

# ORTHOGONAL OVER-PARAMETERIZED TRAINING

Anonymous authors

Paper under double-blind review

## ABSTRACT

The inductive bias of a neural network is largely determined by the architecture and the training algorithm. To achieve good generalization, how to effectively train a neural network is of great importance. We propose a novel orthogonal over-parameterized training (OPT) framework that can provably minimize the hyperspherical energy which characterizes the diversity of neurons on a hypersphere. By maintaining the minimum hyperspherical energy during training, OPT can greatly improve the empirical generalization. Specifically, OPT fixes the randomly initialized weights of the neurons and learns an orthogonal transformation that applies to these neurons. We consider multiple ways to learn such an orthogonal transformation, including unrolling orthogonalization algorithms, applying orthogonal parameterization, and designing orthogonality-preserving gradient descent. For better scalability, we further propose the stochastic OPT which performs orthogonal transformation stochastically for partial dimensions of the neuron. Interestingly, OPT reveals that learning a proper coordinate system for neurons is crucial to generalization. We provide some insights on why OPT yields better generalization. Extensive experiments validate the superiority of OPT.

## 1 INTRODUCTION

The inductive bias encoded in a neural network is generally determined by two major aspects: how the neural network is structured (*i.e.*, network architecture) and how the neural network is optimized (*i.e.*, training algorithm). For the same network architecture, using different training algorithms could lead to a dramatic difference in generalization performance [33, 54] even if the training loss is already close to zero, implying that different training procedures lead to different inductive biases. Therefore, how to effectively train a neural network that can generalize well remains an open challenge.

Recent theories [15, 15, 16, 31, 42] suggest the importance of over-parameterization in linear neural networks. For example, [16] shows that optimizing an underdetermined quadratic objective over a matrix  $M$  with gradient descent on a factorization of  $M$  leads to an implicit regularization that may improve generalization. There is also strong empirical evidence [11, 46] that over-parameterizing the convolutional filters under some regularity is beneficial to generalization. Our paper aims to leverage the power of over-parameterization and explore more intrinsic structural priors in order to train a well-performing neural network.

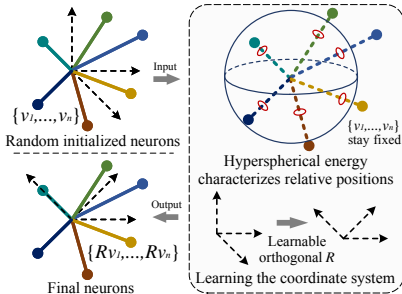


Figure 1: Overview of OPT.

Motivated by this goal, we propose a generic orthogonal over-parameterized training (OPT) framework for neural networks. Different from earlier methods, OPT over-parameterizes a neuron  $w \in \mathbb{R}^d$  with the multiplication of a learnable layer-shared orthogonal matrix  $R \in \mathbb{R}^{d \times d}$  and a fixed randomly initialized weight vector  $v \in \mathbb{R}^d$ , and it follows that the equivalent weight for the neuron is  $w = Rv$ . Once each element of the neuron weight  $v$  has been randomly initialized by a zero-mean Gaussian distribution [14, 19], we fix them throughout the entire training process. Then OPT learns a layer-shared orthogonal transformation  $R$  that is applied to all the neurons (in the same layer). An illustration of OPT is given in Fig. 1. In contrast to standard neural training, OPT decomposes the neuron into an orthogonal transformation  $R$  that learns a proper coordinate system and a weight vector  $v$  that controls the specific position of the neuron. Essentially, the weights  $\{v_1, \dots, v_n \in \mathbb{R}^d\}$  of different neurons determine the relative positions, while the layer-shared orthogonal matrix  $R$  specifies the coordinate system. The decoupled parameterization enables strong modeling flexibility.

Another motivation of OPT comes from an empirical observation [45] that neural networks with lower *hyperspherical energy* generalize better. Hyperspherical energy quantifies the diversity of neurons on a hypersphere, and essentially characterizes the relative positions among neurons via this form of diversity. [45] introduces hyperspherical energy as a regularization in the network but does not guarantee the hyperspherical energy can be effectively minimized (due to the existence of data fitting loss). To address this issue, we leverage the property of hyperspherical energy that it is independent of the coordinate system in which the neurons live and only depends on their relative positions. Specifically, we prove that, if we randomly initialize the neuron weight  $v$  with certain distributions, these neurons are guaranteed to attain minimum hyperspherical energy in expectation. It follows that OPT maintains the minimum energy during training by learning a coordinate system (*i.e.*, layer-shared orthogonal matrix) for the neurons. Therefore, OPT can well minimize the hyperspherical energy.

We consider several ways to learn the orthogonal transformation. First, we unroll different orthogonalization algorithms such as Gram-Schmidt process, Householder reflection and Löwdin’s symmetric orthogonalization. Different unrolled algorithms yield different implicit regularizations to construct the neuron weights. For example, symmetric orthogonalization guarantees that the new orthogonal basis has the least distance in the Hilbert space from the original non-orthogonal basis. Second, we consider to use a special parameterization (*e.g.*, Cayley parameterization) to construct the orthogonal matrix, which is more efficient in training. Third, we consider an orthogonal-preserving gradient descent to ensure that the matrix  $R$  stays orthogonal after each gradient update. Last, we relax the original optimization problem by making the orthogonality constraint a regularization for the matrix  $R$ . Different ways of learning the orthogonal transformation may encode different inductive biases. We note that OPT aims to use orthogonalization as a tool to learn neurons that preserve the hyperspherical energy, rather than to study a specific orthogonalization method. Moreover, we propose a refinement strategy to further reduce the hyperspherical energy for the randomly initialized neuron weights  $\{v_1, \dots, v_n\}$ . In specific, we directly minimize the hyperspherical energy of these random weights as a preprocessing step before training them on actual data.

To improve scalability, we further propose the stochastic OPT that randomly samples neuron dimensions to perform orthogonal transformation. The random sampling process is repeated many times such that each dimension of the neuron is sufficiently learned. Finally, we provide some theoretical insights and discussions to justify the effectiveness of OPT. The advantages of OPT is as follows:

- OPT is a generic neural training framework with strong interpretability and flexibility. There are many ways to learn the orthogonal transformation and each one imposes a unique inductive bias.
- OPT is the first training method where the hyperspherical energy is provably minimized, leading to better empirical generalization. OPT reveals that learning a proper coordinate system is crucial to generalization, and the relative neuron positions are well characterized by hyperspherical energy.
- There is no extra computational cost for OPT-trained neural networks in inference. It has the same inference speed and model size as its standard counterpart. Our experiments also show that OPT performs well on different neural networks and therefore is architecture-agnostic.
- Stochastic OPT can greatly improve the scalability of OPT while having comparable performance.

## 2 RELATED WORK

**Orthogonality in Neural Networks.** Orthogonality has been widely adopted to improve neural networks. [4, 7, 48] uses orthogonality as a regularization for neural network. [3, 24, 28, 39, 52, 65] uses principled orthogonalization methods to guarantee the neurons are orthogonal to each other. In contrast to these works, OPT does not encourage orthogonality among neurons. Instead, OPT utilizes principled orthogonalization methods to learn orthogonal transformations on neurons.

**Parameterization of Neurons.** There are various ways to parameterize a neuron for different applications. [11] over-parameterizes a 2D convolution kernel by combining a 2D kernel of the same size and two additional 1D asymmetric kernels. The resulted convolution kernel has the same effective parameters during testing but more parameters during training. [46] constructs a neuron with a bilinear parameterization and regularizes the bilinear similarity matrix. [68] reparameterizes the neuron matrix with an adaptive fastfood transform to compress model parameters. [27, 44, 63] employ sparse and low-rank structures to construct convolution kernels for a efficient neural network.

**Hyperspherical Learning.** [9, 47, 48] propose to learn representations on hypersphere and show that the angular information in neural networks, in contrast to magnitude information, preserves the most semantic meaning. [45] defines hyperspherical energy that quantifies the diversity of neurons on a hypersphere and empirically shows that the minimum hyperspherical energy improves generalization.

### 3 ORTHOGONAL OVER-PARAMETERIZED TRAINING

#### 3.1 GENERAL FRAMEWORK

OPT parameterizes the neuron as the multiplication of an orthogonal matrix  $\mathbf{R} \in \mathbb{R}^{d \times d}$  and a neuron weight vector  $\mathbf{v} \in \mathbb{R}^d$ , and the equivalent neuron weight becomes  $\mathbf{w} = \mathbf{R}\mathbf{v}$ . The output  $\hat{y}$  of this neuron can be represented by  $\hat{y} = (\mathbf{R}\mathbf{v})^\top \mathbf{x}$  where  $\mathbf{x} \in \mathbb{R}^d$  is the input vector. In OPT, we fix the randomly initialized neuron weight  $\mathbf{v}$  and only learn the orthogonal matrix  $\mathbf{R}$ . In contrast, the standard neuron is directly formulated as  $\hat{y} = \mathbf{v}^\top \mathbf{x}$ , where the weight vector  $\mathbf{v}$  is learned in training.

As an illustrative example, we consider a two-layer linear MLP with a loss function  $\mathcal{L}$  (e.g., the least squares loss:  $\mathcal{L}(e_1, e_2) = (e_1 - e_2)^2$ ). Specifically, the learning objective of the standard training is  $\min_{\{\mathbf{v}_i, u_i, \forall i\}} \sum_{j=1}^m \mathcal{L}(y, \sum_{i=1}^n u_i \mathbf{v}_i^\top \mathbf{x}_j)$ , while differently, our OPT is formulated as

$$\min_{\{\mathbf{R}, u_i, \forall i\}} \sum_{j=1}^m \mathcal{L}(y, \sum_{i=1}^n u_i (\mathbf{R}\mathbf{v}_i)^\top \mathbf{x}_j) \quad \text{s.t. } \mathbf{R}^\top \mathbf{R} = \mathbf{R}\mathbf{R}^\top = \mathbf{I} \quad (1)$$

where  $\mathbf{v}_i \in \mathbb{R}^d$  is the  $i$ -th neuron in the first layer, and  $\mathbf{u} = \{u_1, \dots, u_n\} \in \mathbb{R}^n$  is the output neuron in the second layer. In OPT, each element of  $\mathbf{v}_i$  is usually sampled from a zero-mean Gaussian distribution, and is fixed throughout the entire training process. In general, OPT learns an orthogonal matrix that is applied to all the neurons instead of learning the individual neuron weight. Note that, we usually do not apply OPT to neurons in the output layer (e.g.,  $\mathbf{u}$  in this MLP example, and the final linear classifiers in CNNs), since it makes little sense to fix a set of random linear classifiers. Therefore, the central problem is how to learn these layer-shared orthogonal matrices.

