

---

# SGD with Large Step Sizes Learns Sparse Features

---

Maksym Andriushchenko<sup>1</sup> Aditya Varre<sup>1</sup> Loucas Pillaud-Vivien<sup>1</sup> Nicolas Flammarion<sup>1</sup>

## Abstract

We showcase important features of the dynamics of the Stochastic Gradient Descent (SGD) in the training of neural networks. We present empirical observations that commonly used large step sizes (i) may lead the iterates to jump from one side of a valley to the other causing *loss stabilization*, and (ii) this stabilization induces a hidden stochastic dynamics that *biases it implicitly* toward simple predictors. Furthermore, we show empirically that the longer large step sizes keep SGD high in the loss landscape valleys, the better the implicit regularization can operate and find sparse representations. Notably, no explicit regularization is used: the regularization effect comes solely from the SGD dynamics influenced by the large step sizes schedule. Therefore, these observations unveil how, through the step size schedules, both gradient and noise drive together the SGD dynamics through the loss landscape of neural networks. We justify these findings theoretically through the study of simple neural network models as well as qualitative arguments inspired from stochastic processes. This analysis allows us to shed new light on some common practices and observed phenomena when training deep networks. The code of our experiments is available at <https://github.com/tml-epfl/sgd-sparse-features>.

## 1. Introduction

Deep neural networks have accomplished remarkable achievements on a wide variety of tasks. Yet, the understanding of their remarkable effectiveness remains incomplete. From an optimization perspective, stochastic training procedures challenge many insights drawn from convex models. Notably, large step size schedules used in practice lead to unexpected patterns of stabilizations and sudden drops in

the training loss (He et al., 2016). From a generalization perspective, overparametrized deep nets generalize well while fitting perfectly the data and without any explicit regularizers (Zhang et al., 2017). This suggests that optimization and generalization are tightly intertwined: neural networks find solutions that generalize well *thanks* to the optimization procedure used to train them. This property, known as *implicit bias* or *algorithmic regularization*, has been studied both for regression (Li et al., 2018; Woodworth et al., 2020) and classification (Soudry et al., 2018; Lyu & Li, 2020; Chizat & Bach, 2020). However, for these theoretical results, it is also shown that typical timescales needed to enter the beneficial feature learning regimes are prohibitively long (Woodworth et al., 2020; Moroshko et al., 2020).

In this paper, we aim at staying closer to the experimental practice and consider the SGD schedules from the ResNet paper (He et al., 2016) where the *large step size* is first kept constant and then decayed, potentially multiple times. We illustrate this behavior in Fig. 1 where we reproduce a minimal setting without data augmentation or momentum, and with only one step size decrease. We draw attention to two key observations regarding the large step size phase: (a) quickly after the start of training, the loss remains approximately constant on average and (b) despite no progress on the training loss, running this phase for longer leads to better generalization. We refer to such large step size phase as *loss stabilization*. The better generalization hints at some *hidden dynamics* in the parameter space not captured by the loss curves in Fig. 1. Our main contribution is to unveil the hidden dynamics behind this phase: loss stabilization helps to amplify the noise of SGD that drives the network towards a solution with *sparser features*, meaning that for a feature vector  $\psi(x)$ , only a few unique features are active for a given input  $x$ .

### 1.1. Our Contributions

**The effective dynamics behind loss stabilization.** We characterize two main components of the SGD dynamics with large step sizes: (i) a fast movement determined by the bouncing directions causing loss stabilization, (ii) a slow dynamics driven by the combination of the gradient and the multiplicative noise—which is non-vanishing due to the loss stabilization.

---

<sup>1</sup>EPFL. Correspondence to: Maksym Andriushchenko <maksym.andriushchenko@epfl.ch>.

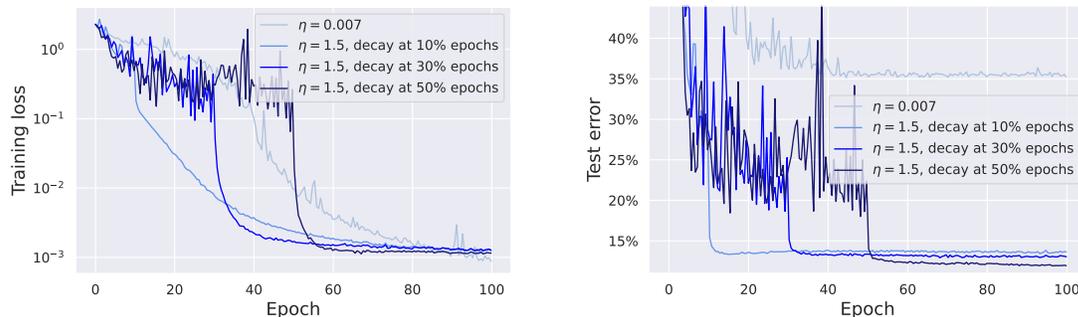


Figure 1. A typical training dynamics for a ResNet-18 trained on CIFAR-10. We use weight decay but no momentum or data augmentation for this experiment. We see a substantial difference in generalization (as large as 12% vs. 35% test error) depending on the step size  $\eta$  and its schedule. When the training loss stabilizes, there is a hidden progress occurring which we aim to characterize.

**SDE model and sparse feature learning.** We model the *effective* slow dynamics during loss stabilization by a stochastic differential equation (SDE) whose multiplicative noise is related to the neural tangent kernel features, and validate this modeling experimentally. Building on the existing theory on diagonal linear networks, which shows that this noise structure leads to sparse predictors, we conjecture a similar “sparsifying” effect on the features of more complex architectures. We experimentally confirm this on neural networks of increasing complexity.

**Insights from our understanding.** We draw a clear general picture: the hidden optimization dynamics induced by large step sizes and loss stabilization enable the transition to a sparse feature learning regime. We argue that after a short initial phase of training, SGD *first* identifies sparse features of the training data and eventually fits the data when the step size is decreased. Finally, we discuss informally how many deep learning regularization methods (weight decay, BatchNorm, SAM) may also fit into the same picture.

## 1.2. Related Work

He et al. (2016) popularized the piece-wise constant step size schedule which often exhibits a clear loss stabilization pattern which was later characterized theoretically in Li et al. (2020) from the optimization point of view. However, the regularization effect of this phase induced by the underlying *hidden stochastic dynamics* is still unclear. Li et al. (2019b) analyzed the role of loss stabilization for a synthetic distribution containing different patterns, but it is not clear how this analysis can be extended to general problems. Jastrzebski et al. (2021) suggest that large step sizes prevent the increase of local curvature during the early phase of training. However, they do not provide an explanation for this phenomenon.

The importance of large step sizes for generalization has been investigated with diverse motivations. Many works conjectured that large step sizes induce minimization of some complexity measures related to the flatness of minima

(Keskar et al., 2016; Smith & Le, 2018; Smith et al., 2021; Yang et al., 2022). Notably, Xing et al. (2018) point out that SGD moves through the loss landscape bouncing between the walls of a valley where the role of large step sizes is to guide the SGD iterates towards a flatter minimum. However, the correct flatness measure is often disputed (Dinh et al., 2017) and its role in understanding generalization is questionable since full-batch GD with large step sizes (unlike SGD) can lead to flat solutions which don’t generalize well (Kaur et al., 2022)

The attempts to explain the effect of large step size on strongly convex models (Nakkiran, 2020; Wu et al., 2021; Beugnot et al., 2022) are inherently incomplete since it is a phenomenon related to the existence of many zero solutions with very different generalization properties. Works based on stability analysis characterize the properties of the minimum that SGD or GD can potentially converge depending on the step size (Wu et al., 2018; Mulayoff et al., 2021; Ma & Ying, 2021; Nacson et al., 2022). However, these approaches *do not capture the entire training dynamics* such as the large step size phase that we consider where SGD converges only after the step size is decayed.

To grasp the generalization of SGD, research has focused on SGD augmented with label noise due to its beneficial regularization properties and resemblance to the standard noise of SGD. Its implicit bias has been first characterized by Blanc et al. (2020) and extended by Li et al. (2022). However, their analysis only holds in the final phase of the training, close to a zero-loss manifold. Our work instead is closer in spirit to Pillaud-Vivien et al. (2022) where the label noise dynamics is analyzed in the *central* phase of the training, i.e., when the loss is still substantially above zero.

The dynamics of GD with large step sizes have received a lot of attention in recent times, particularly the edge-of-stability phenomenon (Cohen et al., 2021) and the catapult mechanism (Lewkowycz et al., 2020; Wang et al., 2022). However, the lack of stochastic noise in their analysis renders them incapable of capturing stochastic training. Note

that it is possible to bridge the gap between GD and SGD by using explicit regularization as in Geiping et al. (2022). We instead focus on the *implicit* regularization of SGD which remains the most practical approach for training deep nets.

Finally, sparse features and low-rank structures in deep networks have been commonly used for model compression, knowledge distillation, and lottery ticket hypothesis (Denton et al., 2014; Hinton et al., 2015; Frankle & Carbin, 2018). A common theme of all these works is the presence of *hidden structure* in the networks learned by SGD which allows one to come up with a much smaller network that approximates well the original one. In particular, Hoefler et al. (2021) note that ReLU activations in deep networks trained with SGD are typically much sparser than 50%. Our findings suggest that the step size schedule can be the key component behind emergence of such sparsity.

## 2. The Effective Dynamics of SGD with Large Step Size: Sparse Feature Learning

In this section, we show that large step sizes may lead the loss to stabilize by making SGD bounce above a valley. We then unveil the effective dynamics induced by this loss stabilization. To clarify our exposition we showcase our results for the mean square error but other losses like the cross-entropy carry the same key properties in terms of the noise covariance (Wojtowytsch, 2021b, Lemma 2.14). We consider a generic parameterized family of prediction functions  $\mathcal{H} := \{x \rightarrow h_\theta(x), \theta \in \mathbb{R}^p\}$ , a setting which encompasses neural networks. In this case, the training loss on input/output samples  $(x_i, y_i)_{1 \leq i \leq n} \in \mathbb{R}^d \times \mathbb{R}$  is equal to

$$\mathcal{L}(\theta) := \frac{1}{2n} \sum_{i=1}^n (h_\theta(x_i) - y_i)^2. \quad (1)$$

We consider the overparameterized setting, i.e.  $p \gg n$ , hence, there shall exist many parameters  $\theta^*$  that lead to zero loss, i.e., perfectly interpolate the dataset. Therefore, the question of which interpolator the algorithm converges to is of paramount importance in terms of generalization. We focus on the SGD recursion with step size  $\eta > 0$ , initialized at  $\theta_0 \in \mathbb{R}^p$ : for all  $t \in \mathbb{N}$ ,

$$\theta_{t+1} = \theta_t - \eta(h_{\theta_t}(x_{i_t}) - y_{i_t})\nabla_\theta h_{\theta_t}(x_{i_t}), \quad (2)$$

where  $i_t \sim \mathcal{U}(\llbracket 1, n \rrbracket)$  is the uniform distribution over the sample indices. In the following, note that SGD with mini batches of size  $B > 1$  would lead to similar analysis but with  $\eta/B$  instead of  $\eta$ .

### 2.1. Background: SGD is GD with Specific Label Noise

To emphasize the combined roles of gradient and noise, we highlight the connection between the SGD dynamics and that of full-batch GD plus a specific label noise. Such

manner of reformulating the dynamics has already been used in previous works attempting to understand the specificity of the SGD noise (HaoChen et al., 2021; Ziyin et al., 2022). We formalize it in the following proposition.

**Proposition 2.1.** *Let  $(\theta_t)_{t \geq 0}$  follow the SGD dynamics Eq.(2) with the random sampling function  $(i_t)_{t \geq 0}$ . For  $t \geq 0$ , define the random vector  $\xi_t \in \mathbb{R}^n$  such that*

$$[\xi_t]_i := (h_{\theta_t}(x_i) - y_i)(1 - n\mathbf{1}_{i=i_t}), \quad (3)$$

for  $i \in \llbracket 1, n \rrbracket$  and where  $\mathbf{1}_A$  is the indicator of the event  $A$ . Then  $(\theta_t)_{t \geq 0}$  follows the full-batch gradient dynamics on  $\mathcal{L}$  with label noise  $(\xi_t)_{t \geq 0}$ , that is

$$\theta_{t+1} = \theta_t - \frac{\eta}{n} \sum_{i=1}^n (h_{\theta_t}(x_i) - y_i^t) \nabla_\theta h_{\theta_t}(x_i), \quad (4)$$

where we define the random labels  $y^t := y + \xi_t$ . Furthermore,  $\xi_t$  is a mean zero random vector with variance such that  $\frac{1}{n(n-1)} \mathbb{E} \|\xi_t\|^2 = 2\mathcal{L}(\theta_t)$ .

This reformulation shows two crucial aspects of the SGD noise: (i) the noisy part at state  $\theta$  always belongs to the linear space spanned by  $\{\nabla_\theta h_\theta(x_1), \dots, \nabla_\theta h_\theta(x_n)\}$ , and (ii) it scales as the training loss. Going further on (ii), we highlight in the following section that the loss can stabilize because of large step sizes: this may lead to a *constant* effective scale of label noise. These two features are of paramount importance when modelling the effective dynamics that take place during loss stabilization.

