

ZerO Initialization: Initializing Neural Networks with only Zeros and Ones

Jiawei Zhao

California Institute of Technology

jiawei@caltech.edu

Florian Schäfer

Georgia Institute of Technology

florian.schaefer@cc.gatech.edu

Anima Anandkumar

California Institute of Technology
NVIDIA

anima@caltech.edu

Reviewed on OpenReview: <https://openreview.net/forum?id=1AxQpKmiTc>

Abstract

Deep neural networks are usually initialized with random weights, with adequately selected initial variance to ensure stable signal propagation during training. However, selecting the appropriate variance becomes challenging especially as the number of layers grows. In this work, we replace random weight initialization with a *fully deterministic* initialization scheme, viz., ZerO, which initializes the weights of networks with only *zeros and ones* (up to a normalization factor), based on identity and Hadamard transforms. Through both theoretical and empirical studies, we demonstrate that ZerO is able to train networks without damaging their expressivity. Applying ZerO on ResNet achieves state-of-the-art performance on various datasets, including ImageNet, which suggests random weights may be unnecessary for network initialization. In addition, ZerO has many benefits, such as training ultra deep networks (without batch-normalization), exhibiting low-rank learning trajectories that result in low-rank and sparse solutions, and improving training reproducibility¹.

1 Introduction

An important question in training deep neural networks is how to initialize the weights. Currently, random weight initialization is the de-facto practice across all architectures and tasks. However, choosing the *variance* of the initial weight distribution is a delicate balance when training deep neural networks. If the variance is too large, it can lead to an excessive amplification of the activations propagating through the network during training, resulting in *exploding gradients*. On the other hand, if the weights are initialized too small, the activations may not propagate at all, resulting in *vanishing gradients*. These issues become more challenging as the number of layers in the network grows.

The above challenges can be avoided if *identity initialization* is used instead. It initializes each layer in the network as an identity mapping, such that the input data can be identically propagated to the network output. In this case, there is no need to introduce any randomness or consider its variance. Identity initialization is well studied theoretically from an optimization perspective. Hardt & Ma (2017) prove the existence of a global minimum close to the identity parameterization in a deep residual network. Bartlett et al. (2019) further prove the rate of convergence of gradient-based optimization under identity initialization.

¹Code repository: <https://github.com/jiaweizhao/Zer0-initialization>.

Table 1: Comparing ZerO initialization with randomly perturbed identity initialization. F represents a transformation¹ of an arbitrary layer l .

Settings	Techniques	Related Works
$\mathbf{x}_{l+1} = \mathbf{x}_l + F(\mathbf{x}_l)$	F is randomly initialized with small variance	Hardt & Ma (2017)
$\mathbf{x}_{l+1} = \mathbf{x}_l + \text{Conv}(F(\mathbf{x}_l))$	F is randomly initialized, kernel in Conv is zero	Zhang et al. (2019)
$\mathbf{x}_{l+1} = \mathbf{x}_l + \text{Norm}(F(\mathbf{x}_l))$	F is randomly initialized, scale in Norm is zero	Goyal et al. (2017)
$\mathbf{x}_{l+1} = \mathbf{x}_l + \alpha(F(\mathbf{x}_l))$	F is randomly initialized, scalar α is zero	Bachlechner et al. (2020)
All settings above	F is a Hadamard/identity transform ²	ZerO (ours)

However, prior theoretical works on identity initialization assume that all layers had the same dimensionality, which does not hold for practical networks. Typically, practical networks have *varying dimensionality* across layers, such as the variations of spatial and channel dimensions in ResNet architectures (He et al., 2016). Directly applying identity initialization to these networks leads to a problem of *training degeneracy*, as our theoretical study will demonstrate later.

To avoid the training degeneracy, previous works (summarized in Table 1) employ identity initialization with *random perturbations* to facilitate escapes from a saddle point or to break feature symmetry (Blumenfeld et al., 2020). Broadly, these approaches satisfy the property of *dynamical isometry*, to preserve the signal propagation and ensure well-behaved gradients at initialization (Saxe et al., 2014). Despite the efficacy of random perturbations, they inevitably introduce additional tuning of variances, which can result in gradient explosion in deep networks without careful tuning.

Our work: we propose ZerO initialization that removes *all randomness in the weight initialization*. As illustrated in Figure 1, ZerO initializes networks with Hadamard and identity transforms, which assigns all the weights to only *zeros and ones*.




Figure 1: Illustrating ZerO in a 3-layer network with input dimension n_x , hidden dimension n_h , and output dimension n_y , where $n_h > n_x, n_y$. \mathbf{H} and \mathbf{I} are $n_h \times n_h$ Hadamard and identity matrix, respectively. The dimension-increasing layer is initialized by columns of the Hadamard matrix. The rest layers are initialized by identity matrix or rows of it.

ZerO is not affected by the problem of training degeneracy or accuracy loss. Compared to random initialization, ZerO provides state-of-the-art performance over various image classification tasks, including ImageNet. We further discover many unique properties and benefits of ZerO:

Stable training without batch normalization. ZerO ensures well-behaved signal propagation, which provides stable training without batch normalization. Testing ResNet with over 500 layers, we find ZerO

¹The transformation may contain linear, nonlinear or convolutional operations.

²Both identity and Hadamard transforms are deterministic parameterizations.

converges faster than carefully designed random methods, such as Fixup and ReZero (Zhang et al., 2019; Bachlechner et al., 2020).

Low-rank learning trajectory. We find that ZerO exhibits a low-rank learning trajectory, where the rank of each matrix gradually increases during training. We believe this is the first time that the greedy low-rank learning (GLRL) trajectory, a theoretical characterization of gradient descent, has been observed in large-scale deep learning applications. GLRL is a consequence of implicit rank regularization by gradient descent under infinitesimal initialization (Li et al., 2021; Razin et al., 2021). It can be viewed as performing a rank-constrained optimization and greedily relaxing the rank restriction by one whenever it fails to reach a global minimizer. GLRL has been used to explain the excellent generalization in gradient-based deep learning, as it converges to a global (or local) minima with the minimum rank. However, the GLRL trajectory has never been observed in practice due to its impractical requirement of infinitesimal initialization.

Sparse and low-rank solutions. We observe that ZerO-initialized networks converge to sparse and low-rank solutions. Compared to randomly initialized networks, the sub-networks obtained in trained ZerO-initialized networks achieve 30% lower (matrix or tensor) rank or 25% higher sparsity without sacrificing accuracy.

Better training reproducibility. Since ZerO does not require any random perturbations, it is a fully deterministic initialization scheme. Unlike determinism in random initialization, which needs fixing pseudorandom number generators in hardware, the weights initialized by ZerO are fixed regardless of how the random seed varies or which hardware is used. ZerO significantly reduces the training variation and thus achieves better training reproducibility (the remaining randomness is only due to batch selection). Compared to random initialization, ZerO produces 20%-40% lower standard deviation of the final accuracy over repeated experiments with different random seeds.

Theoretical analysis of ZerO. Theoretically, we demonstrate that ZerO breaks a training degeneracy that arises when applying identity initialization to networks with varying dimensionality across layers. We prove that the training degeneracy necessarily occurs in standard identity initialization because the rank of any $n_h \times n_h$ matrix in the network is upper bounded by the input and output dimensions n_x and n_y throughout the entire training, no matter how large the size of n_h is. This limits the expressivity of each matrix, resulting in the degeneracy of training.

Our contributions are summarized as follows:

1. We design ZerO initialization, the first fully deterministic initialization that achieves state-of-the-art performance in practice.
2. ZerO is backed with theoretical understanding. As shown in Theorem 1, we prove how ZerO breaks the training degeneracy by applying Hadamard transforms.
3. ZerO has many benefits, such as training ultra deep networks (without batch-normalization), exhibiting low-rank learning trajectory, converging to sparse and low-rank solutions, and improving training reproducibility.