#### 3.2 HYPERSPHERICAL ENERGY PERSPECTIVE

We delve into OPT from the hyperspherical energy perspective. Following [45], the hyperspherical energy of  $n$  neurons is defined as  $\mathbf{E}(\hat{\mathbf{v}}_i|_{i=1}^n) = \sum_{i=1}^n \sum_{j=1, j \neq i}^n \|\hat{\mathbf{v}}_i - \hat{\mathbf{v}}_j\|^{-1}$  in which  $\hat{\mathbf{v}}_i = \frac{\mathbf{v}_i}{\|\mathbf{v}_i\|}$  is the  $i$ -th neuron weight projected onto the unit hypersphere  $\mathbb{S}^{d-1} = \{\mathbf{v} \in \mathbb{R}^d \mid \|\mathbf{v}\| = 1\}$ . Hyperspherical energy is used to characterize the diversity of  $n$  neurons on a unit hypersphere. Assume that we have  $n$  neurons in one layer, and we have learned an orthogonal matrix  $\mathbf{R}$  for these neurons. The hyperspherical energy of these  $n$  OPT-trained neurons is given by

$$\mathbf{E}(\mathbf{R}\hat{\mathbf{v}}_i|_{i=1}^n) = \sum_{i=1}^n \sum_{j=1, j \neq i}^n \|\mathbf{R}\hat{\mathbf{v}}_i - \mathbf{R}\hat{\mathbf{v}}_j\|^{-1} = \sum_{i=1}^n \sum_{j=1, j \neq i}^n \|\hat{\mathbf{v}}_i - \hat{\mathbf{v}}_j\|^{-1} = \mathbf{E}(\hat{\mathbf{v}}_i|_{i=1}^n) \quad (2)$$

which shows that the hyperspherical energy does not change in OPT. Moreover, [45] proves that minimum hyperspherical energy corresponds to the uniform distribution over the hypersphere. As a result, if the initialization of the neurons in the same layer follows the uniform distribution over the hypersphere, then we can guarantee that the hyperspherical energy is minimal in a probabilistic sense.

**Theorem 1.** *For the neuron  $\mathbf{h} = \{h_1, \dots, h_d\}$  where  $h_i, \forall i$  are initialized i.i.d. following a zero-mean Gaussian distribution (i.e.,  $h_i \sim N(0, \sigma^2)$ ), the projections onto a unit hypersphere  $\hat{\mathbf{h}} = \mathbf{h} / \|\mathbf{h}\|$  where  $\|\mathbf{h}\| = \sqrt{\sum_i h_i^2}$  are uniformly distributed on the unit hypersphere  $\mathbb{S}^{d-1}$ . The neurons with minimum hyperspherical energy attained asymptotically approach to the uniform distribution on  $\mathbb{S}^{d-1}$ .*

Theorem 1 implies that, if we initialize the neurons in the same layer with zero-mean Gaussians, the corresponding expected hyperspherical energy is guaranteed to be small. It is because the neurons are uniformly distributed on the unit hypersphere and hyperspherical energy quantifies the uniformity on the hypersphere in some sense. More importantly, prevailing neuron initializations such as [14] and [19] are zero-mean Gaussian distribution. Therefore, our neurons naturally have low hyperspherical energy from the beginning. Appendix L gives geometric properties of the random initialized neurons.

#### 3.3 UNROLLING ORTHOGONALIZATION ALGORITHMS

To learn the orthogonal transformation, we consider to unroll classic orthogonalization algorithms and embed them into the neural network such that the training is still end-to-end. We need to make every step of the orthogonalization algorithm differentiable, and the brief training flow is shown in Fig. 2.




Figure 2: Unrolled orthogonalization.

**Gram-Schmidt Process.** This method takes a linearly independent set and produces an orthogonal set based on it. The Gram-Schmidt Process (GS) usually takes the following steps to orthogonalize

a set of vectors  $\{\mathbf{u}_1, \dots, \mathbf{u}_n\} \in \mathbb{R}^{n \times n}$  and obtain an orthogonal set  $\{e_1, \dots, e_i, \dots, e_n\} \in \mathbb{R}^{n \times n}$ . First, when  $i=1$ , we have  $e_1 = \frac{\tilde{e}_1}{\|\tilde{e}_1\|}$  where  $\tilde{e}_1 = \mathbf{u}_1$ . Then, when  $n \geq i \geq 2$ , we have  $e_i = \frac{\tilde{e}_i}{\|\tilde{e}_i\|}$  where  $\tilde{e}_i = \mathbf{u}_i - \sum_{j=1}^{i-1} \text{Proj}_{e_j}(\mathbf{u}_i)$ .  $\text{Proj}_{\mathbf{b}}(\mathbf{a}) = \frac{\langle \mathbf{a}, \mathbf{b} \rangle}{\langle \mathbf{b}, \mathbf{b} \rangle} \mathbf{b}$  is the projection operator. For better orthogonality, we can unroll an iterative version [23] of GS with multiple steps.

**Householder Reflection.** A Householder reflector is defined as  $\mathbf{H} = \mathbf{I} - 2 \frac{\mathbf{u}\mathbf{u}^\top}{\|\mathbf{u}\|^2}$  where  $\mathbf{u}$  is perpendicular to the reflection hyperplane. In QR factorization, Householder reflection (HR) is used to transform a (non-singular) square matrix into an orthogonal matrix and an upper triangular matrix. Given a matrix  $\mathbf{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_n\} \in \mathbb{R}^{n \times n}$ , we consider the first column vector  $\mathbf{u}_1$ . We use Householder reflector to transform  $\mathbf{u}_1$  to  $e_1 = \{1, 0, \dots, 0\}$ . Specifically, we construct an orthogonal matrix  $\mathbf{H}_1$  with  $\mathbf{H}_1 = \mathbf{I} - 2 \frac{(\mathbf{u}_1 - \|\mathbf{u}_1\| e_1)(\mathbf{u}_1 - \|\mathbf{u}_1\| e_1)^\top}{\|\mathbf{u}_1 - \|\mathbf{u}_1\| e_1\|^2}$ . The first column of  $\mathbf{H}_1 \mathbf{U}$  becomes  $\{\|\mathbf{u}_1\|, 0, \dots, 0\}$ . At the  $k$ -th step, we can view the sub-matrix  $\mathbf{U}_{(k:n, k:n)}$  as a new  $\mathbf{U}$ , and use the same procedure to construct the Householder transformation  $\tilde{\mathbf{H}}_k \in \mathbb{R}^{(n-k) \times (n-k)}$ . We construct the final Householder transformation as  $\mathbf{H}_k = \text{Diag}(\mathbf{I}_k, \tilde{\mathbf{H}}_k)$ . Now we can gradually transform  $\mathbf{U}$  to an upper triangular matrix with  $n$  Householder reflections. Therefore, we have that  $\mathbf{H}_n \dots \mathbf{H}_2 \mathbf{H}_1 \mathbf{U} = \mathbf{R}^{\text{up}}$  where  $\mathbf{R}^{\text{up}}$  is an upper triangular matrix and the obtained orthogonal set is  $\mathbf{Q}^\top = \mathbf{H}_n \dots \mathbf{H}_2 \mathbf{H}_1$ .

**Löwdin’s Symmetric Orthogonalization.** Let the matrix  $\mathbf{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_n\} \in \mathbb{R}^{n \times n}$  be a given set of linearly independent vectors in an  $n$ -dimensional space. A non-singular linear transformation  $\mathbf{A}$  can transform the basis  $\mathbf{U}$  to an orthogonal basis  $\mathbf{R}$ :  $\mathbf{R} = \mathbf{U}\mathbf{A}$ . The matrix  $\mathbf{R}$  will be orthogonal if  $\mathbf{R}^\top \mathbf{R} = (\mathbf{U}\mathbf{A})^\top \mathbf{U}\mathbf{A} = \mathbf{A}^\top \mathbf{M}\mathbf{A} = \mathbf{I}$  where  $\mathbf{M} = \mathbf{U}^\top \mathbf{U}$  is the Gram matrix of the given set  $\mathbf{U}$ . We obtain a general solution to the orthogonalization problem via the substitution:  $\mathbf{A} = \mathbf{M}^{-\frac{1}{2}} \mathbf{B}$  where  $\mathbf{B}$  is an arbitrary unitary matrix. The specific choice  $\mathbf{B} = \mathbf{I}$  gives the Löwdin’s symmetric orthogonalization (LS):  $\mathbf{R} = \mathbf{U}\mathbf{M}^{-\frac{1}{2}}$ . We can analytically obtain the symmetric orthogonalization from the singular value decomposition:  $\mathbf{U} = \mathbf{W}\mathbf{\Sigma}\mathbf{V}^\top$ . Then LS gives  $\mathbf{R} = \mathbf{W}\mathbf{V}^\top$  as the orthogonal set for  $\mathbf{U}$ . LS has a unique property which the other orthogonalizations do not have. The orthogonal set resembles the original set in a nearest-neighbour sense. Specifically, LS guarantees that  $\sum_i \|\mathbf{R}_i - \mathbf{U}_i\|^2$  (where  $\mathbf{R}_i$  and  $\mathbf{U}_i$  are the  $i$ -th column of  $\mathbf{R}$  and  $\mathbf{U}$ , respectively) is minimized. Intuitively, LS indicates the gentlest pushing of the directions of the vectors in order to get them orthogonal.

**Discussion.** These orthogonalization algorithms are fully differentiable and end-to-end trainable. For better orthogonality, these algorithms can be used repeatedly and unrolled with multiple steps. Empirically, one-step unrolling already works well. Givens rotations can also construct the orthogonal matrix, but it requires traversing all lower triangular elements in the original set  $\mathbf{U}$ , which takes  $\mathcal{O}(n^2)$  complexity and is too costly. Interestingly, each orthogonalization encodes a unique inductive bias to the neurons by imposing implicit regularizations (e.g., least distance in Hilbert space for LS). Details about these orthogonalizations are in Appendix A. Unrolling orthogonalization has been considered in different scenarios [24, 50, 62]. Many more advanced orthogonalization methods [38] can also be applied in OPT, but exhaustively applying them to OPT is out of the scope of this paper.

### 3.4 ORTHOGONAL PARAMETERIZATION

A convenient way to ensure orthogonality while learning the matrix  $\mathbf{R}$  is to use a special parameterization that inherently guarantees orthogonality. The exponential parameterization use  $\mathbf{R} = \exp(\mathbf{W})$  (where  $\exp(\cdot)$  denotes the matrix exponential) to represent an orthogonal matrix from a skew-symmetric matrix  $\mathbf{W}$ . The Cayley parameterization (CP) is a Padé approximation of the exponential parameterization, and is a more natural choice due to its simplicity. CP uses the following transform to construct an orthogonal matrix  $\mathbf{R}$  from a skew-symmetric matrix  $\mathbf{W}$ :  $\mathbf{R} = (\mathbf{I} + \mathbf{W})(\mathbf{I} - \mathbf{W})^{-1}$  where  $\mathbf{W} = -\mathbf{W}^\top$ . We note that CP only produces the orthogonal matrices with determinant 1, which belong to the special orthogonal group and thus  $\mathbf{R} \in \text{SO}(n)$ . Specifically, it suffices to learn the upper or lower triangular of the matrix  $\mathbf{W}$  with unconstrained optimization to obtain a desired orthogonal matrix  $\mathbf{R}$ . Cayley parameterization does not cover the entire orthogonal group and is less flexible in terms of representation power, which serves as a explicit regularization for the neurons.