### 2.2. The Effective Dynamics Behind Loss Stabilization

**On loss stabilization.** For generic quadratic costs, e.g.,  $F(\beta) := \|X\beta - y\|^2$ , gradient descent with step size  $\eta$  is convergent for  $\eta < 2/\lambda_{\max}$ , divergent for  $\eta > 2/\lambda_{\max}$  and converges to a bouncing 2-periodic dynamics for  $\eta = 2/\lambda_{\max}$ , where  $\lambda_{\max}$  is the largest eigenvalue of the Hessian. However, the practitioner is not likely to hit perfectly this unstable step size and, almost surely, the dynamics shall either converge or diverge. Yet, non-quadratic costs bring to this picture a particular complexity: it has been shown that, even for non-convex toy models, there exist an open interval of step sizes for which the gradient descent neither converge nor diverge (Ma et al., 2022; Chen & Bruna, 2022). As we are interested in SGD, we complement this result by presenting an example in which loss stabilization occurs almost surely *in the case of stochastic updates*. Indeed, consider a regression problem with quadratic parameterization on one-dimensional data inputs  $x_i$ 's, coming from a distribution  $\hat{\rho}$ , and outputs generated by the linear model  $y_i = x_i\theta^*$ . The loss writes  $F(\theta) := \frac{1}{4} \mathbb{E}_{\hat{\rho}} (y - x\theta)^2$ , and the SGD iterates with step size  $\eta > 0$  follow, for any  $t \in \mathbb{N}$ ,

$$\theta_{t+1} = \theta_t + \eta \theta_t x_{i_t} (y_{i_t} - x_{i_t} \theta_t^2) \quad \text{where} \quad x_{i_t} \sim \hat{\rho}. \quad (5)$$

For the sake of clarity, suppose that  $\theta_* = 1$  and  $\text{supp}(\hat{\rho}) = [a, b]$ , we have the following proposition (a more general result is presented in Proposition B.1 of the Appendix).

**Proposition 2.2.** *For any  $\eta \in (a^{-2}, 1.25 \cdot b^{-2})$  and initialization  $\theta_0 \in (0, 1)$ , for all  $t > 0$ ,*

$$\delta_1 < F(\theta_t) < \delta_2 \text{ almost surely, and} \quad (6)$$

$$\exists T > 0, \forall k > T, \theta_{t+2k} < 1 < \theta_{t+2k+1} \text{ almost surely.} \quad (7)$$

where  $\delta_1, \delta_2, T > 0$  are constant given in the Appendix.

The proposition is divided in two parts: if the step size is large enough, Eq.(6) the loss stabilizes in between level sets  $\delta_1$  and  $\delta_2$  and Eq.(7) shows that after some initial phase, the iterates bounce from one side of the *loss valley* to the other one. Note that despite the stochasticity of the process, the results hold *almost surely*.

**The effective dynamics.** As observed in the prototypical SGD training dynamics of Fig. 1 and proved in the non-convex toy model of Proposition 2.2, large step sizes lead the loss to stabilize around some level set. To further understand the effect of this loss stabilization in parameter space, we shall assume perfect stabilization. Then, from Proposition 2.1, we conjecture the following behaviour

*During loss stabilization, SGD is well modelled by GD with constant label noise.*

Label noise dynamics have been studied recently (Blanc et al., 2020; Damian et al., 2021; Li et al., 2022) thanks to their connection with Stochastic Differential Equations (SDEs). To properly write a SDE model, the drift should match the gradient descent and the noise should have the correct covariance structure (Li et al., 2019a; Wojtowytsch, 2021a). Proposition 2.1 implies that the noise at state  $\theta$  is spanned by the gradient vectors  $\{\nabla_{\theta} h_{\theta}(x_1), \dots, \nabla_{\theta} h_{\theta}(x_n)\}$  and has a constant intensity corresponding to the loss stabilization at a level  $\delta > 0$ . Hence, we propose the following SDE model

$$d\theta_t = -\nabla_{\theta} \mathcal{L}(\theta_t) dt + \sqrt{\eta \delta} \phi_{\theta_t}(X)^{\top} dB_t, \quad (8)$$

where  $(B_t)_{t \geq 0}$  is a standard Brownian motion in  $\mathbb{R}^n$  and  $\phi_{\theta}(X) := [\nabla_{\theta} h_{\theta}(x_i)]_{i=1}^n \in \mathbb{R}^{n \times p}$  referred to as the *Jacobian* (which is also the *Neural Tangent Kernel (NTK) feature matrix* (Jacot et al., 2018)). This SDE can be seen as the *effective slow dynamics* that drives the iterates while they bounce *rapidly* in some directions at the level set  $\delta$ . It highlights the combination of the deterministic part of the full-batch gradient and the noise induced by SGD. Beyond the theoretical justification and consistency of this SDE model, we validate it empirically in Sec. C showing that it indeed captures the dynamics of large step size SGD. In the next section, we leverage the SDE (8) to understand the implicit bias of such learning dynamics.

## 2.3. Sparse Feature Learning

We begin with a simple model of diagonal linear networks that showcase a sparsity inducing dynamics and further disclose our general message about the overall implicit bias promoted by the effective dynamics.

### 2.3.1. A WARM-UP: DIAGONAL LINEAR NETWORKS

An appealing example of simple non-linear networks that help in forging an intuition for more complicated architectures is diagonal linear networks (Vaskevicius et al., 2019; Woodworth et al., 2020; HaoChen et al., 2021; Pesme et al., 2021). They are two-layer linear networks with only diagonal connections: the prediction function writes  $h_{u,v}(x) = \langle u, v \odot x \rangle = \langle u \odot v, x \rangle$  where  $\odot$  denotes *elementwise* multiplication. Even though the loss is convex in the associated linear predictor  $\beta := u \odot v \in \mathbb{R}^d$ , it is not in  $(u, v)$ , hence the training of such simple models already exhibit a rich non-convex dynamics. In this case,  $\nabla_u h_{u,v}(x) = v \odot x$ , and the SDE model Eq.(8) writes

$$du_t = -\nabla_u \mathcal{L}(u_t, v_t) dt + \sqrt{\eta \delta} v_t \odot [X^{\top} dB_t], \quad (9)$$

where  $(B_t)_{t \geq 0}$  is a standard Brownian motion in  $\mathbb{R}^n$ . Equations are symmetric for  $(v_t)_{t \geq 0}$ .

**What is the behaviour of this effective dynamics?** (Pillaud-Vivien et al., 2022) answered this question by analyzing a similar stochastic dynamics and unveiled the sparse nature of the resulting solutions. Indeed, under sparse recovery assumptions, denoting  $\beta^*$  the sparsest linear predictor that interpolates the data, it is shown that the associated linear predictor  $\beta_t = u_t \odot v_t$ : (i) converges exponentially fast to zero outside of the support of  $\beta^*$  (ii) is *with high probability* in a  $\mathcal{O}(\sqrt{\eta \delta})$  neighborhood of  $\beta^*$  in its support after a time  $\mathcal{O}(\delta^{-1})$ .

**Overall conclusion on the model.** During a first phase, SGD with large step sizes  $\eta$  decreases the training loss until stabilization at some level set  $\delta > 0$ . During this loss stabilization, an effective noise-driven dynamics takes place. It shrinks the coordinates outside of the support of the sparsest signal and oscillates in parameter space at level  $\mathcal{O}(\sqrt{\eta \delta})$  on its support. Hence, decreasing later the step size leads to perfect recovery of the sparsest predictor. This behaviour is illustrated in our experiments in Figure 2.

### 2.3.2. THE SPARSE FEATURE LEARNING CONJECTURE FOR MORE GENERAL MODELS

Results for diagonal linear nets recalled in the previous paragraph show that the noisy dynamics (9) induce a *sparsity bias*. As emphasized in HaoChen et al. (2021), this effect is largely due to the multiplicative structure of the noise  $v \odot [X^{\top} dB_t]$  that, in this case, has a shrinking effect *on the coordinates* (because of the coordinate-wise multiplication

with  $v$ ). In the general case, we see, thanks to Eq.(8), that the same multiplicative structure of the noise still happens but this time *with respect to the Jacobian*  $\phi_\theta(X)$ . Hence, this suggests that similarly to the diagonal linear network case, the implicit bias of the noise can lead to a shrinkage effect applied to  $\phi_\theta(X)$ . This effect depends on the noise intensity  $\delta$  and the step size of SGD. Indeed, an interesting property of Brownian motion is that, for  $v \in \mathbb{R}^p$ ,  $\langle v, B_t \rangle = \|v\|_2 W_t$ , where the equality holds in law and  $(W_t)_{t \geq 0}$  is a one-dimensional Brownian motion. Hence, the process Eq.(8) is equivalent to a process whose  $i$ -th coordinate is driven by a noise proportional to  $\|\phi_i\| dW_t^i$ , where  $\phi_i$  is the  $i$ -th column of  $\phi_\theta(X)$  and  $(W_t^i)_{t \geq 0}$  is a Brownian motion. This SDE structure, similar to the geometric Brownian motion, is expected to induce the shrinkage of each multiplicative factor (Oksendal, 2013, Section 5.1), i.e., in our case  $(\|\nabla_\theta h(x_i)\|)_{i=1}^n$ . Thus, we conjecture:

*The noise part of Eq.(8) seeks to minimize the  $\ell_2$ -norm of the columns of  $\phi_\theta(X)$ .*

Note that the *fitting part* of the dynamics prevents the Jacobian to collapse totally to zero, but as soon as they are not needed to fit the signal, *columns* can be reduced to zero. Remarkably, from a stability perspective, Blanc et al. (2020) showed a similar bias: locally around a minimum, the SGD dynamics implicitly tries to minimize the *Frobenius norm*  $\|\phi_\theta(X)\|_F = \sum_{i=1}^n \|\nabla_\theta h_\theta(x_i)\|^2$ . Resolving the above conjecture and characterizing the implicit bias along the trajectory of SGD remains an exciting avenue for future work. Now, we provide a specification of this implicit bias for different architectures:

- **Diagonal linear networks:** For  $h_{u,v}(x) = \langle u \odot v, x \rangle$ , we have  $\nabla_{u,v} h_{u,v}(x) = [v \odot x, u \odot x]$ . Thus, for a generic data matrix  $X$ , minimizing the norm of each column of  $\phi_{u,v}(X)$  amounts to put the maximal number of zero coordinates and hence to minimize  $\|u \odot v\|_0$ .
- **ReLU networks:** We take the prototypical one hidden layer to exhibit the sparsification effect. Let  $h_{a,W}(x) = \langle a, \sigma(Wx) \rangle$ , then  $\nabla_a h_{a,W}(x) = \sigma(Wx)$  and  $\nabla_{w_j} h_{a,W}(x) = a_j x \mathbf{1}_{\langle w_j, x \rangle > 0}$ . Note that the  $\ell_2$ -norm of the column corresponding to the neuron is reduced when it is activated at a *minimal number of training points*, hence the implicit bias enables the learning of *sparse data-active features*. Finally, when some directions are needed to fit the data, similarly activated neurons align to fit, reducing the rank of  $\phi_\theta(X)$ .

**Feature sparsity.** Our main insight is that the Jacobian could be significantly simplified during the loss stabilization phase. Indeed, while the gradient part tries to fit the data and align neurons (see e.g. Fig. 10), the noise part of Eq.(8) intends to minimize the  $\ell_2$ -norm of the columns of  $\phi(X)$ . Hence, in combination, this motivates us to count the average number of *distinct* (i.e., counting a group of aligned

neurons as one), *non-zero* activations over the training set. We refer to this as the *feature sparsity coefficient* (see the next section for a detailed description). Note that the aforementioned sparsity comes both in the number of distinct neurons and their activation.

We show next that the conjectured sparsity is indeed observed empirically for a variety of models. Remark that both the feature sparsity coefficient and the rank of  $\phi_\theta(X)$  can be used as a good proxy to track the hidden progress during the loss stabilization phase.

### 3. Empirical Evidence of Sparse Feature Learning Driven by SGD

Here we present empirical results for neural networks of increasing complexity: from diagonal linear networks to deep DenseNets on CIFAR-10, CIFAR-100, and Tiny ImageNet. We make the following common observations for all these networks trained using SGD schedules with large step sizes:

- (O1) **Loss stabilization:** training loss stabilizes around a high level set until step size is decayed,
- (O2) **Generalization benefit:** longer loss stabilization leads to better generalization,
- (O3) **Sparse feature learning:** longer loss stabilization leads to sparser features.

Importantly, *we use no explicit regularization* (in particular, no weight decay) in our experiments so that the training dynamics is driven purely by SGD and the step size schedule. Additionally, in some cases, we cannot find a single large step size that would lead to loss stabilization. In such cases, whenever explicitly mentioned, we use a *warmup* step size schedule—i.e., increasing step sizes according to some schedule—to make sure that the loss stabilizes around some level set. Warmup is commonly used in practice (He et al., 2016; Devlin et al., 2018) and often motivated purely from the optimization perspective as a way to accelerate training (Agarwal et al., 2021), but we suggest that it is also a way to amplify the regularization effect of the SGD noise which is proportional to the step size.

**Measuring sparse feature learning.** We track the simplification of the Jacobian by measuring both the feature sparsity and the rank of  $\phi_\theta(X)$ . We compute the rank over iterations for each model (except deep networks for which it is prohibitively expensive) by using a fixed threshold on the singular values of  $\phi_\theta(X)$  normalized by the largest singular value. In this way, we ensure that the difference in the rank that we detect is not simply due to different scales of  $\phi_\theta(X)$ . Moreover, we always compute  $\phi_\theta(X)$  on the number of fresh samples equal to the number of parameters  $|\theta|$  to make sure that rank deficiency is not coming from  $n \ll |\theta|$  which is the case in the overparametrized settings we consider. To compute the **feature sparsity coefficient**,

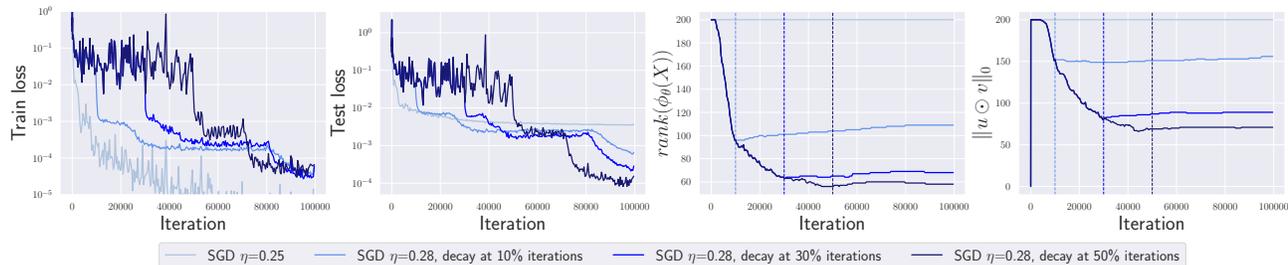


Figure 2. **Diagonal linear networks.** We observe loss stabilization, better generalization for longer schedules, minimization of the rank of  $\phi_\theta(X)$  and sparsity of the predictor  $u \odot v$ .

we count the average fraction of *distinct* (i.e., counting a group of highly correlated activations as one), *non-zero* activations at some layer over the training set. Note that the value of 100% means a completely *dense* feature vector and 0% means a feature vector with all zeros. We count a pair of activations  $i$  and  $j$  as highly correlated if their Pearson’s correlation coefficient is at least 0.95. Unlike  $\text{rank}(\phi_\theta(X))$ , the feature sparsity coefficient scales to deep nets and has an easy-to-grasp meaning.

