THE GEOMETRY OF GROKKING: NORM MINIMIZATION ON THE ZERO-LOSS MANIFOLD

Anonymous authors

Paper under double-blind review

ABSTRACT

Grokking is a puzzling phenomenon in neural networks where full generalization occurs only after a substantial delay following the complete memorization of the training data. Previous research has linked this delayed generalization to representation learning driven by weight decay, but the precise underlying dynamics remain elusive. In this paper, we argue that post-memorization learning can be understood through the lens of constrained optimization: gradient descent effectively minimizes the weight norm on the zero-loss manifold. We formally prove this in the limit of infinitesimally small learning rates and weight decay coefficients. To further dissect this regime, we introduce an approximation that decouples the learning dynamics of a subset of parameters from the rest of the network. Applying this framework, we derive a closed-form expression for the post-memorization dynamics of the first layer in a two-layer network. Experiments confirm that simulating the training process using our predicted gradients reproduces both the delayed generalization and representation learning characteristic of grokking.

1 Introduction

Neural networks have achieved great success, but their mechanisms remain far from being fully understood. Doshi-Velez & Kim (2017) argue that understanding the inner workings of neural networks is crucial for the development of AI systems with increased safety and reliability. Moreover, understanding the learning dynamics of neural networks could also help us improve their performance and efficiency: it is easier to design better learning algorithms when we understand the limitations of the existing ones. Furthermore, insights into artificial neural networks may also enhance our understanding of biological neural networks due to their fundamental similarities (Sucholutsky et al., 2023; Kohoutová et al., 2020).

This work aims to clarify the learning dynamics underlying a particularly puzzling phenomenon termed *grokking*. Under specific training conditions, neural networks achieve generalization on the test data only after an extended period following the complete memorization of the training data. This behavior was first observed in synthetic problems such as modular addition (Power et al., 2022), but was later shown to also happen in real-world datasets (Liu et al., 2022b; Humayun et al., 2024).

In the specific problem of modular addition, interpretability research has revealed that neural networks achieve generalization by placing the embedding vectors on a circle (Gromov, 2023; Zhong et al., 2024). The circular structure of the embedding layer enables the network to perform a symmetric algorithm that generalizes perfectly to unseen data. It is currently known that circular representations emerge gradually during the post-memorization phase (Nanda et al., 2023) and that weight decay is mainly responsible for driving the delayed generalization (Liu et al., 2022b), but the precise dynamics remain unclear.

2 OUR CONTRIBUTIONS

In this work, we aim to answer the following questions:

- Q1. What is the exact role of weight decay in the post-memorization learning dynamics?
- **Q2.** Can we isolate the dynamics of the embedding layer from the rest of the network?

We answer **Q1** in Section 4 by proving that, after memorization is achieved, the learning dynamics approximately follow the minimization of the weight norm, constrained to the zero-loss level set.

We answer **Q2** in Section 5 by proposing an approximation for the isolated learning dynamics of any parameter subset as the minimization of a specific cost function.

We then combine these insights in Section 6 to study the post-memorization learning dynamics of a two-layer network, deriving a closed-form expression for the cost function of the first layer (the embedding layer).

Finally, in Section 7, we validate our theoretical insights on a modular addition task, showing that our approximations reproduce the delayed generalization and circular representations characteristic of grokking.

3 Intuitions from Toy Models

Before presenting our theoretical findings, we give an intuition for our results by discussing and visualizing how they relate to a few highly simplified models.

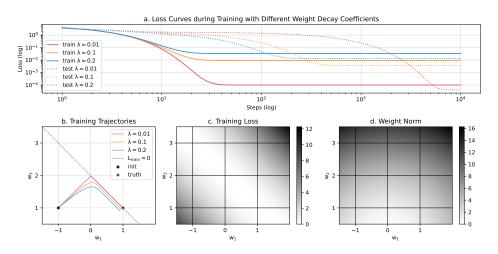


Figure 1: A two-parameter linear model $\hat{y}=w_1x_1+w_2x_2$ groks simple addition when trained with just one sample: $x_1=x_2=1,\ y=2$ (corresponding to 1+1=2). We plot three training runs with different weight decay coefficients λ . After quickly achieving zero loss, learning is entirely driven by the minimization of the weight norm.

3.1 Grokking Addition

We begin by discussing how a linear model can grok addition from just 1+1=2. We use a single-layer linear model with two inputs and two weights: $\hat{y}=w_1x_1+w_2x_2$. We train this model with mean-squared error loss using just one sample: $x_1=x_2=1,\ y=2$. We use three different values of weight decay, $\lambda \in \{0.01, 0.1, 0.2\}$. We initialize our model with $w_1=-1$ and $w_2=1$.

We aim to show that our model can learn to perform standard addition, despite being trained with a single sample. Test accuracy is measured on a set of 100 randomly generated samples, where x_1 and x_2 are sampled from a normal distribution, and $y = x_1 + x_2$.

We show our results in Figure 1. We can see that our model reproduces grokking: training loss becomes very low after just a 10 steps, while test loss takes a few hundred steps. Additionally, the model achieves lower loss with smaller λ , but takes longer to generalize.

Interpretation. From Figure 1 (a, b), we can see that learning follows two phases. In the first phase, driven by loss minimization, the model achieves a low loss by learning $(w_1, w_2) \approx (0, 2)$. In the second phase, learning is entirely driven by weight decay. The model follows norm minimization.

tion, while maintaining zero loss, eventually reaching $(w_1, w_2) \approx (1, 1)$. Note that, for smaller λ , the model remains closer to the zero-loss line, but generalization also takes longer.

3.2 NORM MINIMIZATION

While the previous example illustrates our theoretical framework, it does not capture its full generality. It is unsurprising that applying weight decay encourages a reduction in norm. However, the central claim of this paper is significantly stronger: we argue that the learning dynamics under weight decay do not merely follow *some* norm-decreasing direction, but rather evolve along the direction *that optimally minimizes the norm, subject to remaining on the zero-loss manifold*. Another way to view this is the following: once the model achieves perfect memorization, learning effectively follows gradient descent on the weight norm, constrained to the zero-loss manifold.

To offer a better intuition, we show how a linear model can grok three-number addition. We train a three-parameter linear model $\hat{y}=w_1x_1+w_2x_2+w_3x_3$ with just one sample: $x_1=x_2=x_3=1,\ y=3$. As in the previous section, this model exhibits grokking: after quickly achieving zero training loss, the model slowly reaches the generalizing solution $(w_1,w_2,w_3)\approx (1,1,1)$. We perform four training runs with different initializations and visualize the resulting trajectories in Figure 2. We observe that, for all initializations, the model first converges to the zero-loss plane, then moves directly towards the solution of minimum norm.