2 Is Randomness Necessary in Identity Initialization?

2.1 Background

We wish to train a function $\mathcal{F}(\mathbf{x})$ to learn a particular input-output map given a set of P training samples $(\mathbf{x}^\mu, \mathbf{y}^\mu) \in \mathbb{R}^{n_x \times n_y}$, where $\mu = 1, \dots, P$. Training is accomplished by minimizing the squared error $\mathcal{L} = \frac{1}{2} \sum_{\mu=1}^P \|\mathbf{y}^\mu - \mathcal{F}(\mathbf{x}^\mu)\|_2^2$ using gradient descent with a step size η .

We model $\mathcal{F}(\mathbf{x})$ to be a multilayer perceptron with $L > 2$ hidden layers, such that:

$$\mathbf{x}_l = \mathbf{W}_l \mathbf{z}_{l-1} \quad \mathbf{z}_l = \varphi(\mathbf{x}_l),$$

with $l \in 1, \dots, L$. Let $\mathbf{z}_0 = \mathbf{x}$ and $\mathcal{F}(\mathbf{x}) = \mathbf{z}_L$. φ is an element-wise nonlinearity. We assume that \mathcal{F} has uniform hidden dimension n_h , with $\mathbf{W}_l \in \mathbb{R}^{n_h \times n_h}$ for $l \in 2, \dots, L-1$, $\mathbf{W}_1 \in \mathbb{R}^{n_h \times n_x}$, and $\mathbf{W}_L \in \mathbb{R}^{n_y \times n_h}$.

The input-output Jacobian is a well-studied proxy for estimating the stability of signal propagation at initialization, which is defined as: $\mathbf{J}_{io} = \frac{\partial \mathbf{z}_L}{\partial \mathbf{z}_0}$. Proposed by (Saxe et al., 2014), dynamical isometry is a condition where all singular values of the Jacobian \mathbf{J}_{io} concentrate near 1. If \mathbf{J}_{io} is well-conditioned, the backpropagation error $\frac{\partial \mathcal{L}}{\partial \mathbf{z}_l}$ at any layer l will be well-conditioned as well. This ensures stable signal propagation and well-behaved gradients at initialization.

Consider the case of $n_x = n_y = n_h$. Identity initialization is defined as initializing each matrix to be an identity matrix: $\mathbf{W}_l = \mathbf{I}$. In this case, the dynamical isometry property for a linear \mathcal{F} can be easily verified as $\mathbf{J}_{io} = \mathbf{I}$. It also holds for certain nonlinearities when applying the identity initialization on residual networks: $\mathbf{z}_l = \varphi(\mathbf{x}_l) + \mathbf{x}_{l-1}$ where $\mathbf{W}_l = \mathbf{0}$, such that no signal is passed through the residual branches at initialization. Table 1 lists a few related examples.

From an optimization perspective, Hardt & Ma (2017) suggest that \mathcal{F} has a global minimum very close to its identity initialization, such that $\max_{1 \leq l \leq L} \|\mathbf{W}'_l\| \leq O(1/L)$ for large L , where $\mathbf{W}' = \mathbf{W} - \mathbf{I}$. Bartlett et al. (2019) also proves that under the identity initialization, gradient descent learns an ϵ -approximation of \mathcal{F} within iterations polynomial in $\log(1/\epsilon)$.

2.2 Extending to Large Hidden Dimension

So far, we have discussed identity initialization in the special case of fixed dimensionality. Now we extend our discussion to a more practical setting where \mathcal{F} is equipped with a large hidden dimension, such that $n_h \gg n_x, n_y$. We also focus on the specific case where φ is a Relu nonlinearity.

A straightforward approach to identity initialization in this case is to initialize \mathbf{W}_1 such that it projects the input into a n_x -dimensional subspace of the n_h -dimensional hidden space. This can be achieved by initializing \mathbf{W}_1 with a *partial identity matrix*:

Definition 1 (Partial Identity Matrix). Let $\mathbf{I}^* \in \mathbb{R}^{l \times r}$, the partial identity matrix \mathbf{I}^* is defined as follows:

$$\mathbf{I}^* = \begin{cases} (\mathbf{I}, \mathbf{0}) & \text{where } \mathbf{I} \in \mathbb{R}^{l,l} \text{ and } \mathbf{0} \in \mathbb{R}^{l,r-l} & \text{if } l < r, \\ \mathbf{I} & \text{where } \mathbf{I} \in \mathbb{R}^{l,l} & \text{if } l = r, \\ (\mathbf{I}, \mathbf{0})^\top & \text{where } \mathbf{I} \in \mathbb{R}^{r,r} \text{ and } \mathbf{0} \in \mathbb{R}^{r,l-r} & \text{otherwise.} \end{cases}$$

For a vector $\mathbf{a} \in \mathbb{R}^r$, if $l < r$, then $\mathbf{I}^*(\mathbf{a})$ clips the last few dimension such that $\mathbf{I}^*(\mathbf{a}) = (a_1, a_2, \dots, a_l)^\top$. If $l > r$, then \mathbf{I}^* pads $l - r$ additional dimensions with zero, such that $(a_1, a_2, \dots, a_r, 0 \dots 0)^\top$. This is also known as a zero-padding operator, such as used in channel-expanding layers in ResNet (He et al., 2016). In the network \mathcal{F} , \mathbf{I}^* only needs to be applied in the dimension-varying matrices \mathbf{W}_1 and \mathbf{W}_L , while leaving the remaining $n_h \times n_h$ matrices to be identity matrix \mathbf{I} .

We visualize this process in Figure 2. This may seem like a natural extension of identity initialization to a large width setting, but we will show in Section 2.3 it suffers from a problem we call ‘‘training degeneracy’’. To avoid the problem, we use the Hadamard matrix \mathbf{H} to initialize the dimension-increasing matrices, such that $\mathbf{W}_1 = \mathbf{H}\mathbf{I}^*$. A Hadamard matrix is defined as follows:

Definition 2 (Hadamard matrix). For any Hadamard matrix $\mathbf{H} = \mathbf{H}_m \in \mathbb{R}^{2^m \times 2^m}$ where m is a positive integer, we define $\mathbf{H}_0 = \mathbf{1}$ by the identity, and the matrix with large m is defined recursively:

$$\mathbf{H}_m = \begin{pmatrix} \mathbf{H}_{m-1} & \mathbf{H}_{m-1} \\ \mathbf{H}_{m-1} & -\mathbf{H}_{m-1} \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 & 1 & \dots \\ 1 & -1 & 1 & -1 & \dots \\ 1 & 1 & -1 & -1 & \dots \\ 1 & -1 & -1 & 1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \in \mathbb{R}^{2^m \times 2^m}.$$




Figure 2: A 3-layer network \mathcal{F} ($n_h > n_x = n_y$) where $\mathbf{W}_1, \mathbf{W}_3 = \mathbf{I}^*$ and $\mathbf{W}_2 = \mathbf{I}$ at initialization.

The linear transformation described by the Hadamard matrix, called the Hadamard transform, rotates the coordinate axes to be equally weakly aligned with the standard basis. For example, in a two-dimensional plane, the Hadamard transform rotates the standard basis by an exact angle of 45 degree. This turns out to be an important property for breaking the training degeneracy.

2.3 Identity Initialization limits Network Expressivity

We now present our main result differentiating the training behavior of different initialization methods and describing the problem of training degeneracy.