### 3.5 ORTHOGONALITY-PRESERVING GRADIENT DESCENT

An alternative way to guarantee orthogonality is to modify the gradient update for the matrix  $\mathbf{R}$ . The idea is to initialize  $\mathbf{R}$  with an arbitrary orthogonal matrix and then ensure each gradient update is to apply an orthogonal transformation to  $\mathbf{R}$ . It is essentially conducting gradient descent on the Stiefel manifold [3, 21, 30, 39, 41, 64, 65]. Given a matrix  $\mathbf{U}_{(0)} \in \mathbb{R}^{n \times n}$  that is initialized as an orthogonal

matrix, we aim to construct an orthogonal transformation as the gradient update. We use the Cayley transform to compute a parametric curve on the Stiefel manifold  $\mathcal{M}_s = \{\mathbf{U} \in \mathbb{R}^{n \times n} : \mathbf{U}^\top \mathbf{U} = \mathbf{I}\}$  with a specific metric via a skew-symmetric matrix  $\mathbf{W}$  and use it as the update rule:

$$\mathbf{Y}(\lambda) = (\mathbf{I} - \frac{\lambda}{2}\mathbf{W})^{-1}(\mathbf{I} + \frac{\lambda}{2}\mathbf{W})\mathbf{U}_{(i)}, \mathbf{U}_{(i+1)} = \mathbf{Y}(\lambda) \quad (3)$$

where  $\hat{\mathbf{W}} = \nabla f(\mathbf{U}_{(i)})\mathbf{U}_{(i)}^\top - \frac{1}{2}\mathbf{U}_{(i)}(\mathbf{U}_{(i)}^\top \nabla f(\mathbf{U}_{(i)}\mathbf{U}_{(i)}^\top))$  and  $\mathbf{W} = \hat{\mathbf{W}} - \hat{\mathbf{W}}^\top$ .  $\mathbf{U}_{(i)}$  denotes the orthogonal matrix in the  $i$ -th iteration.  $\nabla f(\mathbf{U}_{(i)})$  denotes the original gradient of the loss function *w.r.t.*  $\mathbf{U}_{(i)}$ . We term this gradient update as orthogonal-preserving gradient descent (OGD). To reduce the computational cost of the matrix inverse in Eq. 3, we use an iterative method [41] to approximate the Cayley transform without matrix inverse. We arrive at the fixed-point iteration from Eq. 3:

$$\mathbf{Y}(\lambda) = \mathbf{U}_{(i)} + \frac{\lambda}{2}\mathbf{W}(\mathbf{U}_{(i)} + \mathbf{Y}(\lambda)) \quad (4)$$

which converges to the closed-form Cayley transform with a rate of  $o(\lambda^{2+n})$  ( $n$  is the iteration number). In practice, we find that two iterations will suffice for a reasonable approximation accuracy.

### 3.6 RELAXATION TO ORTHOGONAL REGULARIZATION

We consider relaxing the original optimization with an orthogonality constraint to an unconstrained optimization with orthogonality regularization (OR). Specifically, we remove the orthogonality constraint, and adopt an orthogonality regularization for  $\mathbf{R}$ , *i.e.*,  $\|\mathbf{R}^\top \mathbf{R} - \mathbf{I}\|_F^2$ . However, OR cannot guarantee the energy stays unchanged. Taking Eq. 1 as an example, the objective becomes

$$\min_{\mathbf{R}, u_i, \forall i} \sum_{j=1}^m \mathcal{L}(y, \sum_{i=1}^n u_i (\mathbf{R}\mathbf{v}_i)^\top \mathbf{x}_j) + \beta \|\mathbf{R}^\top \mathbf{R} - \mathbf{I}\|_F^2 \quad (5)$$

where  $\beta$  is a hyperparameter. This serves as an relaxation of the original OPT objective. Note that, OR is imposed to  $\mathbf{R}$  and is quite different from the existing orthogonal regularization on neurons [4, 48].

### 3.7 REFINING THE RANDOM INITIALIZATION AS PREPROCESSING

**Minimizing hyperspherical energy before training.** Because we randomly initialize the neurons  $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ , there exists a variance that makes the hyperspherical energy deviate from the minima even if the hyperspherical energy is minimal in a probabilistic sense. To further reduce the hyperspherical energy, we propose to refine the random initialization by minimizing its hyperspherical energy as a preprocessing step before the OPT training. Specifically, before feeding these neurons to OPT, we first minimize the hyperspherical energy of the initialized neurons with gradient descent (without fitting the training data). Moreover, since the randomly initialized neurons cannot guarantee to get rid of the collinearity redundancy as shown in [45] (*i.e.*, two neurons are on the same line but have opposite directions), we can perform the half-space hyperspherical energy minimization [45].

**Normalizing the neurons.** The norm of the randomly initialized neurons may have some influence on OPT, serving a role similar to weighting the importance of different neurons. Moreover, the norm makes the hyperspherical energy less expressive to characterize the diversity of neurons, as discussed in Section 5.3. To address this, we propose to normalize the neuron weights such that the weight of each neuron has the unit norm. Because the weights of the neurons  $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$  are fixed throughout the training process and OPT will not change the norm of the final neurons, we only need to normalize the randomly initialized neuron weights as a preprocessing before the OPT training.

## 4 TOWARDS SCALABLE ORTHOGONAL OVER-PARAMETERIZED TRAINING

If the dimension of neurons become is extremely large, then the orthogonal matrix to transform the neurons will also be large. Therefore, it may take large GPU memory and time to train the neural networks with the original OPT. To address this, we propose a scalable OPT variant – stochastic orthogonal over-parameterized training (S-OPT). The key idea of S-OPT is to randomly select some dimensions of neurons in the same layer and construct a small orthogonal matrix to transform these dimensions together. The selection of neuron dimensions is stochastic in each outer iteration, so a small orthogonal matrix is still able to cover all the neuron dimensions. Intuitively, S-OPT aims to approximate the orthogonal transformation of all the neuron dimensions with small orthogonal transformations of random subsets of the neuron dimensions. The approximation will be more accurate if the procedure is randomized over many times. Fig. 2 compares the size of the orthogonal matrix in OPT and S-OPT.




Figure 3: OPT & S-OPT

Most importantly, S-OPT can still preserve the hyperspherical energy of neurons because of the following Theorem 2.

**Theorem 2.** *For  $n$   $d$ -dimensional neurons, selecting any  $p$  ( $p \leq d$ ) dimensions and applying an shared orthogonal transformation (orthogonal matrix of size  $p \times p$ ) to these  $p$  dimensions of all neurons will not change the hyperspherical energy.*

A description of S-OPT is given in Algorithm 1. S-OPT has outer and inner iterations. In each inner iteration, the training is almost the same as OPT, except that the orthogonal matrix transforms a subset of the dimensions and the learnable orthogonal matrix has to be re-initialized as an identity matrix. The selection of neuron dimension is randomized in every outer iteration such that all neuron dimensions are sufficiently covered as the number of outer iteration increases. As a parallel direction to improve the scalability of OPT, we further propose a parameter-efficient OPT in Appendix I. This OPT variant explores structure priors in  $\mathbf{R}$  to improve parameter efficiency.

---

**Algorithm 1** Stochastic OPT

---

```

for  $i = 1, 2, \dots, N_{\text{out}}$  do
  for  $j = 1, 2, \dots, N_{\text{in}}$  do
    1. Randomly select  $p$  dimensions from  $d$ -dimensional neurons in the same layer.
    2. Construct an orthogonal matrix  $\mathbf{R}_p \in \mathbb{R}^{p \times p}$  and initialize it as identity matrix.
    3. Update  $\mathbf{R}_p$  with one iteration of OPT.
  end
  4. Multiply  $\mathbf{R}_p$  back to the  $p$ -dim sub-vectors from the  $d$ -dim neurons for transformation.
end

```

---

## 5 INSIGHTS AND DISCUSSIONS

### 5.1 OPTIMIZATION LANDSCAPE

We follow [40] to visualize the loss landscapes of both standard training and OPT in Fig. 4. For standard training, we perturb the parameter space of all the neurons (*i.e.*, filters). For OPT, we perturb the parameter space of all the trainable matrices (*i.e.*,  $\mathbf{P}$  in Fig. 2), because OPT does not directly learn neuron weights. The general idea is to use two random vectors (*e.g.*, normal distribution) to perturb the parameter space and obtain the loss value with the perturbed network parameters. Details and full results about the visualization are given in Appendix E. The loss landscape of standard training has extremely sharp minima. The red region is very flat, leading to small gradients. In contrast, the loss landscape of OPT is more smooth and convex with flatter minima, matching the finding that flat minimizers generalize well [8, 22, 26]. Moreover, the testing error landscape in Appendix E and additional loss landscape visualization results in Appendix F (with uniform perturbation distributions) well support our argument.

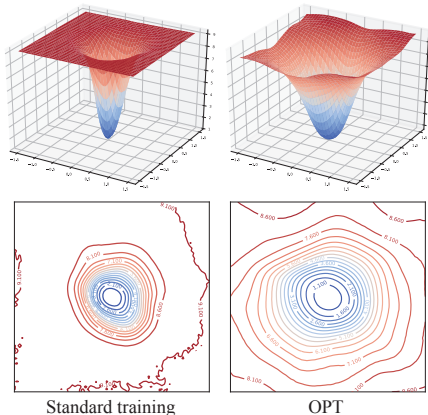


Figure 4: Loss landscape visualization.

### 5.2 OPTIMIZATION AND GENERALIZATION

This subsection discusses why OPT improves optimization and generalization. On one hand, [67] proves that once the neurons are hyperspherically diverse enough in a one-hidden-layer network, the training loss is on the order of the square norm of the gradient and the generalization error will have an additional term  $\tilde{O}(1/\sqrt{m})$  where  $m$  is the number of samples. This suggests that SGD-optimized networks with minimum hyperspherical energy (MHE) attained have no spurious local minima. Since OPT is guaranteed to achieve MHE in expectation, OPT-trained networks enjoy the inductive bias induced by MHE. On the other hand, [1, 12, 16, 31, 42] shows that over-parameterization in neural networks improves the first-order optimization, leads to better generalization, and imposes implicit regularizations. In the light of this, OPT also introduces over-parameterization to each neuron in the neural network, which shares similar spirit to [43]. Specifically, one  $d$ -dimensional neuron has  $d^2 + d$  parameters in OPT, compared to  $d$  parameters in the baseline ( $d^2$  parameters are shared across neurons of the same layer). Although OPT uses more parameters to represent a neuron in training, the equivalent number of parameters for a neuron stays unchanged and will not affect testing speed.