### 3.1. Sparse Feature Learning in Diagonal Linear Networks

**Setup.** The inputs  $x_1, \dots, x_n$  with  $n = 80$  are sampled from  $\mathcal{N}(0, I_d)$  where  $I_d$  is an identity matrix with  $d = 200$ , and the outputs are generated as  $y_i = \langle \beta_*, x_i \rangle$  where  $\beta_* \in \mathbb{R}^d$  is  $r = 20$  sparse. We consider four different SGD runs (started from  $u_i = 0.1, v_i = 0$  for each  $i$ ): one with a small step size and three other with initial large step size decayed after 10%, 30%, 50% iterations, respectively.

**Observations.** We show the results in Fig. 2 and note that (O1)–(O3) hold even in this simple model trained with vanilla SGD without any explicit regularization or layer normalization schemes. We observe that the training loss stabilizes around  $10^{-1.5}$ , the test loss improves for longer schedules, both  $\text{rank}(\phi_\theta(X))$  and  $\|u \odot v\|_0$  decrease during the loss stabilization phase leading to a sparse final predictor. While the training loss has seemingly converged to  $10^{-1.5}$ , a hidden dynamics suggested by Eq.(9) occurs which slowly drifts the iterates to a sparse solution. This implicit sparsification explains the dependence of the final test loss on the time when the large step size is decayed, similarly to what has been observed for deep networks in Fig. 1. Interestingly, we also note that SGD with large step-size schedules encounters saddle points *after* we decay the step size (see the training loss curves in Fig. 2) which resembles the saddle-to-saddle regime described in Jacot et al. (2021) which does not occur in the large-initialization lazy training regime.

**SGD and GD have different implicit biases.** Since we observe from Fig. 2 that for loss stabilization, stochasticity alone does not suffice and large step sizes are necessary, one may wonder if conversely, only large step sizes

can be sufficient to have a sparsifying effect. Even if special instances can be found for which large step sizes are sufficient (such as for non-centered input features as in Nacson et al. (2022)), we answer this negatively showing that gradient descent in general does not go to the sparsest solution as demonstrated in Fig. 11 in the Appendix. Moreover, in Fig. 3, we visualize the difference in trajectory between the two methods taken with large step sizes over a 2D subspace spanned by  $w^* - w_{init}$  and  $w_{flow} - w_{init}$ , where  $w^*$  is the ground truth,  $w_{flow}$  is the result of gradient flow, and  $w_{init}$  is the initialization. This example provides an intuition that loss stabilization alone is not sufficient for sparsification and that the role of noise described earlier is crucial.

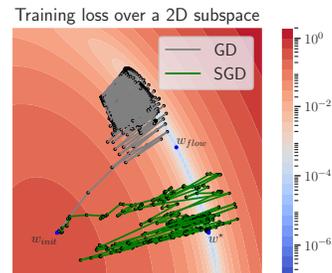


Figure 3. **Diagonal linear networks.** GD and SGD take different trajectories.

### 3.2. Sparse Feature Learning in Simple ReLU Networks

**Two-layer ReLU network in 1D.** We consider the 1D regression task from Blanc et al. (2020) with 12 points, where label noise SGD has been shown to learn a simple model. We show that similar results can be achieved with large-step-size SGD via loss stabilization. We train a ReLU network with 100 neurons with SGD with a long linear warmup (otherwise, we were unable to achieve approximate loss stabilization), directly followed by a step size decay. The two plots correspond to a warmup/decay transition at 2% and 50% of iterations, respectively. The results shown in Fig. 4 confirm that (O1)–(O3) hold: the training loss stabilizes around  $10^{-0.5}$ , the predictor becomes much simpler and is expected to generalize better, and both  $\text{rank}(\phi_\theta(X))$  and the feature sparsity coefficient substantially decrease during the loss stabilization phase. Interestingly, the rank reduction of  $\phi_\theta(X)$  occurs because of zero *activations*, and not because of zero *weights*. For this one-dimensional task, we can directly observe the final predictor which is sparse in terms of the number of distinct ReLU kinks (i.e., having a few piecewise-linear segments) as captured by the feature spar-

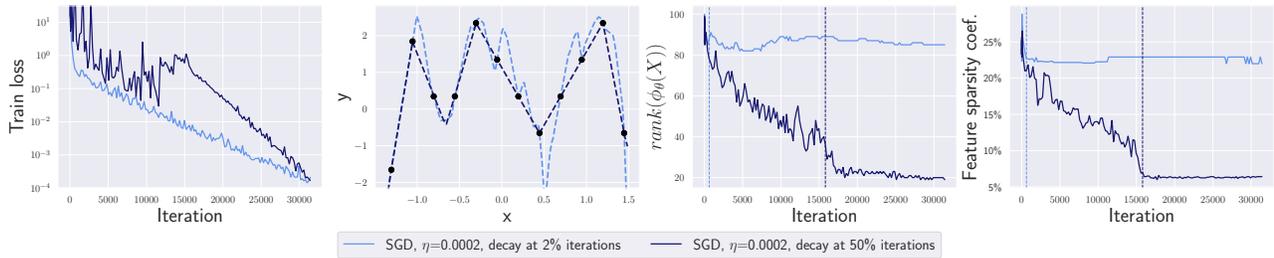


Figure 4. Two-layer ReLU networks for 1D regression. We observe loss stabilization, simplification of the model trained with a longer schedule, lower rank of  $\phi_\theta(X)$ , and much sparser features.

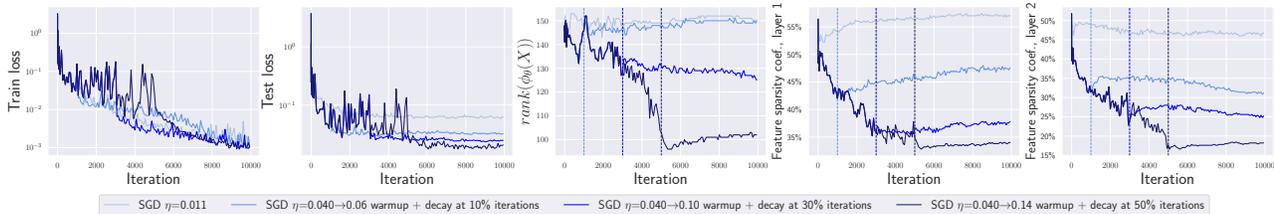


Figure 5. Three-layer ReLU networks in a teacher-student setup. We observe loss stabilization, lower rank of the Jacobian and lower feature sparsity coefficient on both hidden layers.

sity coefficient and the rank of the Jacobian. Interestingly, we also observed *overregularization* for even larger step sizes when we cannot fit all the training points (see Fig. 12 in Appendix). This phenomenon clearly illustrates how the capacity control is induced by the optimization algorithm: *the function class over which we optimize depends on the step size schedule*. Additionally, Fig. 13 in App. shows the evolution of the predictor over iterations. The general picture is confirmed: first the model is simplified during the loss stabilization phase and only then fits the data.

**Deeper ReLU networks.** We use a teacher-student setup with a random *three-layer* teacher ReLU network having 2 neurons on each hidden layer. The student network is over-parametrized with 10 neurons on each layer and is trained on 50 examples. Such teacher-student setup is useful since we know that the student network can implement the ground truth function but might not find it due to the small sample size. We train models using SGD with a medium constant step size and a large step size with warmup decayed after 10%, 30%, 50% iterations, respectively. The results shown in Fig. 5 confirm that (O1)–(O3) hold: the training loss stabilizes around  $10^{-1.5}$ , the test loss is smaller for longer schedules, and both  $\text{rank}(\phi_\theta(X))$  and the feature sparsity coefficient substantially decrease during the loss stabilization phase. All methods have the same value of the training loss ( $10^{-3}$ ) after  $10^4$  iterations but different generalization. Moreover, we see that the feature sparsity coefficient decreases on each layer which makes this metric a promising one to consider for deeper networks.

### 3.3. Sparse Feature Learning in Deep ReLU Networks

**Setup.** We consider here an image classification task and train a DenseNet-100-12 on CIFAR-10, CIFAR-100, and Tiny ImageNet using SGD with batch size 256 and different step size schedules. We use an exponentially increasing warmup schedule with exponent 1.05 to stabilize the training loss. We cannot measure the rank of  $\phi(X)$  here since this matrix is too large ( $\approx (5 \times 10^4) \times (2 \times 10^7)$ ) so we measure only the feature sparsity coefficient taken at two layers: at the end of super-block 3 (i.e., in the middle of the network) and super-block 4 (i.e., right before global average pooling at the end of the network) of DenseNets. We test two settings: a basic setting and a state-of-the-art setting with momentum and standard augmentations.

**Observations.** The results shown in Fig. 6 and 7 confirm that our main findings also hold for deep convolutional networks used in practice: the training loss approximately stabilizes, the test error is becoming progressively better for longer schedules, and the feature sparsity coefficient gradually decreases at both super blocks 3 and 4 until the step size is decayed. We also see that small step sizes consistently lead to suboptimal generalization, e.g., 60% vs. 35% in the basic setting on CIFAR-100. This poor performance confirms that it is crucial to leverage the implicit bias of large step sizes. The difference in the feature sparsity coefficient is also substantial: typically 50%-60% for small step sizes vs. 10%-20% for larger step sizes at block 4. The observations are similar for the state-of-the-art setting as well where we also see a noticeable difference in generalization and feature sparsity depending on the step size and schedule. Finally, we note that while both the feature sparsity coefficient

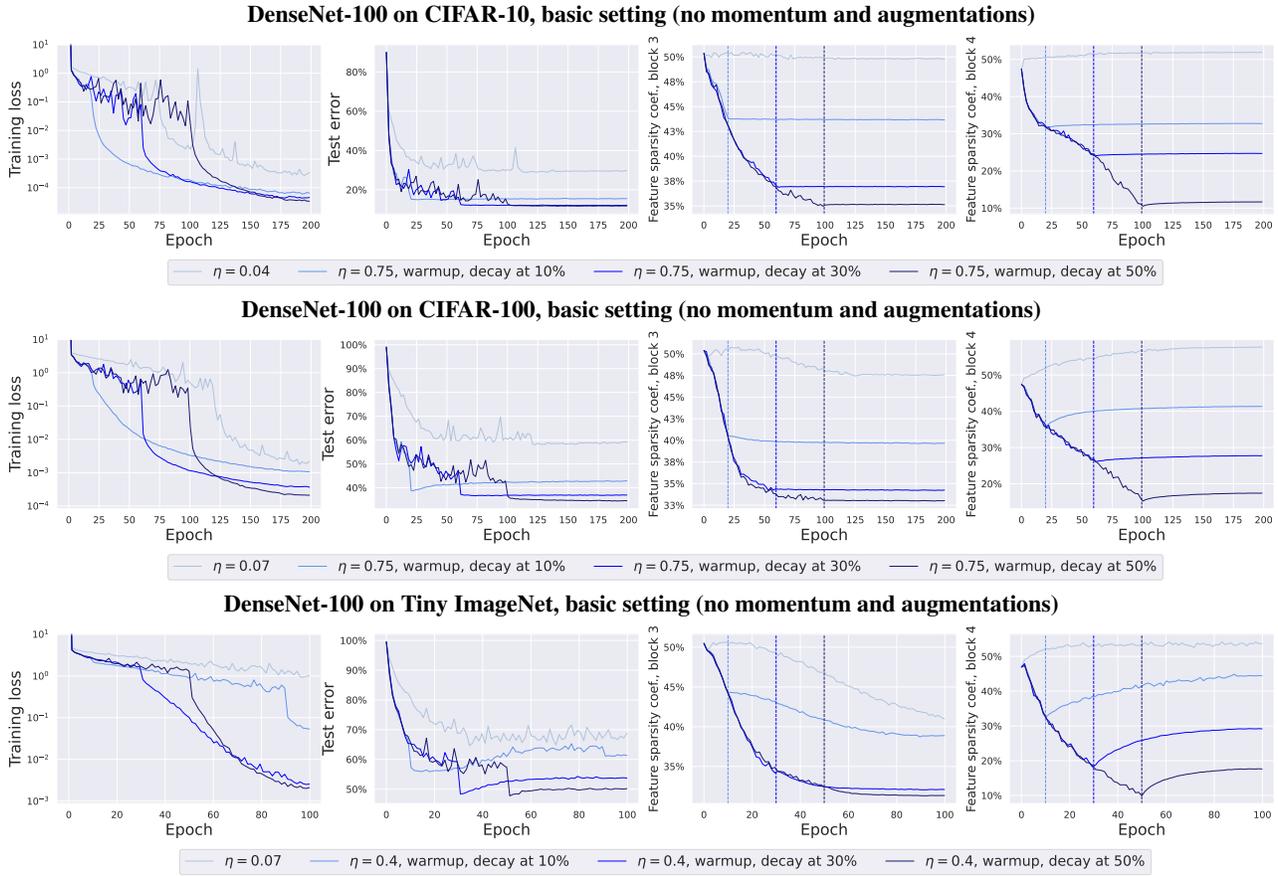


Figure 6. Experiments with DenseNet-100 in the basic setting. We can see that the training loss stabilizes, the test error noticeably depends on the length of the schedule, and the feature sparsity coefficient is minimized during the large step size phase.

cient and test error decrease together, it remains to be seen whether they are causally related on natural datasets.

