

Neuromodulation via Krotov-Hopfield improves accuracy and robustness of restricted Boltzmann machines

Başer Tambaş^{*} and A. Levent Subaşı[†]

Department of Physics, Istanbul Technical University, Istanbul, Türkiye

Alkan Kabakçioğlu[‡]

Department of Physics, Koç University, Istanbul, Türkiye

 (Received 10 April 2025; revised 7 August 2025; accepted 5 December 2025; published 7 January 2026)

In biological systems, neuromodulation tunes synaptic plasticity based on the internal state of the organism, complementing stimulus-driven Hebbian learning. The algorithm recently proposed by Krotov and Hopfield (KH) [D. Krotov and J. J. Hopfield, *Proc. Natl. Acad. Sci. USA* **116**, 7723 (2019)] can be utilized to mirror this process in artificial neural networks, where its built-in intralayer competition and selective inhibition of synaptic updates offer a cost-effective remedy for the lack of lateral connections through a simplified attention mechanism. We demonstrate that KH-modulated restricted Boltzmann machines (RBMs) outperform standard (shallow) RBMs in both reconstruction and classification tasks, offering a superior trade-off between generalization performance and model size, with the additional benefit of robustness to weight initialization as well as to overfitting during training.

DOI: [10.1103/y82t-c34b](https://doi.org/10.1103/y82t-c34b)

Designing artificial neural networks (ANNs) that mimic the essentials of biological learning is a challenging task that gained much attention recently [1]. Most ANNs [2] employ the backpropagation algorithm [3] that broadcasts top-down (TD) signals through the network for end-to-end weight updates. Although highly successful in practice, backpropagation contrasts with the local nature of synaptic plasticity in biological systems. Recent approaches to bypassing backpropagation include methods using alternative feedback mechanisms—such as feedback alignment [4], target propagation [5], and equilibrium propagation [6]—as well as techniques inspired by Hebbian learning. Among the latter are Boltzmann machines (BMs) [7,8], a class of biologically inspired generative neural networks that rely on energy-based dynamics, and more recent methods [9,10] that build on Bienenstock-Cooper-Munro (BCM) theory [11], a refined form of Hebb's rule [12].

Adoption of BMs in large-scale applications is limited by the intractability of sampling and partition function estimation in fully connected architectures [13]. Restricted Boltzmann machines (RBMs) address these challenges by removing lateral connections in the visible and hidden layers (see Fig. 1), which dramatically improves trainability and makes

them amenable to analytical tools from statistical mechanics [14–19], at the expense of reduced representational capacity. Various enhancements—deep Boltzmann machines [13], convolutional RBMs [20], mean-field-inspired augmentations [14], and reintroducing lateral connections [21]—seek to circumvent this issue, but often at the cost of added complexity or computational overhead.

An alternative remedy can be formulated by mimicking neuromodulation in biological systems, where diffusively transmitted chemical signals modulate synaptic plasticity based on the organism's internal state—enabling global coordination beyond the limits of neuronal connectivity [22,23]. In artificial neural networks, this principle is formalized by three-factor learning rules [24], in which a global modulatory signal complements local synaptic updates. Prior implementations along this line have largely employed such signals to dynamically adjust hyperparameters during training [25]. However, neuromodulation can also serve to induce competition among units lacking direct connectivity—such as those within an RBM layer—as we demonstrate below.

In this Letter, we propose the Krotov-Hopfield (KH) learning algorithm [9] as a form of neuromodulatory signaling within RBMs. Although KH was derived from BCM theory [11] and shares Hebbian origins [12], it stands apart by embedding a global competitive mechanism. Specifically, the neuron receiving the strongest input in a layer triggers a broad inhibition that discourages other neurons from becoming coactive. This yields an indirect lateral interaction among hidden units without requiring explicit intralayer connections, which amounts to a rudimentary attention mechanism in a biologically plausible setting that does not rely on backpropagation. As a result, neurons compete for feature selection rather than patterns competing for neuronal attention. KH

^{*}Contact author: tambas19@itu.edu.tr

[†]Contact author: alsubasi@itu.edu.tr

[‡]Contact author: akabakcioglu@ku.edu.tr




FIG. 1. Schematic representation of the KH-modulated RBM with TD and bottom-up (BU) feedback loops for $\delta \mathbf{W}$ weight updates. The graph structure of the RBM has weights $\mathbf{W} = \{W_{ij}\}$ linking N visible $\mathbf{v} = \{v_i\}$ and M hidden $\mathbf{h} = \{h_j\}$ units with associated bias terms $\mathbf{a} = \{a_i\}$ and $\mathbf{b} = \{b_j\}$, respectively. The main text discusses TD feedback only, while results for BU feedback can be found in Sec. B of the Supplemental Material [28].

algorithm has been used successfully to generate competitive filters for convolutional neural nets without employing back-propagation [26].

