Neural Thermodynamics: Entropic Forces in Deep and Universal Representation Learning

Liu Ziyin^{1,3,*}, Yizhou Xu^{2,*}, Isaac Chuang¹

¹Massachusetts Institute of Technology ²École Polytechnique Fédérale de Lausanne ³NTT Research

Abstract

With the rapid discovery of emergent phenomena in deep learning and large language models, understanding their cause has become an urgent need. Here, we propose a rigorous entropic-force theory for understanding the learning dynamics of neural networks trained with stochastic gradient descent (SGD) and its variants. Building on the theory of parameter symmetries and an entropic loss landscape, we show that representation learning is crucially governed by emergent entropic forces arising from stochasticity and discrete-time updates. These forces systematically break continuous parameter symmetries and preserve discrete ones, leading to a series of gradient balance phenomena that resemble the equipartition property of thermal systems. These phenomena, in turn, (a) explain the universal alignment of neural representations between AI models and lead to a proof of the Platonic Representation Hypothesis, and (b) reconcile the seemingly contradictory observations of sharpness- and flatness-seeking behavior of deep learning optimization. Our theory and experiments demonstrate that a combination of entropic forces and symmetry breaking is key to understanding emergent phenomena in deep learning.

1 Introduction

Modern neural networks trained with stochastic gradient descent (SGD) exhibit a complex plethora of emergent behaviors – emergence of capabilities [1, 2, 3], progressive sharpening and flattening [4, 5], phase-transition like behaviors [6, 7], and universal representational alignment across models [8, 9, 10, 11] – that are difficult to explain through loss minimization alone. These behaviors mirror phenomena found in physical systems at finite temperature, suggesting that deep learning dynamics are shaped not just by explicit optimization but also by implicit forces arising from stochasticity and discrete updates. These implicit forces have long been associated with the phenomenon of "implicit bias" in deep learning [12, 13, 14, 15], but their precise mathematical nature remains elusive. In physics, such effects are often captured by *entropic forces*—macroscopic forces that emerge from the system's statistical tendencies rather than its energy landscape alone [16]. The power of this framework lies in the notion of an *effective entropy*, which plays the role of a potential whose gradients define the entropic force. Identifying this effective entropy not only reveals what the system is implicitly optimizing, but also opens the door to leveraging theoretical tools from statistical physics to analyze and improve AI models.

Contributions. We formalize this connection between stochastic learning dynamics and entropic forces through the lens of symmetry and representation learning to:

39th Conference on Neural Information Processing Systems (NeurIPS 2025).

^{*}Equal contribution.

- 1. Derive an entropic loss function and show that the entropic forces of SGD *break continuous* parameter symmetries while preserving discrete ones (Section 3).
- 2. Show that the symmetry breaking due to entropic forces gives rise to a family of *equipartition* theorems that predict the gradient alignment phenomena (Section 4).
- 3. Explain and unify two seemingly disparate but universal observations *progressive sharpening of the loss landscape* and the *emergence of universal representations* as consequences of entropic forces (Section 5).

Our theory establishes a principled framework – akin to a thermodynamics of deep learning – that unifies several universal phenomena under a single formalism. The results suggest that the entropic loss landscape, shaped by both optimization and entropy, plays a foundational role in understanding learning dynamics and emergent phenomena. Full derivations and experimental validations are provided in the appendix.

2 Related Work

Modified Loss and Effective Landscape. The concept of modified or effective losses has emerged as a critical framework for understanding the implicit biases induced by stochastic gradient descent (SGD) in deep learning, which differs from another line of work [17, 18, 19, 20, 21] which leverages the property of stationarity to analyze the stationary distribution of SGD. Ref. [22] introduced the notion of a modified loss to analyze the discrete-time dynamics of SGD, demonstrating how discretization implicitly alters the optimization landscape. Similarly, Refs. [23] and [24] extended the modified loss formulation to where there is a gradient noise due to minibatch sampling. These works conducted numerical simulations to show that training on the effective loss really approximates the original dynamics [23, 25, 26] and leads to similar generalization performances. In this work, we refer to this type of losses as entropic losses for their associations with theoretical physics. Still, these entropic losses remain poorly understood, and their significance for understanding emergent phenomena in deep learning is not yet appreciated. Our work finds the crucial link between the entropic loss and symmetry-breaking dynamics, which is important for understanding the various intriguing nonlinear phenomena of representation learning.

Parameter Symmetry in Neural Networks. Parameter symmetries are shown to play a fundamental role in shaping neural network training dynamics and their emergent properties [27, 28, 29, 7, 30, 31]. A series of works showed that continuous symmetries in the loss function give rise to conservation laws, which imply that the learning result of SGD training is strongly initialization-dependent [32, 28, 33]. More recent works showed how any stochasticity or discretization effect could break the symmetries in a systematic way such that the learned solution is no longer dependent on the initialization, a hint of universality [34, 29, 7, 31]. Particularly, Ref. [29] developed the formalism of exponential symmetries and proved that any loss function with an exponential symmetry leads to a symmetry-breaking dynamics that converges to unique fixed points. This point can be seen as the dynamical equivalence of our Theorem 2, which states that there is essentially no continuous symmetry in the entropic loss. In comparison, our framework takes a different perspective: we study the symmetry from a loss landscape perspective and identify these symmetry-breaking tendencies as entropic forces. This unified perspective enables us to understand the universal learning phenomena with a greatly simplified analysis.

3 Effective Energy for Stochastic Gradient Learning

Define $\ell(x,\theta)$ to be the per-sample loss function. We can define the empirical risk as

$$L(\theta) = \mathbb{E}_{\mathcal{B}} [\mathbb{E}_{x \in \mathcal{B}} \ell_{\gamma}(x, \theta)], \tag{1}$$

where $\ell_{\gamma}(x,\theta) \coloneqq \ell(x,\theta) + \gamma ||\theta||^2$, \mathcal{B} represents the minibatch and γ represents the weight decay. From a dynamical-system perspective, for an infinitesimal learning rate, the loss function coincides with the Bregman Lagrangian of this dynamics, and so one can leverage the Lagrangian formalism to understand the training of gradient flow [35]. This is particularly attractive from a theory perspective because modern theoretical physics are also founded on the Lagrangian formalism and this connection allows one to borrow physics intuitions to understand deep learning.

However, simply studying this loss function is insufficient to understand the learning dynamics of SGD at various learning rates η due to the stochastic discrete-time nature of SGD. This motivates the definition of an entropic loss ϕ_{η} such that running n steps of update on $\phi_{\eta}(\theta)$ with learning rate η/n

is the same as running one step of update on ℓ_{γ} for any x. Taking the limit $n \to \infty$, one can obtain a "renormalized" loss function for which running gradient flow is the same as running gradient descent for the original loss. With this loss, it becomes possible again to leverage Lagrangian formalism to understand SGD training with discrete-time and stochastic learning. Because ϕ_0 coincides with running gradient flow on ℓ_{γ} , one must have that

$$\phi_{\eta} := \ell_{\gamma} + \eta \phi_1 + \eta^2 \phi_2 + O(\eta^3). \tag{2}$$

We can also consider the more general case where the learning rate is a fixed symmetric matrix Λ with $||\Lambda|| = \eta$. The following theorem derives the entropic loss for this case. Many common algorithms, such as Adam, natural gradient descent, and even a wide range of biologically plausible learning rules [36] can be seen as having matrix learning rate.

Theorem 1. (Entropic Loss) For fixed x, starting from θ_0 run one-step gradient descent with Λ on $\ell_{\gamma}(x,\theta)$ to obtain θ_1 . Run n-step gradient descent with Λ/n on $\phi_{\Lambda}(x,\theta) := \ell_{\gamma}(x,\theta) + \phi_{1\Lambda}(x,\theta) + \phi_{1\Lambda}(x,\theta)$ $\phi_{2\Lambda}(x,\theta)$ to obtain θ'_n . Then, assuming $\|\nabla^3 \ell_{\gamma}(x,\theta)\| \leq M$,

- 1. if $\phi_{1\Lambda}(x,\theta) = \frac{1}{4}\nabla\ell_{\gamma}(x,\theta)^{T}\Lambda\nabla\ell_{\gamma}(x,\theta)$, then, $\theta'_{n} = \theta_{1} + O(||\Lambda||^{3} + ||\Lambda||^{2}/n)$; 2. moreover, if $\phi_{2\Lambda}(x,\theta) = \frac{1}{2}\nabla\ell_{\gamma}(x,\theta)^{T}\Lambda\nabla^{2}\ell(x,\theta)\Lambda\nabla\ell_{\gamma}(x,\theta)$, then $\theta'_{n} = \theta_{1} + O(||\Lambda||^{4} + ||\Lambda||^{2}/n + ||\Lambda||^{2}/$ $||\Lambda||^3/n + ||\Lambda||^3M$).

This ϕ_{η} also needs to hold in expectation with respect to the sampling of data points, and so one can define the expected entropic loss $F_{\eta,\gamma}(\theta) = \mathbb{E}[\phi_{\eta}(x,\theta)]$, which is, up to the first order in Λ and γ :

$$F_{\eta,\gamma}(\theta) = \underbrace{\mathbb{E}_{x}\ell(x,\theta)}_{\text{learning, symmetry}} + \underbrace{\gamma||\theta||^{2}}_{\text{regularization}} + \underbrace{\frac{1}{4}\mathbb{E}_{\mathcal{B}}||\sqrt{\Lambda}\mathbb{E}_{x\in\mathcal{B}}\nabla\ell(x,\theta)||^{2}}_{\text{effective entropy due to discretization error, noise, := }S(\theta) + O(||\Lambda||^{2}). \quad (3)$$

Specializing to the first order and with a scalar learning rate, this equation reduces to the "modified loss" previously derived in different contexts [22, 23, 24]. The derivation has a thermodynamic flavor as it essentially computes the degree to which the dynamics is irreversible, and the entropy term is the part that cannot be microscopically reversed. In Theorem 1, the first-order term in Λ encourages the model to have a small gradient fluctuation. The second-order term, on top of gradient regularization, encourages the model to move to flatter solutions; this term has been found to play a role in the edge of stability phenomenon [5]. However, since the first-order term is not yet well-understood, our work focuses on the first-order term in Λ . We also focus on a scalar learning rate $\Lambda = \eta I$ and will comment on the differences when the difference is essential.

Now, treating the original loss plus regularization as an energy, the dynamics of gradient flow on F contains an energy force and and an entropic force: $\dot{\theta} = -\eta(\nabla L + \gamma \theta + \nabla S)$. The ∇S term, the gradient of the effective entropy, will be called the "entropic force." See Figure 1 for an illustration of the entropic force and an example of how entropy evolves during training.

We prove that the entropic force term breaks almost any continuous symmetry of L, a key result that we will leverage to study progressive sharpening and universal representation learning.

Definition 1. A loss function $\ell(x,\theta)$ is said to be K-invariant if items 1-3 are satisfied:

- 1. locality: $K(\theta, \lambda) = \theta + \lambda Q(\theta) + O(\lambda^2)$ for a differentiable Q; 2. consistency: $K(K(\theta, \lambda), \lambda') = K(\theta, \lambda + \lambda')$; 3. invariance: $\ell(x, K(\theta, \lambda)) = \ell(x, \theta)$ for all x, θ and $\lambda \in \mathbb{R}$.