We show the results with similar findings for other architectures (ResNets-18 and ResNets-34) on CIFAR-10 and CIFAR-100 in Fig. 15 and Fig. 16 in Appendix. Additionally, Fig. 17 illustrates that for small step sizes, the early and middle layers stay very close to their random initialization which indicates the absence of feature learning similarly to what is suggested by the neuron movement plot in Fig. 10 in the Appendix for a two-layer network in a teacher-student setup.

#### 4. Conclusions and Insights from our Understanding of the Training Dynamics

Here we provide an extended discussion on the training dynamics of neural networks resulting from our theoretical and empirical findings.

**The multiple stages of the SGD training dynamics.** As analyzed and shown empirically, the training dynamics we considered can be split onto three distinct phases: (i) an

initial phase of reducing the loss down to some level where stabilization can occur, (ii) a loss stabilization phase where noise and gradient directions combine to find architecture-dependent sparse representations of the data, (iii) a final phase when the step size is decreased to fit the training data. This typology clearly disentangles the effect of the stabilization phase (ii) which relies on the implicit bias of SGD to simplify the model. Note that phases (ii) and (iii) can be repeated until final convergence (He et al., 2016). Moreover, in some training schedules, (ii) does not explicitly occur, and the effect of loss stabilization (ii) and data fitting (iii) can occur simultaneously (Loshchilov & Hutter, 2019).

**From lazy training to feature learning.** Similar sparse implicit biases have been shown for regression with infinitely small initialization (Boursier et al., 2022) and for classification (Chizat & Bach, 2020; Lyu & Li, 2020). However, both approaches are not practical from the computational point of view since (i) the origin is a saddle point for regression leading to the vanishing gradient problem (especially, for deep networks), and (ii) max-margin bias for classification is only expected to happen in the asymptotic phase (Moroshko et al., 2020). On the contrary, large step sizes enable

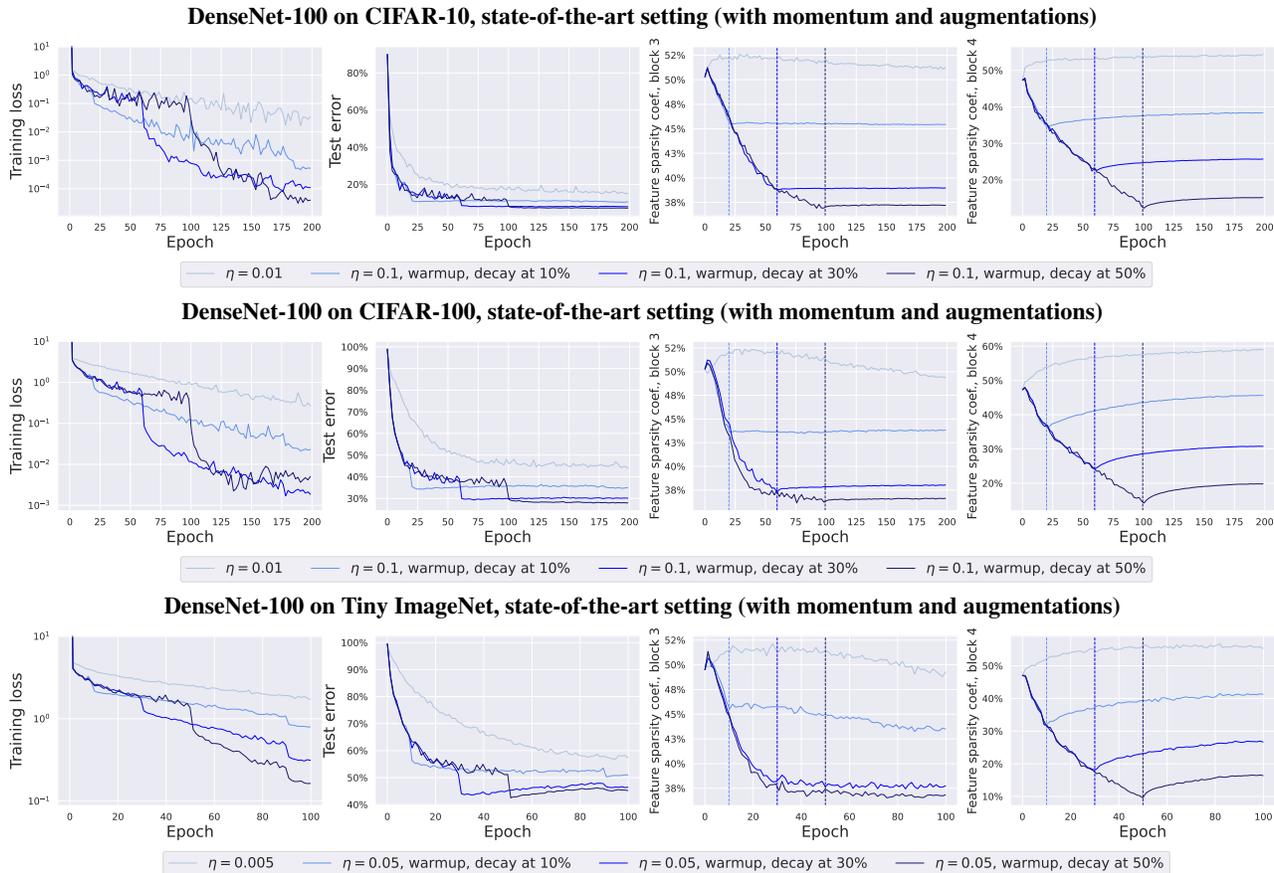


Figure 7. Experiments with DenseNet-100 in the state-of-the-art setting. We can see that the training loss stabilizes, the test error noticeably depends on the length of the schedule, and the feature sparsity coefficient is minimized during the large step size phase.

to initialize far from the origin, while allowing to *efficiently* transition from a regime close to the lazy NTK regime (Jacot et al., 2018) to the rich feature learning regime.

**Common patterns in the existing techniques.** Tuning the step size to obtain loss stabilization can be difficult. To prevent early divergence caused by too large step sizes, we sometimes had to rely on an increasing step size schedule (known as *warmup*). Interpreting such schedules as a tool to favor implicit regularization provides a new explanation to their success and popularity. Additionally, normalization schemes like *batch normalization* or *weight decay*, beyond carrying their own implicit or explicit regularization properties, can be analyzed from a similar lens: they allow to use larger step sizes that boost further the implicit bias effect of SGD while preventing divergence (Bjorck et al., 2018; Zhang et al., 2018; Li & Arora, 2019). Note also that we derived our analysis with batch size equal to one for the sake of clarity, but an arbitrary batch size  $B$  would simply be equivalent to replacing  $\gamma \leftarrow \gamma/B$ . Similarly to the consequence of large step sizes, preferring *smaller batch sizes* (Keskar et al., 2016) while avoiding divergence seem key to benefit from the implicit bias of SGD. Finally, the effect

of large step sizes or small batches is often connected to measures of *flatness* of the loss surface via stability analysis (Wu et al., 2018) and some methods like the Hessian regularization (Damian et al., 2021) or SAM (Foret et al., 2021) explicitly optimize it. Such methods resemble the implicit bias of SGD with loss stabilization implied by the label noise equation (Eq.(8)) where matrix  $\phi_\theta(X)$  is the key component of the Hessian. However, an important practical difference is that the regularization strength in these methods is explicit and decoupled from the step size schedule which may be harder to properly tune since it is simultaneously responsible for optimization *and* generalization.

### Acknowledgements

We thank Maria-Luiza Vladarean and Scott Pesme for insightful discussions and proofs checking. M.A. was supported by the Google Fellowship and Open Phil AI Fellowship. A.V. was supported by the Swiss Data Science Center Fellowship. L.P.V.’s work was partially supported by the Simons Foundation.

## References

- Agarwal, N., Goel, S., and Zhang, C. Acceleration via fractal learning rate schedules. In International Conference on Machine Learning, pp. 87–99. PMLR, 2021.
- Beugnot, G., Mairal, J., and Rudi, A. On the benefits of large learning rates for kernel methods. arXiv preprint arXiv:2202.13733, 2022.
- Bjorck, N., Gomes, C. P., Selman, B., and Weinberger, K. Q. Understanding batch normalization. Advances in neural information processing systems, 31, 2018.
- Blanc, G., Gupta, N., Valiant, G., and Valiant, P. Implicit regularization for deep neural networks driven by an ornstein-uhlenbeck like process. In Conference on learning theory, pp. 483–513. PMLR, 2020.
- Boursier, E., Pillaud-Vivien, L., and Flammarion, N. Gradient flow dynamics of shallow relu networks for square loss and orthogonal inputs. arXiv preprint arXiv:2206.00939, 2022.
- Chen, L. and Bruna, J. On gradient descent convergence beyond the edge of stability. arXiv preprint arXiv:2206.04172, 2022.
- Chizat, L. and Bach, F. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In Conference on Learning Theory, pp. 1305–1338. PMLR, 2020.
- Chizat, L., Oyallon, E., and Bach, F. On lazy training in differentiable programming. Advances in Neural Information Processing Systems, 32, 2019.
- Cohen, J. M., Kaur, S., Li, Y., Kolter, J. Z., and Talwalkar, A. Gradient descent on neural networks typically occurs at the edge of stability. arXiv preprint arXiv:2103.00065, 2021.
- Damian, A., Ma, T., and Lee, J. D. Label noise sgd provably prefers flat global minimizers. Advances in Neural Information Processing Systems, 34:27449–27461, 2021.
- Denton, E. L., Zaremba, W., Bruna, J., LeCun, Y., and Fergus, R. Exploiting linear structure within convolutional networks for efficient evaluation. Advances in neural information processing systems, 27, 2014.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805, 2018.
- Dinh, L., Pascanu, R., Bengio, S., and Bengio, Y. Sharp minima can generalize for deep nets. In International Conference on Machine Learning, pp. 1019–1028. PMLR, 2017.
- Foret, P., Kleiner, A., Mobahi, H., and Neyshabur, B. Sharpness-aware minimization for efficiently improving generalization. In International Conference on Learning Representations, 2021.
- Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. arXiv preprint arXiv:1803.03635, 2018.
- Geiping, J., Goldblum, M., Pope, P. E., Moeller, M., and Goldstein, T. Stochastic training is not necessary for generalization. In International Conference on Learning Representations, 2022.
- HaoChen, J. Z., Wei, C., Lee, J. D., and Ma, T. Shape matters: Understanding the implicit bias of the noise covariance. In Conference on Learning Theory, 2021.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In CVPR, 2016.
- Hinton, G., Vinyals, O., Dean, J., et al. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531, 2(7), 2015.
- Hoefler, T., Alistarh, D., Ben-Nun, T., Dryden, N., and Peste, A. Sparsity in deep learning: Pruning and growth for efficient inference and training in neural networks. JMLR, 22(241):1–124, 2021.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. Advances in neural information processing systems, 31, 2018.
- Jacot, A., Ged, F., Şimşek, B., Hongler, C., and Gabriel, F. Saddle-to-saddle dynamics in deep linear networks: Small initialization training, symmetry, and sparsity. In International Conference on Machine Learning, 2021.
- Jastrzebski, S., Arpit, D., Astrand, O., Kerg, G. B., Wang, H., Xiong, C., Socher, R., Cho, K., and Geras, K. J. Catastrophic fisher explosion: Early phase fisher matrix impacts generalization. In Proceedings of the 38th International Conference on Machine Learning. PMLR, 2021.
- Kaur, S., Cohen, J., and Lipton, Z. C. On the maximum Hessian eigenvalue and generalization. arXiv preprint arXiv:2206.10654, 2022.
- Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., and Tang, P. T. P. On large-batch training for deep

