UBIQUITY OF HEBBIAN DYNAMICS IN COMPLEX LEARNING RULES

Anonymous authors

000

001

002003004

010

011

012

013

014

015

016

017

018

019

021

023

025

026 027

029

031

033

034

036

037

038

039

040

041

042

043

044

046

047

048

051

052

Paper under double-blind review

ABSTRACT

Hebbian and anti-Hebbian plasticity are widely observed in the biological brain, yet their theoretical understanding remains limited. In this work, we find that when a learning method is regularized with L2 weight decay, its update will align with the direction of the Hebbian learning signal as it approaches stationarity. This Hebbian-like behavior is not unique to SGD: almost any learning rule, including random ones, can exhibit the same signature long before learning has ceased. We also provide a theoretical explanation for anti-Hebbian plasticity in regression tasks, demonstrating how it can arise naturally from gradient or input noise, and offering a potential reason for the observed anti-Hebbian effects in the brain. Although the need for pure Hebbian mechanisms is well established in certain brain regions, our results suggest a new possibility: that Hebbian learning may also be auxiliary, emerging from more complex learning mechanisms rather than serving as the predominant learning signal, as has traditionally been assumed.

1 Introduction

Hebbian and anti-Hebbian plasticity are the most commonly observed types of plasticity in the brain (Koch et al., 2013; Zenke & Gerstner, 2017; Lisman, 1989; Lamsa et al., 2007). It is a longstanding belief in neuroscience that Hebbian learning is fundamentally distinct from gradient descent (Hebb, 2005). While Hebbian learning is a simple, unsupervised learning rule that is biologically plausible, gradient-based optimization is widely regarded as requiring access to global error signals and precise coordination across layers—properties not generally supported by neural circuits in the brain. As a result, gradient descent has been largely dismissed as biologically implausible (Rumelhart et al., 1986; Whittington & Bogacz, 2019; Lillicrap et al., 2020), despite its centrality to modern machine learning. Although some recent work has hypothesized algorithms that empirically approximate SGD, which could be implemented through local learning rules in the brain (Lillicrap et al., 2020; Liao et al., 2024), there is limited evidence that our brains are actually learning in any of the proposed mechanisms. At the same time, the mechanistic understanding from the neuroscience side on Hebbian learning and anti-Hebbian learning is poor and often phenomenology-driven. For example, in the spike-timing-dependent-plasticity (STDP) theory (Caporale & Dan, 2008; Froemke et al., 2005; Brzosko et al., 2019), the division between Hebbian and anti-Hebbian learning only depends on the timing of firing, and there is a lack of understanding of why this is the case. However, this separation between artificial and biological learning may be less stark than previously thought. There is some apparent resemblance between Hebbian learning and SGD. Hebbian learning requires weight decay or forms of normalization to ensure convergence, as the core Hebbian principle functions primarily as a learning signal (Oja, 1982). Similarly, in SGD with weight decay, SGD serves as the learning signal, while weight decay acts as a regularization mechanism that promotes robustness and generalization.

In this work, we discover deeper connections between SGD and Hebbian learning. We demonstrate that the standard training routines used in deep learning, especially stochastic gradient descent (SGD) with weight decay and noise, can give rise to learning signals that look Hebbian or anti-Hebbian. Our results demonstrate that:

1. Close to stationarity, almost any learning rule (including SGD) with weight decay will have a learning signal that looks like a Hebbian rule; and the correlation increases monotonically as we use a larger weight decay;

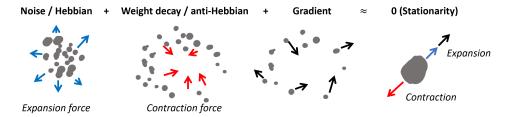


Figure 1: **Balance of contractive and expansive forces.** For deep learning, the noise and weight decay are, respectively, expansion and contraction forces. When they do not balance, the gradient must fill in the gap – if noise outweighs weight decay, the gradient must appear contractive; otherwise, it appears expansive. Similarly, for biology, the Hebbian dynamics is always expansive, and the anti-Hebbian dynamics is always contractive. Thus, to reach a balance, a learning signal will look like, and become aligned with, the Hebbian or anti-Hebbian rule depending on whether it is expansive or contractive.

2. When we inject noise into the learning process, the learning signal aligns anti-Hebbian, and the effect also becomes stronger as the noise gets stronger;

This work is organized as follows. The next section discusses the closely related works and preliminary concepts to understand our result. Section 3 studies when SGD is similar to the Hebbian rule. Section 4 studies when SGD is similar to the anti-Hebbian rule. Section 5 studies the dynamical and nonstationary aspects of these phenomena. Additional figures are presented in the Appendix.

2 RELATED WORKS AND BACKGROUND

Hebbian Learning. As a mathematical model, consider a hidden layer in an arbitrary network:

$$h_b = W h_a(x), \tag{1}$$

where h_a is the postactivation of the previous layer, and h_b is the preactivation of the current layer. In the most conventional form, the simplest homosynaptic update¹ rule states that W is learned according to

$$\Delta W = s\eta h_b h_a^{\mathsf{T}},\tag{2}$$

where $s \in \{-1, 1\}$ is the sign of learning. When s = 1, the rule is Hebbian, which states that if neuron i causes neuron j to fire, then their connection should be strengthened. Similarly, when s = -1, the rule is anti-Hebbian, as it tends to reduce correlation between neurons. η is a positive time constant, which we call the "learning rate." In a neuroscience context, Eq. equation 2 should be regarded a discrete-time approximation to the true underlying continuous-time process, and the rule implies

$$\dot{W} = s\eta h_a h_a^T W, \tag{3}$$

which increases the norm of W when s>0 and decreases it when s<0. Therefore, Hebbian learning in this limit is always expansive, and anti-Hebbian learning is always contractive (see Figure 1 for a visualization). Evidence exists to show that both Hebbian and anti-Hebbian learning exist widely in the brain (Abbott & Nelson, 2000; Magee & Grienberger, 2020). Yet it is not yet clear when the learning is supposed to be Hebbian as opposed to anti-Hebbian. Our theory offers a very simple mechanistic answer to this question.

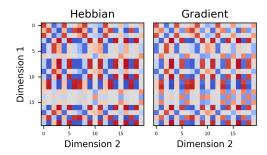
Gradient Descent in the Brain. So far, there has not been any strong evidence that the brain could implement and run any form of gradient descent, despite various theoretical proposals (Kolen & Pollack, 1994; Lillicrap et al., 2020; Whittington & Bogacz, 2019)—and observations of Hebbian plasticity are often implicitly regarded as evidence against gradient descent (e.g., see the criticism of heterosynaptic rules in Porr & Wörgötter (2007)). Our theory shows that gradient descent dynamics can lead to dynamics at stationarity that are consistent with the Hebbian phenomenon, and because of this, observations of Hebbian updates are consistent with the existence of more complicated learning rules in the brain.