An entropic loss $F_{\eta,\gamma}$ is said to have the (**robust**) K-invariance if there exists a neighborhood around η, γ such that $F_{\eta,\gamma}$ is K-invariant.

Theorem 2. (Symmetry Breaking Under the Entropic Loss) Let ℓ be K-invariant. If F is also robustly K-invariant, then, (1) $||K(\theta, \lambda)|| = ||\theta||$ and (2) $\nabla^T \ell \nabla Q(\theta) \nabla \ell = 0$ for all θ .

This means that any symmetry or invariance that $F_{\eta,\gamma}$ has must be norm-preserving transformations. Essentially, this means that any invariance that is not rotation invariance must disappear. The following theorem shows that orthogonal discrete symmetries are preserved.

Theorem 3. Let $OO^T = I$. If $\ell(x, O\theta) = \ell(x, \theta)$ for any θ and x, then $F_{\eta, \gamma}(O\theta) = F_{\eta, \gamma}(\theta)$ for any

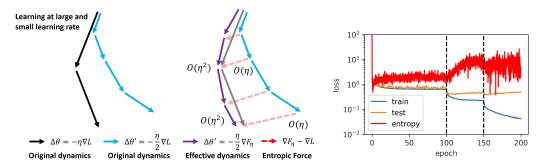


Figure 1: Entropic forces due to discretization error and stochasticity. Left: The learning dynamics of SGD at a large learning rate (LLR) and a small learning rate (SLR) is different. One can view the difference between LLR and SLR training as coming from an entropy term, which is an order η force. After entropic correction, the difference between SLR and LLR is reduced to $O(\eta^2)$ and it becomes possible to analyze LLR SGD with gradient flow. **Right**: An example of entropic effect in neural network training. ResNet18 trained on CIFAR-10 with learning rate decay at 100 and 150 epochs. At the first learning rate drop (black dashed lines), the gradient (entropy) increases. This is unexpected and can only be explained by the entropic loss, where a large learning rate penalizes the entropy, and thus decreasing the learning rate leads to an increase in entropy. The second drop does not create too much effect because the learning rate is too small after the first drop.

Together with the previous result, this shows that when gradient noise or regularization is taken into account, the only relevant remaining symmetries are discrete symmetries. This implies that the results that are based on conservation laws for understanding SGD are questionable and can only hold in the toy setting of an infinitesimal learning rate. The reason is simple: $F_{\eta,\gamma}$ does not have robust invariances at $\eta = \gamma = 0$. Also, note that had we used a generic matrix learning rate (e.g., with Adam), the orthogonal invariances would also be broken. The meaning of these discrete symmetries can be understood through a framework similar to that proposed in Ref. [7] and is left to a future work.

Emergence of Gradient Balance and Equipartition Property

Lie group symmetries exist abundantly in nature and in modern neural networks [27, 28, 29, 7, 32]. In thermodynamics, the existence of symmetries is a crucial fact that leads to the emergence of hierarchical phenomena and phase transitions between them. In a sense, symmetry can be argued to be the "first-order" approximation of the level of hierarchies in the system [37]. The following theorem states that the entropic loss F breaks any nontrivial noncompact Lie group symmetries (also known as exponential symmetries [29]). For formality, we say that ℓ (and L) has a A-exponential **symmetry** if for any $\lambda \in \mathbb{R}$, any x and θ , and any matrix A, $\ell(x,\theta)$ obeys $\ell(x,e^{\lambda A}\theta) = \ell(x,\theta)$.

Theorem 4. (Master Balance Theorem) If $\ell(x,\theta)$ has an A-exponential symmetry, then any local minimum θ^* of Eq.(3) satisfies

$$-\eta \mathbb{E}_{\mathcal{B}} [\mathbb{E}_{x \in \mathcal{B}} (\nabla_{\theta} \ell(x, \theta^*))]^T \tilde{A} [\mathbb{E}_{x \in \mathcal{B}} \nabla_{\theta} \ell(x, \theta^*)] + 4\gamma (\theta^*)^T \tilde{A} \theta^* = 0, \tag{4}$$

 $-\eta \mathbb{E}_{\mathcal{B}} \big[\mathbb{E}_{x \in \mathcal{B}} \big(\nabla_{\theta} \ell(x, \theta^*) \big) \big]^T \tilde{A} \big[\mathbb{E}_{x \in \mathcal{B}} \nabla_{\theta} \ell(x, \theta^*) \big] + 4\gamma (\theta^*)^T \tilde{A} \theta^* = 0, \tag{4}$ where $\tilde{A} := \frac{A + A^T}{2}$. In addition, either (1) $F_{\eta, \gamma}(e^{\lambda A} \theta^*) = F_{\eta, \gamma}(\theta^*)$ for all λ , or (2) there exists no $\lambda \neq 0$ such that $e^{\lambda A} \theta^*$ is a local minimum.

If we had chosen a symmetric A, we would have $\tilde{A} = A$. For an anti-symmetric A, this theorem is trivial, consistent with Theorem 2. Therefore, this theorem is a statement about non-compact Lie group symmetries that extends to infinity. The fact that every point is connected to a point that satisfies Eq (4) means that SGD can reach this condition easily. The meaning of this theorem is a general gradient balance and alignment phenomenon. When $\gamma = 0$, this equation states that the gradient along the positive spectrum of \tilde{A} must balance with the gradient along the negative spectrum. When $\gamma \neq 0$, there is, additionally, a tradeoff between gradient balance and weight balance. We apply this result to various neural networks in this section.

Many works have shown that when training with weight decay or when the model has a small initialization, the weights of the layers become balanced, especially in homogeneous networks [38,

¹See Ref. [31] for a detailed review.

39, 40]. Theorem 4 shows that the SGD in discrete-time or with stochasticity leads to a completely different bias where the gradient noise between all layers must be balanced.

ReLU Layers Consider a deep ReLU network trained on an arbitrary task:

$$f(x) = W_D R_{D-1} ... R_1 W_1 x, (5)$$

where R(x) is a piece-wise constant the zero-one activation matrix functioning as the ReLU activation. The entropy term can be written as

$$S(\theta) = \sum_{i=1}^{D} \text{Tr} \mathbb{E}[g_i g_i^{\mathsf{T}}], \tag{6}$$

where $g_i = \mathbb{E}_{x \in \mathcal{B}} \nabla_{W_i} \ell(\theta, x)$ and \mathbb{E} is a shorthand of $\mathbb{E}_{\mathcal{B}}$. Namely, we can group the gradient covariance according to layer index i. The following theorem states that all layers must have a balanced gradient. The proof shows that while the learning term L is invariant to a class of symmetry transformations, the entropy term is not – and this creates a systematic tendency for the parameters to reduce the entropy.

Theorem 5. (Layer Balance) For all local minimum of Eq. (3),

$$\eta(\mathbb{E}\operatorname{Tr}[g_ig_i^T - g_jg_j^T]) = 4\gamma(\operatorname{Tr}[W_iW_i^T - W_jW_j^T]). \tag{7}$$

Specifically, for $\gamma = 0$ we have gradient balance $\mathbb{E}\mathrm{Tr}[g_ig_i^{\mathsf{T}}] = \mathbb{E}\mathrm{Tr}[g_jg_j^{\mathsf{T}}]$. For $\eta = 0$, we have the standard weight balance $\mathrm{Tr}[W_iW_i^T] = \mathrm{Tr}[W_jW_j^T]$. Otherwise, the solution interpolates between gradient balance and weight balance.

Similarly, within every two neighboring layers, one can group the parameters into neurons:

$$\operatorname{Tr}[g_{i}g_{i}^{\mathsf{T}}] + \operatorname{Tr}[g_{i+1}g_{i+1}^{\mathsf{T}}] = \sum_{j} \left(\operatorname{Tr}[g_{i,j,:}g_{i,j,:}^{\mathsf{T}}] + \operatorname{Tr}[g_{i+1,:,j}g_{i+1,:,j}^{\mathsf{T}}]\right), \tag{8}$$

where $g_{i,j,:}$ is the incoming weights to the *j*-the neuron of the *i*-th layer, and $g_{i+1,:,j}$ is the outgoing weights of the same neuron. The gradient for each neuron must also be balanced because there is a rescaling symmetry in every neuron.

Theorem 6. (Neuron Balance) For all local minimum of Eq. (3) and any i, j,

$$\eta \mathbb{E} \text{Tr}[g_{i,j,:}g_{i,j,:}^{\mathsf{T}} - g_{i+1,:,j}g_{i+1,:,j}^{\mathsf{T}}] = 4\gamma \text{Tr}[w_{i,j,:}w_{i,j,:}^T - w_{i+1,:,j}w_{i+1,:,j}^T].$$
(9)

From a physics perspective, we have proved an equipartition theorem (ET). The elements in the matrix $\mathbb{E}[g_ig_i^T]$ can be seen as the temperature (or, the average energy) felt by each parameter. The trace $\mathrm{Tr}\mathbb{E}[g_ig_i^T]$ thus gives the temperature of the layer. That different layers emerge to have the gradient second momentum is an explicit form of the ET and means that the entropy S must be evenly spread out across every layer. Because the standard ET in physics is a property of thermal equilibrium, our result may be seen as an extension of the physical law to the out-of-equilibrium dynamics of learning. See Figure 2 for the emergence of layer and neuron balances in a ReLU network. We train on the MNIST dataset, but the labels are generated by a teacher ReLU network and trained with an MSE loss. Also, see Appendix A for examples of training trajectories and for an example with a self-attention net.

Polynomial Network Now, consider the case where R(h) is a diagonal matrix such that $R_{ii}(h) = h^d$ corresponds to a polynomial activation. This type of network is also a variant of homogeneous networks [38]. For these networks, one can show that the converge to a state where every layer's gradient norm is d times that of the previous layer, leading to a gradient exploding or vanishing problem (See Appendix B.8).

Self Attention Consider the case when a model has a generic form: $\ell(x, W, U) = \ell(x, WU)$, where W and U are matrices. The loss can contain other trainable parameters, which we ignore. Define $G_W = \mathbb{E}_{x \in \mathcal{B}} \nabla_W \ell(x, W, U)$, $G_U = \mathbb{E}_{x \in \mathcal{B}} \nabla_U \ell(x, W, U)$, and one can prove the following relation:

Theorem 7. (Gradient Alignment) For all local minimum of Eq.(3), we have

$$\eta \mathbb{E}[G_W^T G_W - G_U G_U^T] = 4\gamma (W^T W - U U^T). \tag{10}$$

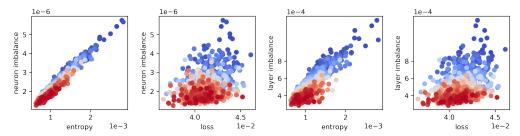


Figure 2: Layer and neuron gradient balance during training of a two-layer ReLU network. Here, every dot is a fixed time during training, where bluer dots are closer to the initialization, and redder dots are closer to convergence. Left 1-2: The entropy is strongly correlated with the neuron balance. As entropy decreases, the neuron balance improves. In contrast, the loss is not correlated to entropic effects at all. Right: Similarly, the layer balance is also correlated with entropy and not with loss function value.