In what follows, we demonstrate that KH-modulated RBMs not only improve reconstruction and classification metrics over standard (shallow) RBMs but also exhibit enhanced robustness against weight initialization and overfitting. Thus, we show that a neuromodulatory-inspired, global inhibition step can effectively offset the absence of lateral interactions to strike a balance between the expressive modeling power of Boltzmann machines and the computational efficiency of RBMs, resulting in an approach that incorporates the “semirestricted Boltzmann machine” [27] functionality without significant additional training cost.

$$\mathbf{v}^{(0)} \sim p_d(\mathbf{v}) \rightarrow \mathbf{h}^{(0)} \sim p_{\theta}(\mathbf{h}^{(0)}|\mathbf{v}^{(0)}) \rightarrow \mathbf{v}^{(1)} \sim p_{\theta}(\mathbf{v}^{(1)}|\mathbf{h}^{(0)}) \rightarrow \dots \rightarrow \mathbf{h}^{(k-1)} \sim p_{\theta}(\mathbf{h}^{(k-1)}|\mathbf{v}^{(k-1)}) \rightarrow \mathbf{v}^{(k)} \sim p_{\theta}(\mathbf{v}^{(k)}|\mathbf{h}^{(k-1)}).$$

Calculating the model expectations in Eq. (4) using the distributions obtained from the k -step Gibbs chain above constitutes the CD_k algorithm. It is noteworthy that even CD_1 works reasonably well in practical applications [33]. In order to check the robustness of our findings to the choice of k , we ran experiments with both CD_1 and CD_{10} .

Krotov-Hopfield updates. Krotov and Hopfield’s unsupervised learning algorithm [9] incorporates a local, Hebbian-like learning rule for synaptic weight updates. A synapse $W_{\mu\nu}$ connecting a presynaptic node μ with activity x_{μ} is modified by taking into account the input current to the postsynaptic node ν given by the Euclidean inner product

$$I_{r_{\nu}} = \langle \mathbf{W}_{*\nu}, \mathbf{x} \rangle = \sum_{\mu} W_{\mu\nu} x_{\mu}. \quad (5)$$

For convenience, the components of $\mathbf{I}$ are indexed by their rank r_{ν} in descending order $I_K \geq I_{K-1} \geq \dots \geq I_1$. The KH

RBM essentials. A shallow RBM consists of binary visible ($\mathbf{v}$) and hidden ($\mathbf{h}$) units connected via a bipartite graph (Fig. 1) with parameters $\theta = \{\mathbf{W}, \mathbf{a}, \mathbf{b}\}$. The energy function is

$$E_{\theta}(\mathbf{v}, \mathbf{h}) = -\mathbf{v}^{\top} \mathbf{W} \mathbf{h} - \mathbf{a}^{\top} \mathbf{v} - \mathbf{b}^{\top} \mathbf{h}, \quad (1)$$

and the associated Boltzmann distribution takes the form $p_{\theta}(\mathbf{v}, \mathbf{h}) = e^{E_{\theta}(\mathbf{v}, \mathbf{h})}/Z_{\theta}$, where $Z_{\theta} = \sum_{\mathbf{v}, \mathbf{h}} e^{-E_{\theta}(\mathbf{v}, \mathbf{h})}$ is the partition function. The predicted distribution of the visible units can be obtained by marginalization as $p_{\theta}(\mathbf{v}) = \sum_{\mathbf{h}} p_{\theta}(\mathbf{v}, \mathbf{h}) \equiv e^{-F_{\theta}(\mathbf{v})}/Z_{\theta}$ and ideally converges to the underlying distribution $p_d(\mathbf{v})$ of the dataset upon training. $F_{\theta}(\mathbf{v})$ is known as the clamped free energy or visible energy. Accordingly, the average *negative log-likelihood*

$$\mathcal{L} \equiv \langle -\log[p_{\theta}(\mathbf{v})] \rangle_{\mathbf{v} \sim p_d(\mathbf{v})} = \langle F_{\theta}(\mathbf{v}) \rangle_{\mathbf{v} \sim p_d(\mathbf{v})} + \log Z_{\theta} \quad (2)$$

serves as a suitable objective function whose minimization by means of gradient descent steps

$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta} \mathcal{L}(\theta_t) \quad (3)$$

for parameters θ_t at time step t can be shown to reduce to

$$\delta W_{ij} = \eta [\langle v_i h_j \rangle_d - \langle v_i h_j \rangle_m], \quad (4)$$

where $\langle \cdot \rangle_d$ and $\langle \cdot \rangle_m$ denote the expectation values over the data and the model distributions, respectively, and η is the learning rate. Updates similar to Eq. (4) can be found for the biases $\{\mathbf{a}, \mathbf{b}\}$ as well [29].