Similarity between learning algorithms. A few works are closely related to ours. Xie & Seung (2003) shows the equivalence of gradient descent to a form of contrastive Hebbian algorithm (CHA). However, CHA is not biologically Hebbian because it is not a homosynaptic rule, required by the

¹We use this term as a synonym of Hebbian learning.

Normalized Weight Update Example Cosine Similarity: 0.930

Normalized Weight Update Example Cosine Similarity: 0.051



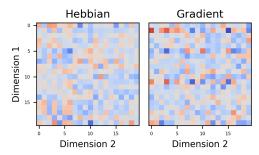


Figure 2: The **left** shows example weight updates with a high alignment between the learning signal $(-\nabla_W \ell)$ and the Hebbian update at the end of training with high weight decay, while the **right** image displays an example update at the end of training with no weight decay which has very low alignment. This figure shows a 20x20 subset of the direction of the Hebbian and learning signal updates for the second layer of an SCE after training with $\eta = 0.1$, and $\gamma = 0.05$, or $\gamma = 0.0$. Dimension 1 can be viewed as the output (post-synaptic) neuron (in the case of SGD, whose incoming weights we are differentiating), and Dimension 2 are input (pre-synaptic) feature/neurons that the weight projects from. We are only visualizing a 20x20 subset of these updates for clarity. Examples of low cosine similarity of the learning signal for $\gamma = 0.05$ at the start and end of training can be seen in Figure 17. In general, we find that stronger weight decay, larger learning rate, and larger batch size lead to better alignment (Figures 8 and 9).

Hebbian principle. There have been several other adjacent ideas to modify the Hebbian rule to lead to learning performance similar to gradient descent or even mathematical equivalences to SGD in certain types of models Scellier et al. (2018); Xiao et al. (2019); Scellier & Bengio (2019); Ernoult et al. (2022). But these works fail to provide a general relationship between arbitrary models trained with SGD and do not identify the key role of regularization and noise.

More recently, it was shown that heterosynaptic circuits such as feedback alignment or SAL (Liao et al., 2024) can lead to dynamics similar to Hebbian dynamics (Ziyin et al., 2025). However, to the best of our knowledge, no paper has successfully shown any robust equivalence between SGD with weight decay and Hebbian learning. On the machine learning side, a recent theoretical work (Ziyin et al., 2024) suggested that the representation learning in neural networks is governed by the expansion and contraction of representations during SGD training. However, the relevance to Hebbian learning and neuroscience is unclear.

3 LEARNING-REGULARIZATION BALANCE PRODUCES HEBBIAN LEARNING

There have been proposals that a balance between Hebbian and anti-Hebbian dynamics must happen for the brain to reach at least some form of homeostasis (stationarity) (Xie & Seung, 2003; Oja, 1982; Bienenstock et al., 1982). There is a similar effect in gradient-based training in neural networks. The use of weight decay contracts the weights to become smaller, but learning can hardly happen if the weights are too small. Therefore, any model that reaches some level of stationarity in training must have a gradient signal that is expansive and opposed to the contractive effect of weight decay.

For the layer defined in Eq. 1, the full weight update is

$$\Delta W \propto \underbrace{-\nabla_{h_b(x)} \ell h_a^T(x)}_{\text{learning signal}} -\gamma W, \tag{4}$$

where and γ is the strength of weight decay. Close to a stationary point, the update should be small in expectation: $\mathbb{E}_x[\nabla_{h_b(x)}\ell h_a^T(x)] + \gamma W \approx 0$. Thus,

$$\operatorname{Tr}\mathbb{E}_{x}[\nabla_{h_{b}(x)}\ell h_{b}^{T}(x)] = \mathbb{E}_{x}[\nabla_{h_{b}(x)}^{T}\ell h_{b}(x)] = -\gamma \operatorname{Tr}[WW^{T}] < 0$$
(5)

Therefore, on average, $\nabla^T_{h_b(x)} \ell h_b(x)$ is negative.

Now, as in Ziyin et al. (2024), we assume a weak decoupling condition: $||h_a||^2 = \mathbb{E}[||h_a||^2]$, which states that norms of all representations are rather close to each other. This is certainly satisfied when, for example, there is a neural collapse (Papyan et al., 2020) or when the representations are

normalized. This means that the expected alignment between the learning signal and the Hebbian update is given by

$$\mathbb{E}\left[\operatorname{Tr}\left[\underbrace{-\nabla_{h_b(x)}\ell h_a^T(x)}_{\text{learning signal}}\underbrace{h_a(x)h_b^T(x)}_{\text{Hebbian update}}\right] = -\mathbb{E}\left[\|h_a\|^2\right]\mathbb{E}_x\left[\nabla_{h_b(x)}^T\ell h_b(x)\right]$$
(6)

$$= \gamma \mathbb{E}[\|h_a\|^2] \operatorname{Tr}[WW^T] > 0. \tag{7}$$

Namely, the learning signal has positive correlation with the Hebbian update, and the alignment becomes stronger as γ increases. See Figure 2 for an example of such alignment.

Perhaps surprisingly, a weaker form of this result generalizes any update rule, precisely because the weight decay term always aligns with the anti-Hebbian update; at stationarity the expected learning signal must align Hebbian. Consider an arbitrary learning signal $g(x, \theta)$; the full weight update is

$$\Delta W = g(x, \theta) - \gamma W,\tag{8}$$

where g is the learning rule and θ is the entirety of all trainable (plastic) parameters. For clarity, η is subsumed into g and γ . Close to stationarity, we have that $\mathbb{E}_x[g(x,\theta)] \approx \gamma W$.