This theorem is thus applicable to matrix factorization, deep linear networks, and, more importantly, self-attention layers. The self-attention logit is computed as $a_{ij} = X_i^T W U X_j$, where W is the key matrix, U is the query matrix. The loss function is a function of a_{ij} viewed as a matrix: $\ell(\{a_{ij}\})$. Let $V = W_2 W_1$. Then, applying this theorem reveals an intriguing relation:

$$W_1 \mathbb{E}[G_V^T G_V] W_1^T = W_2^T \mathbb{E}[G_V G_V^T] W_2. \tag{11}$$

Interpolating Weight Balance and Gradient Balance For all the theorems above, we have also studied how weight decay affects the balance conditions. We see that the weight decay creates something analogous: instead of gradient balance, weight decay encourages weight balance, and this effect often cannot be achieved together with gradient balance. Thus, there is a tradeoff between gradient balance and weight balance. In reality, the network is somewhere in between, where the weight balance and gradient balance has to "balance" with each other. Also, this is a generalized form of an equipartition theorem. If we regard the sum of the regularization and S as a "total" entropic potential Γ , then this means that every layer will contribute an equal amount to Γ .²

5 Implications

Next, we apply these results to study the emergence of universal representations in neural networks, and the progressive sharpening phenomena in deep learning optimization.

5.1 Universal Representation Learning

Recent works found that the representations of learned models are almost universally aligned to different models trained on similar datasets [41, 42], and even to the biological brains [43]. This interesting phenomenon has a rather philosophical undertone and has been termed "Platonic Representation Hypothesis" [8]. Here, we say that the two neural networks have learned a universal representation if for all x_1, x_2 ,

$$h_A(x_1)^T h_A(x_2) = h_B(x_1)^T h_B(x_2),$$
 (12)

where h_A is the activation of network A in one of the hidden layers, and h_B for network B. This is an idealization of what people have observed – and the difference between the two sides is the "degree of alignment." We leverage the entropic force formalism to identify an exact solution to the embedded deep linear network (EDLN) model:

$$\ell(\theta, x) = ||M_1 W_D \cdots W_1 M_2 x - y(x)||^2.$$
(13)

on datasets $\mathcal{D}_{M_3} = \{(M_3x_i, y_i)\}_i$, where M_1, M_2, M_3 are fixed but arbitrary invertible matrices. They have the following meaning:

• M_1 can be seen as a model of the layers coming after the deep linear network;

 $^{^2}$ Now, it is helpful to clarify the role of the L_2 regularization term. Conceptually, it can either be regarded as a part of the loss function ℓ , which we have done so far, or as a part of the entropic term. Treating it as a part of entropy is sometimes conceptually preferred because, like the entropy, it also functions as a regularization of the parameter space, limiting the accessible states.

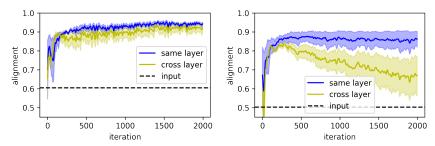


Figure 3: The representations of two 6-layer networks independently trained on randomly transformed MNIST become perfectly aligned for **every** pair of layers. The figure shows the average alignment between the same or different layers of two networks. This alignment does not weaken even if the input is arbitrarily transformed (Theorem 8). The black dashed line shows the average alignment to the input data, which is significantly weaker. **Left**: linear network. **Right**: tanh network.

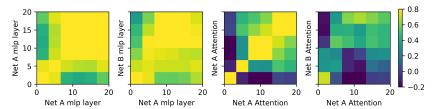


Figure 4: Alignment of representations of two ViT models pretrained on ImageNet. Net A: ViT-B (#param: 86M). Net-B: ViT-H (300M). We see that both mlp layers and the self-attention layers have mutually aligned representations both with itself and the other with each other. In particular, the alignment with itself is slightly better than with the different model, and the alignment of later layers is better than that of the first layers. A similar result with larger models is shown in Appendix A.5.

- M_2 models the layers coming after the embedded network;
- M_3 models different views of the data, which is common in multimodal learning therefore, the two models are trained on two different (but related) datasets.

In Theorem 8, we will train two different models, each with their own and potentially different M_1 , M_2 , M_3 . The arbitrariness of these three matrices implies *universality*.

The data is generated by $y_i = Vx_i + \epsilon_i$ for i.i.d. noise ϵ_i . Assuming that $\mathbb{E}x_i = \mathbb{E}\epsilon_i = 0$ and $\Sigma_{\epsilon} := \mathbb{E}\epsilon_i \epsilon_i^T$, $\Sigma_x := \mathbb{E}x_i x_i^T$, the following theorem characterizes the global minimum of the entropic loss for this network in the case $\gamma = 0$ and $\eta \to 0_+$.

Theorem 8 (Perfect Platonic Representation Hypothesis). Consider two deep linear networks A and B with weights of arbitrary dimensions larger than rank $(\sqrt{\Sigma_{\epsilon}}V\sqrt{\Sigma_{x}})$. Let model A train on $\mathcal{D}_{M_{3}}$ and model B on $\mathcal{D}_{M_{3}'}$. Moreover, the outputs of models are multiplied by M_{1}, M_{1}' , and the inputs are multiplied by M_{2}, M_{2}' , respectively. Then, at the global minimum of Eq. (3), every hidden layer of A is **perfectly** aligned with every hidden layer of B for any x, in the sense that

$$h_A^{L_A}(x) = c_0 R h_B^{L_B}(x)$$
 (14)

for $1 \le L_A < D_A$ and $1 \le L_B < D_B$ and any x, where c_0 is a scalar and $R = U_1U_2^T$ satisfying $U_1^TU_1 = U_2^TU_2 = I$. $h_A^{L_A}(x) \coloneqq \prod_{i=1}^{L_A} W_i^A M_2 M_3 x$, $h_A^{L_A}(x) \coloneqq \prod_{i=1}^{L_A} W_i^A M_2' M_3' x$ denote the output of the L_A, L_B -the layer of network A and B, respectively.

Here, a perfect alignment means that $h_A^{L_A}(x)$ differs from $h_B^{L_B}(x)$ only by a scaling and a rotation. Because of symmetry, SGD converges to a state where all possible pairs of the intermediate layers of two different networks are mutually aligned, independent of the initialization. This is an extraordinary fact because there exist infinitely many solutions that are not perfectly aligned, also due to symmetry. For example, take h_B to be any hidden layer, and transform its incoming weight by A and its outgoing weight by A^{-1} . This remains a global minimum for L, but it is no longer the case that there exists an orthogonal transformation R such that $h_A = h_B$. Therefore, for almost all global minima of L, there is no universal alignment between layers – yet, SGD prefers a universal solution due to the entropy term. The following theorem shows that weight decay will lead to a nonuniversal representation.

From a physics and thermodynamics perspective, it is quite reasonable that the irreversibility of the dynamics leads to the emergence of universal structures. A state is not universal if it contains information about its initial condition. Therefore, the irreversibility of the learning dynamics helps erase information about the parameter initialization, thereby enabling the learning of a universal solution. The following hypothesis can summarize this perspective:

Irreversibility enables universal representation learning.

Now, it is worthwhile to remark on the connection and difference between this result and the original PRH [8]. First, our theory provides strong support for the PRH. The original PRH only hypothesizes a positive similarity between models, and it is unclear whether the alignment score can reach 1 (which implies a perfect alignment) or will only be a small positive value. Our result shows that, in principle, it is possible to reach the perfect alignment limit. Secondly, this result also offers an alternative perspective on the representation alignment phenomenon to that of the original PRH paper. In the old perspective, one regards having no alignment as the default expectation, and positive alignment as something to be explained and understood. In our new perspective, the perfect alignment is the default, and breaking away from it is something to be explained and understood – which is exactly what we will demonstrate in the next theorem. We also refer to Ref. [44] for a more instructive derivation of the proof and for a detailed discussion of different ways to break the perfect PRH.

Theorem 9. Consider the deep linear network (13) with widths larger than d := rank(V). Let $\eta = 0$ and $\gamma \to 0_+$. At the global minimum of (3), we have

$$W_i = U_i P_i \Sigma U_{i-1}^T \tag{15}$$

for $i=1,\cdots,D$. U_D and U_0 are given by the SVD of $M_1^{-1}VM_3^{-1}M_2^{-1}:=U_DSU_0$, where $S\in\mathbb{R}^{d\times d}$ contains the singular values. For $i=2,\cdots,D-1$, U_i are arbitrary matrices satisfying $U_i^TU_i=I_{d\times d}$. Moreover, $\Sigma=S^{1/D}$. $\{P_i\}_{i=1}^D\subset\mathbb{R}^{d\times d}$ are diagonal matrices containing ± 1 and satisfy $\Pi_{i=1}^DP_i=I_d$.

The universal representation property (Theorem 8) does not hold anymore. Therefore, gradient balances lead to universal representations, whereas weight balances do not. See Figure 3 for an experiment with deep linear and nonlinear networks. A surprising aspect is that every layer can be aligned with every other layer. Because any deep nonlinear network is approximated by a linear network for a small weight norm [45], one could say that any nonlinear network is, to first order in $\|\theta\|$, a universal representation learner. In Figure 4, we compare the alignment between different self-attention layers of two differently sized vision transformers pretrained on ImageNet. Note that different layers of the same network and different networks also have significantly positive alignments, consistent with the solution for the embedded deep linear net. That using weight decay does not lead to universal representations is supported by the result in Figure 10.

Theorem 8 is a direct (perhaps the first) proof of the Platonic representation hypothesis, implying that for any x_1, x_2 , Eq. (12) holds. Importantly, the mechanism does not belong to any previously conjectured mechanisms (capacity, simplicity, multitasking [8]). This example has nothing to do with multitasking. The result holds for any deep linear network, all having the same capacity and the same level of simplicity, because all solutions parametrize the same input-output map. Here, the cause of the universal representation is symmetry alone: in the degenerate manifold of solutions, the training algorithm prefers a particular and universal one. This example showcases how symmetry is indeed an overlooked fundamental mechanism in deep learning.

5.2 The Sharpness Paradox

The tendency towards learning universal representations can imply a curse for training. For example, the sharpness of the loss landscape really depends on the distribution of the data, whereas the solution the model finds is independent of these distributions – this could imply that these solutions can be quite bad in terms of, say, optimization properties. Meanwhile, a paradox of the sharpness-seeking behavior of SGD has become explicit. On the one hand, the edge of stability (EOS) states that learning *typically* leads to sharper solutions, whereas a vast majority of works have shown that SGD training leads to flatter solutions [46, 47]. These two cannot happen simultaneously – and the solution must be that the sharpness-seeking behavior of SGD is situation-dependent. This section formalizes this intuition and shows that universal representation learning can be intrinsically related to the edge of stability phenomenon, which is also ubiquitous in deep learning.

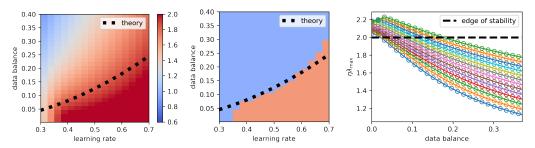


Figure 5: The entropic theory predicts the boundary for the edge of stability (EOS) phenomenon [4]. The theory shows that the imbalance of features and the uncertainty of labels make the model converge to sharper solutions. We run a two-layer linear network trained on a regression task. The **Left** panel plots the quantity $\eta \lambda_{\max}$ at convergence. For stability, $\eta \lambda_{\max}$ must stay (approximately) below 2, and the black dotted line plots the theoretical boundary for $\eta \lambda_{\max} = 2$. **Middle**: The same figure that emphasizes the phase boundary. The blue-red boundary empirically defined by the condition $\eta \lambda_{\max} = 2 - \epsilon$ with $\epsilon = 0.1$ – due to random sampling, the actual edge of stability is slightly smaller than 2 [20]. **Right**: We control the learning rate and balance of the label noise for this experiment. As predicted, as the data noise becomes more balanced, the sharpness metric $\eta \lambda_{\max}$ gets smaller, indicating better dynamical stability during training.