Evaluating Z_{θ} [therefore, the model expectations in Eq. (4)] is hard in practice. Consequently, training of RBMs requires approximate techniques that vary from Markov chain Monte Carlo (MCMC) algorithms, such as contrastive divergence (CD) [8] and persistent CD [30], to mean-field-based methods [31,32]. We here use the CD algorithm proposed by Hinton [8], where an MCMC chain is initiated from a sample drawn from the dataset and iterated with the sequence

weight updates are then expressed as

$$\delta W_{\mu\nu}^{\text{KH}} = \varepsilon \frac{g_{\nu}(\mathbf{I}) \Phi_{\mu\nu}(\mathbf{x}, \mathbf{W})}{\max_{\mu\nu} [g_{\nu}(\mathbf{I}) \Phi_{\mu\nu}(\mathbf{x}, \mathbf{W})]} \quad (6)$$

in terms of the local update function

$$\Phi_{\mu\nu}(\mathbf{x}, \mathbf{W}) \equiv R^2 x_{\mu} - \langle \mathbf{W}_{*\nu}, \mathbf{x} \rangle W_{\mu\nu}, \quad (7)$$

which reduces to Oja’s rule for $g_{\nu}(\mathbf{I}) = \mathbf{I}$. The first term in Eq. (7) then amounts to the Hebbian update rule, while the second term is a homeostatic regularization forcing the incoming synaptic weights onto the sphere $\sum_{\mu} |W_{\mu\nu}|^2 = R^2$ for each postsynaptic unit ν . The hyperparameter $\varepsilon > 0$ and the denominator in Eq. (6) together act as a varying learning rate [34], which can be interpreted as a context-dependent modulatory effect.

Crucially, Krotov and Hopfield proposed choosing $g_v(\mathbf{I})$ as a global modulator [9]

$$g_v(\mathbf{I}) = \begin{cases} 1, & \text{if } r_v = K, \\ -\Delta, & \text{if } r_v = K - \ell, \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

where K is the number of units in the postsynaptic layer, and ℓ and $\Delta > 0$ are hyperparameters. In this scheme, the unit v subject to the highest input current ($r_v = K$) gets full Hebbian reinforcement, while unit with rank $r_v = K - \ell$ receives an anti-Hebbian feedback where the incoming connections from cofiring presynaptic nodes are *weakened* by the factor Δ . The rest of the synaptic weights remain unaltered. $g_v(\mathbf{I})$ thus acts as a neuromodulatory signal, inducing a global competition among postsynaptic neurons by selectively applying Hebbian and anti-Hebbian strategies across units without a direct connection.

KH-modulated RBM. Indirect lateral interactions established in the postsynaptic layer by the KH algorithm can offer a computationally efficient approach to enhancing the limited expressive power of RBMs. This can be achieved simply by placing KH updates as a neuromodulatory signaling stage between RBM updates during training (though alternative forms of integration are also conceivable). Specifically, at each time step t ,

$$\theta_t^{\text{KH}} = \theta_t + \delta\theta_t^{\text{KH}}, \quad \text{then } \theta_{t+1} = \theta_t^{\text{KH}} - \eta \nabla_{\theta} \mathcal{L}(\theta_t^{\text{KH}}), \quad (9)$$

with the intermediary parameter values $\theta_t^{\text{KH}} = \{\mathbf{W}_t + \delta\mathbf{W}_t^{\text{KH}}, \mathbf{a}_t, \mathbf{b}_t\}$ constructed using Eq. (6).

We consider two potential implementations that differ in the choice of the input state vector $\mathbf{x}$ to Eqs. (6) and (7):

(1) **KH_{TD}** (top-down, cognition-driven) uses $\mathbf{h} \sim p_{\theta}(\mathbf{h}|\mathbf{v})$ for the input $\mathbf{x}$ (and $\mathbf{W}^T$ instead of $\mathbf{W}$).

(2) **KH_{BU}** (bottom-up, sensory-driven) uses $\mathbf{v} \sim p_{\theta}(\mathbf{v})$ for $\mathbf{x}$.