Now, consider the cosine similarity between the learning rule and the Hebbian rule:

$$\cos \theta = \frac{\operatorname{Tr}\left[\mathbb{E}_x[g(x,\theta)]\mathbb{E}_x[h_a h_b^T]\right]}{\|\mathbb{E}_x[g(x,\theta)]\|_F \|\mathbb{E}_x[h_a h_b^T]\|_F}$$
(9)

The direction of alignment at stationarity when $\mathbb{E}_x[g(x,\theta)] = \gamma W$ is thus

$$\operatorname{Tr}\left[\mathbb{E}_{x}[g(x,\theta)]\mathbb{E}_{x}[h_{a}h_{b}^{T}]\right] = \gamma \operatorname{Tr}\left[W\mathbb{E}_{x}[h_{a}h_{b}^{T}]\right]$$
(10)

$$= \gamma \mathbb{E}[\|h_b\|^2] > 0. \tag{11}$$

We see that the update must have a positive alignment with the Hebbian rule on average. This shows an intriguing and yet surprising fact: any algorithm with weight decay may look like a Hebbian rule, and the Hebbian rule may just be a "universal" projection of more complicated algorithms. This is a weaker but rather universal result. It is different from Eq.7 in the sense that Eq. 7 predicts that the learning signal and the Hebbian rule are statistically correlated, whereas this result only says that they have the same direction when averaged over all stimuli.

While this theory is only rigorously stated for stationary points of the learning rule, it is also relevant when the dynamics are out-of-stationarity. For example, the alignment is found to be proportional to γ , and to reach this value of alignment, there must be a period of time when the dynamics is nonstationary and is between zero and the stationarized value. We will study the dynamical aspects in Section 5.

Neurobiology. A key feature of this simple theory is that it separately considers the effects of the learning signal and the weight decay, which, in the context of neurobiology, are likely to have different biological substrates. The learning signal is a fast process; it is likely to, for example, come from other neurons and take the form of electric currents and spikes (Lillicrap et al., 2020). The weight decay, however, is a much slower biochemical process and directly corresponds to the changes in the biochemical properties of synapses (such as a spine shrinkage (Stein et al., 2015)). Therefore, the biological realizations of these two processes are likely to take different forms and can be separately measured. This makes it particularly important to have a theory for the learning-signal component of the update, as this can be directly measured through LTD and LTP experiments of Hebbian plasticity (e.g., see (Zenke & Gerstner, 2017)).

Simulations. We empirically find that this trend holds across a wide variety of different learning tasks. We ran simulations performing classification on CIFAR-10 and non-linear regression on synthetic data (Krizhevsky, 2009). We tested both MLPs and transformers, as well as a range of activation functions and optimizers. In some situations, the correlation between the two learning paradigms is very strong (e.g., in Figure 2). In our experiments, we used a default learning rate of $\eta=0.01$ and trained for 50 epochs, which reached convergence. Since this trend only holds near stationarity—a condition achievable in full gradient descent but obscured in SGD by noise—we found it best to use larger batch sizes to compute both the gradient and Hebbian update as suggested

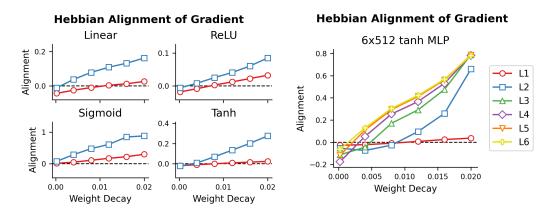


Figure 3: The diagram on the **left** shows that the trend of weight decay increasing Hebbian alignment of the learning signal is robust across different activation functions. The diagram on the **right** shows that the trend can generalize to deeper networks. The SCE MLPs were modified by varying the activation functions across Linear, ReLU, Sigmoid, and Tanh (**left**) and increasing the depth to 6 and layer width to dimension 512 (**right**. L1, L2,...,L6 in this diagram indicate the Hebbian alignment with the learning signal for the corresponding layer. For a small (or zero) weight decay, the learning process sometimes exhibits a weak anti-Hebbian alignment, indicated by a negative alignment with Hebbian learning.

Hebbian Learning SGD Alignment Rolling Window 200 (mean ± SD)

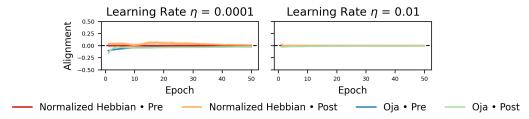


Figure 4: No standard interpretation of Hebbian learning produces alignment with SGD at convergence. The graph shows the mean SGD alignment of the second layer's updates, \pm the standard deviation over a 200-iteration window, when trained with various versions of the Hebbian learning rule for two different learning rates. The SCE was modified to learn with various versions of the Hebbian learning rule. The *Normalized Hebbian* learning rule is the generic Hebbian algorithm with weight standardization after every step. The second algorithm is Oja's rule. We also tested the pre-activation and post-activation versions of both. The average alignment of every combination approaches zero.

in Xu et al. (2023). We found a batch size of 256 to generally show Hebbian phenomena while being small enough to converge to good solutions quickly (Figure 8). To get the alignment between the updates, we compute the cosine similarity of the direction of the gradient update from the loss function (the negative gradient in SGD) and the direction of the Hebbian update. Further, most experiments reported on in this paper followed one of the following setups, and any variations will be reported when relevant.

- (1) Standard Classification Experiment (SCE): In these experiments, we trained a small MLP with 2 layers of 128 dimensions and tanh activation using cross-entropy loss to classify CIFAR-10.
- (2) Standard Regression Experiment (SRE): In these experiments, we trained a small MLP with two hidden layers of 128 units each and tanh activation, using mean squared error to predict the output of a teacher model. The teacher has the same architecture but is initialized with different random parameters. Both the input and output vectors are 32-dimensional. The training dataset consisted of 20,000 randomly generated training examples, and the validation dataset contained 2,000 examples. For the transformer variant of the SRE, we used a transformer with 32-dimensional token embeddings, a vocabulary size of 16, and a maximum sequence length of 32. The encoder consists of 2 layers with 4 attention heads and 32-dimensional feed-forward blocks using ReLU activations. The average of the output token embeddings is passed through the same MLP described above and compared to the teacher output.