- learning: Generalization gap and sharp minima. [arXiv preprint arXiv:1609.04836](#), 2016.
- Lewkowycz, A., Bahri, Y., Dyer, E., Sohl-Dickstein, J., and Gur-Ari, G. The large learning rate phase of deep learning: the catapult mechanism. [arXiv preprint arXiv:2003.02218](#), 2020.
- Li, Q., Tai, C., and Weinan, E. Stochastic modified equations and dynamics of stochastic gradient algorithms i: Mathematical foundations. *The Journal of Machine Learning Research*, 20(1):1474–1520, 2019a.
- Li, Y., Ma, T., and Zhang, H. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Conference On Learning Theory*, pp. 2–47. PMLR, 2018.
- Li, Y., Wei, C., and Ma, T. Towards explaining the regularization effect of initial large learning rate in training neural networks. In *NeurIPS*, 2019b.
- Li, Z. and Arora, S. An exponential learning rate schedule for deep learning. [arXiv preprint arXiv:1910.07454](#), 2019.
- Li, Z., Lyu, K., and Arora, S. Reconciling modern deep learning with traditional optimization analyses: The intrinsic learning rate. In *Advances in Neural Information Processing Systems*, volume 33, pp. 14544–14555, 2020.
- Li, Z., Wang, T., and Arora, S. What happens after sgd reaches zero loss?—a mathematical framework. In *International Conference on Learning Representations*, 2022.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- Lyu, K. and Li, J. Gradient descent maximizes the margin of homogeneous neural networks. In *International Conference on Learning Representations*, 2020.
- Ma, C. and Ying, L. The Sobolev regularization effect of stochastic gradient descent. [arXiv preprint arXiv:2105.13462](#), 2021.
- Ma, C., Wu, L., and Ying, L. The multiscale structure of neural network loss functions: The effect on optimization and origin. [arXiv preprint arXiv:2204.11326](#), 2022.
- Moroshko, E., Woodworth, B. E., Gunasekar, S., Lee, J. D., Srebro, N., and Soudry, D. Implicit bias in deep linear classification: Initialization scale vs training accuracy. In *Advances in Neural Information Processing Systems*, volume 33, pp. 22182–22193, 2020.
- Mulayoff, R., Michaeli, T., and Soudry, D. The implicit bias of minima stability: A view from function space. *Advances in Neural Information Processing Systems*, 34: 17749–17761, 2021.
- Nacson, M. S., Ravichandran, K., Srebro, N., and Soudry, D. Implicit bias of the step size in linear diagonal neural networks. In *International Conference on Machine Learning*, pp. 16270–16295. PMLR, 2022.
- Nakkiran, P. Learning rate annealing can provably help generalization, even for convex problems. [arXiv preprint arXiv:2005.07360](#), 2020.
- Oksendal, B. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- Pesme, S., Pillaud-Vivien, L., and Flammarion, N. Implicit bias of sgd for diagonal linear networks: a provable benefit of stochasticity. In *Advances in Neural Information Processing Systems*, 2021.
- Pillaud-Vivien, L., Reygner, J., and Flammarion, N. Label noise (stochastic) gradient descent implicitly solves the lasso for quadratic parametrisation. In *Conference on Learning Theory*, 2022.
- Smith, S. L. and Le, Q. V. A Bayesian perspective on generalization and stochastic gradient descent. In *International Conference on Learning Representations*, 2018.
- Smith, S. L., Dherin, B., Barrett, D. G., and De, S. On the origin of implicit regularization in stochastic gradient descent. [arXiv preprint arXiv:2101.12176](#), 2021.
- Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- Vaskevicius, T., Kanade, V., and Rebeschini, P. Implicit regularization for optimal sparse recovery. *Advances in Neural Information Processing Systems*, 32, 2019.
- Wang, Y., Chen, M., Zhao, T., and Tao, M. Large learning rate tames homogeneity: Convergence and balancing effect. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=3tbDrs77LJ5>.
- Wojtowysch, S. Stochastic gradient descent with noise of machine learning type. Part II: Continuous time analysis. [arXiv preprint arXiv:2106.02588](#), 2021a.
- Wojtowysch, S. Stochastic gradient descent with noise of machine learning type. Part I: Discrete time analysis. [arXiv preprint arXiv:2105.01650](#), 2021b.

- Woodworth, B., Gunasekar, S., Lee, J. D., Moroshko, E., Savarese, P., Golan, I., Soudry, D., and Srebro, N. Kernel and rich regimes in overparametrized models. In Proceedings of Thirty Third Conference on Learning Theory, volume 125 of Proceedings of Machine Learning Research, pp. 3635–3673. PMLR, 09–12 Jul 2020.
- Wu, J., Zou, D., Braverman, V., and Gu, Q. Direction matters: On the implicit bias of stochastic gradient descent with moderate learning rate. International Conference on Learning Representations, 2021.
- Wu, L., Ma, C., et al. How sgd selects the global minima in over-parameterized learning: A dynamical stability perspective. Advances in Neural Information Processing Systems, 31, 2018.
- Xing, C., Arpit, D., Tsirigotis, C., and Bengio, Y. A walk with sgd. arXiv preprint arXiv:1802.08770, 2018.
- Yang, N., Tang, C., and Tu, Y. Stochastic gradient descent introduces an effective landscape-dependent regularization favoring flat solutions. arXiv preprint arXiv:2206.01246, 2022.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. In International Conference on Learning Representations, 2017.
- Zhang, G., Wang, C., Xu, B., and Grosse, R. Three mechanisms of weight decay regularization. In International Conference on Learning Representations, 2018.
- Ziyin, L., Liu, K., Mori, T., and Ueda, M. Strength of minibatch noise in sgd. In International Conference on Learning Representations, 2022.

## APPENDIX

In Section A, we show Proposition 2.1 on the equivalence between SGD and GD with added noise. In Section B, we provide the proof that loss stabilization occurs as written in Proposition 2.2. In Section C, we show experimentally that the proposed SDE model matches well the SDE dynamics. Finally, we present additional experiments in Section D.

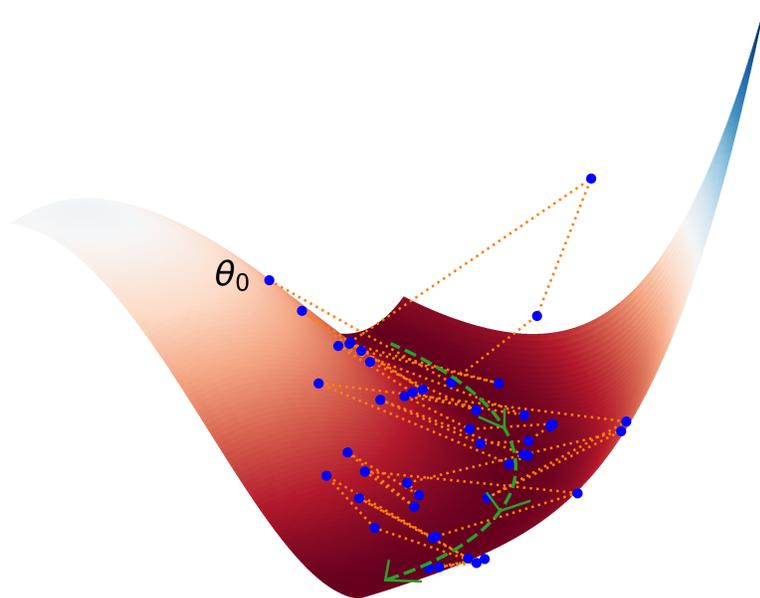


Figure 8. Three-dimensional visualisation of the SGD dynamics in a non-convex loss landscape. The SGD dynamics (blue points) is bouncing side-to-side to the bottom of the valley (the dotted green line). A slow movement occurs pushing the iterates in the direction given by the green arrows.

To begin this appendix, we provide in Figure 8 a toy visualization in which we showcase a typical SGD dynamics when loss stabilization occurs. We run SGD on the diagonal linear network with one sample in two dimensions ( $n = 1, d = 2$ ) adding label noise of the shape given by equation Eq.(9), with balanced layers  $u = v$ . The blue points corresponds to iterates of the dynamics (that are linked with the orange dotted lines). The green line corresponds to the global minimum of the loss, what can be called the “bottom of the valley”. This hopefully will serve the reader for a visual intuition on (i) the bouncing dynamics side-to-side to the bottom of the valley (in green), and (ii) the slow stochastic movement (in the direction of the green arrows).

### A. SGD and Label Noise GD

For the sake of clarity we recall below the statement of the Proposition 2.1 which we prove in this section.

*Proposition 2.1.* Let  $(\theta_t)_{t \geq 0}$  follow the SGD dynamics Eq.(2) with sampling function  $(i_t)_{t \geq 0}$ . Let  $\mathbf{1}_{i=i_t}$  be indicator function, define for  $t \geq 0$ , the random vector  $\xi_t \in \mathbb{R}^n$  such that for all  $i \in \llbracket 1, n \rrbracket$ ,

$$[\xi_t]_i := (h_{\theta_t}(x_i) - y_i)(1 - n\mathbf{1}_{i=i_t}). \quad (10)$$

Then  $(\theta_t)_{t \geq 0}$  follows the full-batch gradient dynamics on  $\mathcal{L}$  with label noise  $(\xi_t)_{t \geq 0}$ , that is

$$\theta_{t+1} = \theta_t - \frac{\eta}{n} \sum_{i=1}^n (h_{\theta_t}(x_i) - y_i^t) \nabla_{\theta} h_{\theta_t}(x_i), \quad (11)$$

where we define the random labels  $y^t := y + \xi_t$ . Furthermore,  $\xi_t$  is a mean zero random vector with variance such that  $\frac{1}{n(n-1)} \mathbb{E} \|\xi_t\|^2 = 2\mathcal{L}(\theta_t)$ .

*Proof.* Note that

$$\sum_{i=1}^n (h_{\theta_t}(x_i) - y_i^t) \nabla_{\theta} h_{\theta_t}(x_i) = \sum_{i=1}^n (h_{\theta_t}(x_i) - y_i - [\xi_t]_i) \nabla_{\theta} h_{\theta_t}(x_i). \quad (12)$$

Using  $[\xi_t]_i := (h_{\theta_t}(x_i) - y_i)(1 - n\mathbf{1}_{i=i_t})$ ,

$$= \frac{1}{n} \sum_{i=1}^n (h_{\theta_t}(x_i) - y_i - (h_{\theta_t}(x_i) - y_i)(1 - n\mathbf{1}_{i=i_t})) \nabla_{\theta} h_{\theta_t}(x_i), \quad (13)$$

$$= \sum_{i=1}^n \mathbf{1}_{i=i_t} (h_{\theta_t}(x_i) - y_i) \nabla_{\theta} h_{\theta_t}(x_i) = (h_{\theta_t}(x_{i_t}) - y_{i_t}) \nabla_{\theta} h_{\theta_t}(x_{i_t}). \quad (14)$$

which is exactly the stochastic gradient wrt to sample  $(x_{i_t}, y_{i_t})$ .

Now we prove the latter part of the proposition regarding the scale of the noise. Recall that, for all  $i \leq n$ , we have  $[\xi_t]_i = (h_{\theta_t}(x_i) - y_i)(1 - n\mathbf{1}_{i=i_t})$ , where  $i_t \sim \mathcal{U}(\llbracket 1, n \rrbracket)$ . Now taking the expectation,

$$\mathbb{E}[\xi_t]_i = \mathbb{E}[(h_{\theta_t}(x_i) - y_i)(1 - n\mathbf{1}_{i=i_t})] = (h_{\theta_t}(x_i) - y_i)(1 - n\mathbb{E}[\mathbf{1}_{i=i_t}]) = 0, \quad (15)$$

as  $\mathbb{E}[\mathbf{1}_{i=i_t}] = 1/n$ . Coming to the variance,

$$\mathbb{E} \|\xi_t\|^2 = \mathbb{E} \left[ \sum_{i=1}^n [\xi_t]_i^2 \right] = \sum_{i=1}^n \mathbb{E} [\xi_t]_i^2 \quad (16)$$

$$= \sum_{i=1}^n (h_{\theta_t}(x_i) - y_i)^2 \mathbb{E} [(1 - n\mathbf{1}_{i=i_t})^2] \quad (17)$$

$$= \sum_{i=1}^n (h_{\theta_t}(x_i) - y_i)^2 \mathbb{E} [(1 - 2n\mathbf{1}_{i=i_t} + n^2\mathbf{1}_{i=i_t})] \quad (18)$$

$$= \sum_{i=1}^n (h_{\theta_t}(x_i) - y_i)^2 (1 - 2 + n) \quad (19)$$

$$= (n-1) \sum_{i=1}^n (h_{\theta_t}(x_i) - y_i)^2 = 2n(n-1)\mathcal{L}(\theta_t), \quad (20)$$

and this concludes the proof of the proposition.  $\square$

## B. Quadratic Parameterization in One Dimension

Again, for the Appendix to be self-contained, we recall the setup of the Proposition 2.2 on loss stabilization. We consider a regression problem with quadratic parameterization on one-dimensional data inputs  $x_i$ 's, coming from a distribution  $\hat{\rho}$ , and outputs generated by the linear model  $y_i = x_i \theta_*^2$ . The loss writes  $F(\theta) := \frac{1}{4} \mathbb{E}_{\hat{\rho}} (y - x\theta^2)^2$ , and the SGD iterates with step size  $\eta > 0$  follow, for any  $t \in \mathbb{N}$ ,

$$\theta_{t+1} = \theta_t + \eta \theta_t x_{i_t} (y_{i_t} - x_{i_t} \theta_t^2) \quad \text{where } x_{i_t} \sim \hat{\rho}. \quad (21)$$

We rewrite the proposition here.

**Proposition B.1. (Extended version of Proposition 2.2)** Assume  $\exists x_{\min}, x_{\max} > 0$  such that  $\text{supp}(\hat{\rho}) \subset [x_{\min}, x_{\max}]$ . Then for any  $\eta \in ((\theta_* x_{\min})^{-2}, 1.25(\theta_* x_{\max})^{-2})$ , any initialization in  $\theta_0 \in (0, \theta_*)$ , for  $t \in \mathbb{N}$ , we have almost surely

$$F(\theta_t) \in (\epsilon_o^2 \theta_*^2, 0.17 \theta_*^2). \quad (22)$$

where  $\epsilon_o = \min\{(\eta(\theta_* x_{\min})^2 - 1)/3, 0.02\}$ . Also, almost surely, there exists  $t, k > 0$  such that  $\theta_{t+2k} \in (0.65 \theta_*, (1 - \epsilon_o) \theta_*)$  and  $\theta_{t+2k+1} \in ((1 + \epsilon_o) \theta_*, 1.162 \theta_*)$ .

*Proof.* Consider SGD recursion Eq.(21) and note that  $y = x\theta_*^2$ .

$$\theta_{t+1} = \theta_t + \eta \theta_t x(x\theta_*^2 - x\theta_t^2) \quad (23)$$

$$\theta_{t+1} = \theta_t + \eta \theta_t x^2 (\theta_*^2 - \theta_t^2) \quad (24)$$

For the clarity of exposition, we consider the rescaled recursion of the original SGD recursion.

$$\theta_{t+1}/\theta_* = \theta_t/\theta_* + \eta \theta_*^2 x^2 \theta_t/\theta_* \left(1 - (\theta_t/\theta_*)^2\right), \quad (25)$$

and, by making the benign change  $\theta_t \leftarrow \theta_t/\theta_*$ , we focus on the stochastic recursion instead,

$$\theta_{t+1} = \theta_t + \gamma \theta_t (1 - \theta_t^2), \quad (26)$$

where  $\gamma \sim \hat{\rho}_\gamma$  the pushforward of  $\hat{\rho}$  under the application  $z \rightarrow \eta \theta_*^2 z^2$ . Let  $\Gamma := \text{supp}(\hat{\rho}_\gamma)$ , the support of the distribution of  $\gamma$ . From the range of  $\eta$ , it can be verified that  $\Gamma \subseteq (1, 1.25)$ . Now the proof of the theorem follows from Lemma B.3.  $\square$

**Lemma B.2** (Bounded Region). *Consider the recursion Eq.(26), for  $\Gamma \subseteq (1, 1.25)$  and  $0 < \theta_0 < 1$ , then for all  $t > 0$ ,  $\theta_t \in (0, 1.162)$ .*

*Proof.* Consider a single step of Eq.(26), for some  $\gamma \in (1, 1.25)$ ,

$$\theta_+ = \theta + \gamma \theta (1 - \theta^2)$$

The aim is to show that  $\theta_+$  stays in the interval  $(0, 1.162)$ . In order to show this, we do a casewise analysis.