Theorem 1. *Let \mathcal{F} be a neural network with L matrices, where $\mathbf{W}_l \in \mathbb{R}^{n_h \times n_h}$ for $l \in 2, \dots, L-1$, $\mathbf{W}_1 \in \mathbb{R}^{n_h \times n_x}$, and $\mathbf{W}_L \in \mathbb{R}^{n_y \times n_h}$. \mathcal{F} has a uniform hidden dimension n_h , input dimension n_x , and output dimension n_y , where $n_h \geq n_x, n_y$. Define residual component $\mathbf{W}'_l = \mathbf{W}_l - \mathbf{I}$. Let $\mathbf{z}_l(\mathbf{x})$ to be the activation in the l -th layer under the input \mathbf{x} . Then we have the following results for different initializations:*

- (i) *Consider a random perturbation $\mu \in \mathbb{R}^{n_h \times n_h}$ where each element is sampled from a Gaussian distribution: $\mu_{ij} \sim \mathcal{N}(0, \sigma^2)$. It is well-known that the randomly perturbed matrices $\mathbf{W}_l = \mathbf{I} + \mu_l$ (for $l \neq 1, L$) are full-rank almost surely:*

$$\text{Prob}(\text{rank}(\mathbf{W}'_l) = n_h) = 1 \quad \text{for } l \in 2, \dots, L-1. \quad (1)$$

- (ii) *When initializing $\mathbf{W}_1, \mathbf{W}_L = \mathbf{I}^*$ and the remaining matrices as $\mathbf{W}_l = \mathbf{I}$ (for $l \neq 1, L$), for any $\mathbf{x} \in \mathbb{R}^{n_x}$, the linear dimension of the set of all possible activations is bounded throughout training as*

$$\dim(\text{span}(\{\mathbf{z}_l(\mathbf{x}) | \mathbf{x} \in \mathbb{R}^{n_x}\})) \leq n_x \quad \text{for } l \in 2, \dots, L-1. \quad (2)$$

As a result, the ranks of the weight matrices remain bounded throughout training as

$$\text{rank}(\mathbf{W}'_l) \leq n_x \quad \text{for } l \in 2, \dots, L-1. \quad (3)$$

- (iii) *When initializing $\mathbf{W}_1 = \mathbf{H}\mathbf{I}^*$, $\mathbf{W}_L = \mathbf{I}^*$, and the remaining matrices as $\mathbf{W}_l = \mathbf{I}$ (for $l \neq 1, L$) it is possible for the activations at an intermediate layer to attain*

$$\dim(\text{span}(\{\mathbf{z}_l(\mathbf{x}) | \mathbf{x} \in \mathbb{R}^{n_x}\})) > n_x, \quad (4)$$

breaking the constraint on the linear dimension described in Equation 2.

Remark 1. *The constraints on linear dimensions and matrix ranks in Equation 2 and 3 suggest that no matter how large hidden dimension n_h is, the network \mathcal{F} is only optimized within a low-dimensional subspace (depending on input dimension n_x) of the full parameter space. This restricts the maximum network expressivity of \mathcal{F} , and thus the training may only converge to an underfitting regime, leading to **training degeneracy**.*

Remark 2. *Under the assumptions of (iii), the breaking of training degeneracy described in Equation 4 appears to be the generic case. As verified empirically in Figure 3, applying the Hadamard transform in (ii) also breaks the rank constraint in Equation 3.*

Remark 3. (i) and (iii) avoid the training degeneracy from different directions. Unlike (i), the proposed (iii) doesn't introduce any randomness with the help of the Hadamard transform.

Almost all existing works use random weight initialization, which largely affects the rank of each matrix as shown in (i). (i) can be proved by showing any column (or row) in a random matrix is linearly independent to the other columns (or rows), almost surely. A detailed proof can be found in the appendix.

Consider identity initialization without any randomness. In (ii), \mathbf{W}_1 identically maps the input \mathbf{x} into a subspace of the hidden layer \mathbf{z}_1 , such that $\mathbf{z}_1 = (\mathbf{x}^\top, 0, \dots, 0)^\top$. Thus, the linear dimension on \mathbf{z}_l (i.e., the linear dimension on activations $\mathbf{z}_l^1, \dots, \mathbf{z}_l^P$) is equivalent to the linear dimension on \mathbf{x} , which is upper bounded by n_x . This result is held for every layer \mathbf{z}_l (where $l \neq 1, L$).




Figure 3: Verify Theorem 1 in practice. We train a network \mathcal{F} with $L = 3$ on MNIST, and visualize $\text{rank}(W_2)$ over the training. The red dash line denotes the rank constraint on W_2 predicted by the theorem, which is $n_x = 784$. **Left:** we verify (ii) in the theorem by varying n_h . No matter how large n_h is, $\text{rank}(W_2)$ follows the rank constraint through the entire training. **Right:** we verify (iii) where applying Hadamard transform breaks the rank constraint introduced in (ii), given $n_h = 2048$. We denote the initializations in (ii) and (iii) as standard and Hadamard-based GI-Init, respectively. As predicted in (i), random initialization achieves its maximum rank immediately after the initialization.

To show the rank constraint in Equation 3, we track a single derivative of the residual component at layer l :

$$\frac{\partial \mathcal{L}}{\partial \mathbf{W}'_l} = \sum_{\mu=1}^P \frac{\partial \mathcal{L}}{\partial \mathbf{x}'_l^\mu} \otimes \mathbf{z}_{l-1}^\mu, \quad (5)$$

where \otimes denotes the outer product. We use the following well-known fact:

Lemma 1. Consider a matrix \mathbf{M} to be a sum of vector outer products: $\mathbf{M} = \sum_{\mu=1}^Q \mathbf{a}^\mu \otimes \mathbf{b}^\mu$, where $\mathbf{a}^\mu \in \mathbb{R}^{n_a}$ and $\mathbf{b}^\mu \in \mathbb{R}^{n_b}$ for $\mu \in 1, \dots, Q$. Let $Q > n_a, n_b$. If linear dimensions $\dim(\text{span}(\mathbf{a}^1, \dots, \mathbf{a}^Q)) \leq U$ and $\dim(\text{span}(\mathbf{b}^1, \dots, \mathbf{b}^Q)) \leq V$, where $U \leq n_a$ and $V \leq n_b$, then:

$$\text{rank}(\mathbf{W}) \leq \min(U, V)$$

By Lemma 1, at initialization, the upper bound n_x on the linear dimension on \mathbf{z}_{l-1} results in a rank constraint on $\text{rank}(\mathbf{W}'_l)$. The rank constraint holds during the entire training as $\frac{\partial \mathcal{L}}{\partial \mathbf{W}'_l}$ has a zero-valued $n_y \times (n_h - n_x)$ sub-matrix at every iteration (as shown in the appendix). Since $\mathbf{W}'_l = \sum_{t=1}^T -\eta \frac{\partial \mathcal{L}}{\partial \mathbf{W}'_l} \Big|_t$ after T weight updates (by gradient descent with a step size η), $\text{rank}(\mathbf{W}'_l)$ is bounded by n_x no matter what T is. This results in the training degeneracy as described in Remark 1. We also verify it empirically in Figure 3.

To avoid the training degeneracy, we need to overcome limitations on the linear dimension of the set of possible activations. This is indeed possible when using the Hadamard matrix as $\mathbf{W}_1 = \mathbf{H}\mathbf{I}^*$, as we will illustrate by means of an example.

Lemma 2. Assume $n_h = 4$ and $n_x = 3$. For any vector $\mathbf{x} \in \text{span}(\mathbf{e}_2, \mathbf{e}_3)$ where \mathbf{e}_2 and \mathbf{e}_3 are coordinate vectors $(0, 1, 0)^\top$ and $(0, 0, 1)^\top$, it holds that:

$$\text{span}(\{\mathbf{z}_1(\mathbf{x}) | \mathbf{x} \in \mathbb{R}^{n_x}\}) = \text{span}(\text{Relu}(\mathbf{H}\mathbf{I}^* \mathbf{e}_2), \text{Relu}(-\mathbf{H}\mathbf{I}^* \mathbf{e}_2), \text{Relu}(\mathbf{H}\mathbf{I}^* \mathbf{e}_3), \text{Relu}(-\mathbf{H}\mathbf{I}^* \mathbf{e}_3)), \quad (6)$$

where $\text{Relu}(\mathbf{H}\mathbf{I}^* \mathbf{e}_2)$, $\text{Relu}(-\mathbf{H}\mathbf{I}^* \mathbf{e}_2)$, $\text{Relu}(\mathbf{H}\mathbf{I}^* \mathbf{e}_3)$, and $\text{Relu}(-\mathbf{H}\mathbf{I}^* \mathbf{e}_3)$ are linearly independent. This indicates that:

$$\dim(\text{span}(\{\mathbf{z}_1(\mathbf{x}) | \mathbf{x} \in \mathbb{R}^{n_x}\})) = 4 = n_h > n_x = 3.$$

When using Hadamard matrix as $\mathbf{W}_1 = \mathbf{H}\mathbf{I}^*$, the breaking of the training degeneracy described in Lemma 2 appears to be the generic case. As verified empirically in Figure 3, this also breaks the rank constraint in Equation 3.