Although both approaches yield qualitatively similar outcomes, we report results for the cognition-driven, top-down modulation, which performed slightly better in our experiments (see Sec. B of the Supplemental Material [28]). In either case, Eq. (9) can be cast into a Langevin-like form

$$\theta_{t+1} = \theta_t - \eta[\nabla_{\theta} \mathcal{L}(\theta_t) + \xi_t], \quad (10)$$

where $\xi_t \simeq (\delta\theta_t^{\text{KH}} \cdot \nabla_{\theta}) \nabla_{\theta} \mathcal{L}(\theta_t) - \eta^{-1} \delta\theta_t^{\text{KH}}$ (obtained by Taylor expansion) encapsulates the total contribution of KH modulation as an additive correction incorporating the local Hessian and the gated Oja's rule in Eq. (6).

Note that Eqs. (9) and (10) repurpose the KH algorithm's competitive, intralayer feedback as a form of neuromodulation integrated alongside the gradient updates to navigate the learning dynamics. In this regard, our approach differs substantially from prior uses of the KH algorithm as a feature extractor in unsupervised learning settings [9,26]. Interestingly, the KH and RBM updates are largely uncorrelated. On the other hand, we found that applying KH updates in a random order does not result in any improvement in performance (see the Appendix for a detailed discussion).

We next demonstrate that KH-modulated RBM improves on the standard version in terms of the test accuracy (or

achieves on-par performance with a fraction of the number of parameters) while simultaneously eliminating overfitting in two classical machine-learning tasks, reconstruction and classification. To this end, we conducted a series of experiments on the Modified National Institute of Standards and Technology (MNIST) [35] dataset whose 70 000 samples were partitioned into 60 000 training samples and 10 000 test samples. Each sample is a 28×28 binary image obtained by thresholding the pixel values $v_i \in [0, 255]$ of the original data using a global threshold of 127. The dataset—being images of handwritten digits—consists of ten classes with a balanced representation in both the training and the test sets, which are otherwise random.

We used two RBMs composed of $M = 100$ and $M = 500$ hidden units, and conducted each experiment on both sizes. The models were trained using contrastive divergence, CD_k , with $k = 1$ and $k = 10$ Gibbs sampling steps and optimized separately. The learning rate and the batch size for all RBMs were set (after a grid search on shallow RBMs) to $\eta = 0.1$ and $n_b = 100$, respectively. KH-related hyperparameters were determined subsequently while keeping RBM parameters fixed. This yields a practically accessible lower bound on the full potential of KH modulation as a comprehensive hyperparameter tuning on KH-RBMs can only improve on our findings, albeit with significant computational cost (see Sec. A of the Supplemental Material [28]).

Weight initialization can make a significant difference in the training performance of standard RBMs. We therefore report the performance on the test dataset for two different and commonly used initializations [36], namely, StdNormal: $W_{\mu\nu} \sim \mathcal{N}(0, 1)$ and LeCun: $W_{\mu\nu} \sim \mathcal{N}(0, 1/N)$, where $\mathcal{N}(0, \sigma^2)$ signifies the normal distribution with zero mean and a standard deviation σ . Upon experimentation, we found that KH updates perform best in an “annealed” setting (see Sec. C of the Supplemental Material [28]), where the KH step size (ε) is gradually reduced according to $\varepsilon = \varepsilon_0(1 - n_{\text{epoch}}/S)^{3/2}$, in which n_{epoch} is the current epoch number and S is the total number of epochs. This scheduling sets a finite timescale for KH modulation. The power $3/2$ was chosen by trial and error, with the intuition that a superlinear decay provides sufficient time for a (almost) pure RBM gradient descent to converge to a minimum in the region targeted by KH modulation.

In all experiments, we set $\ell = 1$ after observing that larger values do not consistently improve model performance. The initial step size ε_0 and the penalty parameter Δ in Eq. (8) were chosen optimally by means of a grid search for each shallow RBM model, for fair comparison. All grid searches used only the 60 000 samples originally reserved for training, 10 000 of which were set apart for hyperparameter validation. Finally, we chose $R = 1$ and 0.1 for StdNormal and LeCun initializations, respectively, making sure that the magnitude of the synaptic weight vector associated with a postsynaptic unit is roughly an order of magnitude larger than the radius of the fixed sphere of Eq. (7) (see Sec. A of the Supplemental Material [28]) [37].

