^{\text{NP}} = 4.64 \times 10^{-1}$.

For the details of the version of NP proposed by Ref. [14], see this article. For the convenience of the reader, here we briefly mention the main differences to the vanilla NP version Eq. (6): For each network neuron, a node perturbation is applied at a simulation time step only with a probability of 0.3% and is drawn from a uniform distribution, $\xi \sim \mathcal{U}(-16, 16)$. The error is given by the absolute difference between output and target. Weight updates are computed via $\Delta w_{ij}^{\text{rec}}(n_k) = -\eta E_k(n_k - 1) [E^{\text{pert}}(n_k) - E_k(n_k - 1)] \sum_{t=1}^T [(x_{it} - \bar{x}_{it}) r_{j,t-1}]^3$ and clipped when they exceed $\pm 3 \times 10^{-4}$ (cf. code accompanying Ref. [14]). t indexes the simulation time step of each trial, T is the total number of simulation time steps per trial, and $\bar{x}_{it} = \bar{x}_{i,t-1} + (1/\tau_x)(x_{it} - \bar{x}_{i,t-1})$ is an exponential average of past activations. Parameter values are $\eta = 0.1$ and $\tau_x = \frac{20}{19}$.

G. MNIST classification task

The input layer of the fully connected feedforward network consists of 784 units encoding the pixel values of the data. The hidden layer consists of 100 neurons with tanh activation function and biases. The output layer consists of ten neurons, one for each single-digit number, with softmax activation function and biases. We use the standard training and test dataset but split the standard training data set into a training dataset of 50000 images and a validation dataset of 10 000 images. No preprocessing is done on the data. We employ vanilla WP Eq. (3), NP Eq. (6), or SGD to train all parameters of the network. The error function is the cross-entropy loss averaged over the batch of length $N_{\text{batch}} = T$. We also try to combine the gradient estimates obtained from WP and NP with Momentum, RMSProp, or Adam [29] but do not find an improvement of performance compared to the vanilla versions with carefully tuned parameters. The same holds for SGD. This may be because of the rather simple network architecture.

To obtain the best-performing parameters (the learning rate for all three algorithms and the standard deviation for WP

and NP), we perform a grid search for each of the considered batch sizes: For each parameter set, we train the network for 50 000 trials (i.e., weight updates) on the training dataset. We then select the best-performing parameter sets based on the final accuracy on the validation dataset and apply them to the test dataset. High final accuracy appears to concur with fast convergence speed, such that a comparison to our analytical results (where learning rate optimizes the convergence speed) seems justified.

The supporting data for this article are openly available from GitHub [94].

ACKNOWLEDGMENTS

We thank the German Federal Ministry of Education and Research (BMBF) for support via the Bernstein Network (Bernstein Award 2014, 01GQ1710).

APPENDIX A: LEARNING MODELS AND TASK PRINCIPLES

1. Mean updates for small perturbations

To calculate the average update of WP, one uses the linear order approximation valid for small perturbations [7,8]:

$$E^{\text{pert}} - E \approx \sum_{m=1}^M \sum_{k=1}^N \frac{\partial E}{\partial w_{mk}} \xi_{mk}^{\text{WP}}, \quad (\text{A1})$$

where the sum is over all weights in the network. With this and $\langle \xi_{ij}^{\text{WP}} \xi_{mk}^{\text{WP}} \rangle = \delta_{im} \delta_{jk} \sigma_{\text{WP}}^2$, where δ is the Kronecker delta, the WP update rule [Eq. (3)] yields

$$\langle \Delta w_{ij}^{\text{WP}} \rangle \approx -\frac{\eta}{\sigma_{\text{WP}}^2} \sum_{m=1}^M \sum_{k=1}^N \frac{\partial E}{\partial w_{mk}} \langle \xi_{mk}^{\text{WP}} \xi_{ij}^{\text{WP}} \rangle = -\eta \frac{\partial E}{\partial w_{ij}}; \quad (\text{A2})$$

i.e., the weight update is along the negative error gradient. The result's independence of σ_{WP} motivates the division by σ_{WP}^2 in Eq. (3).

NP perturbs for temporally extended tasks the total inputs $y_{it} = \sum_{k=1}^N w_{ik} r_{kt}$ of neurons $i = 1, \dots, M$ by the node perturbation vectors ξ_{it}^{NP} . Therefore, the change in the scalar error is, to linear order, given by the projection of ξ_{it}^{NP} onto the error gradient $\partial E / \partial y_{it}$ with respect to y_{it} [10]:

$$E^{\text{pert}} - E \approx \sum_{i=1}^M \sum_{t=1}^T \frac{\partial E}{\partial y_{it}} \xi_{it}^{\text{NP}}. \quad (\text{A3})$$

Since $\partial y_{it} / \partial w_{ij} = r_{jt}$, the gradients with respect to weights and sums of inputs are related via the chain rule by $\partial E / \partial w_{ij} = \sum_t \partial E / \partial y_{it} r_{jt}$. This reveals that the NP weight update [Eq. (6)] is on average along the negative error gradient:

$$\begin{aligned} \langle \Delta w_{ij}^{\text{NP}} \rangle &\approx -\frac{\eta}{\sigma_{\text{NP}}^2} \sum_{m=1}^M \sum_{s,t=1}^T \frac{\partial E}{\partial y_{mt}} \langle \xi_{mt}^{\text{NP}} \xi_{is}^{\text{NP}} \rangle r_{js} = -\eta \sum_{t=1}^T \frac{\partial E}{\partial y_{it}} r_{jt} \\ &= -\eta \frac{\partial E}{\partial w_{ij}}. \end{aligned} \quad (\text{A4})$$

Equations (A2) and (A4) hold for any error function. For corresponding more specific computations for the quadratic error function [Eq. (8)], see Appendix A3.

2. Dependence of weight update noise on error baseline

To show that the choice of E as error baseline in Eq. (3) minimizes the update noise for WP, we compute the variance $\langle \langle \Delta w_{ij}^{\text{WP}} \rangle \rangle$ of the WP weight updates when adding a possibly trial- and synapse-dependent baseline term \tilde{E}_{ij} to the update rule, $\Delta w_{ij}^{\text{WP}} = -(\eta / \sigma_{\text{WP}}^2) (E^{\text{pert}} - E + \tilde{E}_{ij}) \xi_{ij}^{\text{WP}}$. The linear approximation for E^{pert} yields

$$\begin{aligned} \langle \langle \Delta w_{ij}^{\text{WP}} \rangle \rangle &\approx \frac{\eta^2}{\sigma_{\text{WP}}^4} \left\langle \left[\left(\sum_{m=1}^M \sum_{l=1}^N \frac{\partial E}{\partial w_{ml}} \xi_{ml}^{\text{WP}} + \tilde{E}_{ij} \right) \xi_{ij}^{\text{WP}} \right]^2 \right\rangle \\ &\quad - \eta^2 \left(\frac{\partial E}{\partial w_{ij}} \right)^2 \\ &= \eta^2 \sum_{m=1}^M \sum_{l=1}^N \left(\frac{\partial E}{\partial w_{ml}} \right)^2 + \eta^2 \left(\frac{\partial E}{\partial w_{ij}} \right)^2 \\ &\quad + \frac{\eta^2}{\sigma_{\text{WP}}^2} \tilde{E}_{ij}^2. \end{aligned} \quad (\text{A5})$$

To show that the choice of E minimizes the update noise for NP, we analogously compute the variance $\langle \langle \Delta w_{ij}^{\text{NP}} \rangle \rangle$ of the NP weight updates [Eq. (6)] when adding a possibly trial- and neuron-dependent baseline term \tilde{E}_i , $\Delta w_{ij}^{\text{NP}} = -(\eta / \sigma_{\text{NP}}^2) (E^{\text{pert}} - E + \tilde{E}_i) \sum_{t=1}^T \xi_{it}^{\text{NP}} r_{jt}$:

$$\begin{aligned} \langle \langle \Delta w_{ij}^{\text{NP}} \rangle \rangle &\approx \frac{\eta^2}{\sigma_{\text{NP}}^4} \left\langle \left[\left(\sum_{m=1}^M \sum_{s=1}^T \frac{\partial E}{\partial y_{ms}} \xi_{ms}^{\text{NP}} + \tilde{E}_i \right) \sum_{t=1}^T \xi_{it}^{\text{NP}} r_{jt} \right]^2 \right\rangle \\ &\quad - \eta^2 \left(\frac{\partial E}{\partial w_{ij}} \right)^2 \\ &= \eta^2 \sum_{m=1}^M \sum_{s=1}^T \left(\frac{\partial E}{\partial y_{ms}} \right)^2 \sum_{t=1}^T r_{jt}^2 + \eta^2 \left(\frac{\partial E}{\partial w_{ij}} \right)^2 \\ &\quad + \frac{\eta^2}{\sigma_{\text{NP}}^2} \tilde{E}_i^2 \sum_t r_{jt}^2. \end{aligned} \quad (\text{A6})$$

For both algorithms, the variance is minimal if $\tilde{E} = 0$, as apparent from its quadratic occurrence with a positive prefactor.