We point out that the increase of the linear dimension of the set of possible \mathbf{z}_l is only possible due to the nonlinearity. If \mathcal{F} is a linear network, the linear dimension on every \mathbf{z}_l is at most n_x , no matter how the weights are initialized.

Nevertheless, the nonlinearity can not increase the linear dimensionality if we initialize the network with a partial identity matrix. This is because when $\mathbf{W}_1 = \mathbf{I}^*$, $\text{span}(\{\mathbf{z}_l(\mathbf{x})|\mathbf{x} \in \mathbb{R}^{n_x}\})$ is aligned with the standard basis, i.e., each vector in the span at least has $n_h - n_x$ zero coefficients when expressed in the standard basis. Thus, an element-wise nonlinearity can not increase the linear dimension of its input beyond n_x .

To break the alignment $\text{span}(\{\mathbf{z}_l(\mathbf{x})|\mathbf{x} \in \mathbb{R}^{n_x}\})$ with the standard basis, we use the Hadamard transform. This is because it transforms the subspace such that the new basis is equally weakly aligned with the standard basis. We note that other linear transforms may also detach the subspace from the standard basis, but the Hadamard transform is the most natural choice.

3 ZerO Initialization

The initialization analyzed in (iii) of Theorem 1 is based on a network condition in which all hidden spaces $\mathbf{z}_1, \dots, \mathbf{z}_{L-1}$ have the same dimension n_h . Motivated by our theoretical understanding, we propose ZerO initialization, which initializes the weights of any network with arbitrary hidden dimensions. As described in Algorithm 1, ZerO only initializes dimensional-increasing layers with Hadamard matrices to avoid the training degeneracy. Other layers are simply initialized by (partial) identity matrices. We also rescale Hadamard matrices by a normalization factor $2^{-(m-1)/2}$, resulting in an orthonormal Hadamard transform.

Algorithm 1 ZerO Initialization.

Input: a neural network \mathcal{F} with L matrices $\mathbf{W}_l \in \mathbb{R}^{P_l \times Q_l}$ for l in $1, \dots, L$. \mathbf{I}^* is partial identity matrix defined in Definition 1. \mathbf{H}_m is the Hadamard matrix defined in Definition 2.

For l **in** $1, \dots, L$:

If $P_l = Q_l$: $\mathbf{W}_l \leftarrow \mathbf{I}$ ▷ Identity mapping

If $P_l < Q_l$: $\mathbf{W}_l \leftarrow \mathbf{I}^*$ ▷ Propagate the first P_l dimensions

If $P_l > Q_l$: $\mathbf{W}_l \leftarrow c\mathbf{I}^*\mathbf{H}_m\mathbf{I}^*$, where $m = \lceil \log_2(P_l) \rceil$ and $c = 2^{-(m-1)/2}$ ▷ Apply Hadamard matrix

We also apply ZerO to the well-developed ResNet architectures in He et al. (2016). As shown in Algorithm 2, we apply ZerO to convolution in a similar way by considering the variation in channel dimensions. When \mathbf{K} is a 1x1 convolution, \mathbf{K} also can be viewed a $c_{out} \times c_{in}$ matrix, which matches the initialization in Algorithm 1. We note that Algorithm 2 can be applied to every convolution in ResNet, including the first 3x3 convolution, 1x1 convolutions in spatial-downsampling skip connections, and convolutions in basic block and bottleneck block.

To achieve dynamical isometry at initialization, we apply a common technique that initializes the last convolution in each residual block as zero. This helps suppress the signals from residual branches to stabilize signal propagations, as studied in Zhang et al. (2019); Bachlechner et al. (2020).

Algorithm 2 ZerO Initialization on Convolution.

Input: number of input channels c_{in} , number of output channels c_{out} , odd kernel size k .

Return: a $c_{out} \times c_{in} \times k \times k$ convolutional kernel \mathbf{K} .

Let $n \leftarrow \lfloor k/2 \rfloor$

If $c_{out} = c_{in}$: $\mathbf{K}[:, :, n, n] \leftarrow \mathbf{I}$

If $c_{out} < c_{in}$: $\mathbf{K}[:, :, n, n] \leftarrow \mathbf{I}^*$

If $c_{out} > c_{in}$: $\mathbf{K}[:, :, n, n] \leftarrow c\mathbf{I}^*\mathbf{H}_m\mathbf{I}^*$, where $m = \lceil \log_2(P_l) \rceil$ and $c = 2^{-(m-1)/2}$

We also apply ZerO to networks with or without batch normalization. For ResNet with batch normalization, we follow the standard practice to initialize the scale and bias in batch normalization as one and zero,

Dataset	Model	Initialization	Test Error (mean \pm std)
CIFAR-10	ResNet-18	ZerO Init	5.13 \pm 0.08
		Kaiming Init	5.15 \pm 0.13
		Xavier Init	5.23 \pm 0.16
ImageNet	ResNet-50	ZerO Init	23.43 \pm 0.04
		Kaiming Init	23.46 \pm 0.07
		Xavier Init	23.65 \pm 0.11

Table 2: Benchmarking ZerO on CIFAR-10 and ImageNet. We repeat each run 10 times with different random seeds.

respectively. For training without batch normalization, we adopt a technique proposed by Zhang et al. (2019), where the batch normalization is replaced by learnable scalar multipliers and biases.

4 Experiments

In this section, we empirically benchmark ZerO on CIFAR-10 and ImageNet datasets, where we evaluate ResNet-18 on CIFAR-10 and ResNet-50 on ImageNet (Krizhevsky, 2009; Deng et al., 2009). Both ResNet structures follow the design from He et al. (2016), which includes batch normalization by default.

Hyperparameter settings. We find that *ZerO can fully utilize the default hyperparameters*, which include a learning rate of 0.1, a momentum of 0.9, and a weight decay of 0.0001. In addition, we observe the learning rate warmup is essential for ZerO to achieve a large maximal learning rate, as most of the weights start from the exact zero. We warm up the learning rate with 5 and 10 epochs for ImageNet and CIFAR-10, respectively.

We present our main results that compare different initialization schemes. For each dataset, all experiments use the same hyperparameter settings by default. Each experiment is repeated for ten runs with different random seeds. As shown in Table 2, ZerO achieves state-of-the-art accuracy on both datasets compared to other random methods.

In addition, we compare ZerO with a broad range of related works on CIFAR-10 using ResNet-18 and ResNet-50, including ReZerO (Bachlechner et al., 2020), Fixup (Zhang et al., 2019), SkipInit (De & Smith, 2020) and ConstNet (Blumenfeld et al., 2020). As shown in Table 3, ZerO consistently achieves top performance compared to other methods.

We note that the ConstNet proposed by Blumenfeld et al. (2020) is also a deterministic initialization. However, unlike ZerO which preserves feature diversity, ConstNet is designed to eliminate the diversity by averaging the features through layers. The feature symmetric problem in ConstNet causes significant degradation, and additional random noise (e.g., non-deterministic GPU operation and dropout) is needed to break the symmetry.