In Fig. 2, we compare the training progression of KH_{TD}-CD₁ and shallow RBMs with $M = 100$, by monitoring the difference between the predicted and the true distribution of the test data in terms of both the reconstruction




FIG. 2. Comparison of KH_{TD} -RBMs with the shallow RBMs, each with CD_1 and $M = 100$: (a) reconstruction mean-squared error (MSE) and (b) cross-entropy. The final stage of the training where the improved performance of KH modulation over standard RBM on the test dataset is evident even with CD_1 .

mean-squared error

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (v_i - \hat{v}_i)^2 \quad (11)$$

and the cross-entropy

$$\begin{aligned} \mathbb{H} = & - \sum_{i=1}^N v_i \log(p_{\theta}(\hat{v}_i = 1|\mathbf{h})) \\ & + (1 - v_i) \log(1 - p_{\theta}(\hat{v}_i = 1|\mathbf{h})). \end{aligned} \quad (12)$$

Reported quantities are averages over ten independent training runs with different initial random seeds.

Several aspects of KH_{TD} -RBM stand out in Fig. 2: (1) Training is robust to the choice of initialization, unlike the standard RBM; (2) the convergence rate is slow compared to LeCun-initialized standard RBMs but is better than StdNormal's; and (3) for both initializations, the final trained network outperforms the standard RBMs on the test dataset.

In Fig. 2, the high convergence rate of shallow RBMs initialized with LeCun's scheme stands out as a key advantage over KH-RBMs, as faster training may be more desirable than




FIG. 3. Reconstruction MSE on the test dataset during the training of KH_{TD} -RBMs with the MNIST dataset for 500 epochs. We show the average of ten independent runs using $M = 100$, LeCun initialization, CD_1 , and different durations (S) of the scheduling protocol (shown in red).

better accuracy in some applications. We therefore investigated whether KH modulation could be adapted to strike a more favorable balance between speed and performance. We found that it is possible to achieve an almost equally fast convergence to a better validation MSE simply by decaying the KH modulation updates more rapidly. In Fig. 3, we report the MSE on the test dataset for KH_{TD} -RBMs trained with the scheduling hyperparameter $S \in \{50, 100, 200, 300, 400, 500\}$, such that, once the KH modulation fully decays (after S epochs), pure RBM update inherits the dynamics. In all cases, KH_{TD} -RBM outperformed shallow RBM at the end of the KH modulation window and remained so during the rest of the training, irrespective of how small S is.

In order to test whether KH modulation has a positive impact also on a more traditional machine-learning task, we next applied it to cRBMs [38] for MNIST digit classification. To this end, we used graph structures with $M = 100$ and $M = 500$ units with the two different weight initializations, and $k = 1$ and $k = 10$ Gibbs sampling steps, as before. The data presented in Fig. 4 and Table I indicate that KH modulation not only improves the test accuracy of the standard cRBM but also nearly eliminates overfitting—a benefit that is most pronounced in the LeCun-initialized training runs.




FIG. 4. Prediction accuracy vs epochs calculated on the test dataset for the MNIST classification task. The figure compares KH-cRBMs with shallow cRBMs, both using CD_1 and $M = 500$.

TABLE I. The test accuracies for the MNIST classification task achieved using both the proposed and standard cRBM implementations. Accuracies are reported as a function of model size and initialization method, with the maximum test accuracy for each model size highlighted in bold. Training run times per epoch are also given for each model.

	$M = 100$			$M = 500$		
	LeCun	Std	T (s)	LeCun	Std	T (s)
CD ₁	90.6(2)	91.3(2)	1.1	94.6(1)	92.0(3)	2.5
KH _{TD} -CD ₁	91.2(3)	92.1(3)	2.2	95.6(1)	96.5(1)	4.5
CD ₁₀	90.4(3)	90.8(2)	3.9	95.8(3)	93.8(2)	9.0
KH _{TD} -CD ₁₀	90.5(3)	91.2(2)	5.0	96.9(1)	97.0(1)	11.1

With 500 hidden units, StdNormal initialization, and 10 Gibbs sampling steps (CD₁₀), the network reaches its maximum accuracy of 97%. This value is slightly reduced with one-step Gibbs sampling to 96.5%, which is the accuracy reported for a shallow cRBM trained with 6000 hidden units [38]. In other words, KH modulation allows a tenfold reduction in model size without sacrificing from accuracy. Alternatively, comparing models of identical size in Table I, KH modulation yields consistent gains in test accuracy, with improvements exceeding 1% in accuracy (more than 25% reduction in error rate) for $M = 500$. We note that, although higher accuracies have been reported for MNIST classification using specialized RBM variants, these results were derived from a generic implementation, without classification-specific tuning or architectural optimizations.