Table 1: For all models, optimizers and learning rules, Hebbian alignment rises with increasing weight-decay γ . Hebbian alignment (mean \pm SD, n=10) at convergence is shown for the 2nd-layer gradient in a regression MLP and a sequence-to-vector transformer (1st layer for DFA). All experiments were SREs with a few modifications outside of the learning rule and weight decay specified in the table. DFA used $\eta=0.1$ with gradient-norm clip=5 and, as in the original implementation, used biases. RandomNN used gradient-norm clip=1 and a target weight L2 norm of 100 to determine the sign of the update as explained in Section C.3 of the Appendix. Table elements with – indicate the model's weights collapsed to zero.

| Model | Learning Rule | Weight Decay (γ) | | | |
|----------------|--------------------------------|---|---|--|--|
| | | 0 | 5×10^{-5} | 5×10^{-4} | 5×10^{-3} |
| Regression MLP | Adam SGD DFA RandomNN | -0.02 ± 0.00 -0.10 ± 0.01 0.45 ± 0.05 0.00 ± 0.00 | 0.10 ± 0.00 -0.06 ± 0.01 0.45 ± 0.04 0.00 ± 0.00 | 0.66 ± 0.01 0.17 ± 0.01 0.68 ± 0.05 0.05 ± 0.00 | $-\\0.59 \pm 0.01\\0.87 \pm 0.00\\0.50 \pm 0.00$ |
| Transformer | Adam SGD DFA RandomNN | $-0.02 \pm 0.02 0.00 \pm 0.01 0.08 \pm 0.03 0.00 \pm 0.00$ | 0.50 ± 0.24 0.04 ± 0.01 0.07 ± 0.02 0.00 ± 0.00 | 0.99 ± 0.02 0.47 ± 0.06 0.11 ± 0.02 0.01 ± 0.00 | $-0.88 \pm 0.03 \\ 0.12 \pm 0.02 \\ 0.09 \pm 0.01$ |

Classification. We train a series of MLPs using the SCE setup to classify CIFAR-10. Figure 3 shows that as weight decay increases, so does the alignment of the learning signal with the Hebbian update. The trend persists across different activation functions. Although we still detect this trend in larger MLPS (Figure 3), we occasionally observe some layers behaving in an anti-Hebbian direction as the weight increases. Using residual connections and batch normalization can stabilize the network and make it more Hebbian; however, a deeper exploration of this quality is left for future research. We also explore the use of other regularizations in Appendix 16. At the same time, when one trains an identical model using Hebbian-style learning rule, the model performs poorly and there is no alignment at convergence with the SGD learning signal as seen in Figure 4 (Oja, 1982); there is no reason to expect unsupervised learning to produce similar dynamics near convergence as supervised learning.

Regression. We also evaluate the generalization of this trend to student-teacher regression problems as described in SRE. We explored both MLP and Transformer models and evaluated the Hebbian alignment for learning rules outside of SGD. Recall that a key prediction of the theory is that almost any update look like a Hebbian rule when regularized. We test a variety of rules: SGD, Adam, and Direct Feedback Alignment (DFA) (Nøkland, 2016). To demonstrate that this observation is universal, we also run a setting with a randomly initialized neural network whose output is used as a learning signal, based entirely on the input data (Random NN). Notably, Random NN should not be able to *learn anything* given it is effectively only a deterministic random error signal. Results are shown in Table 1. In all cases, alignment to Hebbian learning emerges and becomes stronger as weight decay increases, regardless of the model.

4 Noise-Learning balance leads to Anti-Hebbian learning

We have answered the question of how Hebbian learning can be an emergent and phenomenological byproduct. The second part of the question is when we will see anti-Hebbian learning, as both Hebbian and anti-Hebbian learning are ubiquitous in the brain. Can anti-Hebbian learning also be a byproduct of more complicated learning rules?

The analysis in the previous section does not take into account the existence of noise in learning. In reality, noise is always non-negligible both in biological learning and in artificial learning. That a strong noise leads to an anti-Hebbian learning signal can already be explained by looking at a simple linear regression problem:

$$\ell(w) = (w^T x - y)^2, (12)$$

where $x \in \mathbb{R}^d$, $y \in \mathbb{R}$ are sampled from some underlying distribution at every training step. Here, we inject noise $\epsilon \in \mathcal{N}(0, \sigma I)$ to the weight before every optimization step so that $w = v + \epsilon$, where v is the weight before noise injection. This could be a thermal noise that can exist ubiquitously in the brain (London et al., 2010). It can also be seen as an approximate model of the SGD noise, which

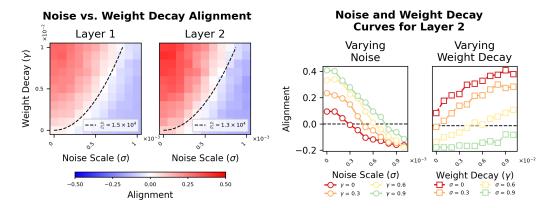


Figure 5: As the noise increases, the Hebbian alignment decreases and higher weight decays lead to higher Hebbian alignment (**right**). The figure on the **left** displays a heatmap of the Hebbian alignment of the learning signal at convergence for a number of different additive noises and weight decays; there is a clear quadratic curve at zero-alignment as predicted by the theory. The SRE was augmented by adding noise to each parameter at the start of each iteration with a mean of zero and the specified standard deviation on the diagram. The trend is even clearer when we follow the behavior of varying the noise of a specific weight decay (Varying Noise) or the weight decay of a specific noise standard deviation (Varying Weight Decay).

causes \boldsymbol{w} to fluctuate around the mean (Liu et al., 2021). The learning signal and Hebbian update are

$$\Delta_{\text{SGD}}w = -x(w^T x - y),\tag{13}$$

$$\Delta_{\text{Hebb}} w = x w^T x. \tag{14}$$

The alignment between the two is

$$\mathbb{E}_{\epsilon}[(\Delta_{\text{SGD}}w)^T(\Delta_{\text{Hebb}}w)] = -\|x\|^2 \mathbb{E}_{\epsilon}[(w^Tx)^2 - w^Txy]$$
(15)

$$= -\|x\|^2 \left[(v^T x)^2 + \sigma^2 \|x\|^2 - v^T xy \right], \tag{16}$$

which is negative for sufficiently large σ^2 and any $||x|| \neq 0$. Thus, large noise leads to anti-Hebbian learning.

An interesting question is how this effect competes and trades off with weight decay. When there is a weight decay, the full weight update is $\Delta_{SGD}w = -x(w^Tx - y) - \gamma w$, and so

$$\mathbb{E}_{\epsilon}[(\Delta_{\text{SGD}}w)^{T}(\Delta_{\text{Hebb}}w)] \approx -\sigma^{2}c_{0} + \gamma c_{1}, \tag{17}$$

where c_0 and c_1 are positive. Thus, one expects a **phase transition** boundary at $\gamma \propto \sigma^2$. When γ is larger than this boundary, the learning is Hebbian-like; when smaller, the learning is anti-Hebbian like. This result provides a straightforward and simple framework to potentially test and understand the Hebbian and anti-Hebbian plasticity in biology. A possible strong biological evidence that would verify this theory is the simultaneous observations of strong noise in the ambient space and anti-Hebbian plasticity. In simulation, this scaling law is verified in the experiments (Figure 5), which justifies this simple analysis.