3. Task setting

We train a linear readout $w \in \mathbb{R}^{M \times N}$ that maps N input traces $r \in \mathbb{R}^{N \times T}$ of time dimension T onto M output traces

$z \in \mathbb{R}^{M \times T}$ of the same time dimension. If not stated otherwise, input and target traces do not change between trials. The readout weights w_{ij} are trained to minimize the mean squared deviation $E = 1/(2T) \sum_{i=1}^M \sum_{t=1}^T (z_{it} - z_{it}^*)^2$ of the output z from a target z^* . The target reads, in general, $z_{it}^* = \sum_{j=1}^N w_{ij}^* r_{jt} + d_{it}$, where w_{ij}^* are target weights—i.e., the output error is certainly minimal for $w_{ij} = w_{ij}^*$ —and d_{it} is an output component that cannot be generated by the network, since it is orthogonal to its inputs, $\sum_{t=1}^T d_{it} r_{jt} = 0 \ \forall i, j$. It is useful to define the (symmetric, positive semidefinite) input correlation matrix

$$S_{ij} = \frac{1}{T} \sum_{t=1}^T r_{it} r_{jt}. \quad (\text{A7})$$

S can be diagonalized by a rotation $O \in \text{SO}(N)$ into $D = O^T S O$ such that $D_{\mu\nu} = \alpha_\mu^2 \delta_{\mu\nu}$ is the diagonal matrix of eigenvalues $\alpha_\mu^2 \geq 0$ (where $\mu, \nu = 1, \dots, N$). We refer to the N -dimensional (spatial) eigenvectors of S as input directions. Another useful characteristic is the autocorrelation of the inputs. We denote the autocorrelation summed over all inputs and normalized by T by

$$C_{ts} = \frac{1}{T} \sum_{j=1}^N r_{jt} r_{js}. \quad (\text{A8})$$

We refer to the T -dimensional (temporal) eigenvectors of C as input components. Reading out from an input direction yields a temporal output vector parallel to the related input component. S and C have the same nonzero eigenvalues α_μ^2 . (This follows, for example, from the more general fact that the products AB and BA of an $N \times T$ matrix A and a $T \times N$ matrix B have the same nonzero eigenvalues [95] by setting $A_{it} = r_{it}$ and $B_{ti} = r_{it}$.) We call the α_μ^2 input “strengths,” since they equal the average strength of the μ th latent input per time step or, equivalently, the average strength of all input activity read out from the μ th input direction. The sum of the eigenvalues, the trace

$$\text{tr}[S] \equiv \alpha_{\text{tot}}^2, \quad (\text{A9})$$

equals the total average input strength per time step α_{tot}^2 . We call α_{tot}^2 the total input strength for short.

Throughout Appendix B and Supplemental Material Secs. I–VI [28], we mainly consider inputs with correlation matrices S that have N_{eff} eigenvalues equal to α^2 and all others zero (Table S1), although intermediate results can

also hold for inputs with general correlations S . The correlation matrix then has the useful properties

$$S^2 = \alpha^2 S, \quad (\text{A10})$$

$$\text{tr}[S] \equiv \alpha_{\text{tot}}^2 = \alpha^2 N_{\text{eff}}. \quad (\text{A11})$$

Varying the effective input dimensionality while keeping the total input strength α_{tot}^2 constant thus implies that the strengths of individual input components scale like $\alpha^2 \sim N_{\text{eff}}^{-1}$.

To measure the strength of the unrealizable target component, we define a quantity α_d^2 analogous to α^2 . Since $\alpha^2 = (1/N_{\text{eff}}) \text{tr}[S] = 1/(N_{\text{eff}} T) \text{tr}[r r^T]$ is, in particular, the average input strength per time and latent input, we set

$$\alpha_d^2 \equiv \frac{1}{MT} \text{tr}[d d^T]; \quad (\text{A12})$$

α_d^2 is the average strength of the unrealizable component of the target per time and output.

Using the weight mismatch $W_{ij} = w_{ij} - w_{ij}^*$ and the correlation matrix S , the error function can be written as

$$E = \frac{1}{2T} \sum_{i=1}^M \sum_{t=1}^T (z_{it} - z_{it}^*)^2 \quad (\text{A13})$$

$$\begin{aligned} &= \frac{1}{2T} \sum_{i=1}^M \sum_{t=1}^T \left(\sum_{j=1}^N W_{ij} r_{jt} - d_{it} \right)^2 \\ &= \frac{1}{2} \text{tr}[W S W^T] + E_{\text{opt}}. \end{aligned} \quad (\text{A14})$$

Here, E_{opt} is the lowest achievable error corresponding to zero weight mismatch:

$$E_{\text{opt}} \equiv \frac{1}{2T} \text{tr}[d d^T] = \frac{1}{2} M \alpha_d^2. \quad (\text{A15})$$

We note that due to the division by T both the correlation matrix and the error no longer scale with the task duration, but the error does scale with the number M of outputs.

The choice of a quadratic error function allows one to compute the evolution of the expected error analytically. For the quadratic error function [Eqs. (8) and (A13)], the average WP and NP weight updates [Eqs. (3) and (6)] follow the gradient exactly; i.e., Eqs. (4), (A2), and (A4) hold exactly for any perturbation size:

$$\begin{aligned} \langle \Delta w_{ij}^{\text{WP}} \rangle &= -\frac{\eta}{\sigma_{\text{WP}}^2} \langle (E^{\text{pert}} - E) \xi_{ij}^{\text{WP}} \rangle = -\frac{\eta}{\sigma_{\text{WP}}^2} \sum_{m=1}^M \sum_{k,l=1}^N \left(W_{mk} S_{kl} \langle \xi_{ml}^{\text{WP}} \xi_{ij}^{\text{WP}} \rangle + \frac{1}{2} S_{kl} \langle \xi_{mk}^{\text{WP}} \xi_{ml}^{\text{WP}} \xi_{ij}^{\text{WP}} \rangle \right) \\ &= -\eta \sum_{k=1}^N W_{ik} S_{kj} = -\eta \frac{\partial E}{\partial w_{ij}}, \end{aligned} \quad (\text{A16})$$

$$\begin{aligned}
\langle \Delta w_{ij}^{\text{NP}} \rangle &= -\frac{\eta}{\sigma_{\text{NP}}^2} \sum_{t=1}^T \langle (E^{\text{pert}} - E) \xi_{it}^{\text{NP}} \rangle r_{jt} \\
&= -\frac{\eta}{\sigma_{\text{WP}}^2 T} \sum_{m=1}^M \sum_{s,t=1}^T \left(\sum_{k=1}^N W_{mk} r_{ks} \langle \xi_{ms}^{\text{NP}} \xi_{it}^{\text{NP}} \rangle + \frac{1}{2} \langle (\xi_{ms}^{\text{NP}})^2 \xi_{it}^{\text{NP}} \rangle - d_{ms} \langle \xi_{ms}^{\text{NP}} \xi_{it}^{\text{NP}} \rangle \right) r_{jt} \\
&= -\eta \sum_{k=1}^N W_{ik} S_{kj} + \frac{\eta}{T} \sum_{t=1}^T \cancel{d_{it}} r_{jt} = -\eta \frac{\partial E}{\partial w_{ij}}. \tag{A17}
\end{aligned}$$

4. Effective perturbation strength

For a fair comparison of WP and NP, we consider the output perturbations $\delta z = z^{\text{pert}} - z$ that they generate. For WP, Eqs. (2) and (7) imply

$$\delta z_{it} = \sum_{j=1}^N \xi_{ij}^{\text{WP}} r_{jt}; \tag{A18}$$

for NP, Eqs. (5) and (7) imply

$$\delta z_{it} = \xi_{it}^{\text{NP}}. \tag{A19}$$

We choose σ_{WP} and σ_{NP} such that weight and node perturbations lead to the same output perturbation strength as measured by $\sigma_{\text{eff}}^2 = 1/(MT) \langle \sum_{it} (\delta z_{it})^2 \rangle$, the total induced output variance per time step and output neuron:

$$\begin{aligned}
\sigma_{\text{eff,WP}}^2 &= \frac{1}{MT} \sum_{i=1}^M \sum_{t=1}^T \left\langle \left(\sum_{j=1}^N \xi_{ij}^{\text{WP}} r_{jt} \right)^2 \right\rangle \\
&= \sigma_{\text{WP}}^2 \cdot \alpha^2 N_{\text{eff}}, \tag{A20}
\end{aligned}$$

$$\sigma_{\text{eff,NP}}^2 = \frac{1}{MT} \sum_{i=1}^M \sum_{t=1}^T \langle (\xi_{it}^{\text{NP}})^2 \rangle = \sigma_{\text{NP}}^2. \tag{A21}$$

Here, we use $\langle \xi_{ij}^{\text{WP}} \xi_{mk}^{\text{WP}} \rangle = \sigma_{\text{WP}}^2 \delta_{im} \delta_{jk}$ and $\langle \xi_{it}^{\text{NP}} \xi_{ms}^{\text{NP}} \rangle = \sigma_{\text{NP}}^2 \delta_{im} \delta_{ts}$. Requiring $\sigma_{\text{eff,WP}}^2 \stackrel{!}{=} \sigma_{\text{eff,NP}}^2 \equiv \sigma_{\text{eff}}^2$ implies

$$\sigma_{\text{NP}}^2 = \sigma_{\text{eff}}^2, \quad \sigma_{\text{WP}}^2 = \frac{1}{\alpha^2 N_{\text{eff}}} \cdot \sigma_{\text{eff}}^2. \tag{A22}$$

We note that, although the induced output perturbations have the same variance, they follow different distributions:

$$\frac{1}{MT} \sum_{it} (\delta z_{it})^2 \sim \begin{cases} \frac{\sigma_{\text{eff}}^2}{MN_{\text{eff}}} \chi_{k=MN_{\text{eff}}}^2 & \text{for WP,} \\ \frac{\sigma_{\text{eff}}^2}{MT} \chi_{k=MT}^2 & \text{for NP,} \end{cases} \tag{A23}$$

where χ_k^2 is the chi-square distribution for k degrees of freedom.