All the experiments described above, including reconstruction and classification tasks, were repeated for the Kuzushiji MNIST (KMIST) [39] dataset, which is considered to be a more challenging version of MNIST. The results reported in Sec. D of the Supplemental Material [28] demonstrate a similar performance boost, increasing the confidence on the validity of our observations across applications.

In conclusion, we presented a proof-of-concept implementation of the KH algorithm as a form of neuromodulatory signaling in RBMs. We showed that the selective inhibition of synaptic updates proposed by Krotov and Hopfield—which can be interpreted as a form of cognition-driven (TD) or sensory-driven (BU) attention mechanism without backpropagation—leads to visible improvements in both reconstruction and classification tasks. Furthermore, KH-modulated RBMs exhibit reduced sensitivity to weight initialization and a marked resilience against overfitting. Typically, the absence of lateral connections in the bipartite structure of RBMs can result in redundant feature learning due to the well-known issue of hidden-unit coadaptation [40]. Incorporating the KH algorithm into the RBM training process appears to partially overcome this problem by promoting more diverse representations through a global competition scheme, thereby enhancing the network’s capacity without need for significant additional resources.

Finally, we emphasize that the reported gains in accuracy and stability are not necessarily confined to RBMs or the KH algorithm specifically. Rather, they reflect a general principle applicable to architectures featuring state-dependent

neuromodulatory feedback guiding gradient descent, as illustrated in Fig. 1. With this insight, the performance of KH-RBMs encourages the use of similar, hybrid frameworks in more complex architectures, which may further bridge the gap between biologically plausible learning and state-of-the-art machine-learning approaches. Another potential future direction is to perform a detailed comparison of the training trajectories with and without KH modulation on the loss landscape to better understand the origins of these improvements.

Acknowledgments. We acknowledge beneficial discussions with A. T. Yıldırım and D. Yuret, as well as a partial support by the Technological and Scientific Research Council of Türkiye (TÜBİTAK) through Grant No. MFAG-119F121 during the initial stages of this work. We also acknowledge the partial support from the Istanbul Technical University Scientific Research Projects Department under Grant No. TDK-2025-47038.

Data availability. The data that support the findings of this article are openly available [37].

Appendix. In the main text, we proposed a hybrid KH-RBM architecture inspired by neuromodulation in biological systems, where unsupervised KH updates are applied alongside the gradient updates, resulting in lower loss values compared to training without KH modulation. In order to elucidate the role of KH modulation in the model’s evolution during KH-RBM training, we also looked at how the stochastic gradient updates compare with the KH updates along the training trajectory. A comparison, in terms of the magnitudes and the directions, of the RBM updates, $\nabla_{\theta} \mathcal{L}(\theta_t)$, and KH modulation, ξ_t , in Eq. (10) during a typical training run is given in Fig. 5. While the decay in the magnitude ratio is by design (the annealing protocol), the existence of a negative directional correlation between the RBM updates and the KH contribution is surprising. Upon inspecting Fig. 5, the KH modulation appears to direct the search primarily perpendicular to the gradient—as a random vector would in high dimensions—while slightly backtracking along the calculated RBM gradient.

In view of the above observations, one might wonder whether the additive contribution of KH in Eq. (10) ultimately plays a role similar to noise injection in the stochastic gradient Langevin dynamics method [41], which can facilitate better posterior sampling and mitigate overfitting. While the KH updates are informed course corrections (with some noise due to minibatching), their minimal overlap with the loss gradient observed in Fig. 5 brings forth the possibility that the observations above can perhaps be attributed to the same mechanism.

To rule out this possibility, we now consider random-KH_{TD}-RBMs, in which we shuffle the KH update vector,

$$\delta\theta_t^{\text{KH}} \longrightarrow \delta\theta_t^{\text{random-KH}},$$

and then apply exactly the same parameter update procedure as described in the main text. Figure 6 compares (1) the original KH_{TD}-RBM setup versus (2) random-KH_{TD}-RBM, showing the cosine similarities of the update vectors to the RBM gradient in Fig. 6(a), and the resulting test MSE in Fig. 6(b). The random-KH_{TD}-RBM quickly loses any non-trivial directional correlation and also fails to improve on the performance of a standard (shallow) RBM.




FIG. 5. Comparison of the RBM gradient and KH modulation weight updates during KH_{TD} -RBM training. The red curve shows the time evolution of the magnitude ratio $|\xi_t| / |\nabla_{\theta} \mathcal{L}(\theta_t)|$ from Eq. (10). The black curve is the directional correlation measured by the cosine similarity, $\nabla_{\theta} \mathcal{L}(\theta_t) \cdot \xi_t / |\nabla_{\theta} \mathcal{L}(\theta_t)| |\xi_t|$. KH modulation is applied with a time-dependent (scheduled) learning rate and is fully phased out at 300 epochs. Each data point is an average of over 20 time steps, and each time step is averaged over 3000 trials.