Simulations We ran experiments to validate the noise prediction using a two-layer MLP with tanh activation. We used a student-teacher model to build a non-linear regression problem and trained until convergence using SGD and varying the variance of the Gaussian noise added at each training step, as well as the weight decay. There is a very smooth alignment trend with SGD, as can be seen in Figure 5. The white region shows the phase boundary between the Hebbian phase and anti-Hebbian phase, and shows a shape in accordance with the quadratic curve $\gamma \approx \sigma^2$.

We observed that at convergence, the Hebbian alignment of the learning signal is higher in low noise environments, and becomes more aligned with anti-Hebbian as the noise increases (Figure 5). Interestingly, we found that solutions with high generalization generally had low Hebbian and anti-Hebbian alignment (Figure 6).

We also observed this trend with other optimizers such as Adam (Figure 13). However, we struggled to robustly reproduce this effect outside of the last few layers of much larger networks or those doing different learning tasks, such as classification. We hypothesize this could be because the magnitude

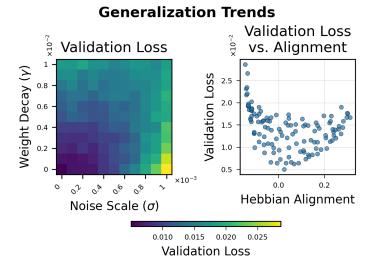


Figure 6: Best performance of the model is achieved when it is not Hebbian or anti-Hebbian on average. The **left** image displays the student validation loss for the experiment in Figure 5, while the **right** image shows a scatter plot of the validation loss vs. Hebbian alignment of the gradient. There seems to be some weak saddle phenomena in loss that occur at the phase transition boundary of Hebbian alignment with respect to noise and scale. The validation loss reduces as both weight decay and noise get smaller.

of the weights does not have as much of an effect on the quality of the learned representations in larger non-linear networks, so the gradient signal does not necessarily need to point in a direction that contracts weights. We also find that adding other types of biologically plausible constraints during learning, such as a sparsification term on layer activations, can lead to a stronger anti-Hebbian alignment of the gradient.

5 Transient phases of Hebbian and Anti-Hebbian Learning

As we mentioned in Section 3, the results are applicable when the dynamics are not yet fully stationary. While the argument we had suggested that one would only observe the Hebbian alignment close to convergence, our empirical results suggest that the alignment is present for much of training. Two key phenomena we discover are the initial alignment bump and the steady state Hebbian oscillations. Outside stationarity, the learning signal often dominates regularization. So it is sufficient to consider the full weight update directly.

Initial Hebbian alignment bump Particularly for networks with ReLU activations, there is a bump in Hebbian alignment of the learning signal that appears to be strongly dependent on initialization scale and learning rate at the beginning of training (Figure 7). During this initial phase of alignment, the full weight update of SGD also increases in alignment to a pure Hebbian update. This early stage of alignment seems to be the result of general feature learning, as the actual scale of weight norms does not change substantially at the start of this period, and with positive weight decays decreases. A higher learning rate makes this process happen faster.

When we examine the Hebbian alignment of the weight updates for individual neurons in the model, a striking pattern appears. During this period, individual neurons seem to take on Hebbian or anti-Hebbian learning roles that can persist for many steps. Like the average behavior of the model, the length of these phases increases as the weight initialization scale magnitude and learning rate decrease. The ratio of anti-Hebbian to Hebbian neurons increases with weight decay.

Hebbian and Anti-Hebbian steady state oscillations There is a phase change that occurs after the initial coherent phase of learning, which is accompanied by a strong shift in the magnitude of the alignment intensity. After this phase change, individual neurons often, though not always, seem to oscillate between strongly Hebbian or anti-Hebbian weight updates (Figure 18). Since near stationarity, the magnitude of the parameters should not increase or decrease on average, we also find the mean of the full weight updates to converge to zero (Figure 15). Often, we find that models with better generalization exhibit strong Hebbian/anti-Hebbian oscillations; however, strong oscillations do not necessarily entail strong generalization.

Initialization Scale vs. Hebbian Alignment of Weight Update (mean ± SD, window=200) & Weight-Norm

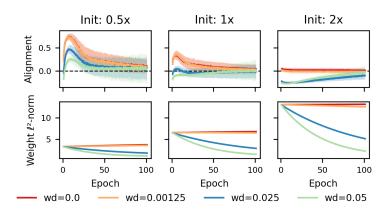


Figure 7: For some activations at low learning rates, there is a sharp jump in Hebbian alignment of the weight update when training with SGD; the size of this jump depends on initial conditions. During this phase, the weight norm decreases monotonically, suggesting the effect is due to feature alignment rather than parameter scale. This experiment used a SCE with $\eta = 0.001$.

6 Discussion

This study suggests that Hebbian and anti-Hebbian plasticity can be understood as emergent regimes of gradient-based optimization rather than fundamentally distinct learning principles. By analyzing the interaction between stochastic gradient descent, weight decay, and stochastic perturbations, we demonstrated that the expected gradient update direction aligns with classic Hebbian plasticity when contraction due to regularization dominates, and switches to an anti-Hebbian alignment when expansion driven by noise prevails. The resulting phase boundary satisfies a simple scaling relation, and the phenomenon was observed across a broad spectrum of architectures, objectives, and alternative update rules.

There are a few limitations that provide interesting areas for future research. We only dealt with smaller models in our experiments. While this decision was reasonable given the scope of this paper, it leaves open the question of whether or not we see Hebbian dynamics in much larger-scale models. As mentioned in the text, as we expanded our models and used different optimizers, we often saw strong average anti-Hebbian alignment of a subset or all of the layers, even at high weight decays. This likely results from our stationarity condition not holding in these models. However, we do not yet have a theory for why and when these regions of anti-Hebbian alignment occur.

Our results have two primary implications. First, our results show that Hebbian and anti-Hebbian plasticity can emerge as regimes of gradient-based optimization, in addition to their conventional role as fundamental learning mechanisms. Second, the presence of Hebbian or anti-Hebbian signatures in neuro-physiological data need not be interpreted as evidence against global error-driven optimization in the brain; such local plasticity patterns may arise as epiphenomena of an underlying optimization process.

Given that unsupervised adaptation and reinforcement are useful and widespread mechanisms, intrinsically Hebbian homosynaptic plasticity likely does exist in the brain. However, much of the existing experimental evidence for Hebbian and anti-Hebbian plasticity is often correlational and phenomenological (e.g., see Lamsa et al. (2007)), so it can be difficult to decide whether the underlying dynamics are actually Hebbian or are more complicated and only appear to be Hebbian. Although some evidence has established the importance of heterosynaptic modulation in memory storage and visual discrimination, its broader role in learning has remained largely speculative due to the challenges of studying it both in vivo and in vitro (Bailey et al., 2000; Chasse et al., 2021). We hope the theory we propose in this paper will serve as a basis for future experimental studies that will validate or challenge the existence of heterosynaptic learning principles in the brain.