APPENDIX B: DERIVATION OF ERROR DYNAMICS

This part starts with a derivation of the error dynamics if our tasks are learned with pure gradient descent (GD) learning. These are comparably simple and a useful benchmark. The sections thereafter provide a full derivation of the error dynamics of WP and NP learning [main text, Eqs. (9)–(14)]. Their analysis and interpretation follow in Supplemental Material Sec. I [28].

1. Error curves for gradient descent

In GD, the updates directly follow the gradient:

$$\Delta w_{ij}^{\text{GD}} = -\eta \frac{\partial E}{\partial w_{ij}} = -\eta \sum_{k=1}^N W_{ik} S_{kj}. \tag{B1}$$

The error after such a deterministic update (which equals the expected error) is

$$\begin{aligned}
E(n) &= \frac{1}{2} \text{tr} \{ [W(n-1) + \Delta w^{\text{GD}}] \tilde{S} [W(n-1) + \Delta w^{\text{GD}}]^T \} \\
&\quad + E_{\text{opt}} \\
&= E(n-1) + \text{tr} [W(n-1) \tilde{S} (\Delta w^{\text{GD}})^T] \\
&\quad + \frac{1}{2} \text{tr} [\Delta w^{\text{GD}} \tilde{S} (\Delta w^{\text{GD}})^T] \\
&= E(n-1) - \eta \text{tr} [W \tilde{S} S W^T] + \frac{1}{2} \eta^2 \text{tr} [W \tilde{S} \tilde{S} W^T], \tag{B2}
\end{aligned}$$

where W stands for $W(n-1)$. To facilitate the tracing of the different terms, we tag the correlation matrix that stems from the cost evaluation at trial n by a tilde, which keeps it distinguishable from those arising from the weight update Eq. (B1); the entries \tilde{S}_{ij} are identical to S_{ij} . For a correlation matrix S that has N_{eff} eigenvalues equal to α^2 and all others zero (see Appendix A 3), we can use Eq. (A10) (and $\tilde{S} \equiv S$) together with $\frac{1}{2} \text{tr} [W S W^T] = E(n-1) - E_{\text{opt}}$ to obtain

$$\begin{aligned}
E(n) &= E(n-1) - \eta \alpha^2 \text{tr} [W S W^T] + \frac{1}{2} \eta^2 \alpha^4 \text{tr} [W S W^T] \\
&= (1 - 2\eta \alpha^2 + \eta^2 \alpha^4) \cdot [E(n-1) - E_{\text{opt}}] + E_{\text{opt}}. \tag{B3}
\end{aligned}$$

This leads to an exponential decay of the error to E_{opt} :

$$E(n) = [E(0) - E_{\text{opt}}]a^n + E_{\text{opt}}, \quad (\text{B4})$$

with

$$a_{\text{GD}} = 1 - 2\eta\alpha^2 + \eta^2\alpha^4, \quad (\text{B5})$$

which becomes zero for the optimal learning rate $\eta^* = \alpha^{-2}$ such that the task is solved in a single trial. As for WP and NP [Eqs. (11) and (B39)], learning diverges once $\eta > 2\eta^*$ where $a > 1$.

2. Error curves for weight perturbation

Since we consider only WP here, for clarity of notation we omit the specifier ‘‘WP’’ in ξ^{WP} , Δw^{WP} , and σ_{WP} . The error E^{pert} as a function of the perturbations ξ applied to the weights then reads

$$\begin{aligned} E^{\text{pert}} &= \frac{1}{2} \text{tr}[(W + \xi)S(W + \xi)^T] + E_{\text{opt}} \\ &= E + \text{tr}[WS\xi^T] + \frac{1}{2} \text{tr}[\xi S\xi^T], \end{aligned} \quad (\text{B6})$$

which yields a weight update [cf. Eq. (3)]

$$\begin{aligned} \Delta w_{ij} &= -\frac{\eta}{\sigma^2} (E^{\text{pert}} - E)\xi_{ij} \\ &= -\frac{\eta}{\sigma^2} \left(\text{tr}[WS\xi^T] + \frac{1}{2} \text{tr}[\xi S\xi^T] \right) \xi_{ij}. \end{aligned} \quad (\text{B7})$$

Averaging over ξ gives the expected error $\langle E(n) \rangle$ after the n th update:

$$\begin{aligned} \Delta E_{\text{update}}^{\text{quad}} &= \frac{1}{2} \langle \text{tr}[\Delta w \tilde{S} \Delta w^T] \rangle \\ &= \frac{\eta^2}{2\sigma^4} \left\langle \left(\text{tr}[WS\xi^T] + \frac{1}{2} \text{tr}[\xi S\xi^T] \right) \cdot \text{tr}[\xi \tilde{S} \xi^T] \cdot \left(\text{tr}[WS\xi^T] + \frac{1}{2} \text{tr}[\xi S\xi^T] \right) \right\rangle \\ &= \frac{\eta^2}{2\sigma^4} \langle \text{tr}[\xi \tilde{S} \xi^T] \text{tr}[WS\xi^T]^2 \rangle + \frac{\eta^2}{8\sigma^4} \langle \text{tr}[\xi \tilde{S} \xi^T] \text{tr}[\xi S\xi^T]^2 \rangle \\ &= \underbrace{\frac{\eta^2}{2\sigma^4} \sum_{imn=1}^M \sum_{jklpqr=1}^N \tilde{S}_{jk} W_{ml} S_{lp} W_{nq} S_{qr} \cdot \langle \xi_{ij} \xi_{ik} \xi_{mp} \xi_{nr} \rangle}_{\equiv \Delta E_{\text{update}}^{\text{quad,a}}} + \underbrace{\frac{\eta^2}{8\sigma^4} \sum_{imn=1}^M \sum_{jklpqr=1}^N \tilde{S}_{jk} S_{lp} S_{qr} \cdot \langle \xi_{ij} \xi_{ik} \xi_{ml} \xi_{mp} \xi_{nq} \xi_{nr} \rangle}_{\equiv \Delta E_{\text{update}}^{\text{quad,b}}}. \end{aligned} \quad (\text{B10})$$

$\Delta E_{\text{update}}^{\text{quad,a}}$ depends on the weight mismatch W and originates from the gradient related part of the error signal $\Delta E_{\text{pert}} - E$, whereas $\Delta E_{\text{update}}^{\text{quad,b}}$ originates from quadratic reward noise (Supplemental Material Sec. I [28]). The moments of ξ can be computed using $\langle \xi_{ij} \xi_{mk} \rangle = \sigma^2 \delta_{im} \delta_{jk}$ and Isserlis’ theorem:

$$\langle \xi_{ik} \xi_{mp} \rangle = \sigma^2 \delta_{im} \delta_{kp}, \quad (\text{B11})$$

$$\langle \xi_{ij} \xi_{ik} \xi_{mp} \xi_{nr} \rangle = \sigma^4 (\delta_{jk} \delta_{mn} \delta_{pr} + \delta_{im} \delta_{jp} \delta_{in} \delta_{kr} + \delta_{in} \delta_{jr} \delta_{im} \delta_{kp}), \quad (\text{B12})$$

$$\begin{aligned} \langle E(n) \rangle &= \frac{1}{2} \langle \text{tr}\{[W(n-1) + \Delta w] \tilde{S} [W(n-1) + \Delta w]^T\} \rangle + E_{\text{opt}} \\ &= \langle E(n-1) \rangle + \underbrace{\langle \text{tr}[W \tilde{S} \Delta w^T] \rangle}_{\equiv \Delta E_{\text{update}}^{\text{lin}}} \\ &\quad + \underbrace{\frac{1}{2} \langle \text{tr}[\Delta w \tilde{S} \Delta w^T] \rangle}_{\equiv \Delta E_{\text{update}}^{\text{quad}}}, \end{aligned} \quad (\text{B8})$$

where W stands for $W(n-1)$. Here, we again tag the correlation matrix with a tilde, since it stems from the cost evaluation at trial n (cf. Appendix B 1). Using Eq. (B7) and that odd moments of ξ vanish, the first highlighted term, $\Delta E_{\text{update}}^{\text{lin}}$, which is linear in the update, gives

$$\begin{aligned} \Delta E_{\text{update}}^{\text{lin}} &= \langle \text{tr}[W \tilde{S} \Delta w^T] \rangle \\ &= -\frac{\eta}{\sigma^2} \left\langle \text{tr} \left\{ W \tilde{S} \left(\text{tr}[WS\xi^T] + \frac{1}{2} \text{tr}[\xi S\xi^T] \right) \xi^T \right\} \right\rangle \\ &= -\frac{\eta}{\sigma^2} \langle \text{tr}[W \tilde{S} \xi^T] \text{tr}[WS\xi^T] \rangle \\ &= -\frac{\eta}{\sigma^2} \sum_{im=1}^M \sum_{jklp=1}^N W_{ij} \tilde{S}_{jk} W_{ml} S_{lp} \cdot \langle \xi_{ik} \xi_{mp} \rangle. \end{aligned} \quad (\text{B9})$$