### 5.3 DISCUSSIONS

**Coordinate system and relative position.** OPT shows that learning the coordinate system yields better generalization than learning neuron weights directly. This implies that the coordinate system is crucial to generalization. However, the relative position does not matter only when the hyperspherical energy is sufficiently low, indicating that hyperspherical energy well characterizes the relative positions among neurons. Lower hyperspherical energy generally leads to better generalization.

**The effects of neuron norm.** Because we will normalize the neuron norm when computing the hyperspherical energy, the effects of neuron norm will not be taken into consideration. Moreover,

simply learning the orthogonal matrices will not change the neuron norm. Therefore, the neuron norm may affect the training. We use an extreme example to demonstrate the effects. Assume that one of the neurons have norm 1000 and the other neurons have norm 0.01. Then no matter what orthogonal matrices we have learned, the final performance will be bad. In this case, the hyperspherical energy can still be minimized to a very low value, but it can not capture the norm distribution. Fortunately, such an extreme case is unlikely to happen, because we are using zero-mean Gaussian distribution to initialize the neuron and every neuron also has the same expected value for the norm. To eliminate the effects of norms, we can normalize the neuron weights in the training, as proposed in Section 3.7.

## 6 EXPERIMENTS AND RESULTS

The section aims to show that OPT can effectively improve all kinds of neural networks. Detailed experimental settings are given in Appendix D. We put many additional experiments in Appendix I,K,M.

### 6.1 ABLATION STUDY AND EXPLORATORY EXPERIMENTS

**Necessity of orthogonality.** We test 6-layer and 9-layer CNN (Appendix D) on CIFAR-100. OPT is compared to unconstrained over-parameterized training (UPT) which learns an unconstrained matrix  $R$  (with only weight decay) using the same network. In Table 1, “FN” denotes whether the randomly initialized neuron weights are fixed in the training. “LR” denotes whether the learnable transformation  $R$  is unconstrained (“U”) or orthogonal (“GS” for Gram-Schmidt process). Table 1 shows that without ensuring orthogonality, UPT performs much worse than OPT.

**Fixed weights vs. learnable weights.** From Table 1, we can see that using fixed neuron weights is consistently better than learnable neuron weights in both UPT and OPT. It indicates that fixing the neuron weights while learning the transformation matrix  $R$  is very beneficial to generalization.

**Refining neuron initialization.** We evaluate two refinement methods (introduced in Section 3.7) for neuron initialization. First, we consider the hyperspherical energy minimization as a preprocessing for the neuron weights. Our experiment uses CNN-6 on CIFAR-100. Specifically, we run gradient descent for 5k iterations to minimize the objective of MHE/HS-MHE [45] before the training starts. Table 2 shows the hyperspherical energy before and after the preprocessing. All methods start with the same random initialization, so all hyperspherical energies start at 3.5109. Testing errors (%) in Table 2 show that the refinement well improves OPT.

Then we evaluate the neuron weight normalization. Section 5.3 justifies the superiority of normalizing the neurons. After random initialization, we normalize the scale of all the neuron weights to 1. These randomly initialized neurons still possess the important property of achieving minimum energy. We conduct classification with CNN-6 on CIFAR-100. Testing errors in Table 3 show that normalizing the neurons greatly improves OPT. Note that, these two refinements are not used by default in other experiments.

### 6.2 EMPIRICAL EVALUATION ON OPT

**Multi-layer perceptrons.** We evaluate OPT on MNIST with a 3-layer MLP. Appendix D gives specific settings. Table 4 shows the testing error with normal initialization (MLP-N) or Xavier initialization [14] (MLP-X). GS/HR/LS denote different orthogonalization unrolling. CP denotes Cayley parameterization. OGD denotes orthogonal-preserving gradient descent. OR denotes relaxed orthogonal regularization. All OPT variants outperform the others by a large margin.

**Convolutional networks.** We evaluate OPT with 6/9-layer plain CNNs and ResNet-20/32 [20] on CIFAR-100. Detailed settings are given in Appendix D. All neurons (*i.e.*, convolution kernels) are initialized by [19]. Batch normalization [25] is used by default. Table 4 shows that all OPT variants outperform both baseline and HS-MHE [45] by a large margin. HS-MHE puts the half-space hyperspherical energy into the loss function and naively minimizes it along with the CNN. We observe that OPT (HR) performs the worse among all OPT variants partially because of its intensive unrolling computation. OPT (GS) achieves the best testing error on CNN-6/9, while OPT (CP) achieves the best testing error on ResNet-20/34, implying that different OPT may encode different inductive biases.

| Method   | FN | LR | CNN-6        | CNN-9        |
|----------|----|----|--------------|--------------|
| Baseline | -  | -  | 37.59        | 33.55        |
| UPT      | ✗  | U  | 48.47        | 46.72        |
| UPT      | ✓  | U  | 42.61        | 39.38        |
| OPT      | ✗  | GS | 37.24        | 32.95        |
| OPT      | ✓  | GS | <b>33.02</b> | <b>31.03</b> |

Table 1: Error (%) on C-100.

“LR” denotes whether the learnable transformation  $R$  is unconstrained (“U”) or orthogonal (“GS” for Gram-Schmidt process). Table 1 shows that without ensuring orthogonality, UPT performs much worse than OPT.

From Table 1, we can see that using fixed neuron weights is consistently better than learnable neuron weights in both UPT and OPT. It indicates that fixing the neuron weights while learning the transformation matrix  $R$  is very beneficial to generalization.

| Method   | Original | MHE    | HS-MHE        |
|----------|----------|--------|---------------|
| OPT (GS) | 33.02    | 32.99  | <b>32.78</b>  |
| OPT (LS) | 34.48    | 34.43  | <b>34.37</b>  |
| OPT (CP) | 33.53    | 33.50  | <b>33.42</b>  |
| Energy   | 3.5109   | 3.5003 | <b>3.4976</b> |

Table 2: Refining energy.

| Method    | w/o Norm | w/ Norm      |
|-----------|----------|--------------|
| Baseline  | 37.59    | <b>36.05</b> |
| OPT (GS)  | 33.02    | <b>32.54</b> |
| OPT (HR)  | 35.67    | <b>35.30</b> |
| OPT (LS)  | 34.48    | <b>32.11</b> |
| OPT (CP)  | 33.53    | <b>32.49</b> |
| OPT (OGD) | 33.37    | <b>32.70</b> |
| OPT (OR)  | 34.70    | <b>33.27</b> |

Table 3: Normalization (%).

| Method         | MNIST       |             | CIFAR-100    |              |              |              |
|----------------|-------------|-------------|--------------|--------------|--------------|--------------|
|                | MLP-N       | MLP-X       | CNN-6        | CNN-9        | ResNet-20    | ResNet-32    |
| Baseline       | 6.05        | 2.14        | 37.59        | 33.55        | 31.11        | 30.16        |
| Orthogonal [7] | 5.78        | 1.93        | 36.32        | 33.24        | 31.06        | 30.05        |
| SRIP [4]       | -           | -           | 34.82        | 32.72        | 30.89        | 29.70        |
| HS-MHE [45]    | 5.57        | 1.88        | 34.97        | 32.87        | 30.98        | 29.76        |
| OPT (GS)       | <b>5.11</b> | <b>1.45</b> | <b>33.02</b> | <b>31.03</b> | 30.49        | 29.34        |
| OPT (HR)       | 5.31        | 1.60        | 35.67        | 32.75        | 30.73        | 29.56        |
| OPT (LS)       | 5.32        | 1.54        | 34.48        | 31.22        | 30.51        | 29.42        |
| OPT (CP)       | 5.14        | 1.49        | 33.53        | 31.28        | <b>30.47</b> | <b>29.31</b> |
| OPT (OGD)      | 5.38        | 1.56        | 33.33        | 31.47        | 30.50        | 29.39        |
| OPT (OR)       | 5.41        | 1.78        | 34.70        | 32.63        | 30.66        | 29.47        |

Table 4: Error (%) of OPT for MLPs and CNNs.

**Training dynamics.** We further look into how hyperspherical energy and testing error changes while training with OPT. Fig. 5 shows that the energy of the baseline will increase dramatically at the beginning and then gradually go down, but it still stays in a high value in the end. HS-MHE can effectively reduce the hyperspherical energy at the end of the training. In contrast, all OPT variants maintain minimal energy in training. OPT with GS, CP and OGD keeps exactly the same energy as the randomly initialized neurons, while OPT (OR) slightly increases the energy due to the relaxation. For convergence, all OPT variants converge efficiently and stably.




Figure 5: Training dynamics on CIFAR-100. Left: Hyperspherical energy vs. iteration. Right: Testing error vs. iteration.

**ImageNet.** We test OPT on the large-scale ImageNet-2012 dataset. Specifically, we use OPT with OGD and CP to train a plain 10-layer CNN (Appendix D) on ImageNet. Our purpose is to validate the superiority of OPT over the corresponding baseline. Table 5 shows that OPT (CP) reduces top-1 and top-5 error by  $\sim 0.7\%$  and  $\sim 0.9\%$ , respectively.

| Method         | Top-1 Err.   | Top-5 Err.   |
|----------------|--------------|--------------|
| Baseline       | 44.32        | 21.13        |
| Orthogonal [7] | 44.13        | 20.97        |
| HS-MHE [45]    | 43.92        | 20.85        |
| OPT (OGD)      | 43.81        | 20.49        |
| OPT (CP)       | <b>43.67</b> | <b>20.26</b> |

Table 5: ImageNet (%).

**Few-shot learning.** To evaluate the cross-task generalization of OPT, we conduct few-shot learning on Mini-ImageNet, following the same experimental setting as [10]. Detailed settings are in Appendix D. Specifically, we apply OPT with CP to train the baseline and baseline++ described in [10], and immediately obtain obvious improvements. Therefore, OPT-trained networks also generalize well in the challenging few-shot case.

| Method            | 5-shot Acc. (%)                    |
|-------------------|------------------------------------|
| MAML [13]         | 62.71 $\pm$ 0.71                   |
| ProtoNet [58]     | 64.24 $\pm$ 0.72                   |
| Baseline [10]     | 62.53 $\pm$ 0.69                   |
| Baseline w/ OPT   | <b>63.27 <math>\pm</math> 0.68</b> |
| Baseline++ [10]   | 66.43 $\pm$ 0.63                   |
| Baseline++ w/ OPT | <b>66.82 <math>\pm</math> 0.62</b> |

Table 6: Few-shot learning.