REFERENCES

- Larry F Abbott and Sacha B Nelson. Synaptic plasticity: taming the beast. *Nature neuroscience*, 3 (11):1178–1183, 2000.
- Craig. Bailey, Maurizio Giustetto, Yan-You Huang, and et al. Is heterosynaptic modulation essential for stabilizing hebbian plasticity and memory. *Nature Reviews Neuroscience*, 1:11–20, 2000. doi: 10.1038/35036191. URL https://doi.org/10.1038/35036191.
 - Elie L Bienenstock, Leon N Cooper, and Paul W Munro. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2(1):32–48, 1982.
 - Zuzanna Brzosko, Susanna B Mierau, and Ole Paulsen. Neuromodulation of spike-timing-dependent plasticity: past, present, and future. *Neuron*, 103(4):563–581, 2019.
 - Natalia Caporale and Yang Dan. Spike timing–dependent plasticity: a hebbian learning rule. *Annu. Rev. Neurosci.*, 31(1):25–46, 2008.
 - Renee Chasse, Alexey Malyshev, Roslyn Holly Fitch, and Maxim Volgushev. Altered heterosynaptic plasticity impairs visual discrimination learning in adenosine a1 receptor knock-out mice. *Journal of Neuroscience*, 41(21):4631–4640, 2021. ISSN 0270-6474. doi: 10.1523/JNEUROSCI. 3073-20.2021. URL https://www.jneurosci.org/content/41/21/4631.
 - Maxence M Ernoult, Fabrice Normandin, Abhinav Moudgil, Sean Spinney, Eugene Belilovsky, Irina Rish, Blake Richards, and Yoshua Bengio. Towards scaling difference target propagation by learning backprop targets. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 5968–5987. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/ernoult22a.html.
 - Robert C Froemke, Mu-ming Poo, and Yang Dan. Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature*, 434(7030):221–225, 2005.
 - Donald Olding Hebb. *The organization of behavior: A neuropsychological theory*. Psychology press, 2005.
 - Giacomo Koch, Viviana Ponzo, Francesco Di Lorenzo, Carlo Caltagirone, and Domenica Veniero. Hebbian and anti-hebbian spike-timing-dependent plasticity of human cortico-cortical connections. *Journal of Neuroscience*, 33(23):9725–9733, 2013.
 - John F Kolen and Jordan B Pollack. Backpropagation without weight transport. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, volume 3, pp. 1375–1380. IEEE, 1994.
 - Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. URL https://api.semanticscholar.org/CorpusID:18268744.
 - Karri P Lamsa, Joost H Heeroma, Peter Somogyi, Dmitri A Rusakov, and Dimitri M Kullmann. Anti-hebbian long-term potentiation in the hippocampal feedback inhibitory circuit. *Science*, 315 (5816):1262–1266, 2007.
 - Qianli Liao, Liu Ziyin, Yulu Gan, Brian Cheung, Mark Harnett, and Tomaso Poggio. Self-assembly of a biologically plausible learning circuit. *arXiv preprint arXiv:2412.20018*, 2024.
- Timothy P Lillicrap, Adam Santoro, Luke Marris, Colin J Akerman, and Geoffrey Hinton. Back-propagation and the brain. *Nature Reviews Neuroscience*, 21(6):335–346, 2020.
 - John Lisman. A mechanism for the hebb and the anti-hebb processes underlying learning and memory. *Proceedings of the National Academy of Sciences*, 86(23):9574–9578, 1989.
 - Kangqiao Liu, Liu Ziyin, and Masahito Ueda. Noise and fluctuation of finite learning rate stochastic gradient descent, 2021.

- Michael London, Arnd Roth, Lisa Beeren, Michael Häusser, and Peter E Latham. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, 466(7302): 123–127, 2010.
- Jeffrey C Magee and Christine Grienberger. Synaptic plasticity forms and functions. *Annual review of neuroscience*, 43(1):95–117, 2020.
 - Arild Nøkland. Direct feedback alignment provides learning in deep neural networks. *Advances in neural information processing systems*, 29, 2016.
 - Erkki Oja. Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 15:267—273, 1982.
 - Vardan Papyan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40): 24652–24663, 2020.
 - Bernd Porr and Florentin Wörgötter. Learning with "relevance": using a third factor to stabilize hebbian learning. *Neural computation*, 19(10):2694–2719, 2007.
 - David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
 - Benjamin Scellier and Yoshua Bengio. Equivalence of equilibrium propagation and recurrent back-propagation. *Neural Computation*, 31(2):312–329, 2019.
 - Benjamin Scellier, Anirudh Goyal, Jonathan Binas, Thomas Mesnard, and Yoshua Bengio. Generalization of equilibrium propagation to vector field dynamics. *arXiv preprint arXiv:1808.04873*, 2018. URL https://arxiv.org/abs/1808.04873.
 - Ivar S Stein, John A Gray, and Karen Zito. Non-ionotropic nmda receptor signaling drives activity-induced dendritic spine shrinkage. *Journal of Neuroscience*, 35(35):12303–12308, 2015.
 - James CR Whittington and Rafal Bogacz. Theories of error back-propagation in the brain. *Trends in cognitive sciences*, 23(3):235–250, 2019.
 - Will Xiao, Honglin Chen, Qianli Liao, and Tomaso Poggio. Biologically-plausible learning algorithms can scale to large datasets. In *International Conference on Learning Representations*, (ICLR 2019), 2019.
 - Xiaohui Xie and H Sebastian Seung. Equivalence of backpropagation and contrastive hebbian learning in a layered network. *Neural computation*, 15(2):441–454, 2003.
 - Mengjia Xu, Tomer Galanti, Akshay Rangamani, Lorenzo Rosasco, and Tomaso Poggio. The janus effects of sgd vs gd: high noise and low rank. 2023.
 - Friedemann Zenke and Wulfram Gerstner. Hebbian plasticity requires compensatory processes on multiple timescales. *Philosophical transactions of the royal society B: biological sciences*, 372 (1715):20160259, 2017.
 - Liu Ziyin, Isaac Chuang, Tomer Galanti, and Tomaso Poggio. Formation of representations in neural networks. *arXiv preprint arXiv:2410.03006*, 2024.
 - Liu Ziyin, Isaac Chuang, and Tomaso Poggio. Heterosynaptic circuits are universal gradient machines, 2025.