Using again Eq. (B7) and that odd moments of ξ vanish, the second highlighted term, $\Delta E_{\text{update}}^{\text{quad}}$, which is quadratic in the updates, can be expanded as

$$\begin{aligned}
\langle \xi_{ij} \xi_{ik} \xi_{ml} \xi_{mp} \xi_{nq} \xi_{nr} \rangle &= \sigma^6 [\delta_{jk} (\delta_{lp} \delta_{qr} + \delta_{mn} \delta_{lq} \delta_{pr} + \delta_{mn} \delta_{lr} \delta_{pq}) && (= \langle \xi_{ij} \xi_{ik} \rangle \langle \xi_{ml} \xi_{mp} \xi_{nq} \xi_{nr} \rangle) \\
&+ \delta_{im} \delta_{jl} (\delta_{kp} \delta_{qr} + \delta_{in} \delta_{kq} \delta_{pr} + \delta_{in} \delta_{kr} \delta_{pq}) && (= \langle \xi_{ij} \xi_{ml} \rangle \langle \xi_{ik} \xi_{mp} \xi_{nq} \xi_{nr} \rangle) \\
&+ \delta_{im} \delta_{jp} (\delta_{kl} \delta_{qr} + \delta_{in} \delta_{kq} \delta_{lr} + \delta_{in} \delta_{kr} \delta_{lq}) && (= \langle \xi_{ij} \xi_{mp} \rangle \langle \xi_{ik} \xi_{ml} \xi_{nq} \xi_{nr} \rangle) \\
&+ \delta_{in} \delta_{jq} (\delta_{im} \delta_{kl} \delta_{pr} + \delta_{im} \delta_{kp} \delta_{lr} + \delta_{kr} \delta_{lp}) && (= \langle \xi_{ij} \xi_{nq} \rangle \langle \xi_{ik} \xi_{ml} \xi_{mp} \xi_{nr} \rangle) \\
&+ \delta_{in} \delta_{jr} (\delta_{im} \delta_{kl} \delta_{pq} + \delta_{im} \delta_{kp} \delta_{lq} + \delta_{kq} \delta_{lp})]. && (= \langle \xi_{ij} \xi_{nr} \rangle \langle \xi_{ik} \xi_{ml} \xi_{mp} \xi_{nq} \rangle) \tag{B13}
\end{aligned}$$

Inserting this into Eqs. (B9) and (B10) and performing the summations is partially lengthy but straightforward. It results in

$$\Delta E_{\text{update}}^{\text{lin}} = -\eta \sum_{i=1}^M \sum_{jkl=1}^N W_{ij} \tilde{S}_{jk} S_{lk} W_{il} = -\eta \text{tr}[W \tilde{S} S W^T], \tag{B14}$$

$$\begin{aligned}
\Delta E_{\text{update}}^{\text{quad.a}} &= \frac{\eta^2}{2} \sum_{imn=1}^M \sum_{jklpqr}^N \tilde{S}_{jk} W_{ml} S_{lp} W_{nq} S_{qr} \cdot (\delta_{jk} \delta_{mn} \delta_{pr} + \delta_{im} \delta_{jp} \delta_{in} \delta_{kr} + \delta_{in} \delta_{jr} \delta_{im} \delta_{kp}) \\
&= \frac{\eta^2}{2} M \text{tr}[W S^2 W^T] \text{tr}[\tilde{S}] + \eta^2 \text{tr}[W S \tilde{S} W^T], \tag{B15}
\end{aligned}$$

$$\Delta E_{\text{update}}^{\text{quad.b}} = \frac{\eta^2 \sigma^2}{8} (M^3 \text{tr}[\tilde{S}] \text{tr}[S]^2 + 2M^2 \text{tr}[\tilde{S}] \text{tr}[S^2] + 4M^2 \text{tr}[\tilde{S} S] \text{tr}[S] + 8M \text{tr}[S \tilde{S} S]). \tag{B16}$$

With this, Eq. (B8) for the evolution of expected error becomes

$$\begin{aligned}
\langle E(n) \rangle &\stackrel{\text{WP}}{=} \langle E(n-1) \rangle + \Delta E_{\text{update}}^{\text{lin}} + \Delta E_{\text{update}}^{\text{quad.a}} + \Delta E_{\text{update}}^{\text{quad.b}} \\
&= \langle E(n-1) \rangle - \eta \text{tr}[W \tilde{S} S W^T] + \frac{\eta^2}{2} M \text{tr}[W S^2 W^T] \text{tr}[\tilde{S}] + \eta^2 \text{tr}[W S \tilde{S} W^T] \\
&\quad + \frac{\eta^2 \sigma^2}{8} (M^3 \text{tr}[\tilde{S}] \text{tr}[S]^2 + 2M^2 \text{tr}[\tilde{S}] \text{tr}[S^2] + 4M^2 \text{tr}[\tilde{S} S] \text{tr}[S] + 8M \text{tr}[S \tilde{S} S]). \tag{B17}
\end{aligned}$$

The result holds for a general input correlation matrix S . In Appendix B 4, we assume same strength latent inputs (Table S1). The resulting properties of S [Eqs. (A10) and (A11)] allow one to reexpress all right-hand-side terms of Eq. (B17) in terms of $E(n-1)$ instead of W , yielding a scalar recurrence relation. We make the assumption of equally strong latent inputs also in Secs. I–III in Supplemental Material [28] (Table S1). The convergence for general input is analyzed in Supplemental Material Sec. IV [28].

3. Error curves for node perturbation

Similar to the previous section, we omit the specifier ‘‘NP’’ in ξ^{NP} , Δw^{NP} , and σ_{NP} in the following for notational clarity. The error E^{pert} then reads as a function of the NP ξ :

$$\begin{aligned}
E^{\text{pert}} &= \frac{1}{2T} \sum_{i=1}^M \sum_{t=1}^T \left(\sum_{j=1}^N W_{ij} r_{jt} + \xi_{it} - d_{it} \right)^2 \\
&= \frac{1}{2T} \sum_{i=1}^M \sum_{t=1}^T \left[\left(\sum_{j=1}^N W_{ij} r_{jt} + \xi_{it} \right)^2 - 2 \sum_{j=1}^N W_{ij} r_{jt} d_{it} - 2 \xi_{it} d_{it} + d_{it}^2 \right] \\
&= E + \frac{1}{T} \text{tr}[W r \xi^T] + \frac{1}{2T} \text{tr}[\xi \xi^T] - \frac{1}{T} \text{tr}[d \xi^T]. \tag{B18}
\end{aligned}$$

The term $\text{tr}[d \xi^T]$ here reflects the interference of learning and unrealizable targets. The other term including d , $1/(2T) \text{tr}[d d^T] = E_{\text{opt}}$, is only an additive constant which also occurs in GD and WP and does not enter the update equation. Equation (B18) yields a weight update [cf. Eq. (6)]:

$$\Delta w_{ij} = -\frac{\eta}{\sigma^2} (E^{\text{pert}} - E) \sum_{t=1}^T \xi_{it} r_{jt} = -\frac{\eta}{\sigma^2 T} \left(\text{tr}[W r \xi^T] + \frac{1}{2} \text{tr}[\xi \xi^T] - \text{tr}[d \xi^T] \right) \sum_{t=1}^T \xi_{it} r_{jt}. \tag{B19}$$