Method	ZerO	ReZero	Fixup	SkipInit	ConstNet	ConstNet*
ResNet-18	5.13	5.20	5.17	5.26	72.39	5.41
ResNet-50	4.53	4.72	4.51	4.63	71.58	4.88

Table 3: Compare ZerO with other initialization methods on CIFAR-10. ConstNet* denotes ConstNet with non-deterministic GPU operations discussed in Blumenfeld et al. (2020). Top-1 test error is reported.

Training ultra deep network without batch normalization Although there are methods attempting to train networks without batch normalization (by achieving dynamical isometry), they inevitably introduce random perturbations at initialization, affecting the training stability when the network is sufficiently deep (Zhang et al., 2019; De & Smith, 2020). We benchmark ZerO with state-of-the-art methods on training without batch normalization. As shown in Figure 4 (left), compared to other methods, ZerO achieves the best training stability for networks with even around 500 layers. It also matches the baseline where batch normalization is enabled.

Improved reproducibility. In addition, as shown in Table 2, ZerO achieves the lowest standard deviation over the repeated runs. On ImageNet, the gap between ZerO and other methods is even more than 40%. Thus, removing the randomness in the weight initialization improves reproducibility, with possible implications for topics such as trustworthy machine learning and network interpretation.




Figure 4: Training extreme deep ResNet on CIFAR-10 over 15 epochs.

4.1 ZerO on Transformer

We also apply ZerO to Transformer and evaluate it on WikiText-2 dataset (Vaswani et al., 2017). In each Transformer layer, we use ZerO to initialize both multi-head attention and feed-forward layers. Because the embedding size is fixed in the multi-head attention, we initialize the projection matrix of queries \mathbf{W}_Q as identity and the projection matrices of keys and values $\mathbf{W}_K, \mathbf{W}_V$ at zero. For the feed-forward layers, we initialize the connection matrices according to their hidden dimensions using Algorithm 1.

We train the Transformer models for 20 epochs with a single learning rate decay at epoch 10^2 . We also vary the number of layers in the model from 2 to 20. As shown in Table 4, ZerO achieves similar performance compared to the standard initialization. In addition, it has better training stability over deeper Transformer models, which is consistent with our previous results on ResNet.

Number of layers	2	4	6	8	10	20
Standard	200.44	168.69	154.67	146.43	diverged	diverged
ZerO	192.34	169.73	151.91	149.27	145.62	141.81

Table 4: Evaluate Transformer on WikiText-2. We vary the number of layers in Transformer, where each layer consists of a multi-head attention and a feed-forward layer. Test perplexity is reported (lower is better).

²We use a transformer architecture (provided by the link here) that was smaller than the transformers typically used for this task, explaining the general degradation of the results.

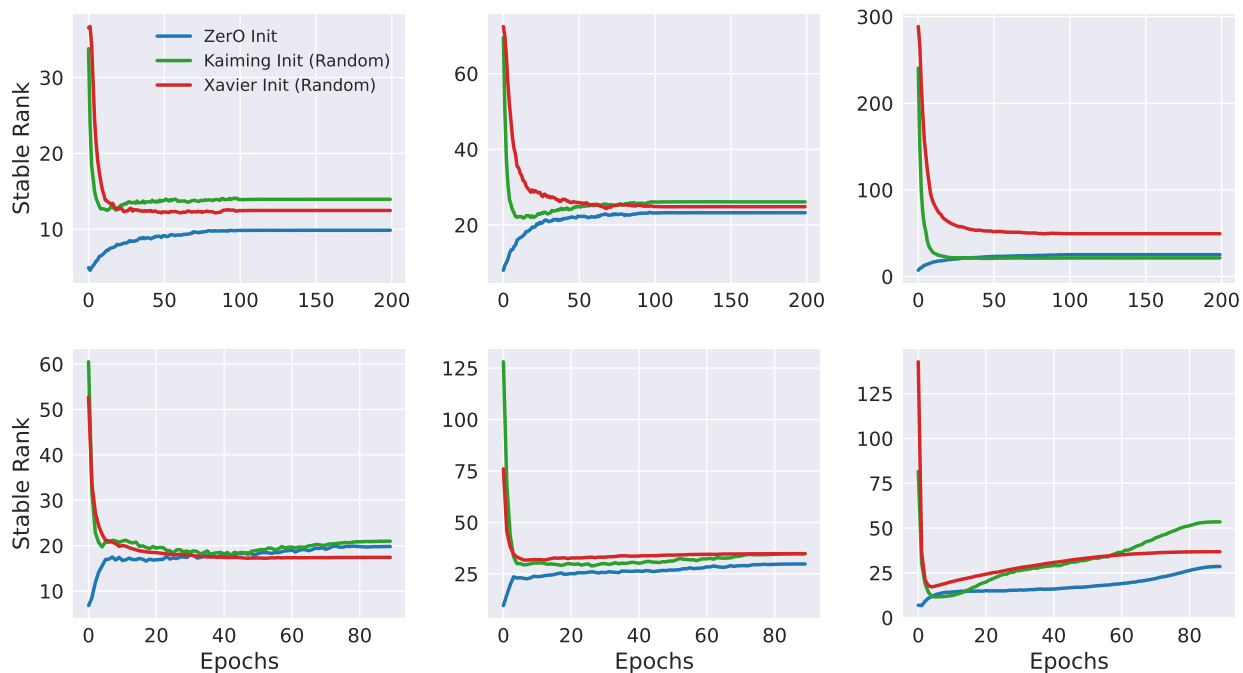


Figure 5: *Low-rank training trajectories in ResNet-18 on CIFAR-10 (top row) and ResNet-50 on ImageNet (bottom row). We visualize trajectories of the first convolutions in second, third, and fourth groups of residual blocks in ResNet.*

5 Low-Rank Learning Trajectory

Although ZerO and random initialization achieve similar test accuracies, their training trajectories differ significantly. In contrast to random initialization, which begins optimization from a complex network (i.e., full-rank weight matrices, as shown in Figure 3), ZerO starts the training from a "simple" network and gradually increases its complexity.

To show the difference in practice, we track the ranks of convolutional kernels in ResNets during training, where the rank of each kernel can reflect its complexity. We measure the stable rank, which is defined as

$$\|\mathbf{W}\|_F^2 / \|\mathbf{W}\|_2^2 = \sum_{i=1}^k \sigma_i^2(\mathbf{W}) / \sigma_{max}^2(\mathbf{W}),$$

for any matrix \mathbf{W} with k singular values σ_i . The stable rank is a soft version of the operator rank, and unlike the operator rank, it is insensitive to small singular values. We compare the stable ranks of various kernels between ZerO and random initialization during training. As shown in Figure 5, in contrast to random methods that begin with extremely high stable ranks, ZerO starts with low stable ranks and gradually increases them during training.

We believe ZerO’s learning trajectory is the first demonstration of greedy low-rank learning (GLRL) in large-scale deep learning applications. GLRL is a theoretical characterization of gradient descent, such that: when matrices are initialized with infinitesimal values, gradient descent performs a rank-constrained optimization and greedily relaxes the rank restriction by one whenever it fails to reach a minimizer (Li et al., 2021; Razin et al., 2021).

For example, when a matrix is initialized sufficiently small (where its rank is approximately zero), gradient descent first searches the solution over all rank-one matrices. If it fails to find a minimizer, it will relax the rank constraint by one and search again over all rank-two matrices. The search is stopped at rank- n if it finds a minimizer among all rank- n matrices.

GLRL suggests that gradient descent implicitly biases the model towards simple solutions by searching through the solution space in an incremental order of the matrix rank. This helps to explain the excellent generalization in gradient-based deep learning, as it converges to a global (or local) minima with the minimum rank.