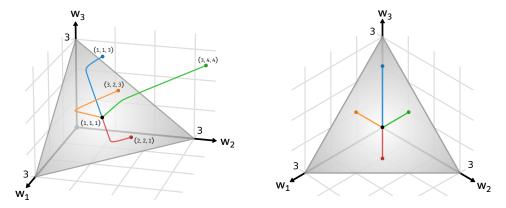


Figure 2: A three-parameter linear model $\hat{y} = w_1x_1 + w_2x_2 + w_3x_3$ groks three-number addition when trained with just one sample: $x_1 = x_2 = x_3 = 1$, y = 3 (corresponding to 1 + 1 + 1 = 3). The gray area shows the zero-loss plane, shaded according to the weight norm, where a lighter shade denotes a lower norm.

3.3 A FEW MATHEMATICAL NUANCES

So far, our examples have shown only flat zero-loss subspaces, but this is not necessarily the case. The zero-loss subspace can more generally be thought of as a manifold: a subspace that locally resembles Euclidean space near each point. For example, we show a curved zero-loss set in Figure 3 (left), along with a few training trajectories.

An important caveat is that the zero-loss set is not necessarily a manifold everywhere: it might contain singular points. Such a singular point is demonstrated in Figure 3 (center). However, such singularities should not worry us too much. If the network is realized by a smooth function, then the singularities will form a null set, i.e. the probability of encountering a singularity during standard training is exactly zero.

Another nuance is that, in practice, neural networks are trained using the ReLU activation function, which is not smooth. This will partition the loss into a finite set of smooth regions with nonsmooth boundaries. The nonsmooth points will also form a null set. We visualize a scenario of this type in Figure 3 (right) with leaky ReLU activation: ReLU(x) = x if x > 0 else x/10.

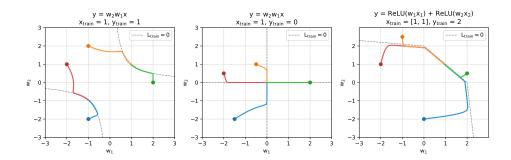


Figure 3: Training trajectories with different data, architectures and initializations. *Left:* a two-layer linear network where the zero-loss set is curved. *Center:* a two-layer linear network where the zero-loss set has a singularity at $(w_1, w_2) = (0, 0)$. *Right:* a single-layer network with leaky ReLU activation groks simple addition.

4 Post-Memorization Dynamics

4.1 ARCHITECTURE

We consider a neural network trained with mean-squared error loss on k samples:

$$\mathcal{L}(\theta) = \sum_{i=1}^{k} \|f(\theta, x_i) - y_i\|^2.$$
 (1)

where $x_i \in \mathbb{R}^n$, $y_i \in \mathbb{R}^m$, and $\theta \in \mathbb{R}^d$ is the parameter vector. The network has d parameters, n inputs, and m outputs. We use $f : \mathbb{R}^{d \times n} \to \mathbb{R}^m$ to denote the network realization function.

We apply a weight decay term (Krogh & Hertz, 1991) with a coefficient $\lambda \in \mathbb{R}^+$ to obtain:

$$\mathcal{L}_{\lambda}(\theta) = \mathcal{L}(\theta) + \lambda \|\theta\|^{2}. \tag{2}$$

We also denote the combined outputs of all samples as using $\mathcal{F}(\theta) \in \mathbb{R}^{km}$, where

$$\mathcal{F}(\theta) = \left[f(\theta, x_1)^\top, f(\theta, x_2)^\top, \dots, f(\theta, x_n)^\top \right]^\top.$$
 (3)

4.2 THEORETICAL SETUP

The following concepts are fundamental to our analysis:

Definition 4.1 (Zero-Loss Set). Let $\mathcal{Z} = \{ \theta \in \mathbb{R}^d \mid \mathcal{L}(\theta) = 0 \}$ denote the zero-loss set.

Definition 4.2 (Singular Points). We say that $\theta \in \mathbb{R}^d$ is a singular point¹ if the Jacobian matrix of \mathcal{F} at θ is not full rank, i.e. $\operatorname{rank}(\nabla \mathcal{F}(\theta)) < \min(d, km)$.

We theoretically study the training dynamics under the following assumptions:

Assumption 4.3 (Over-Parametrization). We assume that $d \ge km$ in order for the model to be able to memorize the entire dataset without learning any representations.

Assumption 4.4 (Perfect Memorization). We study the training dynamics after perfect memorization is achieved. This requires that the zero-loss set is not empty, i.e. $Z \neq \emptyset$.

Assumption 4.5 (Gradient Flow). We assume that \mathcal{F} is smooth and we model the gradient descent trajectory as a continuous-time gradient flow:

$$\frac{\partial \theta(t)}{\partial t} = -\nabla \mathcal{L}_{\lambda}(\theta(t)) \tag{4}$$

where $\theta: \mathbb{R}_{\geq 0} \to \mathbb{R}^d$ is the parameter vector as a function of time.

¹We define singular points in terms of \mathcal{F} , though they correspond to singular points of \mathcal{Z} . With respect to \mathcal{F} , they are more precisely *critical points*.

Assumption 4.6 (Vanishing Weight Decay). We study the learning dynamics in the approximation of a very small weight decay coefficient, i.e. $\lambda \to 0$, motivated by the small values of λ typically used by practitioners (Smith, 2018).

Assumption 4.7 (Well-Behaved Loss). We assume that our training trajectory does not pass through singular points. This is motivated by the fact that singular points form a null set, i.e. a set of Lebesgue measure zero (Morse, 1939; Sard, 1942).

4.3 Constrained to the Zero-Loss

We begin by establishing that the model remains constrained arbitrarily close to the zero-loss set after reaching a memorizing solution.

Definition 4.8 (Distance). Let $\operatorname{dist}_{\mathcal{Z}}(\theta) = \inf_{\phi \in \mathcal{Z}} \|\theta - \phi\|$ be the distance from $\theta \in \mathbb{R}^d$ to \mathcal{Z} .

Theorem 4.9 (Stability of \mathcal{Z}). For every trajectory starting at a zero-loss solution $\theta(0) \in \mathcal{Z}$ and every $\varepsilon > 0$, there exists $\lambda_{\varepsilon} > 0$ such that for all $0 < \lambda < \lambda_{\varepsilon}$ the trajectory under \mathcal{L}_{λ} satisfies

$$\sup_{t\geq 0} \operatorname{dist}_{\mathcal{Z}}(\theta(t)) < \varepsilon. \tag{5}$$

Sketch of the Proof. Our proof is based on the fact that the gradient flow will never increase the optimized quantity $\mathcal{L}_{\lambda}(\theta) = \mathcal{L}(\theta) + \lambda \|\theta\|^2$. Since both terms are non-negative, we can establish any desired bound on $\mathcal{L}(\theta)$ by an appropriate choice of λ . We then use this to obtain the bound on the distance. We give the full proof in Appendix A.