In addition, we found that hidden-unit receptive fields $\mathbf{W}_v$ exhibit reduced overlap when KH modulation is applied. In particular, the average cosine similarity between a receptive field and its maximally overlapping alternative is 0.41 for shallow RBMs, compared to 0.37 for KH_{TD} -RBM and 0.35 for KH_{BU} -RBM models after training. Therefore, rather than acting as a random perturbation, KH modulation




FIG. 6. (a) Cosine similarities of modulatory KH_{TD} (ξ_t) and random- KH_{TD} (ξ_t^{random}) update vectors and RBM gradient ($\nabla_{\theta} \mathcal{L}(\theta_t)$) update vectors, and (b) their performance comparisons with that of the shallow RBM using the reconstruction MSE on the test dataset. The KH_{TD} modulatory and RBM gradient update vectors have a small negative correlation to standard ones and result in improved performance at the end of a scheduling period.

appears to systematically steer the RBM training trajectory toward an alternative minimum on the landscape where a competition-driven weight assignment allows better use of available capacity.

- [1] S. Schmidgall, R. Ziaei, J. Achterberg, L. Kirsch, S. P. Hajiseydrizi, and J. Eshraghian, Brain-inspired learning in artificial neural networks: A review, *APL Mach. Learn.* **2**, 021501 (2024).
- [2] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature (London)* **521**, 436 (2015).
- [3] D. Rumelhart, G. Hinton, and R. Williams, Learning representations by back-propagating errors, *Nature (London)* **323**, 533 (1986).
- [4] 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).
- [5] D.-H. Lee, S. Zhang, A. Fischer, and Y. Bengio, Difference target propagation, in *Machine Learning and Knowledge Discovery in Databases*, edited by A. Appice, P. P. Rodrigues, V. Santos Costa, C. Soares, J. Gama, and A. Jorge (Springer International Publishing, Cham, 2015), pp. 498–515.
- [6] B. Scellier and Y. Bengio, Equilibrium propagation: Bridging the gap between energy-based models and backpropagation, *Front. Comput. Neurosci.* **11**, 24 (2017).
- [7] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, A learning algorithm for Boltzmann machines, *Cognit. Sci.* **9**, 147 (1985).
- [8] G. E. Hinton, Training products of experts by minimizing contrastive divergence, *Neural Comput.* **14**, 1771 (2002).
- [9] D. Krotov and J. J. Hopfield, Unsupervised learning by competing hidden units, *Proc. Natl. Acad. Sci. USA* **116**, 7723 (2019).
- [10] L. Squadrani, N. Curti, E. Giampieri, D. Remondini, B. Blais, and G. Castellani, Effectiveness of biologically inspired neural network models in learning and pattern memorization, *Entropy* **24**, 682 (2022).