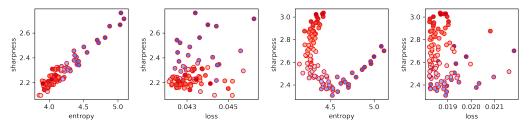


Figure 6: Example of a trajectory of training during the progressive flattening (**Left 1, 2**) and sharpening (**Right 1, 2**) of the experiment in Figure 5. Here, blue dots correspond to earlier in the training, and red dots correspond to later in the training. During progressive flattening, the decrease in sharpness correlates directly with the entropy term (Left 1), whereas the training loss is independent of the sharpness (Left 2). For progressive sharpening (Right), the picture is more complicated. The training trajectory follows three phases. Phase I: flattening correlates with a decrease in entropy; phase II: sharpening correlates with a decreasing entropy; phase III: sharpening correlates with an increasing entropy. The phase III cannot be explained by the leading-order entropic loss because, as the landscape becomes sharper, higher-order effects in η start to dominate training. At the same time, the loss is never correlated with these effects (Right 4).

Definition 2. The total sharpness is defined as $T(\theta) = \text{Tr}\mathbb{E}\nabla^2\ell(x,\theta)$.

This definition is chosen for analytical tractability and has been used in prior works [48, 29]. T upper bounds the largest eigenvalue, and T/d lower bounds it, so it is a good metric of stability and sufficient for the theorem we will prove. The following lemma shows that if there is an exponential symmetry in ℓ , every local minimum of ℓ connects without barrier to a local minimum with arbitrarily high sharpness. It can be seen as a generalization of the result of Ref. [49] to general symmetries.

Lemma 1. (Sharpness Lemma) Assume that A is a symmetric matrix and $\ell(x, e^{\lambda A}\theta) = \ell(x, \theta)$. Moreover, assume that $A\mathbb{E}\nabla^2\ell(x,\theta)\neq 0$. Then, $\limsup_{|\lambda|\to +\infty}|T(e^{\lambda A}\theta)|=+\infty$.

One can analytically solve for the sharpness of two-layer linear networks, and identify a precise cause of the progressive sharpening effect.

Theorem 10. For a two-layer linear network $\ell(x,\theta) = ||y(x) - W_1 W_2 x||^2$ with $y(x) = Vx + \epsilon \in \mathbb{R}^{d_y}$ and $\mathbb{E}x = \mathbb{E}\epsilon = 0$, $\Sigma_x = \mathbb{E}xx^T$, $\Sigma_\epsilon = \mathbb{E}\epsilon\epsilon^T$. Denote $\tilde{U}S'\tilde{V}$ to be the SVD of $V' := \sqrt{\Sigma_\epsilon}V\sqrt{\Sigma_x}$ and assume that the width of the network is larger than rank(V'). Then we have

$$T(\theta) = d_y \operatorname{Tr}\left[\Sigma_x \tilde{V}^T S' \tilde{V}\right] + \operatorname{Tr}\left[\Sigma_{\epsilon}^{-1} \tilde{U} S' \tilde{U}^T\right] \operatorname{Tr}\left[\Sigma_x\right]. \tag{16}$$

at the global optimum of (3). Meanwhile, the minimal sharpness of the global minimum of ℓ is

$$\min T(\theta) = 2\sqrt{d_y \text{Tr}\Sigma_x} \text{Tr}\hat{S}, \tag{17}$$

where \hat{S} is the singular values of $V\Sigma_x$.

Therefore, the minimum of (3) does not correspond to the minimal sharpness unless the data distribution is *fully isotropic* $\Sigma_x = I$, $\Sigma_\epsilon = I$. This result implies that SGD has no inherent preference for flatter minima. See Figure 5-6, where we train a two-layer linear network on a linear regression task with a 2d label $y \in \mathbb{R}^2$. The labels $y = V^*x + \epsilon$, for a ground truth matrix V^* and iid zero-mean noise ϵ such that $\Sigma_\epsilon = \mathrm{diag}(1,\phi_x)$, where $\phi_x \in (0,1)$ is called the "data balance." In Appendix A.6, we also train a deep nonlinear network and we see the same trend where improving balance in the label noise leads to flatter solutions.

As an example, we can choose Σ_x = I, V = I (d_x = d_y = d), which gives V' = $\sqrt{\Sigma_\epsilon}$ = $\tilde{U}S'\tilde{U}^T$ = $\tilde{V}^TS'\tilde{V}$, and thus

$$T(\theta) = d \operatorname{Tr} \left[\sum_{\epsilon}^{1/2} \right] + d \operatorname{Tr} \left[\sum_{\epsilon}^{-1/2} \right], \tag{18}$$

which is dominated by the second term and can be arbitrarily large. Recall that the minimum of $T(\theta)$, on the other hand, does not depend on Σ_{ϵ} , the label noise covariance. Thus, the imbalance of the noise spectrum can lead to arbitrarily high sharpness. This could especially be a problem for language model training because there is a large variation in the randomness of tokens. Some words, like "the" could have a very low conditional entropy, while nouns or verbs can have high entropy, especially when there exist synonyms. Another example is to choose $\Sigma_{\epsilon} = I$, V = I, which gives $V' = \sqrt{\Sigma_{\epsilon}} = \tilde{U}S'\tilde{U}^T = \tilde{V}^TS'\tilde{V}$, and thus

$$T(\theta) = d\text{Tr}[\Sigma_x^{3/2}] + \text{Tr}[\Sigma_x]^2. \tag{19}$$

Meanwhile, the minimal sharpness is given by $\min T(\theta) = 2\sqrt{d} \mathrm{Tr}[\Sigma_x]^{3/2}$. One can check that the sharpness reaches its minimum only if $\Sigma_x = I$. If Σ_x is very imbalanced, e.g. $\Sigma_x = \mathrm{diag}(1,0,\cdots,0)$, we have $T(\theta) = d+1$, significantly larger than $\min T(\theta) = 2\sqrt{d}$ for large d. Thus, an imbalance in the input feature can also lead to sharpness-seeking behaviors. At the same time, if the loss function has scale invariance $(\ell(x,\theta) = \ell(x,\lambda\theta))$ for any $\lambda \in \mathbb{R}$, the learning dynamics leads to a flattening of the curvature during training (See Section B.1). Thus, entropic forces and symmetry are strong factors deciding the sharpness-seeking behavior of SGD.

6 Conclusion

In this work, we have proposed an entropic-force perspective on neural network learning. We derive an entropic loss, which breaks continuous symmetries and preserves discrete ones, leading to universal behaviors such as gradient balance and alignment. Our results provide a unifying framework that explains several emergent phenomena in deep learning, including sharpness-seeking behavior and universal feature structure, as consequences of underlying symmetries and entropic forces. This framework offers predictive and explanatory power across architectures and scales, and points toward a more principled, physics-inspired understanding of learning dynamics and emergent phenomena. Future work may extend this foundation to encompass higher-order corrections, richer architecture structures, and nonequilibrium dynamics of modern training procedures. A major limitation of our work is that we only focused on the problems that have explicit symmetries, and an important future problem is to extend the results to the case where there is only approximate symmetry or no symmetry at all. Our experiments show that these systems without symmetries agree qualitatively with those systems with explicit symmetries, and this suggests that there are underlying, hidden concepts that are yet to be discovered.

An implication of Theorem 2 is that discrete symmetries such as Z_2 still remain in the loss function. This enables the possibilities of spontaneous symmetry breaking and phase transitions in neural networks. Prior works have studied phase transitions [7] without the entropy term and are inherently zero-temperature phase transitions. It may be possible to develop a theory of phase transitions based on entropic loss, similar to the classical Landau theory. Also, a striking aspect of our construction is its similarity to the actual formalism of thermodynamics in physics; our result motivates the development of a robust thermodynamics theory of deep learning.

³Viewed together with the result in Section 5, one reaches an interesting and surprising conjecture: progressive sharpening and universal representation alignment, with entirely different phenomenology, may have the same underlying cause, and could be two sides of the same coin.

Acknowledgment

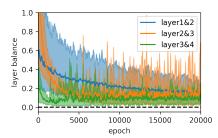
The authors thank Prof. Phillip Isola and Prof. Tomaso Poggio for discussion. ILC acknowledges support in part from the Institute for Artificial Intelligence and Fundamental Interactions (IAIFI) through NSF Grant No. PHY-2019786. This work was also supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF - 1231216.

References

- [1] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022.
- [2] Eric Michaud, Ziming Liu, Uzay Girit, and Max Tegmark. The quantization model of neural scaling. *Advances in Neural Information Processing Systems*, 36:28699–28722, 2023.
- [3] Sanjeev Arora and Anirudh Goyal. A theory for emergence of complex skills in language models. *arXiv preprint arXiv:2307.15936*, 2023.
- [4] Jeremy M Cohen, Simran Kaur, Yuanzhi Li, J Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. *arXiv preprint* arXiv:2103.00065, 2021.
- [5] Jeremy M Cohen, Alex Damian, Ameet Talwalkar, Zico Kolter, and Jason D Lee. Understanding optimization in deep learning with central flows. *arXiv preprint arXiv:2410.24206*, 2024.
- [6] Liu Ziyin and Masahito Ueda. Zeroth, first, and second-order phase transitions in deep neural networks. *Physical Review Research*, 5(4):043243, 2023.
- [7] Liu Ziyin. Symmetry induces structure and constraint of learning. In *Forty-first International Conference on Machine Learning*, 2024.
- [8] Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation hypothesis. *arXiv preprint arXiv:2405.07987*, 2024.
- [9] Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and Saining Xie. Representation alignment for generation: Training diffusion transformers is easier than you think. *arXiv* preprint arXiv:2410.06940, 2024.
- [10] Megan Tjandrasuwita, Chanakya Ekbote, Liu Ziyin, and Paul Pu Liang. Understanding the emergence of multimodal representation alignment. *arXiv preprint arXiv:2502.16282*, 2025.
- [11] Dong Shu, Haiyan Zhao, Jingyu Hu, Weiru Liu, Ali Payani, Lu Cheng, and Mengnan Du. Large vision-language model alignment and misalignment: A survey through the lens of explainability. arXiv preprint arXiv:2501.01346, 2025.
- [12] Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. *Advances in Neural Information Processing Systems*, 32, 2019.
- [13] Tomer Galanti and Tomaso Poggio. Sgd noise and implicit low-rank bias in deep neural networks. arXiv preprint arXiv:2206.05794, 2022.
- [14] Alex Damian, Eshaan Nichani, and Jason D Lee. Self-stabilization: The implicit bias of gradient descent at the edge of stability. arXiv preprint arXiv:2209.15594, 2022.
- [15] Shuo Xie and Zhiyuan Li. Implicit bias of adamw: ℓ_{∞} -norm constrained optimization. In *International Conference on Machine Learning*, pages 54488–54510. PMLR, 2024.
- [16] Rong-Gen Cai, Li-Ming Cao, and Nobuyoshi Ohta. Friedmann equations from entropic force. *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 81(6):061501, 2010.
- [17] Sho Yaida. Fluctuation-dissipation relations for stochastic gradient descent. *arXiv preprint* arXiv:1810.00004, 2018.