**Geometric networks.** We apply OPT to graph convolution network (GCN) [34] and point cloud neural network (PointNet) [51] for graph node and point cloud classification, respectively. The training of GCN and PointNet is conceptually similar to MLP, and the detailed training procedures are given in Appendix D. For GCN, we evaluate OPT on Cora and Pubmed datasets [56]. For PointNet, we conduct experiments on ModelNet-40 dataset [66]. Table 7 shows that OPT effectively improves both GCN and PointNet.

| Method    | GCN         |             | PointNet     |
|-----------|-------------|-------------|--------------|
|           | Cora        | Pubmed      | MN-40        |
| Baseline  | 81.3        | 79.0        | 87.1         |
| OPT (GS)  | 81.9        | 79.4        | 87.23        |
| OPT (CP)  | 82.0        | 79.4        | 87.81        |
| OPT (OGD) | <b>82.3</b> | <b>79.5</b> | <b>87.86</b> |

Table 7: Geometric networks.

### 6.3 EMPIRICAL EVALUATION ON S-OPT

**Convolutional networks.** S-OPT is a scalable OPT variant, and we evaluate its performance in terms of number of trainable parameters and testing error. Training parameters are learnable

variables in training, which are different from model parameters in testing. In testing, all methods have the same number of model parameters. We perform classification on CIFAR-100 with CNN-6 and wide CNN-9. We also evaluate S-OPT with wide ResNet-18 on ImageNet. Detailed settings are given in Appendix D. For S-OPT, we set the sampling dimension as 25% of the original neuron dimension in each layer. Table 8 shows that S-OPT achieves a good trade-off between accuracy and efficiency. Most importantly, S-OPT is able to train on wide neural networks on large-scale datasets, making OPT more useful in practice. As a parallel result, we also propose and experiment an orthogonal transformation with structural parameter sharing to save parameters in OPT in Appendix I.

**Effect of sampling dimension.** Since the sampling dimension  $p$  is a hyperparameter, we study its effect by performing classification with wide CNN-9 on CIFAR-100. In Table 9,  $p = d/4$  means that we randomly sample 1/4 of the original neuron dimension in each layer, so  $p$  may vary in different layer.  $p = 16$  means that we sample 16 dimensions in each layer. 1M trainable parameters come from the last classifier layer, which cannot be saved. Table 9 shows that larger  $p$  yields better accuracy but consumes more trainable parameters.

| Method      | CIFAR-100    |             |              |              | ImageNet     |              |
|-------------|--------------|-------------|--------------|--------------|--------------|--------------|
|             | CNN-6        | Params      | Wide CNN-9   | Params       | ResNet-18    | Params       |
| Baseline    | 37.58        | 249K        | 28.03        | 2.97M        | 32.95        | 11.2M        |
| HS-MHE [45] | 37.58        | 249K        | 25.96        | 2.97M        | 32.50        | 11.2M        |
| OPT (GS)    | <b>33.02</b> | 1.73M       | OOM          | 16.6M        | OOM          | 92.9M        |
| S-OPT (GS)  | 33.70        | <b>170K</b> | <b>25.59</b> | <b>2.02M</b> | <b>32.26</b> | <b>5.77M</b> |

Table 8: OPT vs. S-OPT on CIFAR-100 & ImageNet.

| Dimension $p$ | $d/4$        | $d/8$ | $d/16$ | 16     | 3             |
|---------------|--------------|-------|--------|--------|---------------|
| S-OPT (GS)    | <b>25.59</b> | 28.61 | 32.52  | 33.03  | 45.22         |
| Params        | 2.02M        | 1.29M | 1.11M  | 1.051M | <b>1.049M</b> |

Table 9: Effect of sampling dimension.

## 7 CONCLUDING REMARKS

We propose a novel training framework that effectively minimizes hyperspherical energy for neural networks. OPT over-parameterizes neurons with neuron weights (randomly initialized and fixed) and a layer-shared orthogonal matrix (learnable). Experiments validate OPT’s strong generalizability.



## REFERENCES

- [1] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. *arXiv preprint arXiv:1811.03962*, 2018. 6, 36
- [2] Ramesh Naidu Annavarapu. Singular value decomposition and the centrality of löwdin orthogonalizations. *American Journal of Computational and Applied Mathematics*, 3(1):33–35, 2013. 15
- [3] Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. In *ICML*, 2016. 2, 4
- [4] Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. Can we gain more from orthogonality regularizations in training deep cnns? In *NeurIPS*, 2018. 2, 5, 7
- [5] Johann S Brauchart, Alexander B Reznikov, Edward B Saff, Ian H Sloan, Yu Guang Wang, and Robert S Womersley. Random point sets on the sphere—hole radii, covering, and separation. *Experimental Mathematics*, 27(1):62–81, 2018. 42, 43
- [6] Anna Breger, Martin Ehler, and Manuel Gräf. Points on manifolds with asymptotically optimal covering radius. *Journal of Complexity*, 48:1–14, 2018. 42
- [7] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Neural photo editing with introspective adversarial networks. In *ICLR*, 2017. 2, 7, 8
- [8] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. In *ICLR*, 2017. 6
- [9] Beidi Chen, Weiyang Liu, Animesh Garg, Zhiding Yu, Anshumali Shrivastava, and Animashree Anandkumar. Angular visual hardness. In *ICML*, 2020. 2
- [10] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. *arXiv preprint arXiv:1904.04232*, 2019. 8, 19, 20
- [11] Xiaohan Ding, Yuchen Guo, Guiguang Ding, and Jungong Han. Acnet: Strengthening the kernel skeletons for powerful cnn via asymmetric convolution blocks. In *ICCV*, 2019. 1, 2
- [12] Simon S Du, Jason D Lee, Yuandong Tian, Barnabas Póczos, and Aarti Singh. Gradient descent learns one-hidden-layer cnn: Don’t be afraid of spurious local minima. *arXiv preprint arXiv:1712.00779*, 2017. 6, 36
- [13] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. 8
- [14] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 2010. 1, 3, 7
- [15] Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In *NeurIPS*, 2018. 1
- [16] Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In *NeurIPS*, 2017. 1, 6
- [17] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016. 40
- [18] DP Hardin and EB Saff. Minimal riesz energy point configurations for rectifiable d-dimensional manifolds. *Advances in Mathematics*, 193(1):174–204, 2005. 17
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 2015. 1, 3, 7
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 7, 19, 21
- [21] Mikael Henaff, Arthur Szlam, and Yann LeCun. Recurrent orthogonal networks and long-memory tasks. *arXiv preprint arXiv:1602.06662*, 2016. 4
- [22] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural Computation*, 1997. 6

- [23] Walter Hoffmann. Iterative algorithms for gram-schmidt orthogonalization. *Computing*, 41(4):335–348, 1989. 4, 13
- [24] Lei Huang, Xianglong Liu, Bo Lang, Adams Wei Yu, Yongliang Wang, and Bo Li. Orthogonal weight normalization: Solution to optimization over multiple dependent stiefel manifolds in deep neural networks. In *AAAI*, 2018. 2, 4
- [25] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 7
- [26] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *UAI*, 2018. 6
- [27] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. In *BMVC*, 2014. 2
- [28] Kui Jia, Shuai Li, Yuxin Wen, Tongliang Liu, and Dacheng Tao. Orthogonal deep neural networks. *TPAMI*, 2019. 2
- [29] Xu Jia, Bert De Brabandere, Tinne Tuytelaars, and Luc V Gool. Dynamic filter networks. In *NeurIPS*, 2016. 40
- [30] Li Jing, Yichen Shen, Tena Dubcek, John Peurifoy, Scott Skirlo, Yann LeCun, Max Tegmark, and Marin Soljačić. Tunable efficient unitary neural networks (eunn) and their application to rnns. In *ICML*, 2017. 4
- [31] Kenji Kawaguchi. Deep learning without poor local minima. In *NeurIPS*, 2016. 1, 6, 36
- [32] Kenji Kawaguchi, Bo Xie, and Le Song. Deep semi-random features for nonlinear function approximation. In *AAAI*, 2018. 37
- [33] Nitish Shirish Keskar and Richard Socher. Improving generalization performance by switching from adam to sgd. *arXiv preprint arXiv:1712.07628*, 2017. 1
- [34] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 8
- [35] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012. 21
- [36] N.S. Landkof. *Foundations of modern potential theory*. Springer-Verlag, 1972. 17
- [37] Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent only converges to minimizers. In *COLT*, 2016. 36
- [38] Mario Lezcano-Casado. Trivializations for gradient-based optimization on manifolds. In *NeurIPS*, 2019. 4
- [39] Mario Lezcano-Casado and David Martínez-Rubio. Cheap orthogonal constraints in neural networks: A simple parametrization of the orthogonal and unitary group. In *ICML*, 2019. 2, 4
- [40] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *NeurIPS*, 2018. 6, 22, 29
- [41] Jun Li, Fuxin Li, and Sinisa Todorovic. Efficient riemannian optimization on the stiefel manifold via the cayley transform. In *ICLR*, 2020. 4, 5
- [42] Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *COLT*, 2018. 1, 6
- [43] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013. 6, 40
- [44] Baoyuan Liu, Min Wang, Hassan Foroosh, Marshall Tappen, and Marianna Pinsky. Sparse convolutional neural networks. In *CVPR*, 2015. 2
- [45] Weiyang Liu, Rongmei Lin, Zhen Liu, Lixin Liu, Zhiding Yu, Bo Dai, and Le Song. Learning towards minimum hyperspherical energy. In *NeurIPS*, 2018. 2, 3, 5, 7, 8, 22, 29, 36, 43, 44
- [46] Weiyang Liu, Zhen Liu, James M Rehg, and Le Song. Neural similarity learning. In *NeurIPS*, 2019. 1, 2, 40