Averaging over ξ gives the expected error after the update:

$$\begin{aligned} \langle E(n) \rangle &= \frac{1}{2} \langle \text{tr}\{[W(n-1) + \Delta w] \tilde{S} [W(n-1) + \Delta w]^T\} \rangle + E_{\text{opt}} \\ &= \langle E(n-1) \rangle + \underbrace{\langle \text{tr}[W \tilde{S} \Delta w^T] \rangle}_{\equiv \Delta E_{\text{update}}^{\text{lin}}} + \frac{1}{2} \underbrace{\langle \text{tr}[\Delta w \tilde{S} \Delta w^T] \rangle}_{\equiv \Delta E_{\text{update}}^{\text{quad}}}, \end{aligned} \quad (\text{B20})$$

where W stands for $W(n-1)$ and the correlation matrix that stems from the error evaluation is tagged with a tilde. Using Eq. (B19) and that odd moments of ξ vanish, $\Delta E_{\text{update}}^{\text{lin}}$ and $\Delta E_{\text{update}}^{\text{quad}}$ are

$$\begin{aligned} \Delta E_{\text{update}}^{\text{lin}} &= \langle \text{tr}[W \tilde{S} \Delta w^T] \rangle \\ &= -\frac{\eta}{\sigma^2 T} \left\langle \text{tr} \left\{ W \tilde{S} \left(\text{tr}[W r \xi^T] + \frac{1}{2} \text{tr}[\xi \xi^T] - \text{tr}[d \xi^T] \right) r \xi^T \right\} \right\rangle \\ &= -\frac{\eta}{\sigma^2 T} \langle \text{tr}[W \tilde{S} r \xi^T] \text{tr}[W r \xi^T] - \text{tr}[W \tilde{S} r \xi^T] \text{tr}[d \xi^T] \rangle \\ &= -\frac{\eta}{\sigma^2 T} \sum_{im=1}^M \sum_{jkl=1}^N \sum_{st=1}^T W_{ij} \tilde{S}_{jk} W_{ml} r_{lr} r_{ks} \cdot \langle \xi_{is} \xi_{mt} \rangle + \frac{\eta}{\sigma^2 T} \sum_{im=1}^M \sum_{jk=1}^N \sum_{st=1}^T W_{ij} \tilde{S}_{jk} r_{ks} d_{mt} \cdot \langle \xi_{is} \xi_{mt} \rangle, \end{aligned} \quad (\text{B21})$$

$$\begin{aligned} \Delta E_{\text{update}}^{\text{quad}} &= \frac{1}{2} \langle \text{tr}[\Delta w \tilde{S} \Delta w^T] \rangle \\ &= \frac{\eta^2}{2\sigma^4 T^2} \left\langle \text{tr}[\xi r^T \tilde{S} r \xi^T] \left(\text{tr}[W r \xi^T]^2 + \frac{1}{4} \text{tr}[\xi \xi^T]^2 + \text{tr}[d \xi^T]^2 - 2\text{tr}[W r \xi^T] \text{tr}[d \xi^T] \right) \right\rangle \\ &= \frac{\eta^2}{2\sigma^4 T^2} \sum_{imn=1}^M \sum_{jklp=1}^N \sum_{stuv=1}^T r_{js} \tilde{S}_{jk} r_{kt} W_{ml} r_{lu} W_{np} r_{pv} \cdot \langle \xi_{is} \xi_{it} \xi_{mu} \xi_{nv} \rangle \quad (\Delta E_{\text{update}}^{\text{quad.a}}) \end{aligned} \quad (\text{B22})$$

$$+ \frac{\eta^2}{8\sigma^4 T^2} \sum_{imn=1}^M \sum_{jk=1}^N \sum_{stuv=1}^T r_{js} \tilde{S}_{jk} r_{kt} \cdot \langle \xi_{is} \xi_{it} \xi_{mu} \xi_{mu} \xi_{nv} \xi_{nv} \rangle \quad (\Delta E_{\text{update}}^{\text{quad.b}}) \quad (\text{B23})$$

$$+ \frac{\eta^2}{2\sigma^4 T^2} \sum_{imn=1}^M \sum_{jk=1}^N \sum_{stuv=1}^T r_{js} \tilde{S}_{jk} r_{kt} d_{mu} d_{nv} \cdot \langle \xi_{is} \xi_{it} \xi_{mu} \xi_{nv} \rangle \quad (\Delta E_{\text{update}}^{\text{quad.c}}) \quad (\text{B24})$$

$$- \frac{\eta^2}{\sigma^4 T^2} \sum_{imn=1}^M \sum_{jkl=1}^N \sum_{stuv=1}^T r_{js} \tilde{S}_{jk} r_{kt} W_{ml} r_{lu} d_{nv} \cdot \langle \xi_{is} \xi_{it} \xi_{mu} \xi_{nv} \rangle \quad (\Delta E_{\text{update}}^{\text{quad.d}}). \quad (\text{B25})$$

Like for WP, $\Delta E_{\text{update}}^{\text{quad.a}}$ depends on the weight mismatch W and originates from the gradient-related part of the error signal $\Delta E_{\text{pert}} - E$, whereas $\Delta E_{\text{update}}^{\text{quad.b}}$ originates from quadratic reward noise (Supplemental Material Sec. I [28]). $\Delta E_{\text{update}}^{\text{quad.c}}$ and $\Delta E_{\text{update}}^{\text{quad.d}}$ stem from the reward noise due to coupling to unrealizable target components. The moments of ξ can again be computed from $\langle \xi_{is} \xi_{mt} \rangle = \sigma^2 \delta_{im} \delta_{st}$ and Isserlis' theorem:

$$\langle \xi_{is} \xi_{mt} \rangle = \sigma^2 \delta_{im} \delta_{st}, \quad (\text{B26})$$

$$\langle \xi_{is} \xi_{it} \xi_{mu} \xi_{nv} \rangle = \sigma^4 (\delta_{st} \delta_{mn} \delta_{uv} + \delta_{im} \delta_{su} \delta_{in} \delta_{tv} + \delta_{in} \delta_{sv} \delta_{im} \delta_{tu}), \quad (\text{B27})$$

$$\begin{aligned}
\langle \xi_{is}\xi_{it}\xi_{mu}\xi_{mu}\xi_{nv}\xi_{nv} \rangle &= \sigma^6 [\delta_{st}(1 + 2\delta_{mn}\delta_{uv})] && (= \langle \xi_{is}\xi_{it} \rangle \langle \xi_{mu}\xi_{mu}\xi_{nv}\xi_{nv} \rangle) \\
&+ 2\delta_{im}\delta_{su}(\delta_{im}\delta_{tu} + 2\delta_{in}\delta_{tv}\delta_{mn}\delta_{uv}) && (= 2\langle \xi_{is}\xi_{mu} \rangle \langle \xi_{it}\xi_{mu}\xi_{nv}\xi_{nv} \rangle) \\
&+ 2\delta_{in}\delta_{sv}(\delta_{in}\delta_{tv} + 2\delta_{im}\delta_{tu}\delta_{mn}\delta_{uv}) && (= 2\langle \xi_{is}\xi_{nv} \rangle \langle \xi_{it}\xi_{mu}\xi_{mu}\xi_{nv} \rangle) \\
&= \sigma^6 [\delta_{st}(1 + 2\delta_{mn}\delta_{uv}) + 8\delta_{imn}\delta_{stuv} + 2\delta_{im}\delta_{su}\delta_{tu} + 2\delta_{in}\delta_{sv}\delta_{tv}]. && \text{(B28)}
\end{aligned}$$

Inserting this into Eqs. (B21)–(B25) gives after a partially lengthy but straightforward computation

$$\Delta E_{\text{update}}^{\text{lin}} = -\eta \text{tr}[W\tilde{S}SW^T], \quad \text{(B29)}$$

$$\Delta E_{\text{update}}^{\text{quad,a}} = \frac{\eta^2}{2} M \text{tr}[WSW^T] \text{tr}[\tilde{S}S] + \eta^2 \text{tr}[WS\tilde{S}SW^T], \quad \text{(B30)}$$

$$\Delta E_{\text{update}}^{\text{quad,b}} = \frac{\eta^2 \sigma^2}{8T} \text{tr}[\tilde{S}S] \cdot (M^3 T^2 + 6M^2 T + 8M), \quad \text{(B31)}$$

$$\Delta E_{\text{update}}^{\text{quad,c}} = \frac{\eta^2}{2T} M \text{tr}[\tilde{S}S] \text{tr}[dd^T], \quad \text{(B32)}$$

$$\Delta E_{\text{update}}^{\text{quad,d}} = 0. \quad \text{(B33)}$$

With this, Eq. (B20) for the evolution of expected reward becomes

$$\begin{aligned}
\langle E(n) \rangle &\stackrel{\text{NP}}{=} \langle E(n-1) \rangle + \Delta E_{\text{update}}^{\text{lin}} + \Delta E_{\text{update}}^{\text{quad,a}} + \Delta E_{\text{update}}^{\text{quad,b}} + \Delta E_{\text{update}}^{\text{quad,c}} \\
&= \langle E(n-1) \rangle - \eta \text{tr}[W\tilde{S}SW^T] + \frac{\eta^2}{2} M \text{tr}[WSW^T] \text{tr}[\tilde{S}S] + \eta^2 \text{tr}[WS\tilde{S}SW^T] \\
&\quad + \frac{\eta^2 \sigma_{\text{NP}}^2}{8T} \text{tr}[\tilde{S}S] \cdot (M^3 T^2 + 6M^2 T + 8M) + \eta^2 M \text{tr}[\tilde{S}S] \cdot \frac{1}{2T} \text{tr}[dd^T]. && \text{(B34)}
\end{aligned}$$

As for WP, this result holds for any correlation matrix S .