- [18] Guodong Zhang, Lala Li, Zachary Nado, James Martens, Sushant Sachdeva, George Dahl, Chris Shallue, and Roger B Grosse. Which algorithmic choices matter at which batch sizes? insights from a noisy quadratic model. Advances in neural information processing systems, 32, 2019.
- [19] Liu Ziyin, Kangqiao Liu, Takashi Mori, and Masahito Ueda. Strength of minibatch noise in SGD. In *International Conference on Learning Representations*, 2022.
- [20] Kangqiao Liu, Liu Ziyin, and Masahito Ueda. Noise and fluctuation of finite learning rate stochastic gradient descent. In *International Conference on Machine Learning*, pages 7045–7056. PMLR, 2021.
- [21] Liu Ziyin, Hongchao Li, and Masahito Ueda. Noise balance and stationary distribution of stochastic gradient descent. *Physical Review E*, 111(6):065303, 2025.
- [22] David GT Barrett and Benoit Dherin. Implicit gradient regularization. arXiv preprint arXiv:2009.11162, 2020.
- [23] Samuel L Smith, Benoit Dherin, David GT Barrett, and Soham De. On the origin of implicit regularization in stochastic gradient descent. arXiv preprint arXiv:2101.12176, 2021.
- [24] Yatin Dandi, Luis Barba, and Martin Jaggi. Implicit gradient alignment in distributed and federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6454–6462, 2022.
- [25] Jonas Geiping, Micah Goldblum, Phillip E Pope, Michael Moeller, and Tom Goldstein. Stochastic training is not necessary for generalization. *arXiv preprint arXiv:2109.14119*, 2021.
- [26] Daniel Kunin, Javier Sagastuy-Brena, Lauren Gillespie, Eshed Margalit, Hidenori Tanaka, Surya Ganguli, and Daniel LK Yamins. The limiting dynamics of sgd: Modified loss, phase-space oscillations, and anomalous diffusion. *Neural Computation*, 36(1):151–174, 2023.
- [27] Xingguo Li, Zhaoran Wang, Junwei Lu, Raman Arora, Jarvis Haupt, Han Liu, and Tuo Zhao. Symmetry, saddle points, and global geometry of nonconvex matrix factorization. *arXiv* preprint arXiv:1612.09296, 1:5–1, 2016.
- [28] Bo Zhao, Nima Dehmamy, Robin Walters, and Rose Yu. Symmetry teleportation for accelerated optimization. Advances in Neural Information Processing Systems, 35:16679–16690, 2022.
- [29] Liu Ziyin, Mingze Wang, Hongchao Li, and Lei Wu. Parameter symmetry and noise equilibrium of stochastic gradient descent. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [30] Liu Ziyin, Yizhou Xu, and Isaac L Chuang. Remove symmetries to control model expressivity and improve optimization. The Thirteenth International Conference on Learning Representations, 2025.
- [31] Liu Ziyin, Yizhou Xu, Tomaso Poggio, and Isaac Chuang. Parameter symmetry potentially unifies deep learning theory. *arXiv preprint arXiv:2502.05300*, 2025.
- [32] Hidenori Tanaka and Daniel Kunin. Noether's learning dynamics: Role of symmetry breaking in neural networks. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 25646–25660. Curran Associates, Inc., 2021.
- [33] Sibylle Marcotte, Rémi Gribonval, and Gabriel Peyré. Abide by the law and follow the flow: Conservation laws for gradient flows. 2023.
- [34] Feng Chen, Daniel Kunin, Atsushi Yamamura, and Surya Ganguli. Stochastic collapse: How gradient noise attracts sgd dynamics towards simpler subnetworks. *arXiv* preprint *arXiv*:2306.04251, 2023.

- [35] Andre Wibisono, Ashia C Wilson, and Michael I Jordan. A variational perspective on accelerated methods in optimization. proceedings of the National Academy of Sciences, 113(47):E7351–E7358, 2016.
- [36] Liu Ziyin, Isaac Chuang, and Tomaso Poggio. Heterosynaptic circuits are universal gradient machines, 2025.
- [37] Philip W Anderson. More is different: Broken symmetry and the nature of the hierarchical structure of science. *Science*, 177(4047):393–396, 1972.
- [38] Simon S Du, Wei Hu, and Jason D Lee. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. Advances in neural information processing systems, 31, 2018.
- [39] Andrey Polyakov. Homogeneous artificial neural network. arXiv preprint arXiv:2311.17973, 2023.
- [40] Akshay Rangamani and Andrzej Banburski-Fahey. Neural collapse in deep homogeneous classifiers and the role of weight decay. In *ICASSP* 2022-2022 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4243–4247. IEEE, 2022.
- [41] Yamini Bansal, Preetum Nakkiran, and Boaz Barak. Revisiting model stitching to compare neural representations. *Advances in neural information processing systems*, 34:225–236, 2021.
- [42] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International conference on machine learning*, pages 3519–3529. PMLR, 2019.
- [43] Daniel LK Yamins, Ha Hong, Charles F Cadieu, Ethan A Solomon, Darren Seibert, and James J DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the national academy of sciences*, 111(23):8619–8624, 2014.
- [44] Liu Ziyin and Isaac Chuang. Proof of a perfect platonic representation hypothesis. *arXiv* preprint arXiv:2507.01098, 2025.
- [45] Liu Ziyin, Botao Li, and Xiangming Meng. Exact solutions of a deep linear network. In *Advances in Neural Information Processing Systems*, 2022.
- [46] Alex Damian, Tengyu Ma, and Jason D Lee. Label noise sgd provably prefers flat global minimizers. *Advances in Neural Information Processing Systems*, 34:27449–27461, 2021.
- [47] Lei Wu, Mingze Wang, and Weijie Su. The alignment property of SGD noise and how it helps select flat minima: A stability analysis. *Advances in Neural Information Processing Systems*, 35:4680–4693, 2022.
- [48] Zhiyuan Li, Tianhao Wang, and Sanjeev Arora. What happens after sgd reaches zero loss?—a mathematical framework. In *International Conference on Learning Representations*, 2021.
- [49] L. Dinh, R. Pascanu, S. Bengio, and Y. Bengio. Sharp Minima Can Generalize For Deep Nets. *ArXiv e-prints*, March 2017.



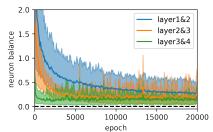
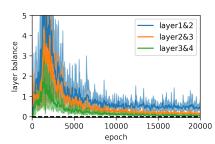


Figure 7: 4-layer linear networks and no weight decay trained on a teacher-student setting. **Left**: Layer imbalance for each layer, which verifies Theorem 5. **Middle**: Neuron balance for each layer, which verifies Theorem 6. The curves are smoothed and averaged over 5 runs for better visualization. **Right**: Loss and entropy.



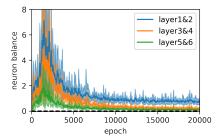


Figure 8: The same setting as Figure 7, but for ReLU activation.

A Experiment

A.1 ResNet

For Figure 1, we train ResNet18 on CIFAR 10 using SGD with momentum 0.9, batchsize 128 and weight decay 5×10^{-4} . The learning rate is 0.1 at the beginning, 0.01 after the 100–th epoch and 0.001 after the 150–th epoch. The entropy is calculated by summing the gradient norm of all parameters. We obtain training accuracy 98% and test accuracy 88% at the end.

A.2 Gradient Balance

For Figure 2, we train on the MNIST dataset but the labels are generated by a teacher ReLU network and trained with an MSE loss. Namely, the loss is

$$||f(\theta) - y(x) - \epsilon_x||^2, \tag{20}$$

where y(x) is a parameterized by a random ReLU teacher network and ϵ_x is an i.i.d. Gaussian noise with 0.2 standard deviation. The training proceeds with SGD for 10^4 steps with a learning rate of 0.01 and batchsize of 200. Layer balance is calculated by $|\mathbb{E}\mathrm{Tr}[g_ig_i^T-g_{i+1}^Tg_{i+1}]|$ and neuron balance is calculated by $\sum_j |\mathbb{E}\mathrm{Tr}[g_{i,j},g_{i,j}^T-g_{i+1,i,j}g_{i+1,i,j}^T]|$ for the i-th layer.

Additional experiments on layer balance and neuron balance are presented in Figures 7, 8 and 9 for 4-layer linear networks, ReLU networks and simple self-attention networks, still for the teacher-student setting. The hidden dimensions are 256, 128, 64 for linear and ReLU networks, and 256 for self-attention. We present the evolution of layer and neuron balance along training. Figures 7 and 8 suggest that layer balance and neuron balance approach zero during training, which verifies Theorems 5 and 6.

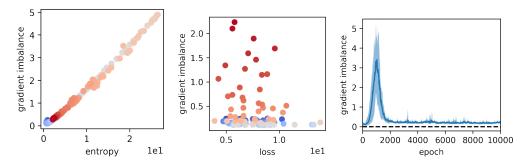


Figure 9: A self attention networks $y = (x^T U V x) w^T x$ and no weight decay trained on a teacher-student setting. The gradient imbalance is evaluated by $||G_U^T G_U - G_V G_V^T||_F$. Left: Gradient imbalance is strongly correlated with the entropy. Middle: Gradient imbalance is weakly correlated with the entropy. Right: The evolution of entropy along training, which is averaged over 5 runs.

A.3 Gradient Balance in Self-Attention Nets

Figure 9 suggests that $G_U^T G_U - G_V G_V^T$ approaches zero during training, and it is correlated with the entropy rather than the loss, which verifies Theorem 7.

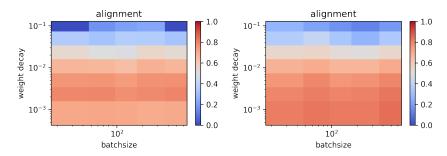


Figure 10: The alignment of two 2-layer ReLU networks independently trained on a teacher-student setting. We measure the alignment for different batchsizes and weight decay. **Left**: SGD. **Right**: Adam.

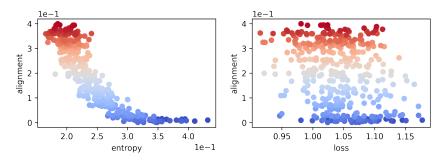


Figure 11: Universal alignment is strongly correlated with the entropy but not the loss. The setting is the same as Figure 10, where we use SGD, weight decay 0 and batchsize 100.

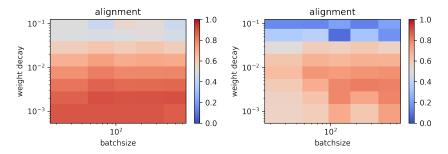


Figure 12: The same setting as Figure 10, but for two-layer linear networks. Left: SGD. Right: Adam.