- [47] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. SpheroFace: Deep hypersphere embedding for face recognition. In *CVPR*, 2017. 2
- [48] Weiyang Liu, Yan-Ming Zhang, Xingguo Li, Zhiding Yu, Bo Dai, Tuo Zhao, and Le Song. Deep hyperspherical learning. In *NeurIPS*, 2017. 2, 5, 21
- [49] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *ECCV*, 2018. 37
- [50] Zakaria Mhammedi, Andrew Hellicar, Ashfaqur Rahman, and James Bailey. Efficient orthogonal parametrization of recurrent neural networks using householder reflections. In *ICML*, 2017. 4
- [51] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 8, 20
- [52] Haozhi Qi, Chong You, Xiaolong Wang, Yi Ma, and Jitendra Malik. Deep isometric learning for visual recognition. *arXiv preprint arXiv:2006.16992*, 2020. 2
- [53] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *NeurIPS*, 2008. 37
- [54] Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. *arXiv preprint arXiv:1904.09237*, 2019. 1
- [55] Matthias Reitzner. Stochastic approximation of smooth convex bodies. *Mathematika*, 51(1-2):11–29, 2004. 42
- [56] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 2008. 8, 20
- [57] Ian H Sloan and Robert S Womersley. Extremal systems of points and numerical integration on the sphere. *Advances in Computational Mathematics*, 21(1-2):107–125, 2004. 43
- [58] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017. 8
- [59] Daniel Soudry and Yair Carmon. No bad local minima: Data independent training error guarantees for multilayer neural networks. *arXiv preprint arXiv:1605.08361*, 2016. 36
- [60] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *JMLR*, 2014. 37
- [61] Vipin Srivastava. A unified view of the orthogonalization methods. *Journal of Physics A: Mathematical and General*, 33(35):6219, 2000. 15
- [62] Yun Tang, Jing Huang, Guangtao Wang, Xiaodong He, and Bowen Zhou. Orthogonal relation transforms with graph context modeling for knowledge graph embedding. In *ACL*, 2020. 4
- [63] Min Wang, Baoyuan Liu, and Hassan Foroosh. Factorized convolutional neural networks. In *ICCV Workshops*, 2017. 2
- [64] Zaiwen Wen and Wotao Yin. A feasible method for optimization with orthogonality constraints. *Mathematical Programming*, 2013. 4
- [65] Scott Wisdom, Thomas Powers, John Hershey, Jonathan Le Roux, and Les Atlas. Full-capacity unitary recurrent neural networks. In *NeurIPS*, 2016. 2, 4
- [66] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, 2015. 8, 20
- [67] Bo Xie, Yingyu Liang, and Le Song. Diverse neural network learns true target functions. In *AISTATS*, 2017. 6, 36, 37
- [68] Zichao Yang, Marcin Moczulski, Misha Denil, Nando de Freitas, Alex Smola, Le Song, and Ziyu Wang. Deep fried convnets. In *ICCV*, 2015. 2

# Appendix

## A DETAILS OF UNROLLED ORTHOGONALIZATION ALGORITHMS

### A.1 GRAM-SCHMIDT PROCESS

**Gram-Schmidt Process.** GS process is a method for orthonormalizing a set of vectors in an inner product space, *i.e.*, the Euclidean space  $\mathbb{R}^n$  equipped with the standard inner product. Specifically, GS process performs the following operations to orthogonalize a set of vectors  $\{\mathbf{u}_1, \dots, \mathbf{u}_n\} \in \mathbb{R}^{n \times n}$ :

$$\text{Step 1: } \tilde{\mathbf{e}}_1 = \mathbf{u}_1, \quad \mathbf{e}_1 = \frac{\tilde{\mathbf{e}}_1}{\|\tilde{\mathbf{e}}_1\|}$$

$$\text{Step 2: } \tilde{\mathbf{e}}_2 = \mathbf{u}_2 - \text{Proj}_{\tilde{\mathbf{e}}_1}(\mathbf{u}_2), \quad \mathbf{e}_2 = \frac{\tilde{\mathbf{e}}_2}{\|\tilde{\mathbf{e}}_2\|}$$

$$\text{Step 3: } \tilde{\mathbf{e}}_3 = \mathbf{u}_3 - \text{Proj}_{\tilde{\mathbf{e}}_1}(\mathbf{u}_3) - \text{Proj}_{\tilde{\mathbf{e}}_2}(\mathbf{u}_3), \quad \mathbf{e}_3 = \frac{\tilde{\mathbf{e}}_3}{\|\tilde{\mathbf{e}}_3\|}$$

$$\text{Step 4: } \tilde{\mathbf{e}}_4 = \mathbf{u}_4 - \text{Proj}_{\tilde{\mathbf{e}}_1}(\mathbf{u}_4) - \text{Proj}_{\tilde{\mathbf{e}}_2}(\mathbf{u}_4) - \text{Proj}_{\tilde{\mathbf{e}}_3}(\mathbf{u}_4), \quad \mathbf{e}_4 = \frac{\tilde{\mathbf{e}}_4}{\|\tilde{\mathbf{e}}_4\|}$$

⋮

$$\text{Step n: } \tilde{\mathbf{e}}_n = \mathbf{u}_n - \text{Proj}_{\tilde{\mathbf{e}}_1}(\mathbf{u}_n) - \text{Proj}_{\tilde{\mathbf{e}}_2}(\mathbf{u}_n) - \text{Proj}_{\tilde{\mathbf{e}}_3}(\mathbf{u}_n) - \dots - \text{Proj}_{\tilde{\mathbf{e}}_{n-1}}(\mathbf{u}_n), \quad \mathbf{e}_n = \frac{\tilde{\mathbf{e}}_n}{\|\tilde{\mathbf{e}}_n\|} \quad (6)$$

where  $\text{Proj}_{\mathbf{a}}(\mathbf{b}) = \frac{\langle \mathbf{a}, \mathbf{b} \rangle}{\langle \mathbf{a}, \mathbf{a} \rangle} \mathbf{a}$  denotes the projection of the vector  $\mathbf{b}$  onto the vector  $\mathbf{a}$ . The set  $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$  denotes the output orthonormal set. The algorithm flowchart can be described as follows:

---

#### Algorithm 2 Gram-Schmidt Process

---

**Input:**  $U = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n\} \in \mathbb{R}^{n \times n}$

**Output:**  $R = \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\} \in \mathbb{R}^{n \times n}$

$R = \mathbf{0}$

**for**  $j = 1, 2, \dots, n$  **do**

$$\left| \begin{array}{l} \mathbf{q}_j = R^\top \mathbf{u}_j \\ \mathbf{t} = \mathbf{u}_j - R \mathbf{q}_j \\ \mathbf{q}_{jj} = \|\mathbf{t}\|_2 \\ \mathbf{e}_j = \frac{\mathbf{t}}{\mathbf{q}_{jj}} \end{array} \right.$$

**end**

---

The vectors  $\mathbf{q}_j, \forall j$  in the algorithm above are used to compute the QR factorization, which is not useful in orthogonalization and therefore does not need to be stored. When the GS process is implemented on a finite-precision computer, the vectors  $\mathbf{e}_j, \forall j$  are often not quite orthogonal, because of rounding errors. Besides the standard GS process, there is a modified Gram-Schmidt (MGS) algorithm which enjoys better numerical stability. This approach gives the same result as the original formula in exact arithmetic and introduces smaller errors in finite-precision arithmetic. Specifically, GS computes the following formula:

$$\begin{aligned} \tilde{\mathbf{e}}_j &= \mathbf{u}_j - \sum_{k=1}^{j-1} \text{Proj}_{\tilde{\mathbf{e}}_k}(\mathbf{u}_j) \\ \mathbf{e}_j &= \frac{\tilde{\mathbf{e}}_j}{\|\tilde{\mathbf{e}}_j\|} \end{aligned} \quad (7)$$

Instead of computing the vector  $\mathbf{e}_j$  as in Eq. 7, MGS computes the orthogonal basis differently. MGS does not subtract the projections of the original vector set, and instead remove the projection of

the previously constructed orthogonal basis. Specifically, MGS computes the following series of formulas:

$$\begin{aligned}
\tilde{e}_j^{(1)} &= \mathbf{u}_j - \text{Proj}_{\tilde{e}_1}(\mathbf{u}_j) \\
\tilde{e}_j^{(2)} &= \tilde{e}_j^{(1)} - \text{Proj}_{\tilde{e}_2}(\tilde{e}_j^{(1)}) \\
&\vdots \\
\tilde{e}_j^{(j-2)} &= \tilde{e}_j^{(j-3)} - \text{Proj}_{\tilde{e}_2}(\tilde{e}_j^{(j-3)}) \\
\tilde{e}_j^{(j-1)} &= \tilde{e}_j^{(j-2)} - \text{Proj}_{\tilde{e}_2}(\tilde{e}_j^{(j-2)}) \\
\mathbf{e}_j &= \frac{\tilde{e}_j^{(j-1)}}{\|\tilde{e}_j^{(j-1)}\|}
\end{aligned} \tag{8}$$

where each step finds a vector  $\tilde{e}_j^{(i)}$  that is orthogonal to  $\tilde{e}_j^{(i-1)}$ . Therefore,  $\tilde{e}_j^{(i)}$  is also orthogonalized against any errors brought by the computation of  $\tilde{e}_j^{(i-1)}$ . In practice, although MGS enjoys better numerical stability, we find the empirical performance of GS and MGS is almost the same in OPT. However, MGS takes longer time to complete since the computation of each orthogonal basis is an iterative process. Therefore, we usually stick to classic GS for OPT.

**Iterative Gram-Schmidt Process.** Iterative Gram-Schmidt (IGS) process is an iterative version of the GS process. It is shown in [23] that GS process can be carried out iteratively to obtain a basis matrix that is orthogonal in almost full working precision. The IGS algorithm is given as follows:

---

**Algorithm 3** Iterative Gram-Schmidt Process

---

**Input:**  $U = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n\} \in \mathbb{R}^{n \times n}$

**Output:**  $R = \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\} \in \mathbb{R}^{n \times n}$

$R = \mathbf{0}$

**for**  $j = 1, 2, \dots, n$  **do**

$\mathbf{q}_j = \mathbf{0}$

$\mathbf{t} = \mathbf{u}_j$

**while**  $\mathbf{t} \perp \text{span}(\mathbf{e}_1, \dots, \mathbf{e}_{j-1})$  **is False do**

$\mathbf{p} = \mathbf{t}$

$\mathbf{s} = R^\top \mathbf{p}$

$\mathbf{v} = R\mathbf{s}$

$\mathbf{t} = \mathbf{p} - \mathbf{v}$

$\mathbf{q}_j \leftarrow \mathbf{q}_j + \mathbf{s}$

**end**

$\mathbf{q}_{jj} = \|\mathbf{t}\|_2$

$\mathbf{e}_j = \frac{\mathbf{t}}{\mathbf{q}_{jj}}$

**end**

---

The vectors  $\mathbf{q}_j, \forall j$  in the algorithm above are used to compute the QR factorization, which is not useful in orthogonalization and therefore does not need to be explicitly computed. The while loop in IGS is an iterative procedure. In practice, we can unroll a fixed number of steps for the while loop in order to improve the orthogonality. The resulting  $\mathbf{q}_j$  in the  $j$ -th step corresponds to the solution of the equation  $\tilde{R}^\top \tilde{R}\mathbf{q}_j = \tilde{R}^\top \mathbf{u}_j$  where  $\tilde{R} = \{\mathbf{e}_1, \dots, \mathbf{e}_{j-1}\}$ . The IGS process corresponds to the Gauss-Jacobi iteration for solving this equation.