For  $\theta \in (0, 1]$ : Since  $0 < \theta \leq 1$ , we have  $\theta_+ \geq \theta > 0$ . To prove the bound above, consider the following quantity,

$$\theta_{max} = \max_{\gamma \in (1, 1.25)} \max_{\theta \in (0, 1]} \theta + \gamma \theta (1 - \theta^2) \quad (27)$$

Say  $h_\gamma(\theta) = \theta + \gamma \theta (1 - \theta^2)$ , note that  $h'_\gamma(\theta) = 1 + \gamma - 3\gamma\theta^2$  and  $h''_\gamma(\theta) = -6\gamma\theta < 0$ . Hence, for any  $\gamma$  in our domain, the maximum is attained at  $\theta_\gamma = \frac{1}{\sqrt{3}} \sqrt{\frac{1}{\gamma} + 1}$  and  $h_\gamma(\theta_\gamma) = \frac{2(1+\gamma)^{3/2}}{3\sqrt{3}\gamma}$ .

$$\max_{\gamma \in (1, 1.25)} \max_{\theta \in (0, 1]} \theta + \gamma \theta (1 - \theta^2) = \max_{\gamma \in (.5, 1.25)} \frac{2(1+\gamma)^{3/2}}{3\sqrt{3}\gamma} \quad (28)$$

It can be verified that  $\frac{2(1+\gamma)^{3/2}}{3\sqrt{3}\gamma}$  is increasing with gamma in the interval  $(1, 1.25)$ . Hence,

$$\max_{\gamma \in (1, 1.25)} \frac{2(1+\gamma)^{3/2}}{3\sqrt{3}\gamma} \leq \left. \frac{2(1+\gamma)^{3/2}}{3\sqrt{3}\gamma} \right|_{\gamma=1.25} < 1.162 \quad (29)$$

Combining them, we get,

$$\theta_+ \leq \max_{\gamma \in (0, 1.25)} \max_{\theta \in (0, 1]} \theta + \gamma \theta (1 - \theta^2) < 1.162 \quad (30)$$

For  $\theta \in (1, 1.162)$ : Since  $\theta > 1$ , we have,  $\theta_+ < \theta < 1.162$ . For lower bound, note that for  $\theta_+$  to be less than 0, we need  $1 + \gamma - \gamma\theta^2 < 0$ . But for  $\gamma \in (1, 1.25)$  and  $\theta \in (1, 1.162)$ ,

$$\gamma(\theta^2 - 1) < 1.25((1.162)^2 - 1) < 1. \quad (31)$$

Hence, it never goes below 0.  $\square$

**Lemma B.3.** *Consider the recursion Eq.(26) with  $\Gamma \subseteq (1, 1.25)$  and  $\theta_0$  initialized uniformly in  $(0, 1)$ . Then, there exists  $\epsilon_0 > 0$ , such that for all  $\epsilon < \epsilon_0$  there exists  $t > 0$  such that for any  $k > 0$ ,*

$$\theta_{t+2k} \in (0.65, 1 - \epsilon) \quad \text{and} \quad \theta_{t+2k+1} \in (1 + \epsilon, 1.162) \quad (32)$$

*almost surely.*

*Proof.* Define  $\gamma_{\min} > 1$  as the infimum of the support  $\Gamma$ . Let  $\epsilon_o = \min\{(\gamma_{\min}-1)/3, 0.02\}$ . Note that  $\epsilon_o > 0$  as  $\gamma_{\min} > 1$ . Now for any  $0 < \epsilon < \epsilon_o$ , we have  $\gamma_{\min}(2 - \epsilon)(1 - \epsilon) > 2$ .

Divide the interval  $(0, 1.162)$  into 4 regions,  $I_0 = (0, 0.65]$ ,  $I_1 = (0.65, 1 - \epsilon)$ ,  $I_2 = [1 - \epsilon, 1)$ ,  $I_3 = (1, 1.162)$ . The strategy of the proof is that the iterates will eventually end up in  $I_1$  and that once it ends up in  $I_1$ , it comes back to  $I_1$  in 2 steps.

Let  $\theta_0$  be initialized uniformly random in  $(0, 1)$ . Consider the sequence  $(\theta_t)_{t \geq 0}$  generated by

$$\theta_{t+1} = h_{\gamma_t}(\theta_t) := \theta_t + \gamma_t \theta_t (1 - \theta_t^2) \quad \text{where} \quad \gamma_t \sim \hat{\rho}_{\gamma}. \quad (33)$$

We prove the following facts **(P1)**-**(P4)**:

**(P1)** There exists  $t \geq 0$  such that the  $\theta_t \in I_1 \cup I_2 \cup I_3$ .

**(P2)** Let  $\theta_t \in I_3$ , then  $\theta_{t+1} \in I_1 \cup I_2$ .

**(P3)** Let  $\theta_t \in I_2$ , there exists  $k > 0$  such that for  $k' < k$ ,  $\theta_{t+2k'} \in I_2$  and  $\theta_{t+2k} \in I_1$ .

**(P4)** When  $\theta_t \in I_1$ , then for all  $k \geq 0$ ,  $\theta_{t+2k} \in I_1$  and  $\theta_{t+2k+1} \in (1 + \epsilon, 1.162)$ .

**Proof of (P1)-(P4):** Let  $t \in \mathbb{N}$ , note first that the event  $\{\theta_t = 1\} = \cup_{k \leq t} \{\theta_k = 1 | \theta_{k-1} \neq 1\}$  and hence a finite union of zero measure sets. Hence  $\{\theta_t = 1\}$  is a zero measure set and therefore we do not consider it below. For any other sequence, from the above four properties, we can conclude that the lemma holds.

**Proof of P1:** Assume that until time  $t > 0$ , the iterates are all in  $I_0$ , then we have

$$\theta_t = \theta_{t-1}(1 + \gamma(1 - \theta_{t-1}^2)) \geq \theta_{t-1}(2 - \theta_{t-1}^2) > 1.5 \theta_{t-1} > 1.5^t \theta_0 \quad (34)$$

Hence, the sequence eventually exits  $I_0$ . We know that it will stay bounded from Lemma B.2, hence it will end up in  $I_1 \cup I_2 \cup I_3$ .

**Proof of P2:** For any  $\theta_t \in (1, 1.162)$ ,  $1 < \gamma < 1.25$ , since  $h_{\gamma}(\cdot)$  is decreasing in  $(1, 1.162)$ , we have  $h_{\gamma}(1.162) < h_{\gamma}(\theta_t) < h_{\gamma}(1)$ . Also  $h_{\gamma}(\theta)$  is linear in gamma with negative coefficient for  $\theta > 1$ . Hence it decreases as  $\gamma$  increases. Using this,

$$.652 = h_{1.25}(1.162) < h_{\gamma}(1.162) < h_{\gamma}(\theta_t) < h_{\gamma}(1) = 1. \quad (35)$$

Hence,  $\theta_{t+1} \in I_1 \cup I_2$ .

**Proof of P3:** The proof of this follows from Lemma B.5.

**Proof of P4:** The proof of this follows from Lemma B.8. □

**Lemma B.4.** For any  $\theta \in I_1 \cup I_2$  and any  $a, b \in \Gamma$ ,  $h_a(h_b(\theta)) \in I_1 \cup I_2$ ,

$$h_{\gamma_{\max}}(h_{\gamma_{\max}}(\theta)) \leq h_a(h_b(\theta)) \leq h_{\gamma_{\min}}(h_{\gamma_{\min}}(\theta)). \quad (36)$$

*Proof.* For any  $\gamma \in \Gamma$ , recall

$$h_{\gamma}(\theta) = \theta + \gamma \theta (1 - \theta^2) = 1 + (1 - \theta)(\gamma \theta (1 + \theta) - 1). \quad (37)$$

Note that for  $\theta \in I_1 \cup I_2$ ,  $\theta(1 + \theta) > 1$ , Hence  $\gamma \theta (1 + \theta) > 1$ . This gives us that  $h_{\gamma}(\theta) > 1$ . Now we will track where  $\theta \in I_1 \cup I_2$  can end up after two stochastic gradient steps.

- For any  $b \in \Gamma$ , as  $\theta \in I_1 \cup I_2$ , we have

$$h_{\gamma_{\max}}(\theta) \geq h_b(\theta) \geq h_{\gamma_{\min}}(\theta) > 1,$$

note  $h_{\gamma_{\max}}(\theta) \geq h_b(\theta) \geq h_{\gamma_{\min}}(\theta)$  holds since  $\theta < 1$ .

- Now for any  $a \in \Gamma$  and  $x > 1$ ,  $h_a(x)$  is a decreasing function in  $x$ . Hence

$$h_a(h_{\gamma_{\max}}(\theta)) \leq h_a(h_b(\theta)) \leq h_a(h_{\gamma_{\min}}(\theta)).$$

Using  $\gamma_{\min} \leq a$ ,  $h_a(h_{\gamma_{\min}}(\theta)) \leq h_{\gamma_{\min}}(h_{\gamma_{\min}}(\theta))$ , Similarly using  $\gamma_{\max} > a$ , we have,  $h_{\gamma_{\max}}(h_{\gamma_{\max}}(\theta)) \leq h_a(h_{\gamma_{\max}}(\theta))$ . Combining them we get,

$$h_{\gamma_{\max}}(h_{\gamma_{\max}}(\theta)) \leq h_a(h_b(\theta)) \leq h_{\gamma_{\min}}(h_{\gamma_{\min}}(\theta)). \quad (38)$$

Similar argument can extend it to,

$$h_{1.25}(h_{1.25}(\theta)) < h_a(h_b(\theta)) < h_1(h_1(\theta)). \quad (39)$$

□

**Lemma B.5.** *Let  $\theta_t \in I_2$ , there exists  $k > 0$  such that  $\theta_{t+2k} \in I_1$ .*

*Proof.* For any  $\gamma \in \Gamma$ , let  $\theta_+ = h_\gamma(\theta)$ , then we have

$$h_\gamma(h_\gamma(\theta)) - \theta = h_\gamma(\theta_+) - \theta = \gamma\theta(1 - \theta^2) + \gamma\theta_+(1 - \theta_+^2). \quad (40)$$

Furthermore,

$$\theta_+ = \theta + \gamma\theta(1 - \theta^2) = \theta(1 + \gamma(1 - \theta^2)), \quad (41)$$

$$1 + \theta_+ = 1 + \theta + \gamma\theta(1 - \theta^2) = (1 + \theta)(1 + \gamma\theta(1 - \theta)), \quad (42)$$

$$1 - \theta_+ = 1 - \theta - \gamma\theta(1 - \theta^2) = (1 - \theta)(1 - \gamma\theta(1 + \theta)). \quad (43)$$

And multiplying the above three terms and adding  $\theta(1 - \theta^2)$ , we get,

$$\theta_+(1 - \theta_+^2) + \theta(1 - \theta^2) = \theta(1 - \theta^2) \underbrace{\{1 + [(1 + \gamma(1 - \theta^2))(1 + \gamma\theta(1 - \theta))(1 - \gamma\theta(1 + \theta))]\}}_{P(\theta)} \quad (44)$$

For  $\theta \in I_2$ , using  $\gamma_{\min}(2 - \epsilon)(1 - \epsilon) > 2$ , we have the inequalities

$$(1 + \gamma(1 - \theta^2))(1 + \gamma\theta(1 - \theta)) > 1, \quad (45)$$

$$(1 - \gamma\theta(1 + \theta)) < 1 - \gamma_{\min}(2 - \epsilon)(1 - \epsilon) < -1, \quad (46)$$

$$P(\theta) < -1. \quad (47)$$

Hence,

$$h_\gamma(h_\gamma(\theta)) - \theta = \gamma(1 - \theta^2)(1 + P(\theta)) < 0. \quad (48)$$

Therefore, for  $[1 - \epsilon, 1)$ , for any  $\gamma \in \Gamma$ ,  $h_\gamma(h_\gamma(\theta)) < \theta$ . Hence for any two stochastic gradient step with  $a, b \in \Gamma$ , from Eq.(36),  $\theta_{t+2} = h_a(h_b(\theta_t)) \leq h_{\gamma_{\min}}(h_{\gamma_{\min}}(\theta_t)) < \theta_t$ . From any point in  $I_2$ , we have  $|\theta_{t+2} - 1| > |\theta_t - 1|$ , for any  $a, b \in \Gamma$ . Intuitively this means that in *two gradient steps* the iterates move *further away from 1* until it eventually leaves the interval  $I_2$  as the sequence  $\{\theta_{t+2k}\}_{k \geq 0}$  is strictly decreasing with no limit point in  $I_2$ . From Lemma B.7, we know that in two steps the iterates will never leave  $I_1 \cup I_2$ . Hence they will eventually end up in  $I_1$  leaving  $I_2$ . □

*Property B.6.* Define  $g_\gamma(\theta) := h_\gamma(h_\gamma(\theta))$  for the sake of brevity. The followings properties hold for  $\theta \in I_1 \cup I_2$ ,  $\gamma \in \Gamma$  and  $\theta_\gamma$  the root of  $h'_\gamma(\theta)$ :

**Q1**  $g_\gamma(\theta) \geq g_\gamma(\theta_\gamma)$ .

**Q2** The function  $g_\gamma(\cdot)$  is decreasing in  $[0.65, \theta_\gamma)$  and increasing in  $(\theta_\gamma, 1]$ .