A.4 Unversal Representation Learning in MLP

For Figure 3, we train two independent 6-layer networks on MNIST. The networks have linear or tanh activation and 128 neurons in each hidden layer, and for the second network, the input MNIST data is transformed by a random Gaussian matrix. We train the networks with Adam optimizer, learning rate 10^{-4} for 5 epochs. During training, we measure the representation alignment between every pair of layers, defined as the cosine similarity between the two sides of (12), averaged over the test set. We then plot the average alignment between the same or different layers of two networks. The input alignment denotes the average alignment between every layer representation and the input data.

In Figures 10 and 12, we test the influence of batchsize and weight decay on universal representation in a teacher-student setting. In Figure 10, both the teachers and students are two-layer ReLU networks. In Figure 12, the student is replaced by a linear network. Their hidden dimensions are 100. Similar to Figure 3, we measure the representation alignment between the middle layers of two independently trained networks, and the input of the second network is rotated by a Gaussian matrix. We use random Gaussian data, and the labels are generated by the teacher network. We train the student networks with SGD, learning rate 5×10^{-2} or with Adam, learning rate 10^{-4} . For both SGD and Adam optimizers, Figures 10 and 12 suggest that universal alignment does not rely on the batchsize as predicted, but disappears for large weight decay, which verifies Theorem 9. In Figure

11, we show that the increase of alignment is more correlated with the decrease of the entropy rather than the loss.

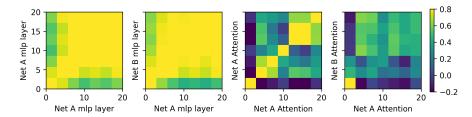


Figure 13: Alignment of representations of two larger ViT models pretrained on ImageNet. Net A: ViT-L (#param: 304M). Net-B: ViT-H (633M). This is similar to Figure 4

A.5 Universal Representation in ViT

See Figure 13 for the alignment in Vision Transformer. The pretrained weights are taken from https://docs.pytorch.org/vision/main/models.html. We measure the CKA alignment between the two models or with itself with a minibatch size of 300 images from the ImageNet dataset.

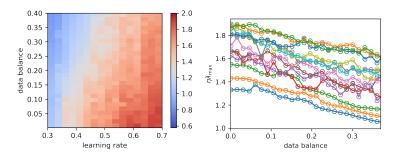


Figure 14: Sharpness at convergence for a two-hidden-layer ReLU network. The setting is identical to that of Figure 5. Again, we see that higher imbalance in the label leads to a sharper solution.

A.6 Edge of Stability

See Figure 5 and 6, where we train a two-layer linear network on a linear regression task with a 2d label $y \in \mathbb{R}^2$. The labels $y = V^*x + \epsilon$, for a ground truth matrix V^* and iid zero-mean noise ϵ such that $\Sigma_{\epsilon} = \operatorname{diag}(1, \phi_x)$, where $\phi_x \in (0, 1)$ is called the "data balance."

We train with different learning rates and data balance. The training proceeds with SGD with a batchsize of 32 for 4×10^4 iterations. Input $x \in \mathbb{R}^2$ is drawn from a standard Gaussian distribution, and the model has dimensions $2 \to 10 \to 2$.

See Figure 14 for an experiment with ReLU. Here, the architecture is has dimensions $2 \rightarrow 10 \rightarrow 10 \rightarrow 2$.

B Theory

B.1 Scale Invariance Leads to Flattening

The symmetry generator is A = I. We have

$$\frac{d}{d\lambda}F_{\eta,0}(e^{\lambda}\theta) = -\eta \mathbb{E}_x(\nabla_{\theta}\ell(x,\theta)^T \nabla_{\theta}\ell(x,\theta)) = -S < 0$$
(21)

by Theorem 4. As $\frac{d}{d\lambda}F_{\eta,\gamma}(e^{\lambda}\theta)|_{\lambda=0} = \theta^T \nabla F_{\eta,0}$. Therefore, when we do gradient descent along $-\nabla F_{\eta,0}$, λ monotonously increases. Meanwhile,

$$T(e^{\lambda}\theta) = \text{Tr}[e^{-2\lambda}\mathbb{E}\nabla^{2}\ell(x,\theta)]$$
 (22)

decreases with λ , and thus the sharpness decreases along training.

B.2 Proof of Theorem 1

Proof. For notational simplicity we drop the subscript γ in the proof. When running gradient descent on $\ell(x,\theta)$, we have

$$\theta_1 = \theta_t - \Lambda \nabla \ell(x, \theta_0). \tag{23}$$

When running gradient descent on ϕ_{Λ} , we have

$$\theta_1' = \theta_0 - \frac{\Lambda}{n} (\nabla \ell(x, \theta_0) + \nabla \phi_{1\Lambda}(x, \theta_0) + \nabla \phi_{2\Lambda}(x, \theta_0)) + O(||\Lambda||^3)$$
(24)

and

$$\theta_2' = \theta_1' - \frac{\Lambda}{n} (\nabla \ell(x, \theta_1') + \nabla \phi_1(x, \theta_1') + \nabla \phi_2(x, \theta_1')) + O(||\Lambda||^4)$$

$$= \theta_0 - \frac{2\Lambda}{n} \nabla \ell(x, \theta_0) + \frac{\Lambda}{n^2} \nabla^2 \ell(x, \theta_0) \Lambda \nabla \ell(x, \theta_0) - \frac{\Lambda}{n} \nabla \phi_1(x, \theta_0)$$

$$+ \frac{\Lambda}{n^2} \nabla^2 \phi_1(x, \theta_0) \Lambda \nabla \ell(x, \theta_0) - \frac{\Lambda}{n} \nabla \phi_2(x, \theta_0) + O(||\Lambda||^4).$$
(25)

Similarly we can obtain

$$\theta_n' = \theta_0 - \Lambda \nabla \ell(x, \theta_0) + \frac{\Lambda}{2} \nabla^2 \ell(x, \theta_0) \Lambda \nabla \ell(x, \theta_0) - \Lambda \nabla \phi_1(x, \theta_0) + \frac{\Lambda}{2} \nabla^2 \phi_1(x, \theta_0) \Lambda \nabla \ell(x, \theta_0) - \Lambda \nabla \phi_2(x, \theta_0) + O(\|\Lambda\|^4 + \|\Lambda\|^2/n + \|\Lambda\|^3/n).$$
(26)

Therefore, we have $\theta_n' = \theta_1 + O(||\Lambda||^3 + ||\Lambda||^2/n)$ if we choose $\phi_1(x, \theta) = \frac{1}{4}\nabla \ell(x, \theta)^T \Lambda \nabla \ell(x, \theta)$.

For small $\nabla^3 \ell(x,\theta)$, we have

$$\nabla^{2}\phi_{1}(x,\theta_{0})\Lambda\nabla\ell(x,\theta_{0}) = \nabla^{2}\ell(x,\theta_{0})\Lambda\nabla^{2}\ell(x,\theta_{0})\Lambda\nabla\ell(x,\theta_{0}) + O(\|\Lambda\|^{2}\|\nabla^{3}\ell(x,\theta)\|)$$

$$= \frac{1}{2}\nabla[\nabla\ell(x,\theta_{0})^{T}\Lambda\nabla^{2}\ell(x,\theta_{0})\Lambda\nabla\ell(x,\theta_{0})] + O(\|\Lambda\|^{2}\|\nabla^{3}\ell(x,\theta)\|).$$
(27)

Therefore, we can choose

$$\phi_2(x,\theta) := \frac{1}{2} \nabla \ell(x,\theta)^T \Lambda \nabla^2 \ell(x,\theta) \Lambda \nabla \ell(x,\theta)$$
 (28)

to obtain
$$\theta'_n = \theta_1 + O(||\Lambda||^4 + ||\Lambda||^2/n + ||\Lambda||^3/n + ||\Lambda||^3||\nabla^3 \ell(x, \theta)||).$$

B.3 Proof of Theorem 4

Proof. By the definition of the exponential symmetry,

$$\ell(x, e^{\lambda A}\theta) = \ell(x, \theta), \tag{29}$$

Taking derivative w.r.t. λ on (29), we have that $\nabla_{\theta}\ell(x,\theta)^T A\theta = 0$. Then taking derivative w.r.t. θ , we have that $A^T \nabla_{\theta}\ell(x,\theta) + \nabla_{\theta}^2\ell(x,\theta) A\theta = 0$.

Let $I(\lambda) := \frac{d}{d\lambda} F_{\eta,\gamma}(e^{\lambda A} \theta^*)$ and $\theta_{\lambda} := e^{\lambda A} \theta^*$. Then we have

$$I(\lambda) = \frac{\eta}{2} (\theta_{\lambda})^{T} A^{T} \mathbb{E}_{\mathcal{B}} [\mathbb{E}_{x \in \mathcal{B}} \nabla^{2} \ell(x, \theta_{\lambda})] [\mathbb{E}_{x \in \mathcal{B}} \nabla \ell(x, \theta_{\lambda})] + 2\gamma (\theta_{\lambda})^{T} A \theta_{\lambda}$$

$$= -\frac{\eta}{2} \mathbb{E}_{\mathcal{B}} (\mathbb{E}_{x \in \mathcal{B}} \nabla_{\theta} \ell(x, \theta_{\lambda}))^{T} A \mathbb{E}_{x \in \mathcal{B}} \nabla_{\theta} \ell(x, \theta_{\lambda}) + 2\gamma (\theta_{\lambda})^{T} A \theta_{\lambda}$$

$$= -\frac{\eta}{2} \text{Tr} [\Sigma(\theta_{\lambda}) A] + 2\gamma (\theta_{\lambda})^{T} A \theta_{\lambda}$$

$$= -\frac{\eta}{2} \text{Tr} [\Sigma(\theta_{\lambda}) \tilde{A}] + 2\gamma (\theta_{\lambda})^{T} \tilde{A} \theta_{\lambda},$$
(30)

where $\Sigma(\theta_{\lambda}) := \mathbb{E}_{\mathcal{B}}[\mathbb{E}_{x \in \mathcal{B}} \nabla_{\theta} \ell(x, \theta_{\lambda})][\mathbb{E}_{x \in \mathcal{B}} \nabla_{\theta} \ell(x, \theta_{\lambda})]^T$ is positive semi-definite. We also use $\operatorname{Tr}[\Sigma(\theta_{\lambda}) \frac{A - A^T}{2}] = 0$ and $(\theta_{\lambda})^T \frac{A - A^T}{2} \theta_{\lambda} = 0$. By [29, Lemma B.1], we have $\operatorname{Tr}[\Sigma(\theta_{\lambda}) \tilde{A}] = \operatorname{Tr}[e^{-2\lambda} \Sigma(\theta^*) \tilde{A}]$, and thus

$$I(\lambda) = -\frac{\eta}{2} \text{Tr} \left[\tilde{A} e^{-2\lambda \tilde{A}} \Sigma(\theta^*) \right] + 2\gamma (\theta^*)^T e^{2\lambda \tilde{A}} \tilde{A} \theta^*$$

$$= \sum_i -\frac{\eta}{2} \mu_i e^{-2\lambda \mu_i} (n_i^T \Sigma(\theta^*) n_i) + 2\gamma \mu_i e^{2\lambda \mu_i} (n_i^T \theta_i^*)^2,$$
(31)

where μ_i , n_i are eigenvalues of eigenvectors of the symmetric matrix \tilde{A} . Therefore,

$$I'(\lambda) = \sum_{i} \mu_i^2 \left(\eta e^{-2\lambda\mu_i} n_i^T \Sigma(\theta^*) n_i + 4\gamma e^{2\lambda\mu_i} (n_i^T \theta_i^*)^2 \right) \ge 0.$$
 (32)

We have $I(\lambda) \equiv 0$ (which happens only if \tilde{A} is not full rank) or $I(\lambda)$ strictly monotonic. As θ^* is a local minimum, we have I(0) = 0, which gives (4). Then we have $I(\lambda) \equiv 0$ or $I(\lambda) = 0$ iff $\lambda = 0$, which finishes the proof.