4.4 Loss Gradient Orthogonality

We will now provide our main theoretical result, which states that $\nabla \mathcal{L}(\theta)$ around \mathcal{Z} becomes orthogonal to any available direction.

Definition 4.10 (Available Direction). We say that $v \in \mathbb{R}^d$ is an available direction at $\theta \in \mathcal{Z}$ if there exists a smooth trajectory $s : \mathbb{R} \to \mathbb{R}^d$ such that $s(0) = \theta$, s'(0) = v, and $\mathcal{L}(s(t)) = 0$ for all $t \in \mathbb{R}$.

Definition 4.11 (Tangent Space). We denote by \mathcal{T}_{θ} the set of all available directions at $\theta \in \mathcal{Z}$.

Definition 4.12 (Projection). Let $\operatorname{proj}_{\mathcal{Z}}(\theta) = \operatorname{arg\,inf}_{\theta' \in \mathcal{Z}} \|\theta - \theta'\|$ be the projection of θ onto \mathcal{Z} .

Theorem 4.13 (Gradient Orthogonality). Let $S \subseteq \mathbb{R}^d$ be a compact space with $\operatorname{proj}_{\mathcal{Z}}(S \cap \mathcal{Z}) \subseteq S$. Assume that S contains no singular points. Then, there exist constants C > 0 and $x_0 > 0$ such that

$$\left| \frac{v^{\top}}{\|v\|} \frac{\nabla \mathcal{L}(\theta)}{\|\nabla \mathcal{L}(\theta)\|} \right| < C \operatorname{dist}_{\mathcal{Z}}(\theta)$$
 (6)

holds for all $\theta \in S \setminus \mathcal{Z}$ with $\operatorname{dist}_{\mathcal{Z}}(\theta) < x_0$ and $v \in \mathcal{T}_{\operatorname{proj}_{\mathcal{Z}}(\theta)}$.

Sketch of the Proof. We approximate the loss gradient at θ using the Taylor expansion around $\operatorname{proj}_{\mathcal{Z}}(\theta)$ to obtain $\nabla \mathcal{L}(\theta) = Hx + O(\|x\|^2)$, where $x = \theta - \operatorname{proj}_{\mathcal{Z}}(\theta)$ and $H = \nabla^2 \mathcal{L}(\operatorname{proj}_{\mathcal{Z}}(\theta))$.

Using $v \in \mathcal{T}_{\text{proj}_{\mathcal{Z}}(\theta)}$ and the absence of singular points, we are able to show that Hv = 0 and $\|\nabla \mathcal{L}(\theta)\| = \Theta(\|x\|)$. Therefore, the normalized dot product will be $O(\|x\|)$. We give the full proof in Appendix B.

Remark 4.14. In other words, the loss gradient does not induce any movement near \mathcal{Z} . After a memorizing solution is reached, learning will be driven entirely by weight decay. The loss will only serve to keep the model near \mathcal{Z} , while weight decay will be free to push the model towards norm minimization along any of the available directions.

5 ISOLATED DYNAMICS OF A NETWORK COMPONENT

5.1 MOTIVATION

It often happens that generalization in neural networks hinges on representation learning within specific network components. A striking example is the role of the embedding layer in the modular

addition task (Zhong et al., 2024). In such cases, it can be highly beneficial to simplify the complex learning dynamics of a deep network by isolating and analyzing only the component of interest.

Specifically, we desire to approximate the learning dynamics of individual components, independent from the rest of the network. Recent work has already begun to explore such approximations, also referred to as *effective theories* (van Rossem & Saxe, 2024; Liu et al., 2022a; Musat, 2024; Mehta et al.). In this section, we establish the theoretical foundations for such approximations and introduce a new one of our own.

5.2 ISOLATED DYNAMICS

The parameter vector can be decomposed into two orthogonal parameter subsets $\theta = [\theta_1, \theta_2]$, where $\theta_1 \in \mathbb{R}^{d_1}$, $\theta_2 \in \mathbb{R}^{d_2}$, and $d_1 + d_2 = d$. We are interested in the learning dynamics of θ_1 :

$$\dot{\theta}_1 = -\nabla_{\theta_1} \mathcal{L}_{\lambda}(\theta_1, \theta_2) \tag{7}$$

The gradient flow assumption means that the trajectory of θ_1 is a one-dimensional curve in a d_1 -dimensional space. This suggests that it is highly unlikely that our trajectory will pass through the same θ_1 twice. If we assume that a training trajectory only goes through unique values of θ_1 , then it is possible to parametrize θ_2 as a function of θ_1 :

$$\theta_2 = \phi(\theta_1) \tag{8}$$

where $\phi : \mathbb{R}^{d_1} \to \mathbb{R}^{d_2}$ is a function specific to the loss function and the initial parameters. This parametrization allows us to isolate the dynamics of θ_1 by expressing them as a function of θ_1 alone:

$$\dot{\theta}_1 = -\nabla_{\theta_1} \mathcal{L}_{\lambda}(\theta_1, \phi(\theta_1)) \tag{9}$$

While the function ϕ is generally intractable, working with reasonable approximations can provide valuable insights into the learning dynamics of θ_1 .

5.3 APPROXIMATE COST FUNCTION

We propose approximating ϕ by assuming that parameters θ_2 are optimal for the current value of θ_1 :

$$\phi(\theta_1) = \arg\min_{\theta_2} \mathcal{L}_{\lambda}(\theta_1, \theta_2) \tag{10}$$

This approximation can also be understood as treating θ_1 as the slow learning component, while θ_2 is the fast learning component that quickly adapts to the current value of θ_1 .