*Proof.* Note  $h'_\gamma(\theta) = 1 + \gamma - \gamma 3\theta^2$  has at most one root  $\theta_\gamma \in (0, 1)$ . Note that for all  $\gamma \in \Gamma$ ,  $\theta_\gamma \in I_1 \cup I_2$ . For any  $\gamma$ ,  $g'_\gamma(\theta) = h'_\gamma(h_\gamma(\theta))h'_\gamma(\theta)$ . For any  $\theta \in I_1 \cup I_2$ , we have,  $h_\gamma(\theta) > 1 \implies h'_\gamma(h_\gamma(\theta)) < 0$ . Therefore,  $g'_\gamma(\theta)$  has only one root in  $I_1 \cup I_2$ . Since  $\theta_\gamma \in I_1 \cup I_2$ , note  $g''_\gamma(\theta_\gamma) = h'_\gamma(h_\gamma(\theta_\gamma))h''_\gamma(\theta_\gamma) > 0$ . Therefore,  $g_\gamma(\cdot)$  attains its minimum at  $\theta_\gamma$  and this shows the desired properties. □

**Lemma B.7.** For any  $\theta \in I_1 \cup I_2$  and any  $a, b \in \Gamma$ ,  $h_a(h_b(\theta)) \in I_1 \cup I_2$ .

*Proof.* **Lower Bound:** From Eq.(39), we know

$$h_{1.25}(h_{1.25}(\theta)) < h_a(h_b(\theta)) \quad (49)$$

We know that from property **Q1** that  $g_\gamma(\theta) \geq g_\gamma(\theta_\gamma)$ . Hence

$$g_{1.25}(\theta_{1.25}) < g_{1.25}(\theta) < h_a(h_b(\theta)) \quad (50)$$

It can be quickly checked that  $.65 < g_{1.25}(\theta_{1.25})$ . Hence the lower bound holds.

**Upper Bound:** From Eq.(39), we know

$$h_a(h_b(\theta)) < h_1(h_1(\theta)) \quad (51)$$

We know that from property **Q2** that  $g_1(\theta) \leq \max\{g_1(1), g_1(0.65)\}$ . It can be easily verified that  $g_1(0.65) < 0.98$ . Hence  $g_1(\theta) < 1$ . □

**Lemma B.8.** For any  $\theta \in I_1$  and any  $a, b \in \Gamma$ ,  $h_a(h_b(\theta)) \in I_1$  and  $h_a(\theta) \in (1 + \epsilon, 1.162)$ .

*Proof.* The lower bound in Lemma B.7 holds here. For the upper bound, from and Eq.(36),

$$h_a(h_b(\theta)) \leq h_{\gamma_{\min}}(h_{\gamma_{\min}}(\theta)). \quad (52)$$

Using property **Q2**,

$$h_{\gamma_{\min}}(h_{\gamma_{\min}}(\theta)) \leq \max\{g_{\gamma_{\min}}(1 - \epsilon), g_{\gamma_{\min}}(0.65)\} \quad (53)$$

From Eq.(48),  $g_{\gamma_{\min}}(1 - \epsilon) < 1 - \epsilon$ . From Eq.(39),  $g_{\gamma_{\min}}(0.65) < g_1(0.65) < 0.98 < 1 - \epsilon$ . In  $I_1$ , the function  $h_a(\cdot)$  first increases reaches maximum and decreases. Hence for  $\theta \in I_1$ ,  $h_a(\theta) \geq \min\{h_a(0.65), h_a(1 - \epsilon)\}$ .

$$h_a(1 - \epsilon) \geq 1 - \epsilon + a(1 - (1 - \epsilon)^2)(1 - \epsilon), \quad (54)$$

$$= 1 - \epsilon + a(2\epsilon - \epsilon^2)(1 - \epsilon), \quad (55)$$

$$\geq 1 - \epsilon + \gamma_{\min}(2\epsilon - \epsilon^2)(1 - \epsilon), \quad (56)$$

$$= 1 + \epsilon + \epsilon(\gamma_{\min}(2 - \epsilon)(1 - \epsilon) - 2) > 1 + \epsilon. \quad (57)$$

Also  $h_a(0.65) > h_1(0.65) > 1.02 > 1 + \epsilon$ , therefore  $h_a(\theta) > 1 + \epsilon$  and this completes the proof. □

## C. Empirical Validation of the SDE Modeling

In this section, we experimentally check the validity of the SDE modeling of SGD in Eq.(8) in terms of the key metrics: training loss, test loss, rank of the Jacobian, and feature sparsity.

**SDE discretization.** Let  $\gamma_t$  be the SDE discretization step size,  $\eta_t$  the step size of the corresponding SGD that we aim to validate,  $\delta_t$  the noise intensity level, and  $Z_t \sim \mathcal{N}(0, I_n)$ . Then we discretize the SDE from Eq.(8) as follows:

$$\theta_{t+1} = \theta_t - \gamma_t \nabla_{\theta} \mathcal{L}(\theta_t) + \sqrt{\gamma_t} \sqrt{\eta_t \delta_t} \phi_{\theta_t}(X)^\top Z_t. \quad (58)$$

To approximate continuous time, we use a small discretization step size  $\gamma_t := \eta_t/10$  and run the discretization for  $10 \times$  longer than the corresponding SGD run. We use  $\eta_t := \eta_{\lfloor t/10 \rfloor}^{SGD}$  and  $\delta_t := c \cdot \mathcal{L}(\theta_{\lfloor t/10 \rfloor}^{SGD})$  where  $c$  is a constant that we select for each setting separately to match the training dynamics of the corresponding SGD run. In addition, we also evaluate a discretization of gradient flow (i.e., Eq.(58) without the noise term) which helps to draw conclusions about the role of the noise term.

**Experimental results.** We present the discretization results in Fig. 9 for all models considered in the paper except deep networks for which computing the Jacobian  $\phi_{\theta_t}$  on each iteration of the SDE discretization is too costly. In all cases, the

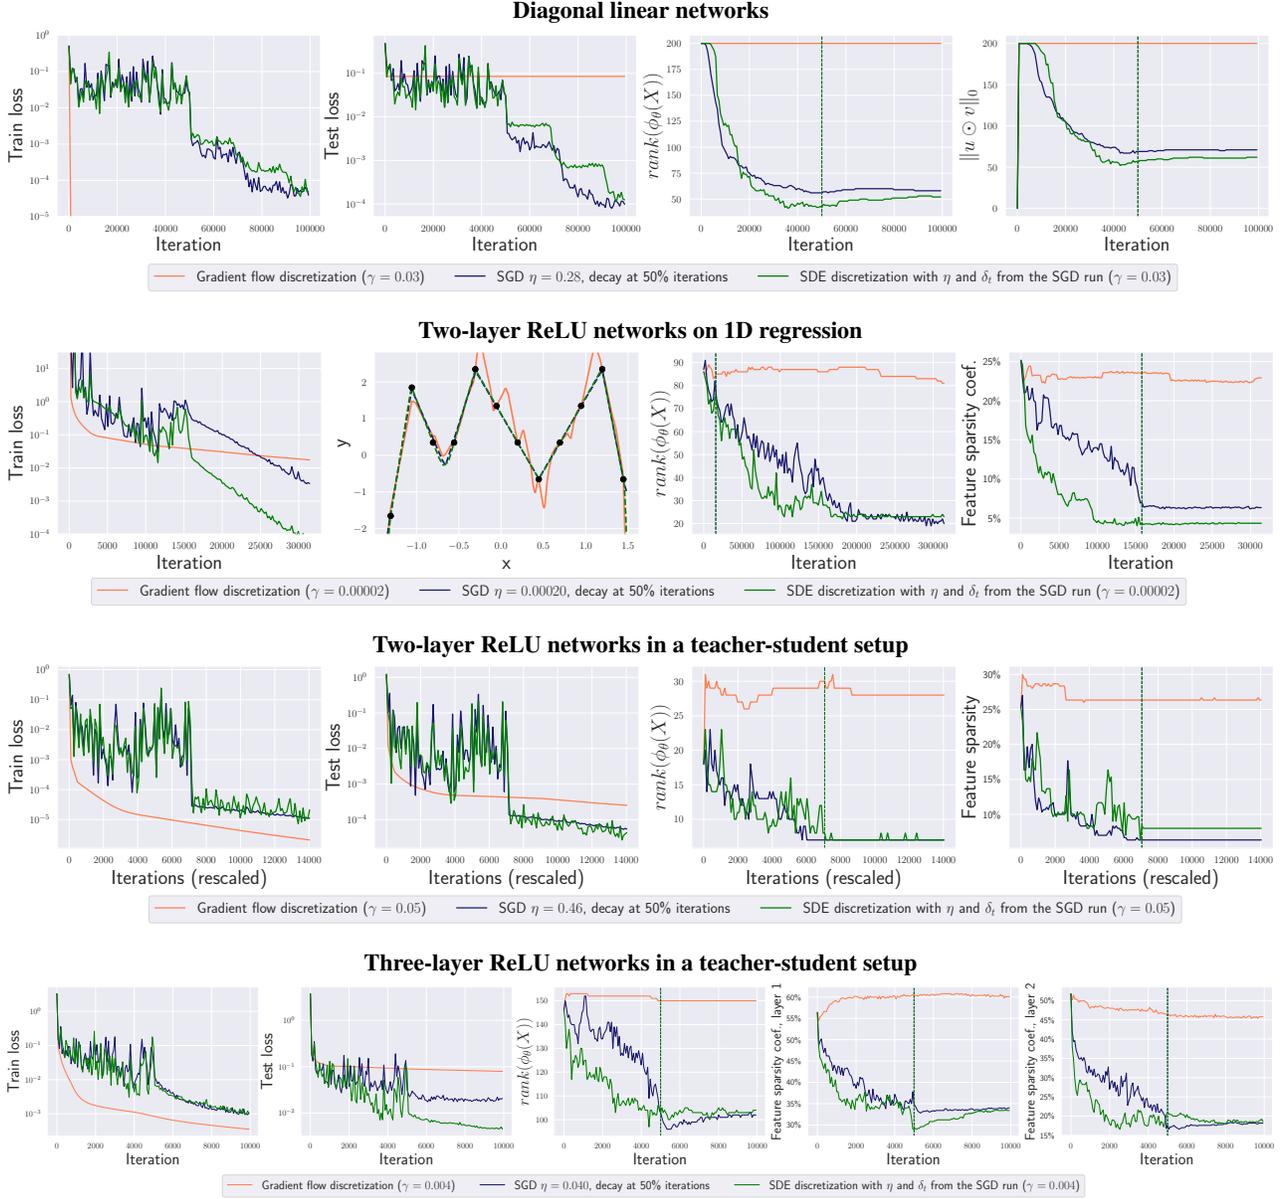


Figure 9. Empirical validation of the SDE modeling. In all cases, the dynamics of the SDE discretization qualitatively matches the dynamics of the corresponding SGD run. Moreover, gradient flow discretization exhibits no rank minimization or feature sparsity which suggests that the presence of the noise plays a key role in learning sparse features.

dynamics of the SDE discretization qualitatively matches the dynamics of the corresponding SGD run. In particular, we observe similar levels of decrease in the rank of the Jacobian and feature sparsity coefficient. We note that the match between SDE and SGD curves is not expected to be precise due to the inherent randomness of the process. Finally, we observe that gradient flow discretization exhibits no rank minimization or feature sparsity which suggests that the presence of the noise (either from the original SGD or its SDE discretization) plays a key role in learning sparse features.

## D. Additional Experimental Results

This section of the appendix presents additional experiments complementing the ones presented in the main text.

**Illustration of neuron dynamics.** We illustrate the change of neurons during training of two-layer ReLU networks (without biases) in the teacher-student setup of Chizat et al. (2019) (see Fig. 1 therein) using a large initialization scale for which small step sizes of GD or SGD lead to lazy training. We illustrate (O1)–(O3) in Fig. 14 and show neuron dynamics in Fig. 10. We see that for SGD with a small step size, the neurons  $w_i$  stay close to their initialization, while for a large step size, there is a clear clustering of directions  $w_i$  along the teacher directions  $w_i^*$ . The overall picture is very similar to Fig. 1 of Chizat et al. (2019) where the same feature learning effect is achieved via gradient flow from a small initialization which is, however, much more computationally expensive due to the saddle point at zero. Finally, we note that the clustering phenomenon of neurons  $w_i$  motivates the removal of highly correlated activations in the feature sparsity coefficient: although the corresponding activations are often non-zero, many of them in fact implement *the same feature* and thus should be counted only once.

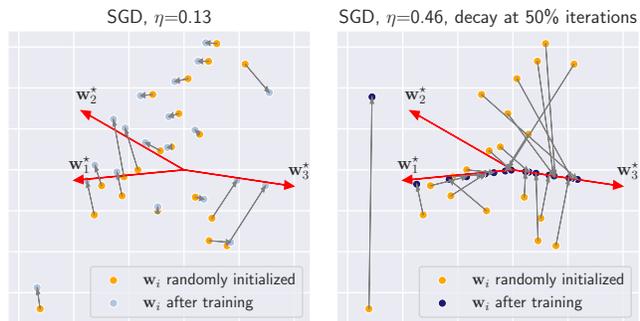


Figure 10. Only for a large step size, the neurons  $w_i$  cluster along the teacher neurons  $w_i^*$  leading to a model that uses a sparse set of features.

**Further results.** We give a short overview of additional figures referred to in the main text. More details can be found in the captions.