4. Error curves for equally strong input components

In this section, we consider the case that there are N_{eff} latent inputs of equal strength; see Appendix A 3. Using Eqs. (A10) and (A11) and $\tilde{S} = S$, we first simplify the evolution equation of WP's expected error [Eq. (B17)]. By identifying occurrences of $\frac{1}{2} \text{tr}[WSW^T]$ and replacing them with $\langle E(n-1) \rangle - E_{\text{opt}}$ [Eq. (A14)], we obtain a linear recurrence relation:

$$\begin{aligned}
\langle E(n) \rangle &\stackrel{\text{WP}}{=} \langle E(n-1) \rangle - \eta \text{tr}[W\tilde{S}SW^T] + \frac{\eta^2}{2} M \text{tr}[WS^2W^T] \text{tr}[\tilde{S}] + \eta^2 \text{tr}[WS\tilde{S}SW^T] \\
&\quad + \frac{\eta^2 \sigma_{\text{WP}}^2}{8} (M^3 \text{tr}[\tilde{S}] \text{tr}[S]^2 + 2M^2 \text{tr}[\tilde{S}] \text{tr}[S^2] + 4M^2 \text{tr}[\tilde{S}S] \text{tr}[S] + 8M \text{tr}[S\tilde{S}S]) \\
&= [1 - 2\eta\alpha^2 + \eta^2\alpha^4(MN_{\text{eff}} + 2)] \cdot [\langle E(n-1) \rangle - E_{\text{opt}}] \\
&\quad + \frac{1}{8} \eta^2 \sigma_{\text{WP}}^2 \alpha^6 \cdot (M^3 N_{\text{eff}}^3 + 6M^2 N_{\text{eff}}^2 + 8MN_{\text{eff}}) + E_{\text{opt}}. && \text{(B35)}
\end{aligned}$$

The evolution equation of NP's expected error [Eq. (B34)] similarly simplifies to a linear recurrence relation:

$$\begin{aligned}
\langle E(n) \rangle &\stackrel{\text{NP}}{=} \langle E(n-1) \rangle - \eta \text{tr}[W\tilde{S}SW^T] + \frac{\eta^2}{2} M \text{tr}[WSW^T] \text{tr}[\tilde{S}S] + \eta^2 \text{tr}[WS\tilde{S}SW^T] \\
&\quad + \frac{\eta^2 \sigma_{\text{NP}}^2}{8T} \text{tr}[\tilde{S}S] \cdot (M^3 T^2 + 6M^2 T + 8M) + \eta^2 M \text{tr}[\tilde{S}S] \cdot \frac{1}{2T} \text{tr}[dd^T] \\
&= [1 - 2\eta\alpha^2 + \eta^2\alpha^4(MN_{\text{eff}} + 2)] \cdot [\langle E(n-1) \rangle - E_{\text{opt}}] \\
&\quad + \frac{1}{8} \eta^2 \sigma_{\text{NP}}^2 \alpha^4 \cdot \left(M^3 N_{\text{eff}} T + 6M^2 N_{\text{eff}} + 8M \frac{N_{\text{eff}}}{T} \right) + \eta^2 \alpha^4 MN_{\text{eff}} \cdot E_{\text{opt}} + E_{\text{opt}}. && \text{(B36)}
\end{aligned}$$

Both Eqs. (B35) and (B36) have the form

$$\langle E(n) \rangle = [\langle E(n-1) \rangle - E_{\text{opt}}] \cdot a + b + E_{\text{opt}}, \quad (\text{B37})$$

with (different) constant parameters a and b . The convergence factor a characterizes the convergence speed and b the per-update increase in error. The recurrence relation can be solved straightforwardly by iteration, by expanding b to $[b/(1-a)](1-a)$, shifting $b/(1-a) + E_{\text{opt}}$ to the equation's left-hand side, and considering $\langle E(n) \rangle - E_{\text{opt}} - b/(1-a)$ as a recurrently specified variable:

$$\langle E(n) \rangle = \left(E(0) - \frac{b}{1-a} - E_{\text{opt}} \right) \cdot a^n + \frac{b}{1-a} + E_{\text{opt}}. \quad (\text{B38})$$

Another possibility is to consider $\langle E(n) \rangle - E_{\text{opt}}$ as a recurrently specified variable and to observe that the iteration

gives rise to a finite geometric series that yields $(1-a^n)/(1-a)b$. The values of the constants are

$$a = 1 - 2\eta\alpha^2 + \eta^2\alpha^4(MN_{\text{eff}} + 2), \quad (\text{B39})$$

$$b_{\text{WP}} = \frac{1}{8}\eta^2\sigma_{\text{WP}}^2\alpha^6 \cdot (M^3N_{\text{eff}}^3 + 6M^2N_{\text{eff}}^2 + 8MN_{\text{eff}}), \quad (\text{B40})$$

$$b_{\text{NP}} = \frac{1}{8}\eta^2\sigma_{\text{NP}}^2\alpha^4 \cdot \left(M^3N_{\text{eff}}T + 6M^2N_{\text{eff}} + 8M\frac{N_{\text{eff}}}{T} \right) + \eta^2\alpha^4MN_{\text{eff}} \cdot E_{\text{opt}}. \quad (\text{B41})$$

Finite b due to finite perturbation size σ^2 causes a finite residual error $b/(1-a)$ in Eq. (B38) even if E_{opt} is zero. To enable a fair comparison, σ_{WP} and σ_{NP} are chosen such that they lead to output perturbations δz of the same strength; see Appendix A4. Expressing σ_{WP}^2 and σ_{NP}^2 through the strength σ_{eff}^2 of the output perturbation that they generate [Eq. (A22)], the constants b become

$$b_{\text{WP}} = \frac{1}{8}\eta^2\sigma_{\text{eff}}^2\alpha^4 \cdot (M^3N_{\text{eff}}^2 + 6M^2N_{\text{eff}} + 8M), \quad (\text{B42})$$

$$b_{\text{NP}} = \frac{1}{8}\eta^2\sigma_{\text{eff}}^2\alpha^4 \cdot \underbrace{\left(M^3N_{\text{eff}}T + 6M^2N_{\text{eff}} + 8M\frac{N_{\text{eff}}}{T} \right)}_{b_{\text{NP}}^{\text{quad}}} + \underbrace{\eta^2\alpha^4MN_{\text{eff}} \cdot E_{\text{opt}}}_{b_{\text{NP}}^{\text{lin}}} \quad (\text{B43})$$

$$\equiv b_{\text{NP}}^{\text{quad}} + b_{\text{NP}}^{\text{lin}}, \quad (\text{B44})$$

where the splitting of b_{NP} into $b_{\text{NP}}^{\text{quad}}$ and $b_{\text{NP}}^{\text{lin}}$ is motivated in the next part. a is independent of the perturbation size and the same for WP and NP.

5. Optimal learning rate

The optimal learning rate and convergence factor are obtained by minimizing the quadratic function a as a function of η with respect to η :

$$\eta^* = \underset{\eta}{\text{argmin}} a(\eta) = \frac{1}{(MN_{\text{eff}} + 2)\alpha^2}, \quad (\text{B45})$$

$$a^* = \min_{\eta} a(\eta) = 1 - \frac{1}{MN_{\text{eff}} + 2} = a(\eta^*). \quad (\text{B46})$$

Learning diverges for $\eta \rightarrow 2\eta^*$ because then $a \rightarrow 1$:

$$a(2\eta^*) = 1 - \frac{4}{MN_{\text{eff}} + 2} + 4\frac{MN_{\text{eff}} + 2}{(MN_{\text{eff}} + 2)^2} = 1. \quad (\text{B47})$$

[1] D. M. Wolpert, J. Diedrichsen, and J. R. Flanagan, *Principles of Sensorimotor Learning*, *Nat. Rev. Neurosci.* **12**, 739 (2011).

[2] J. W. Krakauer and P. Mazzoni, *Human Sensorimotor Learning: Adaptation, Skill, and beyond*, *Curr. Opin. Neurobiol.* **21**, 636 (2011).

[3] T. P. Lillicrap, A. Santoro, L. Marris, C. J. Akerman, and G. Hinton, *Backpropagation and the Brain*, *Nat. Rev. Neurosci.* **21**, 335 (2020).

[4] C. Klos, Y. F. K. Kossio, S. Goedeke, A. Gilra, and R.-M. Memmesheimer, *Dynamical Learning of Dynamics*, *Phys. Rev. Lett.* **125**, 088103 (2020).

[5] P. Dayan and L. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (MIT, Cambridge, MA, 2001).

[6] R. S. Sutton, *Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning Series)* (Bradford, Cambridge, MA, 2018).

[7] A. Dembo and T. Kailath, *Model-Free Distributed Learning*, *IEEE Trans. Neural Networks* **1**, 58 (1990).

[8] G. Cauwenberghs, *A Fast Stochastic Error-Descent Algorithm for Supervised Learning and Optimization*, in *Advances in Neural Information Processing Systems* (Morgan Kaufmann, Burlington, 1993), Vol. 5, pp. 244–251.

[9] J. Werfel, X. Xie, and H. S. Seung, *Learning Curves for Stochastic Gradient Descent in Linear Feedforward Networks*, *Neural Comput.* **17**, 2699 (2005).

[10] I. R. Fiete and H. S. Seung, *Gradient Learning in Spiking Neural Networks by Dynamic Perturbation of Conductances*, *Phys. Rev. Lett.* **97**, 048104 (2006).