Additionally, optimizing θ under this approximation is equivalent to optimizing the following cost function:

$$\mathcal{R}(\theta_1) = \min_{\theta_2} \mathcal{L}_{\lambda}(\theta_1, \theta_2) \tag{11}$$

Theorem 5.1. The learning dynamics of θ_1 under Equations (9) and (10) follow the gradient flow of \mathcal{R} when ϕ is differentiable:

$$\dot{\theta}_1 = -\nabla_{\theta_1} \mathcal{R}(\theta_1) \tag{12}$$

Proof. Note that $\mathcal{R}(\theta_1) = \mathcal{L}_{\lambda}(\theta_1, \phi(\theta_1))$. By differentiating it with respect to θ_1 we obtain that $\nabla_{\theta_1} \mathcal{R}(\theta_1) = \nabla_{\theta_1} \mathcal{L}_{\lambda}(\theta_1, \phi(\theta_1)) + \nabla_{\theta_2} \mathcal{L}_{\lambda}(\theta_1, \phi(\theta_1)) \nabla_{\theta_1} \phi(\theta_1)$. However, since $\phi(\theta_1)$ is a minimum of \mathcal{L}_{λ} , we have that $\nabla_{\theta_2} \mathcal{L}_{\lambda}(\theta_1, \phi(\theta_1)) = 0$, giving us the desired result.

6 Two-Layer Networks

6.1 SETUP

We turn our attention to the learning dynamics of a two-layer neural network trained with mean squared error loss and weight decay:

$$\mathcal{L} = \| \sigma(XW_1) W_2 - Y \|_F^2 \tag{13}$$

where $X \in \mathbb{R}^{n \times d_{in}}$ is the input data, $Y \in \mathbb{R}^{n \times d_{out}}$ is the target output, $\sigma : \mathbb{R} \to \mathbb{R}$ is the activation function, $W_1 \in \mathbb{R}^{d_{in} \times d_h}$ is the first layer weights, $W_2 \in \mathbb{R}^{d_h \times d_{out}}$ is the second layer weights, and $\|\cdot\|_F$ denotes the Frobenius norm.

After applying a weight decay coefficient $\lambda > 0$, we get:

$$\mathcal{L}_{\lambda} = \mathcal{L} + \lambda \left(\|W_1\|_F^2 + \|W_2\|_F^2 \right) \tag{14}$$

6.2 ISOLATED DYNAMICS OF THE FIRST LAYER

Using the approximation from Section 5.3, we can isolate the learning dynamics of the first layer by assuming that the second layer weights are optimal for the current value of the first layer weights:

$$W_2 \approx \phi(W_1) = \arg\min_{\tilde{W}_2} \mathcal{L}_{\lambda}(W_1, \tilde{W}_2)$$
(15)

Since the second layer is just a linear transformation of the hidden layer activations $H = \sigma(XW_1)$, finding the optimal second layer weights is equivalent to the classic problem of ridge regression (Hoerl & Kennard, 1970). The solution is given by:

$$\phi(W_1) = (H^{\top}H + \lambda I)^{-1}H^{\top}Y \tag{16}$$

By combining Equations (13), (14) and (16), we can obtain the cost function for the isolated learning dynamics of the first layer:

$$\mathcal{R}(W_1) = \mathcal{L}_{\lambda}(W_1, \phi(W_1)) \tag{17}$$

This cost function is not particuarly simple, but it is fully differentiable, allowing us to approximate the learning dynamics of the first layer:

$$\dot{W}_1 \approx -\nabla \mathcal{R}(W_1) \tag{18}$$

6.3 Zero-Loss Approximation

Following the theoretical framework developed in Section 4, we can further simplify equation Equation (16) by working in the limit of very small weight decay $\lambda \to 0$:

$$\phi(W_1) = (H^\top H)^{-1} H^\top Y \tag{19}$$

The expression $H^+ = (H^\top H)^{-1} H^\top$ is also known as the Moore-Penrose pseudo-inverse of H. If H has full column rank, then the pseudo-inverse is also given by $H^+ = H^\top (HH^\top)^{-1}$. This allows us to further simplify the cost function down to:

$$\mathcal{R}(W_1) = \lambda \|W_1\|_F^2 + \lambda \operatorname{Tr}(Y^\top (HH^\top)^{-1}Y)$$
(20)

By differentiating this cost function, we can obtain a closed-form expression for the isolated learning dynamics of the first layer in the overparameterized zero-loss approximation:

$$\dot{W}_1 \approx X^T \Big((AYY^T AH) \odot \sigma'(XW_1) \Big) - W_1$$
 (21)

where $H = \sigma(XW_1)$, $A = (HH^T)^{-1}$, σ is the activation function, and \odot denotes the Hadamard product. We provide a detailed derivation in Appendix C.

7 SIMULATED DYNAMICS

In this section, we empirically validate our combined theoretical insights on isolated dynamics and post-memorization dynamics. By applying equation Equation (21) to a network trained on the modular addition task, we show that our approximations reproduce the delayed generalization and circular representations characteristic of grokking.

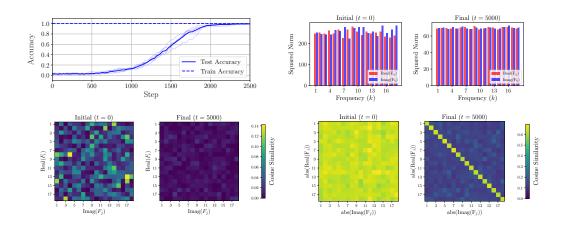


Figure 4: Simulated dynamics according to Equation (21) reproduce the phenomena of delayed generalization and representation learning. *Top left:* generalization emerges after about 1000 steps, despite training loss being exactly zero throughout. *Top right:* Fourier features norms equalize, suggesting the presence of equally-sized circles. *Bottom left:* Fourier features become orthogonal, suggesting that circles are located in orthogonal planes. *Bottom right:* Fourier features absolute values become dissimilar, suggesting that each circle leverages a different subset of hidden activations.

7.1 Dataset

We train the network to perform modular addition modulo a fixed number p. The dataset consists of k = p(p+1)/2 unique input pairs and their sum, $D = \{(a,b,c) \mid 0 \le a \le b < p, c = (a+b) \mod p\}$. We construct the input data $X \in \mathbb{R}^{k \times p}$ and the target output $Y \in \mathbb{R}^{k \times p}$ as $X_i = e_{a_i} + e_{b_i}$ and $Y_i = e_{c_i}$ for all $i = 1, \ldots, k$, where e_i is the i-th unit vector in \mathbb{R}^p and $(a_i, b_i, c_i) \in D$ is the i-th sample in the dataset.

We split the dataset into $(X_{\text{train}}, Y_{\text{train}})$ and $(X_{\text{test}}, Y_{\text{test}})$, using a fraction f_s of the dataset for training and the remaining $1 - f_s$ for testing.

7.2 ARCHITECTURE

We train a two-layer neural network with input dimension p, hidden dimension d_h , output dimension p, and a non-linear activation $\sigma: \mathbb{R} \to \mathbb{R}$.

Since the inputs are sums of one-hot vectors, we refer to the first layer weights as the embedding matrix $E \in \mathbb{R}^{p \times d_h}$. We refer to the second layer weights as simply the weights matrix $W \in \mathbb{R}^{d_h \times p}$.