B.4 Proof of Theorem 2

Proof. By the definition of K-invariance, and taking derivative ∇_{θ} of both sides of $\ell(x, K(\theta, \lambda)) = \ell(x, \theta)$, we have

$$\nabla_{\theta} K(\theta, \lambda)^{T} \nabla_{K(\theta, \lambda)} \ell(x, K(\theta, \lambda)) = \nabla_{\theta} \ell(x, \theta), \tag{33}$$

where the l.h.s. follows from the chain rule. This imlies that

$$(I + \lambda \nabla Q + O(\lambda^2)) \nabla_{K(\theta, \lambda)} \ell(x, K(\theta, \lambda)) = \nabla_{\theta} \ell(x, \theta), \tag{34}$$

and so

$$\nabla_{K(\theta,\lambda)}\ell(x,K(\theta,\lambda)) = (I - \lambda\nabla Q + O(\lambda^2))\nabla_{\theta}\ell(x,\theta), \tag{35}$$

If F is K-invariant, the following Equation holds:

$$\ell(x, K(\theta, \lambda)) + \|\nabla_{K(\theta, \lambda)}\ell(x, K(\theta, \lambda))\|^{2} + \|K(\theta, \lambda)\|^{2} = \ell(x, \theta) + \|\nabla_{\theta}\ell(x, \theta)\|^{2} + \|\theta\|^{2}.$$
 (36)

By the assumption, $\ell(x, K(\theta, \lambda)) = \ell(x, \theta)$, and (36), we have that

$$\|\nabla_{\theta}\ell(x,\theta)\|^{2} + \|\theta\|^{2} = \|\nabla_{K(\theta,\lambda)}\ell(x,K(\theta,\lambda))\|^{2} + \|K(\theta,\lambda)\|^{2}$$
(37)

$$= \|\nabla_{\theta} \ell(x, \theta)\|^2 + \|\theta\|^2 + 2\lambda (\nabla^T \ell \nabla Q^T \nabla \ell - \gamma Q^T \theta) + O(\lambda^2). \tag{38}$$

Thus,

$$\eta \nabla^T \ell \nabla Q^T \nabla \ell - \gamma Q^T \theta = 0. \tag{39}$$

There are two cases: (1) $\eta \nabla^T \ell \nabla Q^T \nabla \ell - \gamma Q^T \theta = 0$, and (2) $\eta \nabla^T \ell \nabla Q^T \nabla \ell - \gamma Q^T \theta \neq 0$.

For case (1), we are done. For case (2), the equation cannot hold for $\gamma + d\gamma$ because the first term is independent of γ . Thus, we can only have case (1).

This means that for any θ

$$||K(\theta,\lambda)||^2 = ||\theta||^2 + O(\lambda^2),$$
 (40)

which is only possible if $||K(\theta, \lambda)||^2 = ||\theta||^2$. This completes the proof.

B.5 Proof of Theorem 3

Proof. By definition we have $||O\theta||^2 = ||\theta||^2$. Take derivative on both sides of $\ell(x, O\theta) = \ell(x, \theta)$, we have

$$O^{T}\nabla_{O\theta}\ell(x,O\theta) = \nabla_{\theta}\ell(x,\theta). \tag{41}$$

Thus we have

$$\nabla_{O\theta}\ell(x, O\theta) = O^{-T}\nabla_{\theta}\ell(x, \theta), \tag{42}$$

which gives $\|\nabla_{O\theta}\ell(x,O\theta)\|^2 = \|O^{-T}\nabla_{\theta}\ell(x,\theta)\|^2 = \|\nabla_{\theta}\ell(x,\theta)\|^2$. Combining the above results we have $F_{\eta,\gamma}(O\theta) = F_{\eta,\gamma}(\theta)$.

B.6 Proof of Theorem 5

Proof. Rescaling symmetry implies that if we make the transform

$$W_i \to e^{\lambda} W_i, \ W_j \to e^{-\lambda} W_j,$$
 (43)

L does not change.

This corresponds to the choice of

$$A_{klm}^{\tilde{k}\tilde{l}\tilde{m}} = \begin{cases} 1 & k = \tilde{k} = i \\ -1 & k = \tilde{k} = j \\ 0 & \text{otherwise} \end{cases}$$
 (44)

in Theorem 4, where the index klm corresponds to the m-th element of the l-th unit of the k-th layer. Then we have

$$\eta(\mathbb{E}\operatorname{Tr}[g_ig_i^T - g_jg_i^T]) = 4\gamma(\operatorname{Tr}[W_iW_i^T - W_jW_i^T]). \tag{45}$$

This finishes the proof by setting $\gamma = 0$.

B.7 Proof of Theorem 6

Proof. We can choose A to be a rescaling matrix w.r.t. the j-th neuron of the i-th layer. Specifically, we choose

$$A_{klm}^{\tilde{k}\tilde{l}\tilde{m}} = \begin{cases} 1 & k = \tilde{k} = i, l = \tilde{l} = j \\ -1 & k = \tilde{k} = i + 1, m = \tilde{m} = j \\ 0 & \text{otherwise} \end{cases}$$
 (46)

in Theorem 4, which gives

$$\eta \mathbb{E} \text{Tr}[g_{i,j,:}g_{i,j,:}^{\mathsf{T}} - g_{i+1,:,j}g_{i+1,:,j}^{\mathsf{T}}] = 4\gamma \text{Tr}[w_{i,j,:}w_{i,j,:}^T - w_{i+1,:,j}w_{i+1,:,j}^T]. \tag{47}$$

B.8 Gradient Imbalance in Polynomial Networks

Theorem 11. (Neuron Balance) For all local minima of Eq. (3) and any i, j,

$$\eta \mathbb{E} \text{Tr}[g_{i,j,:}g_{i,j,:}^{\mathsf{T}} - dg_{i+1,:,j}g_{i+1,:,j}^{\mathsf{T}}] = 4\gamma \text{Tr}[w_{i,j,:}w_{i,j,:}^{T} - dw_{i+1,:,j}w_{i+1,:,j}^{T}]. \tag{48}$$

This means that unless d=1, these networks will either have a noise or weight explosion problem. If $\gamma=0$, the gradient fluctuation grows like d^D , exponential in depth D. When d<1, later layers will have an exploding noise; when d>1, earlier layers will have an exploding noise. When both η and $\gamma\neq 0$, the sum of the noise and gradient norm will explode exponentially. In some sense, this implies that linear or sublinear types of activations are the only stable activations for deep neural networks.

Proof. We can still choose A to be a rescaling matrix, but this time we should rescale the i + 1-th layer more

$$w_{i,j,:} \to e^{\lambda} w_{i,j,:}, \ w_{i+1,:,j} \to e^{-d\lambda} w_{i+1,:,j}.$$
 (49)

This corresponds to

$$A_{klm}^{\tilde{k}\tilde{l}\tilde{m}} = \begin{cases} 1 & k = \tilde{k} = i, l = \tilde{l} = j \\ -d & k = \tilde{k} = i + 1, m = \tilde{m} = j \\ 0 & \text{otherwise} \end{cases}$$
 (50)

in Theorem 4, which gives

$$\eta \mathbb{E} \text{Tr}[g_{i,j,:}g_{i,j,:}^{\mathsf{T}} - dg_{i+1,:,j}g_{i+1,:,j}^{\mathsf{T}}] = 4\gamma \text{Tr}[w_{i,j,:}w_{i,j,:}^T - dw_{i+1,:,j}w_{i+1,:,j}^T].$$
 (51)

This gives

$$\mathbb{E}\operatorname{Tr}[g_{i,j,:}g_{i,i:}^{\mathsf{T}}] = d\mathbb{E}\operatorname{Tr}[g_{i+1,:,j}g_{i+1,:,j}^{\mathsf{T}}]$$
(52)

by choosing $\gamma = 0$.

B.9 Proof of Theorem 7

Proof. The double rotation symmetry can be written as

$$U \to e^{\lambda A} U, \ W \to e^{-\lambda A} W,$$
 (53)

where A is an arbitrary matrix. We can thus choose the following generator

$$A_{klm}^{\tilde{k}\tilde{l}\tilde{m}} = \begin{cases} 1 & k = \tilde{k} = 1, \ l = i, \tilde{l} = j, \text{ or } l = j, \tilde{l} = i \\ -1 & k = \tilde{k} = 2, \ l = i, \tilde{l} = j, \text{ or } l = j, \tilde{l} = i \\ 0 & \text{otherwise} \end{cases}$$
(54)

in Theorem 4, where k = 1 corresponds to W and k = 2 corresponds to U. This gives

$$\sum_{k} \eta [G_{W_{ki}} G_{W_{kj}} - G_{U_{ik}} G_{U_{jk}}] = \sum_{k} 4\gamma [W_{ki} W_{kj} - U_{ik} U_{jk}], \tag{55}$$

which finishes the proof.

B.10 Proof of Lemma 1

Proof. By definition, we have

$$e^{\lambda A} \nabla_{e^{\lambda A} \theta}^{2} \ell(x, e^{\lambda A} \theta) e^{\lambda A} = \nabla^{2} \ell(x, \theta), \tag{56}$$

and thus

$$T(e^{\lambda A}\theta) = \text{Tr}[e^{-2\lambda A}\mathbb{E}\nabla^2\ell(x,\theta)]. \tag{57}$$

Let $A := \sum_{i} \mu_{i} n_{i} n_{i}^{T}$, and thus

$$T(e^{\lambda A}\theta) = \sum_{i} e^{-2\lambda\mu_{i}} (n_{i}^{T} \mathbb{E} \nabla^{2} \ell(x,\theta) n_{i}).$$
 (58)

As $A\mathbb{E}\nabla^2\ell(x,\theta)\neq 0$, there exists i such that $\mu_i\neq 0$ and $n_i^T\mathbb{E}\nabla^2\ell(x,\theta)n_i\neq 0$. Therefore, we have $\lim_{\lambda\to+\infty}|T(e^{\lambda A}\theta)|=+\infty$ if $\mu_i<0$, and $\lim_{\lambda\to-\infty}|T(e^{\lambda A}\theta)|=+\infty$ if $\mu_i>0$.

B.11 Proof of Theorem 8

We first prove a theorem, which we will leverage to prove Theorem 8.