Validation Accuracy and Alignment for Layer 2

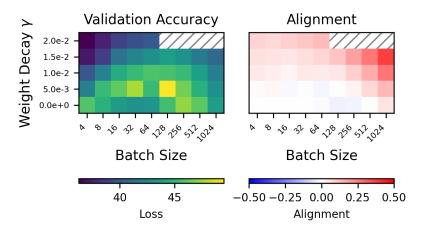


Figure 8: The optimal performance seems to be at a critical position between Hebbian and anti-Hebbian gradient alignment when varying batch size and weight decay. This shows the accuracy (left) and the Hebbian alignment of gradient update (right) for SCEs with a variety of weight decays and batch sizes. The striped background indicates NaN values.

REPRODUCTION

All experiments were run on MIT's OpenMind cluster using Quadro RTX 6000 GPUs and cumulatively took under 50 hours of compute time.

LLM USAGE В

The authors used LLMs to assist in editing the manuscript and writing experimental code.

EXPERIMENTS

In the following document, we provide additional figures and explanations that were referenced in the main text.

ADDITIONAL INFLUENCES ON HEBBIAN ALIGNMENT AND GENERALIZATION

C.1.1 BATCH SIZE

See Figure 8.

Validation Accuracy and Alignment for Layer 2 Validation Accuracy **Alignment** Weight Decay y 5e-3 √e.3 ze.2 ze.a √e.3 sen ze's Learning Rate η Learning Rate η -0.5 0.0 -1.0 0.5 1.0

Figure 9: The optimal performance seems to be at a critical position between Hebbian and anti-Hebbian gradient alignment when varying learning weight and weight decay. This shows the accuracy (**left**) and the Hebbian alignment of gradient update (**right**) for SCEs with a variety of weight decays and learning rates.

Accuracy

Alignment

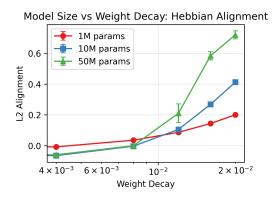


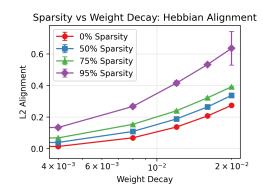
Figure 10: This diagram shows the effect of model size on Hebbian alignment and weight decay. Each point represents the mean of the alignment for the final 200 steps of the run ± the std across 10 seeds. The trend of weight decay leading to increased Hebbian alignment of the learning signal holds with larger models as well. The diagram above shows the alignment of the second layer of the respective MLPs. The MLPs had the following number of total layers: 3, 7, and 9 for the 1M, 10M, and 50M models, respectively. All hidden dimensions were assigned to reach the target parameter count as closely as possible. See equation 18 for how the exact hidden dimensions were computed.

C.1.2 LEARNING RATE

See Figure 9.

C.1.3 MODEL SCALE

See Figure 10.



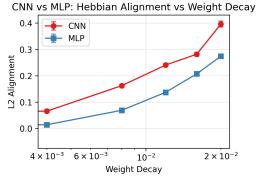
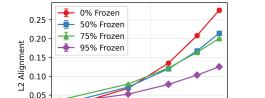


Figure 11: Empirically, the Hebbian alignment of the learning signal increases with sparsity (**left**). We also see that the linear layers in a convolutional neural network, which are highly sparse, have increased Hebbian alignment. Each point represents the mean of the alignment for the final 200 steps of the run \pm the std across 10 seeds. The models on the left were the standard MLPs with varying sparsity. The MLP on the right was a standard MLP and the CNN had the following architecture: Conv Layers: $(c_{in} = 3, c_{out} = 32, s = 3, p = 1)$, (32, 64, 3, 1), (64, 128, 3, 1) MaxPool: (2, 2) Linear hidden dimensions: 2048, 512, 256.

Limited Plasticity vs Weight Decay: Hebbian Alignment



0.00 4×10^{-3} 6×10^{-3} 10^{-2} 2×10^{-2} Weight Decay

Figure 12: The alignment decreases as the fraction of parameters of the standard MLP that are frozen increases; however, the trend still persists. Each point represents the mean of the alignment for the final 200 steps of the run \pm the std across 10 seeds.

$$h = \begin{cases} \frac{-(i+o) + \sqrt{(i+o)^2 + 4t}}{2}, & t = 10^6 \\ \frac{-(i+o) + \sqrt{(i+o)^2 + 20t}}{10}, & t = 10^7 \\ \frac{-(i+o) + \sqrt{(i+o)^2 + 28t}}{14}, & t = 5 \cdot 10^7 \end{cases}$$
(18)

Where i is the dimension of the input, o is the dimension of the output and t is the target parameter count.

C.1.4 MODEL SPARSITY

See Figure 11.

C.1.5 FROZEN PARAMATERS

See Figure 12







Hebbian Alignment of Gradient Heatmap

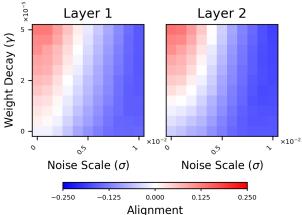


Figure 13: Again, there is a clear trend that even for the Adam optimizer, as noise increases, alignment of the learning signal decreases, and as weight decay increases, so too does alignment. Adam was very sensitive to the parameter ranges for which we'd see the trend, so we used a different weight decay and standard deviation range than the prior experiment. However, the rest of the architecture and experimental setup are identical to that described in Figure 5.

Input Noise vs. Weight Decay Alignment

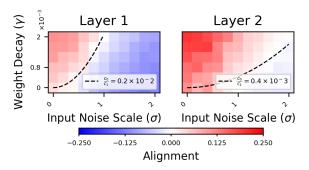


Figure 14: Additive noise to the input can also lead to anti-Hebbian learning. Since noise is only added to the input of the network, the exact phase boundary changes with depth. The results depicted above are from a SRE with input noise injected.

C.1.6TRAINING DURATION

C.1.7Noise

See Figure 13 and 14.

C.1.8 FULL UPDATE VS. LEARNING SIGNAL

As defined in the terminology section, the learning signal $g(x,\theta) \equiv -\nabla_W \ell(\cdot)$ represents the gradient contribution to weight updates, while the full weight update $\Delta W = \eta(g(x,\theta) - \gamma W)$ includes both the learning signal and regularization terms. While we see that the learning signal aligns on average with the Hebbian update, the full weight update can not, otherwise the weight would explode. See Figure 15 for a visualization of the alignment of the learning signal and the full weight update over the course of training. Still, we see a very interesting trend where the full weight update often has strongly Hebbian or anti-Hebbian updates that, on average, cancel.