The network output is given by $\hat{Y} = \sigma(XE)W$. To emphasize the role of the first layer as an embedding, we can also write the output as $\hat{Y}_i = \sigma(E_{a_i} + E_{b_i})W$.

We define the test accuracy as the percentage of correctly predicted test samples. We say that a test sample is correctly predicted if the index of the maximum value in the predicted output \hat{Y}_i matches c_i .

7.3 SIMULATED OPTIMIZATION

Our goal is not to train the network, but to validate that our approximate learning dynamics reproduce the phenomena observed during standard training.

We simulate the evolution of the embedding matrix E under the isolated dynamics given by equation Equation (21). We start from a random initialization $E \sim \mathcal{N}(0, p^{-1/2})$ and update it for T steps as $E \leftarrow E + \eta \Delta E$, where $\eta > 0$ is the step size and $\Delta E = X^T \left(\left(AYY^TAH \right) \odot \sigma'(XE) \right) - E$.

The isolated dynamics assume that $W = (H^T H)^{-1} H^T Y$ is optimal for the current value of E, which guarantees zero loss and perfect accuracy on the training data throughout training. This also ensures that predicted outputs perfectly match the target outputs on the training data, which is less principled for a classification task, but generally performs well in practice (Rifkin et al., 2003).

7.4 DETAILS

We use p = 37, $d_h = 512$, $\sigma(x) = \max(0, x)$, $f_s = 0.7$, $\eta = 10^{-3}$, T = 5000.

7.4.1 DELAYED GENERALIZATION

We simulate 5 runs starting from different random initializations and plot the test accuracy in Figure 4. Despite the fact that the training loss is exactly zero throughout, the test accuracy is not better than random guessing for the first 500 steps. However, the network eventually achieves perfect generalization on the test data after about 2000 steps, reproducing the delayed generalization phenomena Power et al. (2022).

7.4.2 FOURIER FEATURES

Using a discrete Fourier transform, we decompose the embedding matrix E into a linear combination of circles with different frequencies:

$$F_k = \frac{1}{p} \sum_{j=0}^{p-1} e^{-i2\pi jk/p} E_j \qquad \forall k \in \{1, \dots, (p-1)/2\}$$

Projecting the embeddings onto the plane spanned by $Re(F_k)$ and $Im(F_k)$ gives us a circle where the embeddings appear in the order $\{0, k, 2k, 3k, \ldots, (p-1)k\} \mod p$. Note that a circle of frequency k is equivalent to a circle of frequency p-k, so we only need to consider frequencies up to (p-1)/2.

We visualize several comparisons of the Fourier features of the initial and final embedding matrices for a single run Figure 4. First, the norms of the real and imaginary parts of the Fourier features equalize, suggesting the presence of *equally sized circles with perfect aspect ratios*. Second, the real and imaginary parts of the Fourier features become orthogonal, indicating that *circles are located in orthogonal planes*. Third, by taking the absolute value of the real and imaginary parts of the Fourier features, we obtain vectors very similar for the same frequency, but very different for different frequencies. This suggests that *each circle leverages a different subset of hidden units*.

8 CONCLUSION

We have formally established that the learning dynamics of neural networks in the grokking regime approximate as the minimization of the weight norm within the zero-loss set. Additionally, we have established a theoretical basis for approximating the learning dynamics of individual network components.

Limitations. This work does not cover cross-entropy loss, which is commonly used in practice. With regard to isolated dynamics, our work is limited to the case of two-layer networks. Exciting challenges lie ahead in understanding the grokking dynamics of more complex settings and architectures.

Impact Statement. We believe that understanding the learning dynamics of neural networks is essential for the design of more efficient and accurate AI systems. However, the development and deployment of such systems should be approached with caution.

LLM Usage. Large Language Models (LLMs) were used in standard ways throughout this work to polish the writing, assist with coding, and support brainstorming of mathematical proofs.

Reproducibility. We provide the full code for training and plotting used for the experiments with simulated dynamics from Section 7.

REFERENCES

- Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.
- Andrey Gromov. Grokking modular arithmetic. arXiv preprint arXiv:2301.02679, 2023.
- Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
 - Ahmed Imtiaz Humayun, Randall Balestriero, and Richard Baraniuk. Deep networks always grok and here is why. *arXiv preprint arXiv:2402.15555*, 2024.
 - Lada Kohoutová, Juyeon Heo, Sungmin Cha, Sungwoo Lee, Taesup Moon, Tor D Wager, and Choong-Wan Woo. Toward a unified framework for interpreting machine-learning models in neuroimaging. *Nature protocols*, 15(4):1399–1435, 2020.
 - Anders Krogh and John Hertz. A simple weight decay can improve generalization. *Advances in neural information processing systems*, 4, 1991.
 - Ziming Liu, Ouail Kitouni, Niklas S Nolte, Eric Michaud, Max Tegmark, and Mike Williams. Towards understanding grokking: An effective theory of representation learning. *Advances in Neural Information Processing Systems*, 35:34651–34663, 2022a.
 - Ziming Liu, Eric J Michaud, and Max Tegmark. Omnigrok: Grokking beyond algorithmic data. In *The Eleventh International Conference on Learning Representations*, 2022b.
 - Rohan Mehta, Ziming Liu, and Max Tegmark. Neural embeddings evolve as interacting particles.
 - Anthony P Morse. The behavior of a function on its critical set. *Annals of Mathematics*, 40(1): 62–70, 1939.
 - Tiberiu Musat. Clustering and alignment: Understanding the training dynamics in modular addition. *arXiv preprint arXiv:2408.09414*, 2024.
 - Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. Progress measures for grokking via mechanistic interpretability. *arXiv preprint arXiv:2301.05217*, 2023.
 - Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*, 2022.
 - Ryan Rifkin, Gene Yeo, Tomaso Poggio, et al. Regularized least-squares classification. *Nato Science Series Sub Series III Computer and Systems Sciences*, 190:131–154, 2003.
 - Arthur Sard. The measure of the critical values of differentiable maps. 1942.
 - Leslie N Smith. A disciplined approach to neural network hyper-parameters: Part 1–learning rate, batch size, momentum, and weight decay. *arXiv* preprint arXiv:1803.09820, 2018.
 - Ilia Sucholutsky, Lukas Muttenthaler, Adrian Weller, Andi Peng, Andreea Bobu, Been Kim, Bradley C Love, Erin Grant, Iris Groen, Jascha Achterberg, et al. Getting aligned on representational alignment. *arXiv preprint arXiv:2310.13018*, 2023.
 - Loek van Rossem and Andrew M Saxe. When representations align: Universality in representation learning dynamics. *arXiv* preprint arXiv:2402.09142, 2024.
 - Ziqian Zhong, Ziming Liu, Max Tegmark, and Jacob Andreas. The clock and the pizza: Two stories in mechanistic explanation of neural networks. *Advances in Neural Information Processing Systems*, 36, 2024.