Although the previous works have proved the existence of GLRL trajectory, it has never been observed in practice due to its impractical requirement of infinitesimal initialization. ZerO’s learning trajectory we observed suggests that GLRL not only exists under infinitesimal initialization, but also under initialization around the identity. If a matrix \mathbf{W} is initialized as \mathbf{I} , the low-rank structure is actually inside its residual component: $\mathbf{W}'_l = \mathbf{W}_l - \mathbf{I}$. To be noted, every convolutional kernel $\text{conv}(x)$ we measured in Figure 5 can be viewed as the residual component of $\text{conv}(x) + \mathbf{I}$, where the skip connection is included.

Figure 5 also suggests that the kernels never reach their maximal stable ranks during training under ZerO initialization. This implies that the searching over the space of full-rank weight matrices may be unnecessary, suggesting new avenues towards improved computational efficiency. We hope to explore this direction in future work.

We also observe that ZerO-based networks converge to low-complexity solutions. As shown in both Figure 5 and 6 (left), the convolutional kernels trained by ZerO usually have lower ranks than the kernels trained by random initialization. We further measure model complexity through both network pruning and low-rank approximation.

For network pruning, we use a standard magnitude-based pruning that prunes a portion of weights with the lowest magnitudes in each layer (Frankle & Carbin, 2019). For low-rank approximation, we apply Tucker-2 decomposition over channel dimensions in convolutions to select the most significant components (Kim et al., 2016).

As shown in Figure 6, compared to randomly initialized networks, the sub-networks obtained in trained ZerO-initialized networks achieve 25% higher sparsity or 30% lower (matrix or tensor) rank without sacrificing accuracy. This suggests ZerO encourages the networks to converge to low-complexity solutions, which improves the computational efficiency for inference.




Figure 6: **Left:** comparing kernel rank in ResNet-18 trained by ZerO and Kaiming methods. **Middle:** a magnitude-based network pruning on ResNet-18. **Right:** a Tucker-2 decomposition for a particular convolution with 512 channels in ResNet-18.

6 Related Works

To ensure stable training with random weight initialization, previous works such as Glorot & Bengio (2010); He et al. (2015) study the propagation of variance in the forward and backward pass under different activations. Several studies provide a more detailed characterization of the signal propagation with dynamical isometry (Saxe et al., 2014; Pennington et al., 2017; Xiao et al., 2018).

Inspired by the dynamical isometry property, various initialization methods are proposed to increase the convergence speed and stabilize the signal propagation, including Saxe et al. (2014); Bachlechner et al.

(2020); Gehring et al. (2017); Balduzzi et al. (2017). De & Smith (2020); Hoffer et al. (2018) study the reason behind the success of batch normalization (Ioffe & Szegedy, 2015) , and Zhang et al. (2019); De & Smith (2020) propose initialization methods to train residual networks without batch normalization.

As summarized in Table 1, many of these methods can be categorized as identity initialization with random perturbations. However, ZerO eliminates all the randomness using Hadamard and identity transforms. In another related work, Blumenfeld et al. (2020) discusses whether random initialization is needed from the perspective of feature diversity. They propose networks with identical features at initialization, which still require random perturbations to avoid the symmetric problem and improve performance.

Gradient descent biasing models towards low-rank solutions has been well studied in matrix factorization (Arora et al., 2019). Recent works also demonstrate the existence of greedy low-rank learning trajectory induced by gradient descent (Li et al., 2021; Razin et al., 2021; Jacot et al., 2021). However, no prior work demonstrates the greedy low-rank learning trajectory in large-scale applications of deep learning, as most only consider the problems of matrix factorization or applications of shallow neural networks.

7 Conclusion

In this work, we propose a simple and fully deterministic initialization called ZerO. Extensive experiments demonstrate that ZerO achieves state-of-the-art performance, suggesting that random weight initialization may not be necessary for initializing deep neural networks. ZerO has many benefits, such as training ultra deep networks (without batch-normalization), exhibiting low-rank learning trajectories that result in low-rank and sparse solutions, and improving training reproducibility.

We believe that ZerO opens up many new possibilities given its various benefits. It can be applied to networks and tasks sensitive to the variances in weight initialization. Its low-rank learning trajectories enable the development of rank-constrained training methods that improve computational efficiency. Finally, the improved training reproducibility can aid model interpretability. We hope our results will inspire other researchers to consider deterministic initialization schemes and to rethink the role of weight initialization in training deep neural networks.

Acknowledgments

We are grateful to the anonymous reviewers for their helpful comments and NVIDIA for the computational support. Dr. Anandkumar is supported by the Bren chair professorship at Caltech.

References

- Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 7411–7422, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/c0c783b5fc0d7d808f1d14a6e9c8280d-Abstract.html>.
- Thomas Bachlechner, Bodhisattwa Prasad Majumder, Huanru Henry Mao, Garrison W. Cottrell, and Julian McAuley. ReZero is All You Need: Fast Convergence at Large Depth. *ArXiv preprint*, abs/2003.04887, 2020. URL <https://arxiv.org/abs/2003.04887>.
- David Balduzzi, Marcus Frean, Lennox Leary, J. P. Lewis, Kurt Wan-Duo Ma, and Brian McWilliams. The shattered gradients problem: If resnets are the answer, then what is the question? In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 342–350. PMLR, 2017. URL <http://proceedings.mlr.press/v70/balduzzi17b.html>.
- Peter L. Bartlett, David P. Helmbold, and Philip M. Long. Gradient Descent with Identity Initialization Efficiently Learns Positive-Definite Linear Transformations by Deep Residual Networks. *Neural*

- Computation*, 31(3):477–502, 2019. ISSN 0899-7667, 1530-888X. doi: 10.1162/neco_a_01164. URL <https://direct.mit.edu/neco/article/31/3/477-502/8456>.
- Yaniv Blumenfeld, Dar Gilboa, and Daniel Soudry. Beyond signal propagation: Is feature diversity necessary in deep neural network initialization? In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 960–969. PMLR, 2020. URL <http://proceedings.mlr.press/v119/blumenfeld20a.html>.
- Soham De and Samuel L. Smith. Batch normalization biases residual blocks towards the identity function in deep networks. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/e6b738eca0e6792ba8a9cbbcba6c1881d-Abstract.html>.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, pp. 248–255. IEEE Computer Society, 2009. doi: 10.1109/CVPR.2009.5206848. URL <https://doi.org/10.1109/CVPR.2009.5206848>.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=rJl-b3RcF7>.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. Convolutional sequence to sequence learning. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1243–1252. PMLR, 2017. URL <http://proceedings.mlr.press/v70/gehring17a.html>.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In Yee Whye Teh and Mike Titterton (eds.), *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pp. 249–256, Chia Laguna Resort, Sardinia, Italy, 2010. PMLR. URL <https://proceedings.mlr.press/v9/glorot10a.html>.
- Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour. *ArXiv preprint*, abs/1706.02677, 2017. URL <https://arxiv.org/abs/1706.02677>.
- Moritz Hardt and Tengyu Ma. Identity matters in deep learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=ryxBORttx>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pp. 1026–1034. IEEE Computer Society, 2015. doi: 10.1109/ICCV.2015.123. URL <https://doi.org/10.1109/ICCV.2015.123>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- Elad Hoffer, Ron Banner, Itay Golan, and Daniel Soudry. Norm matters: efficient and accurate normalization schemes in deep networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018*,