- [11] R. J. Williams, *Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning*, *Mach. Learn.* **8**, 229 (1992).
- [12] R. Legenstein, S. M. Chase, A. B. Schwartz, and W. Maass, *A Reward-Modulated Hebbian Learning Rule Can Explain Experimentally Observed Network Reorganization in a Brain Control Task*, *J. Neurosci.* **30**, 8400 (2010).
- [13] G. M. Hoerzer, R. Legenstein, and W. Maass, *Emergence of Complex Computational Structures from Chaotic Neural Networks through Reward-Modulated Hebbian Learning*, *Cereb. Cortex* **24**, 677 (2014).
- [14] T. Miconi, *Biologically Plausible Learning in Recurrent Neural Networks Reproduces Neural Dynamics Observed during Cognitive Tasks*, *eLife* **6**, e20899 (2017).
- [15] M. Jabri and B. Flower, *Weight Perturbation: An Optimal Architecture and Learning Technique for Analog VLSI Feedforward and Recurrent Multilayer Networks*, *IEEE Trans. Neural Networks* **3**, 154 (1992).
- [16] H. Saito, K. Katahira, K. Okanoya, and M. Okada, *Statistical Mechanics of Structural and Temporal Credit Assignment Effects on Learning in Neural Networks*, *Phys. Rev. E* **83**, 051125 (2011).
- [17] I. R. Fiete, M. S. Fee, and H. S. Seung, *Model of Birdsong Learning Based on Gradient Estimation by Dynamic Perturbation of Neural Conductances*, *J. Neurophysiol.* **98**, 2038 (2007).
- [18] R. Mooney, J. Prather, and T. Roberts, *Neurophysiology of Birdsong Learning*, in *Learning and Memory: A Comprehensive Reference* (Elsevier, New York, 2008), pp. 441–474.
- [19] P. Gao, E. Trautmann, B. Yu, G. Santhanam, S. Ryu, K. Shenoy, and S. Ganguli, *A Theory of Multineuronal Dimensionality, Dynamics and Measurement*, bioRxiv.
- [20] M. C. Tresch and A. Jarc, *The Case for and against Muscle Synergies*, *Curr. Opin. Neurobiol.* **19**, 601 (2009).
- [21] J. A. Gallego, M. G. Perich, L. E. Miller, and S. A. Solla, *Neural Manifolds for the Control of Movement*, *Neuron* **94**, 978 (2017).
- [22] K. Doya and T. J. Sejnowski, *A Computational Model of Birdsong Learning by Auditory Experience and Auditory Feedback*, in *Central Auditory Processing and Neural Modeling* (Springer, New York, 1998), pp. 77–88.
- [23] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal Dynamics—From Single Neurons to Networks and Models of Cognition* (Cambridge University Press, Cambridge, England, 2014).
- [24] B. Widrow and M. A. Lehr, *30 Years of Adaptive Neural Networks: Perceptron, Madaline, and Backpropagation*, *Proc. IEEE* **78**, 1415 (1990).
- [25] M. London and M. Häusser, *Dendritic Computation*, *Annu. Rev. Neurosci.* **28**, 503 (2005).
- [26] R.-M. Memmesheimer, *Quantitative Prediction of Intermittent High-Frequency Oscillations in Neural Networks with Supralinear Dendritic Interactions*, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 11092 (2010).
- [27] D. Breuer, M. Timme, and R.-M. Memmesheimer, *Statistical Physics of Neural Systems with Non-additive Dendritic Coupling*, *Phys. Rev. X* **4**, 011053 (2014).
- [28] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevX.13.021006> for detailed derivations, further analysis, and accompanying simulations.
- [29] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT, Cambridge, MA, 2016), <http://www.deeplearningbook.org>.
- [30] F. Zenke, G. Hennequin, and W. Gerstner, *Synaptic Plasticity in Neural Networks Needs Homeostasis with a Fast Rate Detector*, *PLoS Comput. Biol.* **9**, e1003330 (2013).
- [31] N. E. Ziv and N. Brenner, *Synaptic Tenacity or Lack Thereof: Spontaneous Remodeling of Synapses*, *Trends Neurosci.* **41**, 89 (2018).
- [32] R. Kawai, T. Markman, R. Poddar, R. Ko, A. L. Fantana, A. K. Dhawale, A. R. Kampff, and B. P. Ölveczky, *Motor Cortex Is Required for Learning but not for Executing a Motor Skill*, *Neuron* **86**, 800 (2015).
- [33] K. Katahira, T. Cho, K. Okanoya, and M. Okada, *Optimal Node Perturbation in Linear Perceptrons with Uncertain Eligibility Trace*, *Neural Netw.* **23**, 219 (2010).
- [34] M. Beyeler, E. L. Rounds, K. D. Carlson, N. Dutt, and J. L. Krichmar, *Neural Correlates of Sparse Coding and Dimensionality Reduction*, *PLoS Comput. Biol.* **15**, e1006908 (2019).
- [35] R. Hahnloser, A. Kozhevnikov, and M. Fee, *An Ultra-sparse Code Underlies the Generation of Neural Sequences in a Songbird*, *Nature (London)* **419**, 65 (2002).
- [36] R. Q. Quiroga, L. Reddy, G. Kreiman, C. Koch, and I. Fried, *Invariant Visual Representation by Single Neurons in the Human Brain*, *Nature (London)* **435**, 1102 (2005).
- [37] D. V. Buonomano and M. M. Merzenich, *Temporal Information Transformed into a Spatial Code by a Neural Network with Realistic Properties*, *Science* **267**, 1028 (1995).
- [38] P. F. Dominey, *Complex Sensory-Motor Sequence Learning Based on Recurrent State Representation and Reinforcement Learning*, *Biol. Cybern.* **73**, 265 (1995).
- [39] W. Maass, T. Natschläger, and H. Markram, *A Model for Real Time Computation in Generic Microcircuits*, in *Advances in Neural Information Processing Systems*, edited by S. Becker, S. Thrun, and K. Obermayer (MIT, Cambridge, MA, 2003).
- [40] D. Sussillo and L. F. Abbott, *Generating Coherent Patterns of Activity from Chaotic Neural Networks*, *Neuron* **63**, 544 (2009).
- [41] D. Thalmeier, M. Uhlmann, H. J. Kappen, and R.-M. Memmesheimer, *Learning Universal Computations with Spikes*, *PLoS Comput. Biol.* **12**, e1004895 (2016).
- [42] R. Pyle and R. Rosenbaum, *A Reservoir Computing Model of Reward-Modulated Motor Learning and Automaticity*, *Neural Comput.* **31**, 1430 (2019).
- [43] G. Hennequin, T. P. Vogels, and W. Gerstner, *Optimal Control of Transient Dynamics in Balanced Networks Supports Generation of Complex Movements*, *Neuron* **82**, 1394 (2014).
- [44] M. F. Bear, B. W. Connors, and M. A. Paradiso, *Neuroscience—Exploring the Brain* (Wolters Kluwer, Philadelphia, 2016).
- [45] N. Y. Masse, G. R. Yang, H. F. Song, X.-J. Wang, and D. J. Freedman, *Circuit Mechanisms for the Maintenance*