PROOF OF THEOREM 4.9 (STABILITY OF \mathcal{Z}) We begin by establishing the following intermediate result: **Lemma A.1.** For every initialization $\theta(0) \in \mathcal{Z}$ and every $\varepsilon > 0$, there exists $\lambda_{\varepsilon} > 0$ such that for all $0 < \lambda < \lambda_{\varepsilon}$ the corresponding trajectory $\theta(t)$ under \mathcal{L}_{λ} satisfies $\sup_{t\geq 0} \mathcal{L}(\theta(t)) < \varepsilon.$ (22)*Proof.* From Equation (4), we see that gradient flow will never increase the optimized quantity $\mathcal{L}_{\lambda}(\theta) = \mathcal{L}(\theta) + \lambda \|\theta\|^2$. By choosing $\lambda_{\varepsilon} < \varepsilon / \|\theta(0)\|^2$, we ensure that $\mathcal{L}(\theta(t)) < \varepsilon$ for all $t \ge 0$. \square We recall Theorem 4.9: **Theorem 4.9** (Stability of \mathcal{Z}). For every trajectory starting at a zero-loss solution $\theta(0) \in \mathcal{Z}$ and every $\varepsilon > 0$, there exists $\lambda_{\varepsilon} > 0$ such that for all $0 < \lambda < \lambda_{\varepsilon}$ the trajectory under \mathcal{L}_{λ} satisfies $\sup_{t\geq 0} \operatorname{dist}_{\mathcal{Z}}(\theta(t)) < \varepsilon.$ (5) *Proof.* Our training trajectory will not reach any $\theta \in \mathbb{R}^d$ with $\|\theta\| > \|\theta(0)\|$. This is because any such configuration is unreachable by gradient flow from $\theta(0)$ for any $\lambda > 0$ since $\mathcal{L}_{\lambda}(\theta) > \mathcal{L}_{\lambda}(\theta_0)$. We are left to show unreachability of the set $\Phi = \{\theta \in \mathbb{R}^d : \mathcal{D}(\theta) \geq \varepsilon \text{ and } \|\theta\| \leq \|\theta(0)\|\}$. Since Φ is compact, $m = \min_{\Phi} \mathcal{L}(\theta)$ exists and is positive. Applying Lemma A.1, there exists $\lambda > 0$ such that optimizing \mathcal{L}_{λ} starting from $\theta(0)$ is guaranteed to maintain $\mathcal{L}(\theta(t)) < m$, thus making Φ unreachable.

B PROOF OF THEOREM 4.13 (GRADIENT ORTHOGONALITY)

We begin with a few intermediate results.

Lemma B.1. At a non-singular point of the zero-loss set, the tangent space is exactly the null space of the Hessian matrix, i.e. $\mathcal{T}_{\theta} = \{ v \in \mathbb{R}^d \mid \nabla^2 \mathcal{L}(\theta) \ v = 0 \}$ for any non-singular $\theta \in \mathcal{Z}$.

Proof. Part I: $\nabla^2 \mathcal{L}(\theta) \ v = 0 \implies v \in \mathcal{T}_{\theta}$.

We write the loss function as

$$\mathcal{L}(\theta) = \|\mathcal{F}(\theta) - y_{all}\|^2$$

using $y_{all} \in \mathbb{R}^{km}$ as the concatenation of all target outputs:

$$y_{all} = \begin{bmatrix} y_1^\top, \ y_2^\top, \ \dots, \ y_k^\top \end{bmatrix}^\top$$

Differentiating $\mathcal{L}(\theta)$ with respect to θ , we obtain the gradient:

$$\nabla \mathcal{L}(\theta) = 2 \, \nabla \mathcal{F}(\theta)^{\top} \big(\mathcal{F}(\theta) - y \big),$$

where $\nabla \mathcal{F}(\theta) \in \mathbb{R}^{km \times d}$ is the Jacobian of the network output with respect to the parameters.

Differentiating again, we obtain the Hessian:

$$\nabla^2 \mathcal{L}(\theta) = 2 \, \nabla \mathcal{F}(\theta)^\top \nabla \mathcal{F}(\theta) + 2 \sum_{i=1}^{km} (\mathcal{F}_i(\theta) - (y_{all})_i) \nabla^2 \mathcal{F}_i(\theta),$$

where $\nabla^2 \mathcal{F}_i(\theta) \in \mathbb{R}^{d \times d}$ is the Hessian of the *i*-th output component.

At $\theta \in \mathcal{Z}$, where $\mathcal{F}(\theta) = y_{all}$, the second term vanishes, simplifying the Hessian to:

$$\nabla^2 \mathcal{L}(\theta) = 2 \, \nabla \mathcal{F}(\theta)^\top \nabla \mathcal{F}(\theta).$$

For any direction $v \in \mathbb{R}^d$, this yields:

$$v^{\top} \nabla^2 \mathcal{L}(\theta) v = 2 \|\nabla \mathcal{F}(\theta) v\|^2.$$

Therefore,

$$\boldsymbol{v}^{\top} \, \nabla^2 \mathcal{L}(\boldsymbol{\theta}) \, \boldsymbol{v} = \boldsymbol{0} \quad \Leftrightarrow \quad \nabla \mathcal{F}(\boldsymbol{\theta}) \, \boldsymbol{v} = \boldsymbol{0}.$$

Moreover, every $\theta \in \mathcal{Z}$ is a local minimum where the Hessian matrix is symmetric and positive semi-definite. This gives the us equivalence

$$\nabla^2 \mathcal{L}(\theta) \ v = 0 \quad \Leftrightarrow \quad v^\top \ \nabla^2 \mathcal{L}(\theta) \ v = 0.$$

Therefore,

$$\nabla^2 \mathcal{L}(\theta) \, v = 0 \quad \Leftrightarrow \quad \nabla \mathcal{F}(\theta) \, v = 0.$$

Because $km \geq d$ and $\nabla \mathcal{F}(\theta)$ has full rank, the *inverse function theorem* implies that the preimage of \mathcal{F} locally has the structure of a smooth manifold whose tangent space is exactly the null space of $\nabla \mathcal{F}(\theta)$. To simplify our analysis, we directly restate the inverse function theorem below in a form that is slightly non-standard, but perfectly equivalent:

Theorem B.2 (Inverse Function Theorem). Assume that $\mathcal{F}: \mathbb{R}^d \to \mathbb{R}^{km}$ is a smooth function with $d \geq km$ and $\operatorname{rank}(\nabla \mathcal{F}(\theta)) = km$ for some $\theta \in \mathbb{R}^d$. Then, for all $v \in R^d$ such that $\nabla \mathcal{F}(\theta) v = 0$, there exists a smooth trajectory $s: R \to R^d$ such that $s(0) = \theta$, s'(0) = v, and $\mathcal{F}(s(t)) = \mathcal{F}(\theta)$ for all $t \in \mathbb{R}$.