Theorem 12. Let $V' = \sqrt{\sum_{\epsilon} V} \sqrt{\sum_{x}}$ such that $V' = \tilde{U}S'\tilde{V}$ is its SVD and rank(V') = d. Assume that every layer has more than d hidden units. Then at the global minimum of (3), we have

$$\sqrt{\Sigma_{\epsilon}} M_1 W_D = \tilde{U} \Sigma_D U_{D-1}^T, \ W_i = U_i \Sigma_i U_{i-1}^T, \ W_1 M_2 M_3 \sqrt{\Sigma_x} = U_1 \Sigma_1 \tilde{V}, \tag{59}$$

for $i = 2, \dots, D-1$, where U_i are arbitrary matrices satisfying $U_i^T U_i = I_{d \times d}$, and $\Sigma_x = \mathbb{E}[xx^T]$, $\Sigma_{\epsilon} = \mathbb{E}[\epsilon \epsilon^T]$. Moreover,

$$\Sigma_1 = \Sigma_D = \left(\frac{d}{\operatorname{Tr}S'}\right)^{(D-2)/2D} \sqrt{S'}, \ \Sigma_i = \left(\frac{\operatorname{Tr}S'}{d}\right)^{1/D} I_{d \times d}.$$
 (60)

Proof. Consider two consecutive layers W_i and W_{i+1} . By the same argument as Theorem 7, we have

$$\eta \mathbb{E}[G_{W_{i+1}}^T G_{W_{i+1}} - G_{W_i} G_{W_i}^T] = 0.$$
(61)

By the MSE loss $\ell(x,y) = ||y - M_1 W_D \cdots W_1 M_2 M_3 x||^2$, this gives

$$W_{i}h_{i}\mathbb{E}[\|\xi_{i+1}^{T}\tilde{r}\|^{2}\tilde{x}\tilde{x}^{T}]h_{i}^{T}W_{i}^{T} = W_{i+1}^{T}\xi_{i+1}^{T}\mathbb{E}[\|h_{i}\tilde{x}\|^{2}\tilde{r}\tilde{r}^{T}]\xi_{i+1}W_{i+1}, \tag{62}$$

where $\tilde{x} = \mathbb{E}_{x \in \mathcal{B}} x$ and $\tilde{r} := \mathbb{E}_{x \in \mathcal{B}} [y - M_1 W_D \cdots W_1 M_2 M_3 x]$ satisfy $\mathbb{E} \tilde{x} = \mathbb{E} \tilde{r} = 0$, and thus $\mathbb{E} \tilde{x} \tilde{x}^T = \frac{\Sigma_x}{|\mathcal{B}|}$ and $\mathbb{E} \tilde{\epsilon} \tilde{\epsilon}^T = \frac{\Sigma_x}{|\mathcal{B}|}$. We use the fact that at the global minimum we have $M_1 W_D \cdots W_1 M_2 M_3 = V$, and thus $y - M_1 W_D \cdots W_1 M_2 M_3 x$ is independent of x. We also denote $\xi_{i+1} := M_1 W_D \cdots W_{i+2}$ and $h_i := W_{i-1} \cdots W_1 M_2 M_3$.

Finally denote $W_1' = W_1 M_2 M_3 \sqrt{\Sigma_x}$ and $W_D' = \sqrt{\Sigma_\epsilon} M_1 W_D$, which gives $W_D' \cdots W_1' = V'$ and

$$W_{i+1}^{\mathsf{T}} \frac{W_{i+2}^{\mathsf{T}} \cdots W_D^{\mathsf{T}} W_D^{\mathsf{T}} \cdots W_{i+2}}{\operatorname{Tr} \left[W_{i+2}^{\mathsf{T}} \cdots W_D^{\mathsf{T}} W_D^{\mathsf{T}} \cdots W_{i+2} \right]} W_{i+1} = W_i \frac{W_{i-1} \cdots W_1^{\mathsf{T}} W_1^{\mathsf{T}} \cdots W_{i-1}^{\mathsf{T}}}{\operatorname{Tr} \left[W_{i-1} \cdots W_1^{\mathsf{T}} W_1^{\mathsf{T}} \cdots W_{i-1}^{\mathsf{T}} \right]} W_i^{\mathsf{T}}.$$
(63)

The following steps are the same as [29, Theorem 5.4]. We can verify that the solution satisfies $M_1[\prod_{i=1}^{D_A}W_i^A]M_2M_3 = V$, or $h_A^{D_A}(x) = Vx$ as expected.

Now, we prove the actual theorem.

Proof. According to Theorem 12, at the global minimum of (3), the solution of a D_A -layer network for the dataset \mathcal{D}_M is given by

$$\sqrt{\sum_{\epsilon}} M_1 W_{D_A}^A = \tilde{U} \Sigma_D U_{D_A - 1}^T, \ W_i^A = U_i \Sigma_i U_{i - 1}^T, \ W_1^A M_2 M_3 \sqrt{\Sigma_x} = U_1 \Sigma_1 \tilde{V}$$
 (64)

for $i = 2, \dots, D-1$, where U_i are arbitrary matrices satisfying $U_i^T U_i = I_{d \times d}$, and

$$\Sigma_1 = \Sigma_D = \left(\frac{d}{\operatorname{Tr}S'}\right)^{(D_A - 2)/2D_A} \sqrt{S'}, \ \Sigma_i = \left(\frac{\operatorname{Tr}S'}{d}\right)^{1/D_A} I_{d \times d}.$$
 (65)

The solution suggests that

$$h_A^{L_A}(x) = \prod_{i=1}^{L_A} W_i^A M_2 M_3 = U_{L_A} \left(\frac{\text{Tr}S'}{d}\right)^{\frac{2L_A - D_A}{2D_A}} \sqrt{S'} \tilde{V} \Sigma_x^{-1/2} x.$$
 (66)

Similarly

$$h_B^{L_B}(x) = \prod_{i=1}^{L_B} W_i^B M_2' M_3' = U_{L_B} \left(\frac{\text{Tr}S'}{d}\right)^{\frac{2L_B - D_B}{2D_B}} \sqrt{S'} \tilde{V} \Sigma_x^{-1/2} x.$$
 (67)

The proof is complete by comparing (66) and (67).

One might also consider the case where the minimal width of network B is $d_B < d$. In this case, we denote $\bar{S}' \in \mathbb{R}^{d_B \times d_B}$ containing top d_B values of S'. Then we have

$$h_B^{L_B}(x) = \prod_{i=1}^{L_B} W_i^B M_2' M_3' = U_{L_B} \left(\frac{\text{Tr}\bar{S}'}{d}\right)^{\frac{2L_B - D_B}{2D_B}} \sqrt{\bar{S}'} \tilde{V} \Sigma_x^{-1/2} x.$$
 (68)

It is now not fully aligned with $h_A(x)$. To calculate the alignment, the corresponding kernel is

$$K_A(x_1, x_2) = h_A^{L_A}(x_1)^T h_A^{L_A}(x_2) = \left(\frac{\text{Tr}S'}{d}\right)^{\frac{2L_A - D_A}{D_A}} x_1^T \Sigma_x^{-1/2} \tilde{V}^T S' \tilde{V} \Sigma_x^{-1/2} x_2 \tag{69}$$

and

$$K_B(x_1, x_2) = h_B^{L_B}(x_1)^T h_B^{L_B}(x_2) = \left(\frac{\text{Tr}\bar{S}'}{d}\right)^{\frac{2L_B - D_B}{D_B}} x_1^T \Sigma_x^{-1/2} \tilde{V}^T \bar{S}' \tilde{V} \Sigma_x^{-1/2} x_2. \tag{70}$$

We then have

$$\langle K_A, K_A \rangle_F = \mathbb{E} K_A(x_1, x_2)^2 = \left(\frac{\text{Tr}S'}{d}\right)^{\frac{2(2L_A - D_A)}{D_A}} \text{Tr}[(S')^2].$$
 (71)

Similarly we have

$$\langle K_B, K_B \rangle_F = \left(\frac{\text{Tr}\bar{S}'}{d}\right)^{\frac{2(2L_B - D_B)}{D_B}} \text{Tr}[(\bar{S}')^2]$$
(72)

and

$$\langle K_A, K_B \rangle_F = \mathbb{E} K_A(x_1, x_2) K_B(x_1, x_2) = \left(\frac{\text{Tr}S'}{d}\right)^{\frac{2L_A - D_A}{D_A}} \left(\frac{\text{Tr}\bar{S}'}{d}\right)^{\frac{2L_B - D_B}{D_B}} \text{Tr}[\bar{S}'S'].$$
 (73)

Therefore, the alignment is given by

$$\frac{\langle K_A, K_B \rangle_F}{\sqrt{\langle K_A, K_A \rangle_F \langle K_B, K_B \rangle_F}} = \frac{\text{Tr}[\bar{S}'S']}{\sqrt{\text{Tr}[(S')^2]\text{Tr}[(\bar{S}')^2]}} = \sqrt{\frac{\text{Tr}[(\bar{S}')^2]}{\text{Tr}[(S')^2]}},$$
(74)

which is some value between 0 and 1.

B.12 Proof of Theorem 9

Proof. By Theorem 7 we have

$$W_{i+1}^T W_{i+1} = W_i W_i^T (75)$$

for $i=1,\cdots,D-1$, which suggests that $W_i=U_i\Sigma_iU_{i-1}^T$ with $\Sigma_i^2=\Sigma_{i+1}^2$. At the global minimum we have $U_D(\Pi_{i=1}^D\Sigma_i)U_0^T=\Pi_{i=1}^DW_i=M_1^{-1}VM_2^{-1}$, which shows that the left side is the SVD of $M_1^{-1}VM_2^{-1}$. This finishes the proof.

B.13 Proof of Theorem 10

Proof. By Theorem 12 we have

$$\sqrt{\Sigma_{\epsilon}}U = \tilde{U}\sqrt{S'}U_1^T, \ W\sqrt{\Sigma_x} = U_1\sqrt{S'}\tilde{V}. \tag{76}$$

By [31, Proposition 5.3] we have $S(\theta) = d_y ||W\Sigma_x||_F^2 + ||U||_F^2 \text{Tr}[\Sigma_x]$. This finishes the proof of (17).

Meanwhile we can also calculate

$$U^*, W^* = \arg\min_{U,W} d_y ||W\Sigma_x||_F^2 + ||U||_F^2 \text{Tr}[\Sigma_x].$$
 (77)

As $UW\Sigma_x = V\Sigma_x := \hat{U}\hat{S}\hat{V}$, the minimum is given by

$$U^* = \left(\frac{\text{Tr}\Sigma_x}{d_u}\right)^{1/4} \hat{U}\sqrt{\hat{S}}\hat{U}_1, \ W^*\Sigma_x = \left(\frac{\text{Tr}\Sigma_x}{d_u}\right)^{-1/4} \hat{U}_1\sqrt{\hat{S}}\hat{V}$$
 (78)

and

$$\min S(\theta) = 2\sqrt{d_y \text{Tr} \Sigma_x} \text{Tr} \hat{S}$$
 (79)

where \hat{U}_1 is an arbitrary orthogonal matrix.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Our abstract explicitly states that out main contribution is a theoretical framework for understanding the learning dynamics and emergent phenomena of neural networks trained with stochastic gradient descent (SGD) based an entropic loss landscape.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discussed this in the conclusion section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All theorems are stated with the full set of assumptions.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We present all experiment details in Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The experiments are only for demonstration and are straightforward to reproduce following the description.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We disclose all hyperparameters of experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We plot the standard error as the shaded area.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Personal computers are sufficient for all our experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: We closely follow the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There is no societal impact of the work performed.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We use pytorch. We use MNIST and CIFAR datasets, and the publically available pretrained weights from the pytorch website.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

• If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: Our paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can
 either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: This research does not involve LLMs as any important, original, or non-standard components.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.