- and Manipulation of Information in Working Memory, *Nat. Neurosci.* **22**, 1159 (2019).
- [46] A. K. Dhawale, M. A. Smith, and B. P. Ölveczky, *The Role of Variability in Motor Learning*, *Annu. Rev. Neurosci.* **40**, 479 (2017).
- [47] J. P. Cunningham and B. M. Yu, *Dimensionality Reduction for Large-Scale Neural Recordings*, *Nat. Neurosci.* **17**, 1500 (2014).
- [48] L. F. Abbott, K. Rajan, and H. Sompolinsky, *The Dynamic Brain: An Exploration of Neuronal Variability and Its Functional Significance* (Oxford University, New York, 2011), Chap. Interactions between Intrinsic and Stimulus-Evoked Activity in Recurrent Neural Networks, pp. 65–82.
- [49] C. J. Cueva, A. Saez, E. Marcos, A. Genovesio, M. Jazayeri, R. Romo, C. D. Salzman, M. N. Shadlen, and S. Fusi, *Low-Dimensional Dynamics for Working Memory and Time Encoding*, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 23021 (2020).
- [50] M. Russo, M. D’Andola, A. Portone, F. Lacquaniti, and A. d’Avella, *Dimensionality of Joint Torques and Muscle Patterns for Reaching*, *Front. Comput. Neurosci.* **8**, 00024 (2014).
- [51] A. Destexhe, M. Rudolph, and D. Paré, *The High-Conductance State of Neocortical Neurons In Vivo*, *Nat. Rev. Neurosci.* **4**, 739 (2003).
- [52] J. D. Murray, A. Bernacchia, D. J. Freedman, R. Romo, J. D. Wallis, X. Cai, C. Padoa-Schioppa, T. Pasternak, H. Seo, D. Lee, and X.-J. Wang, *A Hierarchy of Intrinsic Timescales across Primate Cortex*, *Nat. Neurosci.* **17**, 1661 (2014).
- [53] T. Teşileanu, B. Ölveczky, and V. Balasubramanian, *Rules and Mechanisms for Efficient Two-Stage Learning in Neural Circuits*, *eLife* **6**, 20944 (2017).
- [54] J. M. Murray and G. S. Escola, *Learning Multiple Variable-Speed Sequences in Striatum via Cortical Tutoring*, *eLife* **6**, 26084 (2017).
- [55] A. K. Dhawale, R. Poddar, S. B. Wolff, V. A. Normand, E. Kopelowitz, and B. P. Ölveczky, *Automated Long-Term Recording and Analysis of Neural Activity in Behaving Animals*, *eLife* **6**, 27702 (2017).
- [56] S. Cheng, *The CRISP Theory of Hippocampal Function in Episodic Memory*, *Front. Neural Circuits* **7**, 88 (2013).
- [57] A. Maes, M. Barahona, and C. Clopath, *Learning Spatio-temporal Signals Using a Recurrent Spiking Network that Discretizes Time*, *PLoS Comput. Biol.* **16**, e1007606 (2020).
- [58] H. S. Seung, *Learning in Spiking Neural Networks by Reinforcement of Stochastic Synaptic Transmission*, *Neuron* **40**, 1063 (2003).
- [59] Ł. Kuśmiercz, T. Isomura, and T. Toyozumi, *Learning with Three Factors: Modulating Hebbian Plasticity with Errors*, *Curr. Opin. Neurobiol.* **46**, 170 (2017).
- [60] R. L. Redondo and R. G. M. Morris, *Making Memories Last: The Synaptic Tagging and Capture Hypothesis*, *Nat. Rev. Neurosci.* **12**, 17 (2011).
- [61] J. N. Reynolds and J. R. Wickens, *Dopamine-Dependent Plasticity of Corticostriatal Synapses*, *Neural Netw.* **15**, 507 (2002).
- [62] Q. Zhou and M. Poo, *Reversal and Consolidation of Activity-Induced Synaptic Modifications*, *Trends Neurosci.* **27**, 378 (2004).
- [63] J. Luboevski and C. Tetzlaff, *Memory Consolidation and Improvement by Synaptic Tagging and Capture in Recurrent Neural Networks*, *Commun. Biol.* **4**, 275 (2021).
- [64] E. M. Izhikevich, *Solving the Distal Reward Problem through Linkage of STDP and Dopamine Signaling*, *Cereb. Cortex* **17**, 2443 (2007).
- [65] M. A. Farries and A. L. Fairhall, *Reinforcement Learning with Modulated Spike Timing-Dependent Synaptic Plasticity*, *J. Neurophysiol.* **98**, 3648 (2007).
- [66] X. Yao, *Evolving Artificial Neural Networks*, *Proc. IEEE* **87**, 1423 (1999).
- [67] D. Floreano, P. Dürr, and C. Mattiussi, *Neuroevolution: From Architectures to Learning*, *Evolut. Intell.* **1**, 47 (2008).
- [68] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever, *Evolution Strategies as a Scalable Alternative to Reinforcement Learning*, [arXiv:1703.03864](https://arxiv.org/abs/1703.03864).
- [69] J. T. Trachtenberg, B. E. Chen, G. W. Knott, G. Feng, J. R. Sanes, E. Welker, and K. Svoboda, *Long-Term In Vivo Imaging of Experience-Dependent Synaptic Plasticity in Adult Cortex*, *Nature (London)* **420**, 788 (2002).
- [70] A. J. Holtmaat, J. T. Trachtenberg, L. Wilbrecht, G. M. Shepherd, X. Zhang, G. W. Knott, and K. Svoboda, *Transient and Persistent Dendritic Spines in the Neocortex In Vivo*, *Neuron* **45**, 279 (2005).
- [71] N. Yasumatsu, M. Matsuzaki, T. Miyazaki, J. Noguchi, and H. Kasai, *Principles of Long-Term Dynamics of Dendritic Spines*, *J. Neurosci.* **28**, 13592 (2008).
- [72] A. Minerbi, R. Kahana, L. Goldfeld, M. Kaufman, S. Marom, and N. E. Ziv, *Long-Term Relationships between Synaptic Tenacity, Synaptic Remodeling, and Network Activity*, *PLoS Biol.* **7**, e1000136 (2009).
- [73] A. Knoblauch, *The Role of Structural Plasticity and Synaptic Consolidation for Memory and Amnesia in a Model of Cortico-Hippocampal Interplay*, in *Connectionist Models of Behaviour and Cognition II* (World Scientific, Singapore, 2009).
- [74] M. Deger, M. Helias, S. Rotter, and M. Diesmann, *Spike-Timing Dependence of Structural Plasticity Explains Cooperative Synapse Formation in the Neocortex*, *PLoS Comput. Biol.* **8**, e1002689 (2012).
- [75] A. Knoblauch, E. Körner, U. Körner, and F. T. Sommer, *Structural Synaptic Plasticity Has High Memory Capacity and Can Explain Graded Amnesia, Catastrophic Forgetting, and the Spacing Effect*, *PLoS One* **9**, e96485 (2014).
- [76] M. Fauth, F. Wörgötter, and C. Tetzlaff, *The Formation of Multi-synaptic Connections by the Interaction of Synaptic and Structural Plasticity and Their Functional Consequences*, *PLoS Comput. Biol.* **11**, e1004031 (2015).
- [77] A. Knoblauch, *Impact of Structural Plasticity on Memory Formation and Decline*, in *The Rewiring Brain* (Elsevier, New York, 2017), pp. 361–386.
- [78] M. J. Fauth and M. C. van Rossum, *Self-Organized Reactivation Maintains and Reinforces Memories despite Synaptic Turnover*, *eLife* **8**, 43717 (2019).
- [79] D. Kappel, R. Legenstein, S. Habenschuss, M. Hsieh, and W. Maass, *A Dynamic Connectome Supports the Emergence of Stable Computational Function of Neural Circuits through Reward-Based Learning*, *eNeuro* **5**, 0301 (2018).
- [80] P. Suszynski and P. Wawrzynski, *Learning Population of Spiking Neural Networks with Perturbation of*

- Conductances*, in *Proceedings of the 2013 International Joint Conference on Neural Networks (IJCNN)* (IEEE, New York, 2013).
- [81] X. Xie and H. S. Seung, *Learning in Neural Networks by Reinforcement of Irregular Spiking*, *Phys. Rev. E* **69**, 041909 (2004).
- [82] T. Cho, K. Katahira, K. Okanoya, and M. Okada, *Node Perturbation Learning without Noiseless Baseline*, *Neural Netw.* **24**, 267 (2011).
- [83] P. Mazzoni, R. A. Andersen, and M. I. Jordan, *A More Biologically Plausible Learning Rule for Neural Networks*, *Proc. Natl. Acad. Sci. U.S.A.* **88**, 4433 (1991).
- [84] R. Darshan, A. Leblois, and D. Hansel, *Interference and Shaping in Sensorimotor Adaptations with Rewards*, *PLoS Comput. Biol.* **10**, e1003377 (2014).
- [85] E. Vasilaki, S. Fusi, X.-J. Wang, and W. Senn, *Learning Flexible Sensori-motor Mappings in a Complex Network*, *Biol. Cybern.* **100**, 147 (2009).
- [86] K. Takiyama and M. Okada, *Maximization of Learning Speed in the Motor Cortex due to Neuronal Redundancy*, *PLoS Comput. Biol.* **8**, e1002348 (2012).
- [87] M. N. Abdelghani, T. P. Lillicrap, and D. B. Tweed, *Sensitivity Derivatives for Flexible Sensorimotor Learning*, *Neural Comput.* **20**, 2085 (2008).
- [88] B. B. Vladimirov, E. Vasilaki, R. Urbanczik, and W. Senn, *Stimulus Sampling as an Exploration Mechanism for Fast Reinforcement Learning*, *Biol. Cybern.* **100**, 319 (2009).
- [89] J. Friedrich, R. Urbanczik, and W. Senn, *Code-Specific Learning Rules Improve Action Selection by Populations of Spiking Neurons*, *Int. J. Neural Syst.* **24**, 1450002 (2014).
- [90] A. Payeur, J. Guerguiev, F. Zenke, B. A. Richards, and R. Naud, *Burst-Dependent Synaptic Plasticity Can Coordinate Learning in Hierarchical Circuits*, *Nat. Neurosci.* **24**, 1010 (2021).
- [91] T. P. Lillicrap, D. Cownden, D. B. Tweed, and C. J. Akerman, *Random Synaptic Feedback Weights Support Error Backpropagation for Deep Learning*, *Nat. Commun.* **7**, 13276 (2016).
- [92] K. Doya, *Bifurcations in the Learning of Recurrent Neural Networks*, in *Proceedings of the 1992 IEEE International Symposium on Circuits and Systems, San Diego, CA* (IEEE, New York, 1992), Vol. 6, pp. 2777–2780.
- [93] M. D. Giudice, *Effective Dimensionality: A Tutorial*, *Multivar. Behav. Res.* **56**, 527 (2021).
- [94] P. Züge and C. Klos, *Source code for simulations supplementing the article “Weight versus Node Perturbation Learning in Temporally Extended Tasks: Weight Perturbation Often Performs Similarly or Better”*, GitHub, <https://github.com/PaulZOnGit/WP-versus-NP> (2023).
- [95] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. (Cambridge University Press, Cambridge, England, 2012).