Note that any such trajectory will also have $\mathcal{L}(s(t)) = 0$ since $\mathcal{F}(s(t)) = \mathcal{F}(\theta) = y_{all}$. This implies the desired result.

Part II: $v \in \mathcal{T}_{\theta} \Rightarrow \nabla^2 \mathcal{L}(\theta) \ v = 0.$

Using the equation $\mathcal{L}(f(t))=0$ from Definition 4.10 and differentiating it twice with respect to t, we obtain:

 $f'(t)^{\top} \nabla^2 \mathcal{L}(f(t)) f'(t) + \nabla \mathcal{L}(f(t))^{\top} f''(t) = 0.$ (23)

Since any point $\theta \in \mathcal{Z}$ is a local minimum, we have that $\nabla \mathcal{L}(\theta) = 0$ and $\nabla^2 \mathcal{L}(\theta)$ is symmetric and positive semi-definite (PSD).

By evaluating Equation (23) at t=0, we obtain that $v^{\top} \nabla^2 \mathcal{L}(\theta) v=0$. Since $\nabla^2 \mathcal{L}(\theta)$ is symmetric and PSD, this implies the desired result.

Definition B.3 (Normal Space). Let $\mathcal{N}_{\phi} = \{ \alpha \in \mathbb{R}^d \mid \alpha^{\top} v = 0, \ \forall v \in \mathcal{T}_{\phi} \}$ denote the normal space at a point $\phi \in \mathcal{Z}$.

Proposition B.4. The displacement of a point from the zero-loss set belongs to the normal space at the point's projection, i.e. $\theta - \operatorname{proj}_{\mathcal{Z}}(\theta) \in \mathcal{N}_{\operatorname{proj}_{\mathcal{Z}}(\theta)}$ for all $\theta \in \mathbb{R}^d$.

Proof. We give a proof by contradiction. Assume that there exists $v \in \mathcal{T}_{\mathrm{proj}_{\mathcal{Z}}(\theta)}$ such that $v^{\top}(\theta - \mathrm{proj}_{\mathcal{Z}}(\theta)) \neq 0$. Then, by Definition 4.10, there must exist a smooth trajectory $f : \mathbb{R} \to \mathbb{R}^d$ with $f(0) = \mathrm{proj}_{\mathcal{Z}}(\theta)$, f'(0) = v, and $f(t) \in \mathcal{Z}$ for all $t \in R$. By moving $\mathrm{proj}_{\mathcal{Z}}(\theta)$ along this trajectory, we can get closer to θ . However, this should not be possible according to Definition 4.12.

We now establish our main result:

Theorem 4.13 (Gradient Orthogonality). Let $S \subseteq \mathbb{R}^d$ be a compact space with $\operatorname{proj}_{\mathcal{Z}}(S \cap \mathcal{Z}) \subseteq S$. Assume that S contains no singular points. Then, there exist constants C > 0 and $x_0 > 0$ such that

$$\left| \frac{v^{\top}}{\|v\|} \frac{\nabla \mathcal{L}(\theta)}{\|\nabla \mathcal{L}(\theta)\|} \right| < C \operatorname{dist}_{\mathcal{Z}}(\theta)$$
 (6)

holds for all $\theta \in S \setminus \mathcal{Z}$ with $\operatorname{dist}_{\mathcal{Z}}(\theta) < x_0$ and $v \in \mathcal{T}_{\operatorname{proj}_{\mathcal{Z}}(\theta)}$.

Proof. By parameterizing θ as

$$\theta = \operatorname{proj}_{\mathcal{Z}}(\theta) + \|\theta - \operatorname{proj}_{\mathcal{Z}}(\theta)\| \frac{\theta - \operatorname{proj}_{\mathcal{Z}}(\theta)}{\|\theta - \operatorname{proj}_{\mathcal{Z}}(\theta)\|}$$

we can denote the quantity of interest as

$$\frac{\boldsymbol{v}^{\top} \nabla \mathcal{L}(\boldsymbol{\theta})}{\|\boldsymbol{v}\| \|\nabla \mathcal{L}(\boldsymbol{\theta})\|} = g \bigg(\mathrm{proj}_{\mathcal{Z}}(\boldsymbol{\theta}), \ \frac{\boldsymbol{\theta} - \mathrm{proj}_{\mathcal{Z}}(\boldsymbol{\theta})}{\|\boldsymbol{\theta} - \mathrm{proj}_{\mathcal{Z}}(\boldsymbol{\theta})\|}, \ \|\boldsymbol{\theta} - \mathrm{proj}_{\mathcal{Z}}(\boldsymbol{\theta})\|, \ \boldsymbol{v} \bigg)$$

where

$$g(\phi, \alpha, x, v) = \frac{v^{\top} \nabla \mathcal{L}(\phi + x\alpha)}{\|v\| \|\nabla \mathcal{L}(\phi + x\alpha)\|}$$

with $\phi \in \mathcal{Z}$, $\alpha \in \mathcal{N}_{\phi}$, $x \in \mathbb{R}^+$, and $v \in \mathcal{T}_{\phi}$.

To obtain the desired result, it suffices to show that there exists C > 0 such that

$$q(\phi, \alpha, x, v) < Cx$$

for all $\phi \in S \cap \mathcal{Z}$, $\alpha \in U(\mathcal{N}_{\phi})$, x > 0, $v \in \mathcal{T}_{\phi}$, where $U(\mathcal{N}_{\phi}) = \{ v \in \mathcal{N}_{\phi} \mid ||v|| = 1 \}$.