- Montréal, Canada, pp. 2164–2174, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/a0160709701140704575d499c997b6ca-Abstract.html>.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Francis R. Bach and David M. Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 448–456. JMLR.org, 2015. URL <http://proceedings.mlr.press/v37/ioffe15.html>.
- Arthur Jacot, François Ged, Berfin Şimşek, Clément Hongler, and Franck Gabriel. Saddle-to-Saddle Dynamics in Deep Linear Networks: Small Initialization Training, Symmetry, and Sparsity. *ArXiv preprint*, abs/2106.15933, 2021. URL <https://arxiv.org/abs/2106.15933>.
- Yong-Deok Kim, Eunhyeok Park, Sungjoo Yoo, Taelim Choi, Lu Yang, and Dongjun Shin. Compression of deep convolutional neural networks for fast and low power mobile applications. In Yoshua Bengio and Yann LeCun (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1511.06530>.
- Alex Krizhevsky. Learning multiple layers of features from tiny images, 2009.
- Zhiyuan Li, Yuping Luo, and Kaifeng Lyu. Towards resolving the implicit bias of gradient descent for matrix factorization: Greedy low-rank learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=AH0s7Sm5H7R>.
- Jeffrey Pennington, Samuel S. Schoenholz, and Surya Ganguli. Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 4785–4795, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/d9fc0cdb67638d50f411432d0d41d0ba-Abstract.html>.
- Noam Razin, Asaf Maman, and Nadav Cohen. Implicit regularization in tensor factorization. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8913–8924. PMLR, 2021. URL <http://proceedings.mlr.press/v139/razin21a.html>.
- Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In Yoshua Bengio and Yann LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6120>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5998–6008, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- Lechao Xiao, Yasaman Bahri, Jascha Sohl-Dickstein, Samuel S. Schoenholz, and Jeffrey Pennington. Dynamical isometry and a mean field theory of cnns: How to train 10, 000-layer vanilla convolutional neural networks. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5389–5398. PMLR, 2018. URL <http://proceedings.mlr.press/v80/xiao18a.html>.
- Hongyi Zhang, Yann N. Dauphin, and Tengyu Ma. Fixup initialization: Residual learning without normalization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=H1gsz30cKX>.

A Additional Proofs in Theorem 1

A.1 Proof of (i)

We want to show that for any matrix $\mathbf{W} \in \mathbb{R}^{M \times N}$, if each entry is sampled from a Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$, then \mathbf{W} will achieve full-rank (i.e., $\text{rank}(\mathbf{W}) = \min(M, N)$) with probability of one.

Assume $M \geq N$ without losing generality. We want to prove the following statement:

$$\text{Prob}(\text{rank}(\mathbf{W}) < N) = 0.$$

Let $\{\mathbf{v}_i\}$ be columns of \mathbf{W} , where $\mathbf{v}_i \in \mathbb{R}^M$ for $i \in \{1, \dots, N\}$. We denote P_i as the probability of the event where a column \mathbf{v}_i is linearly dependent on the rest of the columns $\{\mathbf{v}_j\}$, for $j \neq i$. If \mathbf{W} does not reach the full-rank, then there exists at least a column that is linearly dependent on the rest of the columns. Thus, by the union bound, we know that:

$$\text{Prob}(\text{rank}(\mathbf{W}) < N) \leq \sum_{i=1}^N P_i.$$

For each P_i , it is equivalent to the probability of the event where $\mathbf{v}_i \in \mathbb{S} = \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_j)$, for $j \neq i$. Since \mathbb{S} at most has $N - 1$ dimension, it is a subspace of \mathbb{R}^M . Because each entry of $\mathbf{v}_i \in \mathbb{R}^M$ is sampled from a continuous Gaussian distribution, the probability that \mathbf{v}_i falls into a low-dimensional subspace \mathbb{S} is zero. Thus, $P_i = 0$ for any $i \in \{1, \dots, N\}$, and $\text{Prob}(\text{rank}(\mathbf{W}) < N) = 0$. We can conclude that \mathbf{W} reaches the full-rank almost surely.

A.2 Additional Proofs of (ii)

Here we prove that the rank constraint in (ii) is persistent through the entire training. Assume we perform gradient descent during training. The initial derivatives become:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{W}'_1} = \begin{pmatrix} \Lambda \\ \mathbf{0} \end{pmatrix} \in \mathbb{R}^{n_h \times n_x}, \quad \frac{\partial \mathcal{L}}{\partial \mathbf{W}'_l} = \begin{pmatrix} \Lambda & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \in \mathbb{R}^{n_h \times n_h}, \quad \frac{\partial \mathcal{L}}{\partial \mathbf{W}'_L} = (\Lambda \quad \mathbf{0}) \in \mathbb{R}^{n_y \times n_h},$$

where $\Lambda = \sum_{\mu=1}^P \text{Relu}'(\mathbf{z}^\mu) \odot (\mathbf{z}^\mu - \mathbf{y}^\mu) \mathbf{x}^{\mu T} \in \mathbb{R}^{n_y \times n_x}$, and $\mathbf{z}^\mu = \mathbf{I}_L \circ \text{Relu} \circ \mathbf{I}_1 \mathbf{x}^\mu$ and $l \in 2, \dots, L - 1$. $\mathbf{0}$ is a zero vector or a zero matrix depending on the dimensionality. \odot represents an element-wise multiplication.

As shown above, in the initial derivative of each matrix, the elements are zero except for a sub-matrix determined by Λ . Because the parameters with zero initial derivatives stall at zero during the entire training, the rank of each derivative in $\frac{\partial \mathcal{L}}{\partial \mathbf{W}'_2}, \dots, \frac{\partial \mathcal{L}}{\partial \mathbf{W}'_{L-1}}$ is upper bounded by n_x . Thus, the rank constraint in (ii) is satisfied during the entire training.

A.3 Proof of Lemma 1

We know that:

$$\mathbf{M} = \sum_{\mu=1}^Q \mathbf{a}^\mu \otimes \mathbf{b}^\mu = \mathbf{A} \mathbf{B}^\top,$$

where we define $\mathbf{A} = (\mathbf{a}^1, \dots, \mathbf{a}^Q)$ and $\mathbf{B} = (\mathbf{b}^1, \dots, \mathbf{b}^Q)$. Since $\dim(\text{span}(\mathbf{a}^1, \dots, \mathbf{a}^Q)) \leq U$, there are at most U independent columns in \mathbf{A} , and thus $\text{rank}(\mathbf{A}) \leq U$. Similarly, we can prove that $\text{rank}(\mathbf{B}) \leq V$. By the linear algebra, we can show that:

$$\text{rank}(\mathbf{M}) = \text{rank}(\mathbf{A} \mathbf{B}^\top) \leq \min(\text{rank}(\mathbf{A}), \text{rank}(\mathbf{B}^\top)) = \min(U, V)$$

A.4 Proof of Lemma 2

Let $\mathbf{x} = \alpha \cdot \mathbf{e}_2 + \beta \cdot \mathbf{e}_3$ where α and β are scalars. Then we have:

$$\text{Relu } \mathbf{HI}^* \mathbf{x} = \begin{cases} \alpha \text{Relu}(\mathbf{HI}^* \mathbf{e}_2) + \beta \text{Relu}(\mathbf{HI}^* \mathbf{e}_3) & \text{for } \alpha > 0 \text{ and } \beta > 0, \\ \alpha \text{Relu}(-\mathbf{HI}^* \mathbf{e}_2) + \beta \text{Relu}(\mathbf{HI}^* \mathbf{e}_3) & \text{for } \alpha < 0 \text{ and } \beta > 0, \\ \alpha \text{Relu}(\mathbf{HI}^* \mathbf{e}_2) + \beta \text{Relu}(-\mathbf{HI}^* \mathbf{e}_3) & \text{for } \alpha > 0 \text{ and } \beta < 0, \\ \alpha \text{Relu}(-\mathbf{HI}^* \mathbf{e}_2) + \beta \text{Relu}(-\mathbf{HI}^* \mathbf{e}_3) & \text{for } \alpha < 0 \text{ and } \beta < 0. \end{cases}$$

Thus, it holds that:

$$\text{span}(\{\text{Relu } \mathbf{HI}^* \mathbf{x} | \mathbf{x} \in \mathbb{R}^{n_x}\}) = \text{span}(\text{Relu}(\mathbf{HI}^* \mathbf{e}_2), \text{Relu}(-\mathbf{HI}^* \mathbf{e}_2), \text{Relu}(\mathbf{HI}^* \mathbf{e}_3), \text{Relu}(-\mathbf{HI}^* \mathbf{e}_3)).$$