Both GS and IGS are easy to be embedded in the neural networks, since they are both differentiable. In our experiments, we find that the performance gain of unrolling multiple steps in IGS over GS is not very obvious (partially because GS has already achieved nearly perfect orthogonality), but IGS costs longer training time. Therefore, we unroll the classic GS process by default.

## A.2 HOUSEHOLDER REFLECTION

Let  $\mathbf{v} \in \mathbb{R}^n$  be a non-zero vector. A matrix  $\mathbf{H} \in \mathbb{R}^{n \times n}$  of the form

$$\mathbf{H} = \mathbf{I} - \frac{2\mathbf{v}\mathbf{v}^\top}{\mathbf{v}^\top\mathbf{v}} \quad (9)$$

is a Householder reflection. The vector  $\mathbf{v}$  is the Householder vector. If a vector  $\mathbf{x}$  is multiplied by the matrix  $\mathbf{H}$ , then it will be reflected in the hyperplane  $\text{span}(\mathbf{v})^\perp$ . Householder matrices are symmetric and orthogonal.

For a vector  $\mathbf{x} \in \mathbb{R}^n$ , we let  $\mathbf{v} = \mathbf{x} \pm \|\mathbf{x}\|_2 \mathbf{e}_1$  where  $\mathbf{e}_1$  is a vector of  $\{1, 0, \dots, 0\}$  (the first element is 1 and the remaining elements are 0). Then we construct the Householder reflection matrix with  $\mathbf{v}$  and multiply it to  $\mathbf{x}$ :

$$\mathbf{H}\mathbf{x} = \left( \mathbf{I} - 2\frac{\mathbf{v}\mathbf{v}^\top}{\mathbf{v}^\top\mathbf{v}} \right) \mathbf{x} = \mp \|\mathbf{x}\|_2 \mathbf{e}_1 \quad (10)$$

which indicates that we can make any non-zero vector become  $\alpha\mathbf{e}_1$  where  $\alpha$  is some constant by using Householder reflection. By left-multiplying a reflection we can turn a dense vector  $\mathbf{x}$  into a vector with the same length and with only a single nonzero entry. Repeating this  $n$  times gives us the Householder QR factorization, which also orthogonalizes the original input matrix. Householder reflection orthogonalizes a matrix  $\mathbf{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$  by triangularizing it:

$$\mathbf{U} = \mathbf{H}_1\mathbf{H}_2 \cdots \mathbf{H}_n\mathbf{R} \quad (11)$$

where  $\mathbf{R}$  is a upper-triangular matrix in the QR factorization.  $\mathbf{H}_j, j \geq 2$  is constructed by  $\text{Diag}(\mathbf{I}_{j-1}, \tilde{\mathbf{H}}_{n-j+1})$  where  $\tilde{\mathbf{H}}_{n-j+1} \in \mathbb{R}^{(n-j+1) \times (n-j+1)}$  is the Householder reflection that is performed on the vector  $\mathbf{U}_{(j:n,j)}$ . The algorithm flowchart is given as follows:

---

**Algorithm 4** Householder Reflection Orthogonalization
 

---

**Input:**  $\mathbf{U} = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n\} \in \mathbb{R}^{n \times n}$

**Output:**  $\mathbf{U} = \mathbf{Q}\mathbf{R}$ , where  $\mathbf{Q} = \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\} \in \mathbb{R}^{n \times n}$  is the orthogonal matrix and  $\mathbf{R} \in \mathbb{R}^{n \times n}$  is a upper triangular matrix

**for**  $j = 1, 2, \dots, n - 1$  **do**

$\{\mathbf{v}, \beta\} = \text{Householder}(\mathbf{U}_{j:n,j})$   
      $\mathbf{U}_{j:n,j:n} \leftarrow \mathbf{U}_{j:n,j:n} - \beta\mathbf{v}(\mathbf{v}^\top\mathbf{U}_{j:n,j:n})$   
      $\mathbf{U}_{j+1:n,j} \leftarrow \mathbf{v}(2:\text{end})$

**end**

**function**  $\{\mathbf{v}, \beta\} = \text{Householder}(\mathbf{x})$

$\sigma^2 = \|\mathbf{x}_{2:\text{end}}\|_2^2$

$\mathbf{v} \leftarrow \begin{bmatrix} 1 \\ \mathbf{x}_{2:\text{end}} \end{bmatrix}$

**if**  $\sigma^2 = 0$  **then**

$\beta = 0$

**else**

**if**  $x_1 \leq 0$  **then**

$v_1 = x_1 - \sqrt{x_1^2 + \sigma^2}$

**else**

$v_1 = -\frac{\sigma^2}{x_1 + \sqrt{x_1^2 + \sigma^2}}$

**end**

$\beta = \frac{2v_1^2}{\sigma^2 + v_1^2}$

$\mathbf{v} \leftarrow \frac{\mathbf{v}}{v_1}$

**end**

**end function**

---

The algorithm follows the Matlab notation where  $\mathbf{U}_{j:n,j:n}$  denotes the submatrix of  $\mathbf{U}$  from the  $j$ -th column to the  $n$ -th column and from the  $j$ -th row to the  $n$ -th row. Note that, there are a number of

variants for the Householder reflection orthogonalization, such as the implicit variant where we do not store each reflection  $H_j$  explicitly. Here  $Q$  is the final orthogonal matrix we need.

### A.3 LÖWDIN’S SYMMETRIC ORTHOGONALIZATION

Let  $U = \{u_1, u_2, \dots, u_n\}$  be a set of linearly independent vectors in a  $n$ -dimensional space. We define a general non-singular linear transformation  $A$  that can transform the basis  $U$  to a new basis  $R$ :

$$R = UA \quad (12)$$

where the basis  $R$  will be orthonormal if (the transpose will become conjugate transpose in complex space)

$$R^\top R = (UA)^\top(UA) = A^\top U^\top UA = A^\top MA = I \quad (13)$$

where  $M = U^\top U$  is the gram matrix of the given basis  $U$ .

A general solution to this orthogonalization problem can be obtained via the substitution:

$$A = M^{-1}B \quad (14)$$

in which  $B$  is an arbitrary orthogonal (or unitary) matrix. When  $B = I$ , we will have the symmetric orthogonalization, namely

$$R := \Phi = UM^{-\frac{1}{2}} \quad (15)$$

When  $B = V$  in which  $V$  diagonalizes  $M$ , then we have the canonical orthogonalization, namely

$$\Lambda = UVd^{-\frac{1}{2}}. \quad (16)$$

Because  $V$  diagonalizes  $M$ , we have that  $M = VdV^\top$ . Therefore, we have the  $M^{-\frac{1}{2}}$  transformation as  $M^{-\frac{1}{2}} = Vd^{-\frac{1}{2}}V^\top$ . This is essentially an eigenvalue decomposition of the symmetric matrix  $M = U^\top U$ .

In order to compute the Löwdin’s symmetric orthogonalized basis sets, we can use singular value decomposition. Specifically, SVD of the original basis set  $U$  is given by

$$U = W\Sigma V^\top \quad (17)$$

where both  $W \in \mathbb{R}^{n \times n}$  and  $V \in \mathbb{R}^{n \times n}$  are orthogonal matrices.  $\Sigma$  is the diagonal matrix of singular values. Therefore, we have that

$$\begin{aligned} R &= UM^{-\frac{1}{2}} \\ &= W\Sigma V^\top Vd^{-\frac{1}{2}}V^\top \\ &= W\Sigma d^{-\frac{1}{2}}V^\top \end{aligned} \quad (18)$$

where we have  $\Sigma = d^{\frac{1}{2}}$  due to the connections between eigenvalue decomposition and SVD. Therefore, we end up with

$$R = WV^\top \quad (19)$$

which is the output orthogonal matrix for Löwdin’s symmetric orthogonalization.

An interesting feature of the symmetric orthogonalization is to ensure that

$$R = \arg \min_{P \in \text{orth}(U)} \sum_i \|P_i - U_i\| \quad (20)$$

where  $P_i$  and  $U_i$  are the  $i$ -th column vectors of  $P \in \mathbb{R}^{n \times n}$  and  $U$ , respectively.  $\text{orth}(U)$  denotes the set of all possible orthonormal sets in the range of  $U$ . This means that the symmetric orthogonalization functions  $R_i$  (or  $\Phi_i$ ) are the least distant in the Hilbert space from the original functions  $U_i$ . Therefore, symmetric orthogonalization indicates the gentlest pushing of the directions of the vectors in order to make them orthogonal.

More interestingly, the symmetric orthogonalized basis sets has unique geometric properties [2, 61] if we consider the Schweinler-Wigner matrix in terms of the sum of squared projections.

## B PROOF OF THEOREM 1

To be more specific, neurons with each element initialized by a zero-mean Gaussian distribution are uniformly distributed on a hypersphere. We show this argument with the following theorem.

**Theorem 3.** *The normalized vector of Gaussian variables is uniformly distributed on the sphere. Formally, let  $x_1, x_2, \dots, x_n \sim \mathcal{N}(0, 1)$  and be independent. Then the vector*

$$\mathbf{x} = \left[ \frac{x_1}{z}, \frac{x_2}{z}, \dots, \frac{x_n}{z} \right] \quad (21)$$

follows the uniform distribution on  $\mathbb{S}^{n-1}$ , where  $z = \sqrt{x_1^2 + x_2^2 + \dots + x_n^2}$  is a normalization factor.

*Proof.* A random variable has distribution  $\mathcal{N}(0, 1)$  if it has the density function

$$f(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2}. \quad (22)$$

A  $n$ -dimensional random vector  $\mathbf{x}$  has distribution  $\mathcal{N}(0, 1)$  if the components are independent and have distribution  $\mathcal{N}(0, 1)$  each. Then the density of  $\mathbf{x}$  is given by

$$f(\mathbf{x}) = \frac{1}{(\sqrt{2\pi})^n} e^{-\frac{1}{2}\langle \mathbf{x}, \mathbf{x} \rangle}. \quad (23)$$

Then we introduce the following lemma (Lemma 1) about the orthogonal-invariance of the normal distribution.