Hebbian Alignment and Loss of Learning Signal and Full Weight Update During Training

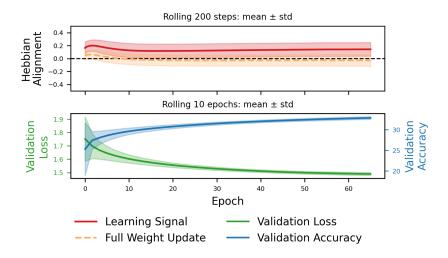


Figure 15: Comparison of Hebbian alignment for learning signal vs. full weight update during training (top) and the corresponding validation loss and accuracy (bottom). The learning signal alignment shows the characteristic patterns described in the main text, while the full weight update alignment approaches zero near stationarity as expected, since the mean of the full update must be zero at convergence. Individual neuron-wise signals can still oscillate between Hebbian and anti-Hebbian phases even when the mean full update is zero. Learning Signal alignment begins and persists far before learning has stopped.

Hebbian Alignment of Gradient Update vs. Regularizers During Training (Rolling Window 200, mean ± SD)

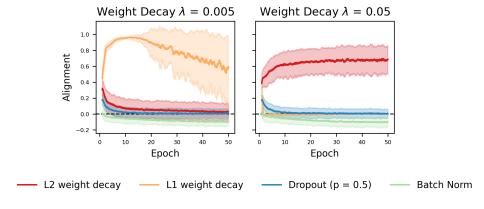


Figure 16: Other regularization techniques have a variety of effects on the Hebbian alignment of the learning signal. While we only developed a theory for L2 weight decay, the alignment seems to exist for some other regularizers as well, such as L1 weight decay and an anti-Hebbian alignment for batch normalization when used to augment SCEs. Batch normalization seems to have an anti-Hebbian effect, while both L1 and L2 weight decay can have a Hebbian effect.

C.1.9 OTHER REGULARIZATION TECHNIQUES

See Figure 16.

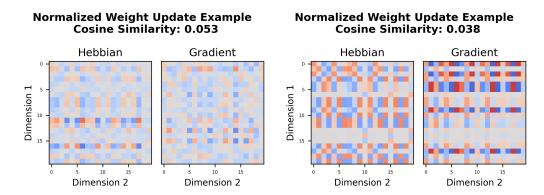


Figure 17: With weight decay, even after the first epoch (**left**), there starts to be an alignment of the directions; at convergence (**right**), even when specific steps have low cosine similarity, there is still clearly a lot of similar structure. At the end of training, many learning signals with low Hebbian alignment still share a surprising amount of structure. The plots above are from a SCE with $\eta = 0.1$ and $\gamma = 0.05$.

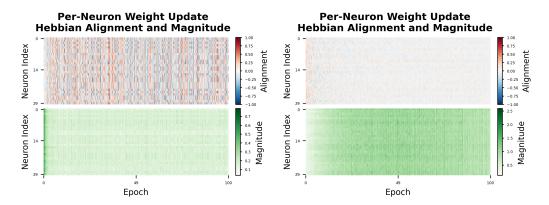


Figure 18: The full weight update of neurons in the neural network strongly oscillates between positively and negatively aligned to the Hebbian update late into training with weight decay (**left**). As can be seen by the blue and red stripes, there is some form of global coherent oscillation in alignment at higher weight decays. This phenomenon becomes much weaker without weight decay (**right**). Both diagrams show 30 example neurons from the second hidden layer of an MLP trained on the standard regression experiment over the course of training. The diagram on the left has a $\gamma = 0.05$ while the one on the right has a gamma of 0.0, both used tanh activations. Our experiments suggest that the oscillation in alignment for individual neurons is not related to the magnitude of the weight update that neuron is receiving, though the experiment run without weight decay does have higher magnitude weight updates since the weights are larger.

C.2 Example alignments during training

C.2.1 LOW ALIGNMENT UPDATE AT END OF TRAINING

C.2.2 FULL WEIGHT UPDATE HEBBIAN OSCILLATION

C.3 RANDOMNN FORMULATION

See Figure 17.

See Figure 18.

The RandomNN was a MLP with 3 hidden vectors of size 128 and tanh activations. The MLP took the same input as the student model but outputted a vector of length 4. The output was averaged across the batch and then and then multiplied by a random projection matrix unique to each parameter and reshaped to be the dimensions of that parameter. No parameters of RandomNN change after initialization. The resulting learning signal for W is a deterministically random low-rank matrix, W^* .

The full weight update is given by:

$$\Delta W = \eta \left(q(x, \theta) - \gamma W \right)$$

where

$$g(x,\theta) = p(W^*)s_{dir}(W)s_{red}(W,W^*)W^*$$

and where,

$$s_{red}(W, W^*) = \text{sign}(\|W\|_2 - \|W - W^*\|_2)$$

$$s_{dir}(W) = \text{sign} (100 - ||W||_2)$$

$$p(W^*) = \begin{cases} 1 & \text{if } ||W^*||_2 \le 1\\ \frac{1}{||W^*||_2 + \epsilon} & \text{otherwise} \end{cases}$$

The minimal requirements to have non-zero weights and reach stationarity require $g(x,\theta)$ to be some forcing function that wants to make the weights larger than zero. This is the case with any descent learning algorithm, as with zero weights, one can not learn or express anything besides 0. However, it is not only true of learning algorithms.

There are a number of trivial constructions that satisfy this condition, such as setting $f(x,\theta) = A$ where A is a random matrix defined at initialization. This will naturally be an expanding force and become aligned with the Hebbian rule, however, it will do this even without regularization. But is there a way to make a non-learning model that does not behave Hebbian at all without regularization, but does with regularization?

RandomNN is one such construction. In it, we produce random weight update vectors in a subspace of the possible directions of the student model's weight updates. This means that after some number of updates, the value of the weight is not orthogonal to the random update vectors, and in fact becomes highly aligned to them. Thus, for a given weight update, the norm of the weights will either increase or decrease, not strictly increase. We can make an attractor to push the norm of the weights to a specific non-zero value by choosing to either add or subtract the random update, depending on which one will move it closer to the target value. Thus, without any regularization, the model's weights will converge to have the target norm and will, on average, not increase or decrease, resulting in no Hebbian alignment. However, once a weight decay term is added, the attractor will try to strictly increase to approach the target, and thus align with the Hebbian update. We also apply a weight update norm clip for stability.