We write the gradient of the loss function around $\phi \in \mathcal{Z}$ using the Taylor expansion:

$$\nabla \mathcal{L}(\phi + x\alpha) = x \nabla^2 \mathcal{L}(\phi) \alpha + x h_{\phi,\alpha}(x)$$

where $h_{\phi,\alpha}(x)$ is a remainder term that vanishes as $x \to 0$. In other words, there exist constants $M_{\phi,\alpha}$, $a_{\phi,\alpha} > 0$ such that

$$||h_{\phi,\alpha}(x)|| < M_{\phi,\alpha} x \qquad \forall x, x < a_{\phi,\alpha}.$$

Since S is closed and \mathcal{L} is continuous, $S \cap \mathcal{Z}$ will also be closed. The set of normal vectors with unit norm at any point is also closed. This implies that the following quantities exist and are positive:

$$M_{\sup} = \sup_{\substack{\phi \in S \cap \mathcal{Z} \\ \alpha \in U(\mathcal{N}_{\psi})}} M_{\phi,\alpha} \qquad a_{\inf} = \inf_{\substack{\phi \in S \cap \mathcal{Z} \\ \alpha \in U(\mathcal{N}_{\psi})}} a_{\phi,\alpha}$$

We express $g(\theta, \alpha, x, v)$ as:

$$g(\phi, \alpha, x, v) = \frac{x \, v^\top \, \nabla^2 \mathcal{L}(\phi) \, \alpha + x \, v^\top \, h_{\phi, \alpha}(x)}{\|v\| \, \|x \, \nabla^2 \mathcal{L}(\phi) \, \alpha + x \, h_{\phi, \alpha}(x)\|} = \frac{v^\top \, \nabla^2 \mathcal{L}(\phi) \, \alpha + v^\top \, h_{\phi, \alpha}(x)}{\|v\| \, \|\nabla^2 \mathcal{L}(\phi) \, \alpha + h_{\phi, \alpha}(x)\|}$$

Since $v \in \mathcal{T}_{\phi}$, from Lemma B.1, we have that $v^{\top} \nabla^2 \mathcal{L}(\phi) \alpha = 0$. This gives

$$g(\phi, \alpha, x, v) = \frac{v^{\top} h_{\phi, \alpha}(x)}{\|v\| \|\nabla^{2} \mathcal{L}(\phi) \alpha + h_{\phi, \alpha}(x)\|}$$
$$\leq \frac{\|h_{\phi, \alpha}(x)\|}{\|\nabla^{2} \mathcal{L}(\phi) \alpha + h_{\phi, \alpha}(x)\|}$$

Since $\nabla^2 \mathcal{L}(\phi) \alpha \neq 0$, the following also exists and is positive:

$$\lambda_{\inf} = \inf_{\substack{\phi \in S \cap \mathcal{Z} \\ \alpha \in U(\mathcal{N}_{\psi})}} \|\nabla^{2} \mathcal{L}(\phi) \alpha\|$$

Assuming that $x < \lambda_{\inf}/M_{\sup}$ guarantees that $\|\nabla^2 \mathcal{L}(\phi) \alpha\| > \|h_{\phi,\alpha}(x)\|$, which gives

$$g(\phi, \alpha, x, v) \le \frac{\|h_{\phi,\alpha}(x)\|}{\|\nabla^2 \mathcal{L}(\phi) \alpha\| - \|h_{\phi,\alpha}(x)\|}$$
$$\le \frac{M_{\sup} x}{\lambda_{\inf} - M_{\sup} x}$$

for any $x < a_{\inf}$. An appropriate choice of C and x_0 gives the desired bound, for example:

$$x_0 = \min\left(a_{\rm inf}, \frac{\lambda_{\rm inf}}{2M_{\rm sup}}\right) \qquad C = \frac{2M_{\rm sup}}{\lambda_{\rm inf}}$$

C COST FUNCTION GRADIENT IN TWO-LAYER NETWORKS

We want to derive the gradient of the following cost function:

$$\mathcal{R}(W_1) = \lambda \|W_1\|_F^2 + \lambda \|\phi(W_1)\|_F^2$$

where

$$\phi(W_1) = H^{\top}(HH^{\top})^{-1}Y \qquad H = \sigma(XW_1) \qquad X \in \mathbb{R}^{n \times d_{\text{in}}} \qquad W_1 \in \mathbb{R}^{d_{\text{in}} \times d_h} \qquad Y \in \mathbb{R}^{n \times d_{\text{out}}}$$

The activation function $\sigma: \mathbb{R} \to \mathbb{R}$ is applied is applied elementwise. Note that, since we work in the zero-loss approximation, there is no explicit error term.

We decompose $\mathcal{R}(W_1)$ into two terms:

$$\mathcal{R}(W_1) = \underbrace{\lambda \|W_1\|_F^2}_{\text{Term 1}} + \underbrace{\lambda \|\phi(W_1)\|_F^2}_{\text{Term 2}}$$

Since the first term is a standard Frobenius norm squared, its gradient is:

$$\nabla_{W_1} \left[\lambda \|W_1\|_F^2 \right] = 2 \lambda W_1$$

To find the gradient of the second term, we analyze:

$$f(W_1) = \|\phi(W_1)\|_F^2 = \|H^\top (HH^\top)^{-1} Y\|_F^2$$

Since the Frobenius norm squared satisfies $||M||_F^2 = \text{Tr}(M^\top M)$, we write:

$$f(W_1) = \operatorname{Tr}\left(\left(H^{\top}(HH^{\top})^{-1}Y\right)^{\top}H^{\top}(HH^{\top})^{-1}Y\right)$$
$$= \operatorname{Tr}\left(Y^{\top}(HH^{\top})^{-1}HH^{\top}(HH^{\top})^{-1}Y\right)$$
$$= \operatorname{Tr}\left(YY^{\top}(HH^{\top})^{-1}\right)$$

Using the following known result from matrix calculus:

$$\frac{\partial \operatorname{Tr}(MP^{-1})}{\partial P} = -P^{-1}MP^{-1}$$

with $M = YY^{\top}$ and $P = HH^{\top}$, we obtain:

$$\frac{\partial f}{\partial H} = -2(HH^{\top})^{-1}YY^{\top}(HH^{\top})^{-1}H$$

Propagating the gradient using the chain rule, we get:

$$\frac{\partial f}{\partial W_1} = -2X^{\top} \left[\left((HH^{\top})^{-1} Y Y^{\top} (HH^{\top})^{-1} H \right) \odot \sigma'(XW_1) \right]$$

where \odot denotes the Hadamard product.

Finally, multiplying by λ , we obtain the gradient of the second term:

$$\nabla_{W_1} \left[\lambda \| \phi(W_1) \|_F^2 \right] = -2\lambda X^{\top} \left[(AYY^{\top}AH) \odot \sigma'(XW_1) \right]$$

where $A = (HH^{\top})^{-1}$.

Thus, the final expression is:

$$\nabla \mathcal{R}(W_1) = -2\lambda W_1 + 2\lambda X^{\top} \Big[(AYY^{\top}AH) \odot \sigma'(XW_1) \Big]$$