- [11] E. Bienenstock, L. Cooper, and P. Munro, Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex, *J. Neurosci.* **2**, 32 (1982).
- [12] D. O. Hebb, *The Organization of Behavior: A Neuropsychological Theory* (John Wiley and Sons, New York, 1949).
- [13] R. Salakhutdinov and G. E. Hinton, Deep Boltzmann machines, in *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*, edited by D. van Dyk and M. Welling (Machine Learning Research, Florida, 2009), Vol. 5, pp. 448–455.
- [14] M. Mézard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, Oxford, 2009).
- [15] A. Decelle, G. Fissore, and C. Furtlehner, Thermodynamics of restricted Boltzmann machines and related learning dynamics, *J. Stat. Phys.* **172**, 1576 (2018).
- [16] E. W. Tramel, M. Gabrié, A. Manoel, F. Caltagirone, and F. Krzakala, Deterministic and generalized framework for unsupervised learning with restricted Boltzmann machines, *Phys. Rev. X* **8**, 041006 (2018).
- [17] E. Agliari, A. Barra, and B. Tirozzi, Free energies of Boltzmann machines: Self-averaging, annealed and replica symmetric approximations in the thermodynamic limit, *J. Stat. Mech.* (2019) 033301.
- [18] A. Decelle and C. Furtlehner, Restricted Boltzmann machine: Recent advances and mean-field theory, *Chin. Phys. B* **30**, 040202 (2021).
- [19] A. Fachechi, E. Agliari, M. Aquaro, A. Coolen, and M. Mulder, Fundamental operating regimes, hyper-parameter fine-tuning and glassiness: Towards an interpretable replica-theory for trained restricted Boltzmann machines, *J. Phys. A: Math. Theor.* **58**, 065004 (2025).
- [20] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng, Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations, in *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)* (Association for Computing Machinery, Montreal, Canada, 2009), pp. 609–616.
- [21] G. Shi, J. Zhang, N. Ji, and C. Wang, A new variant of restricted Boltzmann machine with horizontal connections, *Neural Comput. Appl.* **31**, 6521 (2019).
- [22] F. Nadim and D. Bucher, Neuromodulation of neurons and synapses, *Curr. Opin. Neurobiol.* **29**, 48 (2014).
- [23] J. M. Shine, Neuromodulatory influences on integration and segregation in the brain, *Trends Cognit. Sci.* **23**, 572 (2019).
- [24] N. Fremaux and W. Gerstner, Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules, *Front. Neural Circuits* **9**, 85 (2016).
- [25] J. Mei, R. Meshkinnejad, and Y. Mohsenzadeh, Effects of neuromodulation-inspired mechanisms on the performance of deep neural networks in a spatial learning task, *iScience* **26**, 106026 (2023).
- [26] L. Grinberg, J. Hopfield, and D. Krotov, Local unsupervised learning for image analysis, [arXiv:1908.08993](https://arxiv.org/abs/1908.08993).
- [27] S. Osindero and G. Hinton, Modeling image patches with a directed hierarchy of Markov random fields, in *Proceedings of the 21st International Conference on Neural Information Processing Systems, NIPS'07* (Curran Associates Inc., Red Hook, NY, 2007), pp. 1121–1128.
- [28] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/y82t-c34b> for details on hyperparameter selection, TD/BU-RBM comparison, and experiments on scheduling and KMNIST (a more challenging learning task), which includes Refs. [35,37,39].
- [29] A. Fischer and C. Igel, An introduction to restricted Boltzmann machines, in *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Application Lecture Notes in Computer Science*, edited by L. Alvarez, M. Mejail, L. Gomez, and J. Jacobo (Springer, Berlin, 2012), pp. 14–36.
- [30] T. Tieleman, Training restricted Boltzmann machines using approximations to the likelihood gradient, in *Proceedings of the 25th International Conference on Machine Learning (ICML'08)* (ACM Press, Helsinki, Finland, 2008), pp. 1064–1071.
- [31] M. Welling and G. E. Hinton, A new learning algorithm for mean-field Boltzmann machines, in *Artificial Neural Networks—ICANN 2002*, edited by G. Goos, J. Hartmanis, J. Van Leeuwen, and J. R. Dorronsoro (Springer, Berlin, 2002), Vol. 2415, pp. 351–357.
- [32] M. Gabrié, E. W. Tramel, and F. Krzakala, Training restricted Boltzmann machine via the Thouless-Anderson-Palmer free energy, in *Advances in Neural Information Processing Systems*, edited by C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, Inc., Montreal, Canada, 2015), Vol. 28.
- [33] G. E. Hinton, A practical guide for training restricted Boltzmann machines, in *Neural Networks: Tricks of the Trade*, 2nd ed., Lecture Notes in Computer Science, edited by G. Montavon, G. B. Orr, and K.-R. Müller (Springer, Berlin, 2012), pp. 599–619.
- [34] D. Krotov, Biological learning, GitHub repository, 2020, https://github.com/DimaKrotov/Biological_Learning.
- [35] Y. LeCun, C. Cortes, and C. Burges, MNIST handwritten digit database (1998), <http://yann.lecun.com/exdb/mnist/>.
- [36] Y. A. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller, Efficient backprop, in *Neural Networks: Tricks of the Trade: Second Edition*, edited by G. Montavon, G. B. Orr, and K.-R. Müller (Springer, Berlin, 2012), pp. 9–48.
- [37] The code implementations are made available on GitHub: <https://github.com/basertambas/KH-RBM>.
- [38] H. Larochelle, M. Mandel, R. Pascanu, and Y. Bengio, Learning algorithms for the classification restricted Boltzmann machine, *J. Mach. Learn. Res.* **13**, 643 (2012).
- [39] T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, and D. Ha, Deep learning for classical Japanese literature, <https://doi.org/10.20676/00000341>.
- [40] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov, Improving neural networks by preventing co-adaptation of feature detectors, [arXiv:1207.0580](https://arxiv.org/abs/1207.0580).
- [41] M. Welling and Y. W. Teh, Bayesian learning via stochastic gradient Langevin dynamics, in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)* (Omnipress, Washington, 2011), pp. 681–688.