After a simple computation, we also know that:

$$\begin{aligned} \mathbf{HI}^* \mathbf{e}_2 &= \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \end{pmatrix} & -\mathbf{HI}^* \mathbf{e}_2 &= \begin{pmatrix} 0 \\ 1 \\ 0 \\ 1 \end{pmatrix} \\ \mathbf{HI}^* \mathbf{e}_3 &= \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} & -\mathbf{HI}^* \mathbf{e}_3 &= \begin{pmatrix} 0 \\ 0 \\ 1 \\ 1 \end{pmatrix}. \end{aligned}$$

Therefore, the four vectors are linearly independent, and thus:

$$\dim(\text{span}(\{\text{Relu } \mathbf{HI}^* \mathbf{x} | \mathbf{x} \in \mathbb{R}^{n_x}\})) = 4.$$

B Details of the experiments on MNIST

We train a network \mathcal{F} defined in Theorem 1 with depth $L = 3$ on MNIST, where $\mathbf{W}^1 \in \mathbb{R}^{2048 \times 784}$, $\mathbf{W}^2 \in \mathbb{R}^{2048 \times 2048}$, and $\mathbf{W}^3 \in \mathbb{R}^{10 \times 2048}$. We initialize $\mathbf{W}^1 = \mathbf{H}\mathbf{I}^*$, $\mathbf{W}^2 = \mathbf{I}$, and $\mathbf{W}^3 = \mathbf{I}^*$. The network is trained for 14 epochs using SGD with a learning rate of 0.1. We compare ZerO with Kaiming and Xavier initialization under the same setting. All candidates achieve test accuracy above 98% after the training. Figure 7 shows the weight distributions at different training iterations.

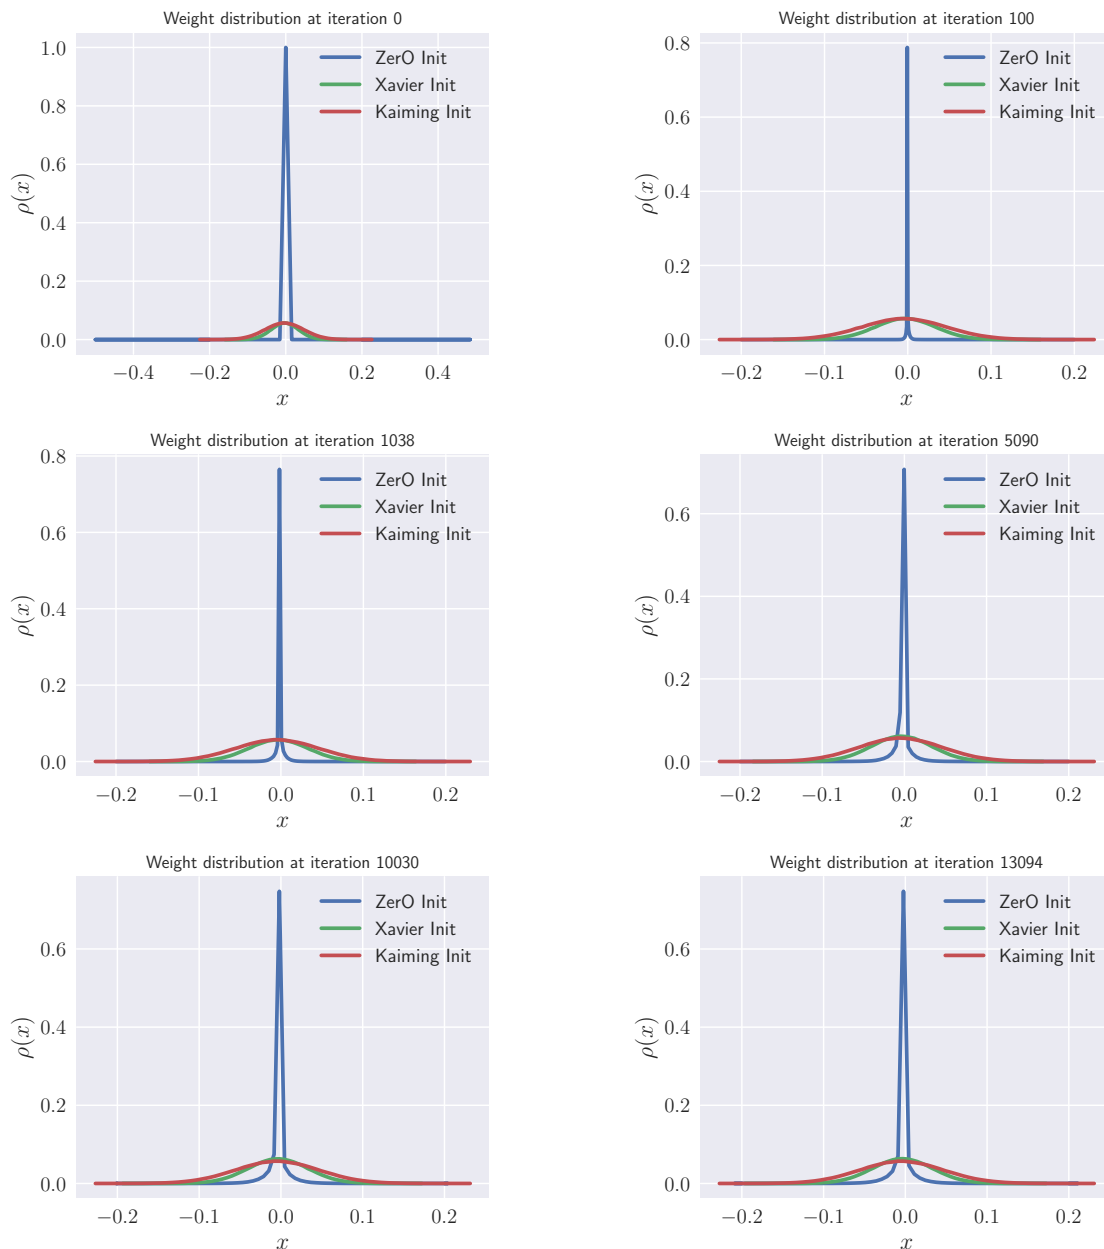


Figure 7: Weight distributions at different training iterations.

C Learning rate warmup

We observe that learning rate warmup is essential for ZeRO to achieve a large maximal learning rate. This is because a large initial learning rate leads to gradient explosion, as most of the weights start from the exact zero at the beginning. To demonstrate it, we measure the gradient norms during training for the settings with or without learning rate warmup.

As shown in Figure 8, we present the gradient norms wrt. the weights of the first convolutional layer in ResNet-18, which is trained on CIFAR-10 for the first 400 iterations. The results indicate that a learning rate warmup is needed for ZeRO because the smaller initial learning rate prevents the gradients from the signal explosion at the beginning of training.

Interestingly, we observe that the warmup has a much larger effect on ZeRO than other random initialization methods. For Kaiming and Xavier initialization, gradient norms are nearly unchanged after applying warmup. However, for ZeRO initialization, the gradient norms become significantly smaller (even smaller than the baselines) after applying the warmup.




Figure 8: Measuring gradient norms of ResNet-18 for the first 400 iterations.

For random initialization methods, the small changes in gradient norms are because the norms are largely dominated by the initial weight variance, which is controlled by the initializers. Choosing inappropriate initial weight variance usually leads to gradient explosion. To avoid the explosion, previous studies on random initialization propose various principles to control the variances for different architectures and nonlinearities. However, as ZeRO initialization directly trains the weights from zeros and ones (without pre-defined weight variances), the learning rate becomes the dominant factor that controls the norm of the initial gradients. Hence, choosing an appropriate initial learning rate is sufficient for avoiding exploding gradients under ZeRO initialization.