- Figure 11 shows that even if loss stabilization occurs in diagonal linear networks, the implicit bias towards sparsity is largely weaker than that of SGD and generalization is poor.
- Figures 12 and 13 demonstrate that the implicit bias resulting from high-loss stabilization makes the neural nets learn *first* a simple model *then eventually* fits the data.
- Figure 14 presents the sparsifying effect corresponding to the neurons’ movements exhibited in Figure 10.
- Figures 15 and 16 exhibit the feature sparsity in ResNet-18 / ResNet-34 architectures on CIFAR-10 and CIFAR-100 in the basic and state-of-the-art settings.
- Figure 17 showcases the features learning induced by large step sizes for different layers of ResNets-18 when trained on CIFAR-10.

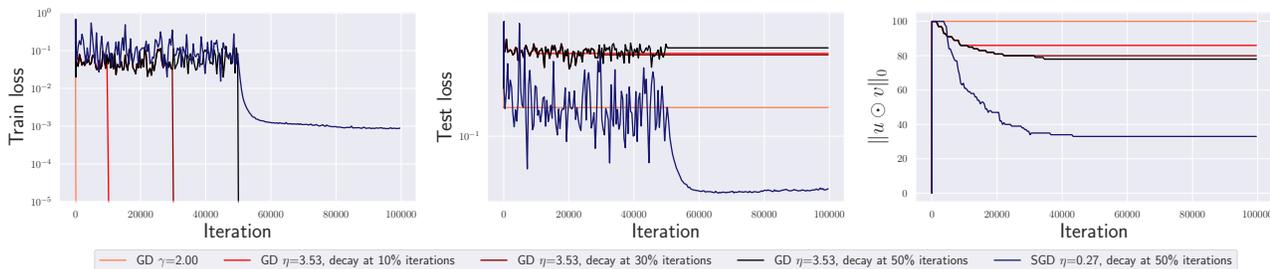


Figure 11. **Diagonal linear networks.** Loss stabilization also occurs for *full-batch gradient descent* but does not lead to a similar level of sparsity as SGD and also does not improve the test loss.

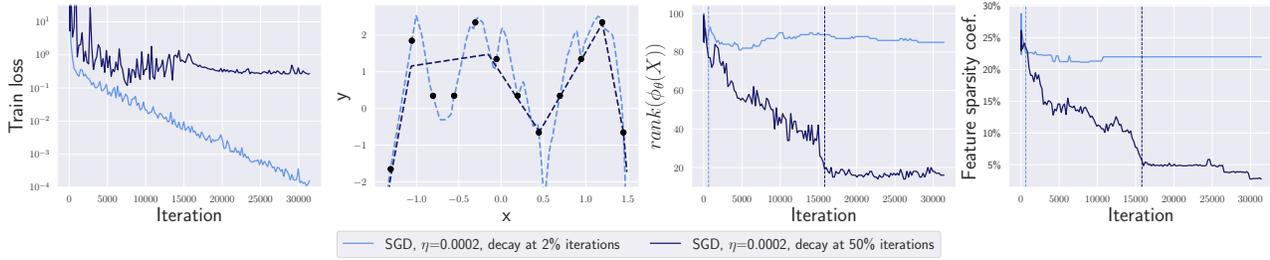


Figure 12. **Two-layer ReLU networks for 1D regression.** Unlike for Fig. 4, here we use a larger warmup coefficient ( $500\times$  vs.  $400\times$ ) which leads to overregularization such that the 50%-schedule run fails to fit all the training points and gets stuck at a too high value of the training loss ( $\approx 10^{-0.5}$ ).

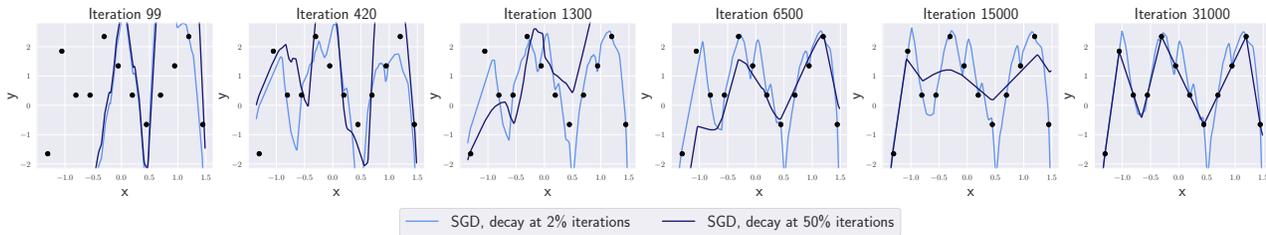


Figure 13. **Two-layer ReLU networks for 1D regression.** Illustration of the resulting models from Fig. 4 over training iterations. We can see that first the model is simplified and only then it fits the training data.

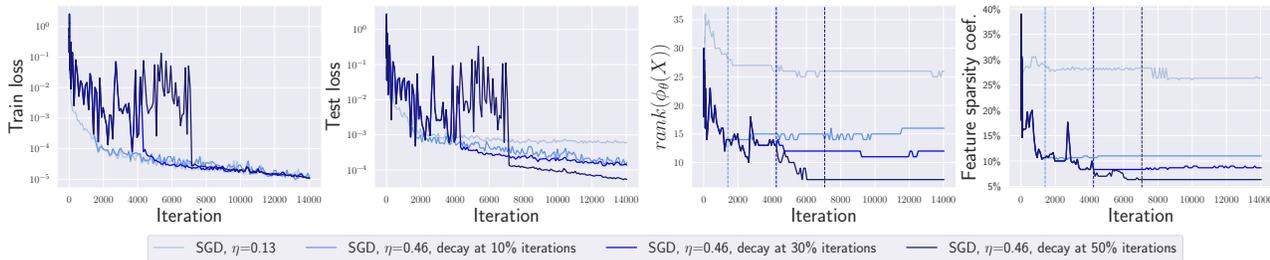


Figure 14. **Two-layer ReLU networks in a teacher-student setup.** Loss stabilization for two-layer ReLU nets in the teacher-student setup with input dimension  $d = 2$ . We observe loss stabilization, better test loss for longer schedules and sparser features due to simplification of  $\phi(X)$ .

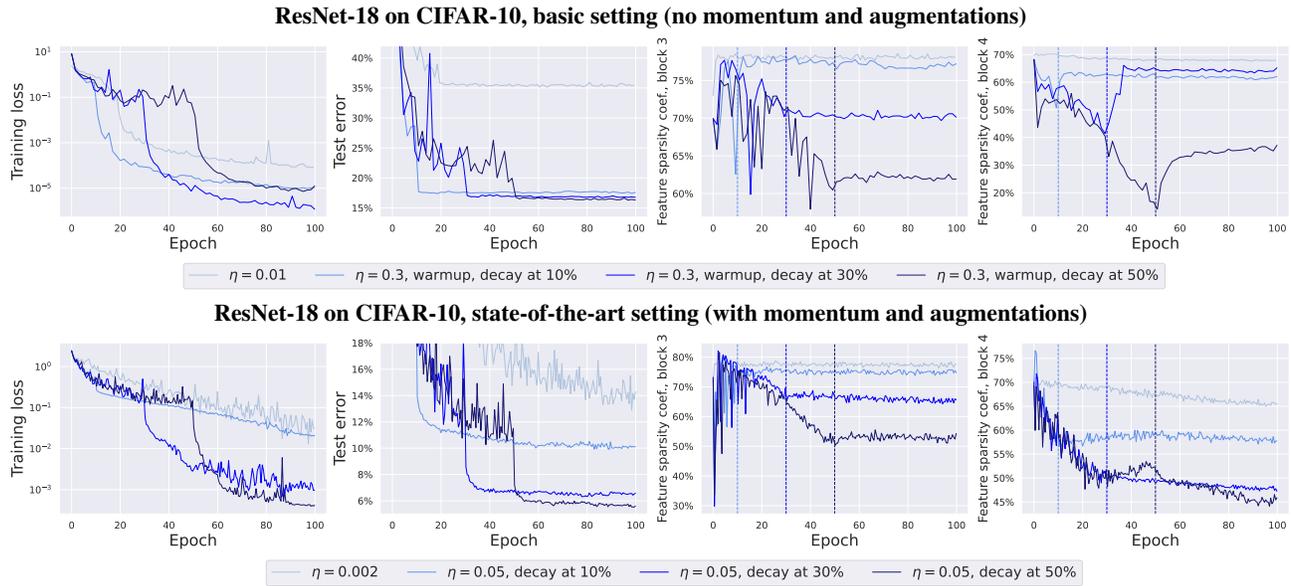


Figure 15. ResNet-18 trained on CIFAR-10. Both in the basic and state-of-the-art settings, the training loss stabilizes, the test loss noticeably depends on the length of the schedule, and the feature sparsity coefficient is minimized over iterations.

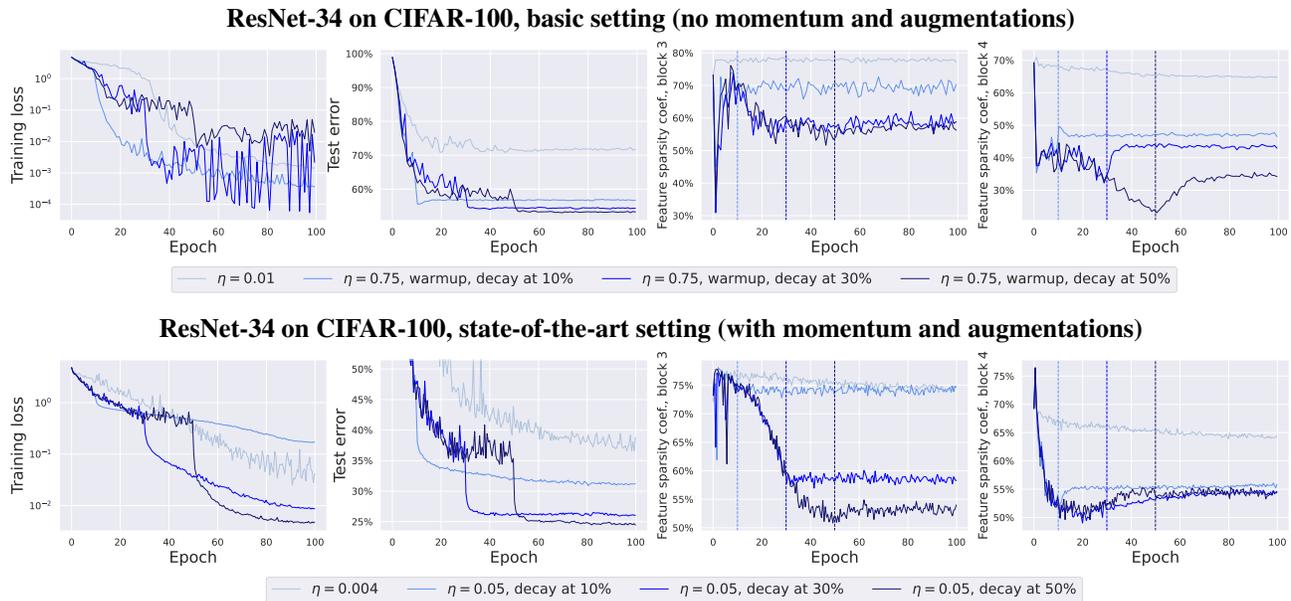


Figure 16. ResNet-34 trained on CIFAR-100. Both in the basic and state-of-the-art settings, the training loss stabilizes, the test loss significantly depends on the length of the schedule, and feature sparsity is minimized over iterations. However, differently from the plots on CIFAR-10, here without explicit regularization we observe oscillating behavior after the step size decay (although at a very low level between  $10^{-4}$  and  $10^{-2}$ ).

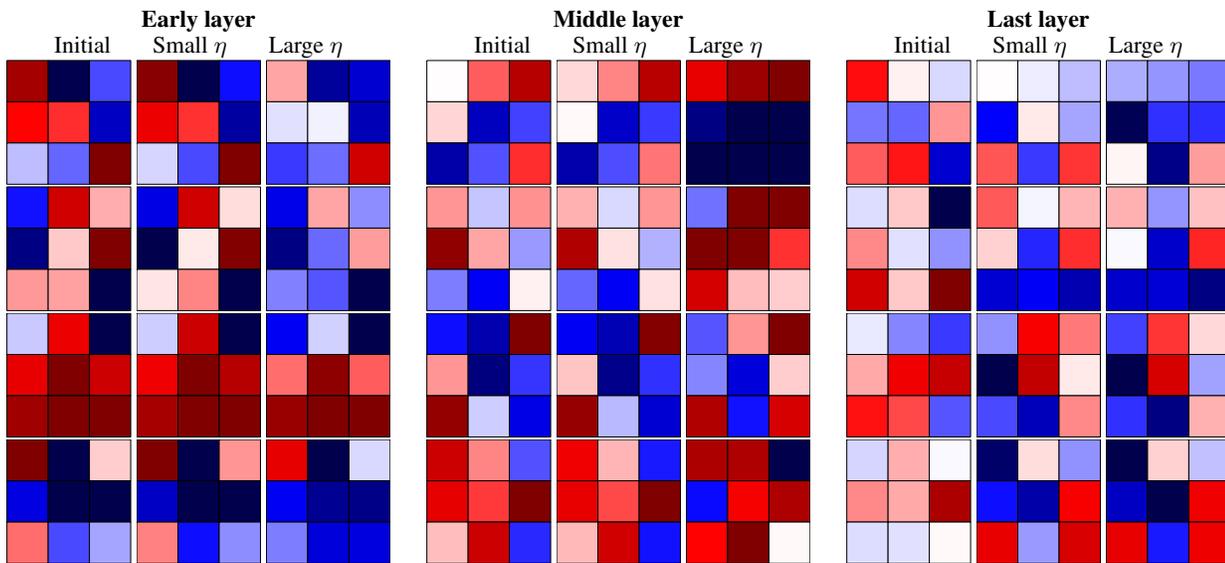


Figure 17. Visualization on four sets of convolutional filters taken from different layers of ResNets-18 trained on CIFAR-10 with small vs. large step size  $\eta$  (the 50% decay schedule). For small step sizes, the early and middle layers stay very close to randomly initialized ones which indicates the absence of feature learning.