**Lemma 1.** *Let  $\mathbf{x}$  be a  $n$ -dimensional random vector with distribution  $\mathcal{N}(0, 1)$  and  $\mathbf{U} \in \mathbb{R}^{n \times n}$  be an orthogonal matrix ( $\mathbf{U}\mathbf{U}^\top = \mathbf{U}^\top\mathbf{U} = \mathbf{I}$ ). Then  $\mathbf{Y} = \mathbf{U}\mathbf{x}$  also has the distribution of  $\mathcal{N}(0, 1)$ .*

*Proof.* For any measurable set  $A \subset \mathbb{R}^n$ , we have that

$$\begin{aligned} P(\mathbf{Y} \in A) &= P(\mathbf{x} \in \mathbf{U}^\top A) \\ &= \int_{\mathbf{U}^\top A} \frac{1}{(\sqrt{2\pi})^n} e^{-\frac{1}{2}\langle \mathbf{x}, \mathbf{x} \rangle} \\ &= \int_A \frac{1}{(\sqrt{2\pi})^n} e^{-\frac{1}{2}\langle \mathbf{U}\mathbf{x}, \mathbf{U}\mathbf{x} \rangle} \\ &= \int_A \frac{1}{(\sqrt{2\pi})^n} e^{-\frac{1}{2}\langle \mathbf{x}, \mathbf{x} \rangle} \end{aligned} \quad (24)$$

because of orthogonality of  $\mathbf{U}$ . Therefore the lemma holds.  $\square$

Because any rotation is just a multiplication with some orthogonal matrix, we know that normally distributed random vectors are invariant to rotation. As a result, generating  $\mathbf{x} \in \mathbb{R}^n$  with distribution  $\mathcal{N}(0, 1)$  and then projecting it onto the hypersphere  $\mathbb{S}^{n-1}$  produces random vectors  $\mathbf{U} = \frac{\mathbf{x}}{\|\mathbf{x}\|}$  that are uniformly distributed on the hypersphere. Therefore the theorem holds.  $\square$

Then we show the normalized vector  $\mathbf{y}$  where each element follows a zero-mean Gaussian distribution with some constant variance  $\sigma^2$ :

$$\mathbf{y} = \left[ \frac{y_1}{r}, \frac{y_2}{r}, \dots, \frac{y_n}{r} \right] \quad (25)$$

where  $r = \sqrt{y_1^2 + y_2^2 + \dots + y_n^2}$ . Because we have that  $\frac{y_i}{\sigma} \sim \mathcal{N}(0, 1)$ , we can rewrite  $\mathbf{y}$  as the following random vector:

$$\mathbf{y} = \left[ \frac{y_1/\sigma}{r/\sigma}, \frac{y_2/\sigma}{r/\sigma}, \dots, \frac{y_n/\sigma}{r/\sigma} \right] \quad (26)$$

where  $r/\sigma = \sqrt{(y_1/\sigma)^2 + (y_2/\sigma)^2 + \dots + (y_n/\sigma)^2}$ . Therefore, we directly can apply Theorem 3 and conclude that  $\mathbf{y}$  also follows the uniform distribution on  $\mathbb{S}^{n-1}$ . Now we obtain that any random vector with each element following a zero-mean Gaussian distribution with some constant variance follows the uniform distribution on  $\mathbb{S}^{n-1}$ .



Then we show that the minimum hyperspherical energy asymptotically corresponds to the uniform distribution over the unit hypersphere. We first write down the hyperspherical energy of  $N$  neurons  $\{\mathbf{w}_1, \dots, \mathbf{w}_N \in \mathbb{R}^{d+1}\}$  (we also define that  $\hat{\mathbf{w}}_i = \frac{\mathbf{w}_i}{\|\mathbf{w}_i\|} \in \mathbb{S}^d$ ):

$$E_{s,d}(\hat{\mathbf{w}}_{i=1}^N) = \sum_{i=1}^N \sum_{j=1, j \neq i}^N f_s(\|\hat{\mathbf{w}}_i - \hat{\mathbf{w}}_j\|) = \begin{cases} \sum_{i \neq j} \|\hat{\mathbf{w}}_i - \hat{\mathbf{w}}_j\|^{-s}, & s > 0 \\ \sum_{i \neq j} \log(\|\hat{\mathbf{w}}_i - \hat{\mathbf{w}}_j\|^{-1}), & s = 0 \end{cases} \quad (27)$$

where  $s$  is a hyperparameter that controls the behavior of hyperspherical energy. We then define a  $N$ -point minimal hyperspherical  $s$ -energy over  $\mathbf{A}$  with

$$\varepsilon_{s,d}(\mathbf{A}, \hat{\mathbf{W}}_N) := \inf_{\hat{\mathbf{W}}_N \subset \mathbf{A}} E_{s,d}(\hat{\mathbf{w}}_{i=1}^N) \quad (28)$$

where we denote that  $\hat{\mathbf{W}}_N = \{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_N\}$ . Typically, we will assume that  $\mathbf{A}$  is compact. Based on [18], we discuss the asymptotic behavior (as  $N \rightarrow \infty$ ) of  $\varepsilon_{s,d}(\mathbf{A}, \hat{\mathbf{W}}_N)$  in three different scenarios: (1)  $0 < s < d$ ; (2)  $s = d$ ; and  $s > d$ . The reason behind is the behavior of the following energy integral:

$$I_s(\mu) = \iint_{\mathbb{S}^d \times \mathbb{S}^d} \|\mathbf{u} - \mathbf{v}\|^{-s} d\mu(\mathbf{u}) d\mu(\mathbf{v}), \quad (29)$$

is quite different under these three scenarios. In scenario (1), Eq. 29 that is taken over all probability measures  $\mu$  supported on  $\mathbb{S}^d$  will be minimal for normalized Lebesgue measure  $\frac{\mathcal{H}_d(\cdot)|_{\mathbb{S}^d}}{\mathcal{H}_d(\mathbb{S}^d)}$  on  $\mathbb{S}^d$ . In the case of  $s \geq d$ , we will have that  $I_s(\mu)$  is positive infinity for all such measures  $\mu$ . Therefore, the behaviour of the minimum hyperspherical energy is different in these three cases. In general, as the parameter  $s$  increases, there is a transition from the global effects to the more local influences (from nearest neighbors). The transition happens when  $s = d$ . However, we typically have  $0 < s < d$  in the neural networks. Therefore, we will mostly study the case of  $0 < s < d$  and the theoretical asymptotic behavior is quite standard results from the potential theory [36]. From the classic potential theory, we have the following known lemma:

**Lemma 2.** *If  $0 < s < d$ , we have that*

$$\lim_{N \rightarrow \infty} \frac{\varepsilon_{s,d}(\mathbb{S}^d, \hat{\mathbf{W}}_N)}{N^2} = I_s\left(\frac{\mathcal{H}_d(\cdot)|_{\mathbb{S}^d}}{\mathcal{H}_d(\mathbb{S}^d)}\right) \quad (30)$$

Moreover, any sequence of  $s$ -energy configuration of minimal hyperspherical energy  $(\hat{\mathbf{W}}_N^*)_{N=2}^\infty \subset \mathbb{S}^d$  is asymptotically uniformly distributed in the sense that for the weak-star topology of measures,

$$\frac{1}{N} \sum_{\mathbf{v} \in \hat{\mathbf{W}}_N^*} \delta_{\mathbf{v}} \rightarrow \frac{\mathcal{H}_d(\cdot)|_{\mathbb{S}^d}}{\mathcal{H}_d(\mathbb{S}^d)} \quad \text{as } N \rightarrow \infty \quad (31)$$

where  $\delta_{\mathbf{v}}$  denotes the unit point mass at  $\mathbf{v}$ .

The lemma above concludes that the neuron configuration with minimal hyperspherical energy asymptotically corresponds to the uniform distribution on  $\mathbb{S}^d$  when  $0 < s < d$ . From [18], we also have the following lemma that shows the same conclusion holds for the the case of  $s = d$  and  $s > d$ :

**Lemma 3.** *Let  $\mathcal{B}^d := \bar{B}(0, 1)$  denote the closed unit ball in  $\mathbb{R}^d$ . For the case of  $s = d$ , we have that*

$$\lim_{N \rightarrow \infty} \frac{\varepsilon_{s,d}(\mathbb{S}^d, \hat{\mathbf{W}}_N)}{N^2 \log N} = \frac{\mathcal{H}_d(\mathcal{B}^d)}{\mathcal{H}_d(\mathbb{S}^d)} = \frac{1}{d} \frac{\Gamma(\frac{d+1}{2})}{\sqrt{\pi} \Gamma(\frac{d}{2})} \quad (32)$$

and any sequence  $(\hat{\mathbf{W}}_N^*) \subset \mathbb{S}^d$  of minimal  $s$ -energy configurations satisfies Eq. 31.

The lemma above shows that the same conclusion holds for  $s = d$ . For the case of  $s > d$ , the theoretical analysis is more involved, but the conclusion that the neuron configuration with minimal hyperspherical energy asymptotically corresponds to the uniform distribution on  $\mathbb{S}^d$  still holds. Note that, we usually will not have the case of  $s > d$  in our applications.

## C PROOF OF THEOREM 2

We consider a set of  $n$   $d$ -dimensional neurons  $\mathbf{W} = \{\mathbf{w}_1, \dots, \mathbf{w}_n\} \in \mathbb{R}^{d \times n}$ . The hyperspherical energy of the original set of neurons can be written as:

$$E(\mathbf{w}_i|_{i=1}^n) = \sum_{i=1}^n \sum_{j=1, j \neq i}^n \left\| \frac{\mathbf{w}_i}{\|\mathbf{w}_i\|_2} - \frac{\mathbf{w}_j}{\|\mathbf{w}_j\|} \right\|^{-1} \quad (33)$$

which means that if the pairwise angle between any two neurons stays unchanged, then the hyperspherical energy will also stay unchanged. Now we consider the cosine value of the angle  $\theta_{(\mathbf{w}_i, \mathbf{w}_j)}$  between any two neuron  $\mathbf{w}_i$  and  $\mathbf{w}_j$ :

$$\cos(\theta_{(\mathbf{w}_i, \mathbf{w}_j)}) = \frac{\mathbf{w}_i^\top \mathbf{w}_j}{\|\mathbf{w}_i\| \cdot \|\mathbf{w}_j\|} = \frac{\sum_{k=1}^d w_{ik} \cdot w_{jk}}{\|\mathbf{w}_i\| \cdot \|\mathbf{w}_j\|} \quad (34)$$

where  $w_{ik}$  is the  $k$ -th element of the neuron  $\mathbf{w}_i$ . From the equation above, we can observe that permuting the order of the elements in the neurons together will not change the angle. For example, switching the  $i$ -th and  $j$ -th element in all the neurons will not change the hyperspherical energy. Assume that we randomly select  $p$  dimensions from the  $d$  dimensions and denote the set of  $p$  dimension as  $\mathbf{s} = \{s_1, \dots, s_p\} \in \mathbb{R}^p$ . Therefore we can construct a new set of neurons  $\tilde{\mathbf{W}} = \{\tilde{\mathbf{w}}_1, \dots, \tilde{\mathbf{w}}_n\}$  by permuting the  $p$  dimensions in  $\mathbf{s}$  to become the first  $p$  elements for all the neurons. Essentially, we use permutation to make  $\tilde{\mathbf{w}}_i = \mathbf{w}_{s_i}$  for  $i \in [1, p]$ . Therefore, we can have the following equation:

$$E(\mathbf{w}_i|_{i=1}^n) = E(\tilde{\mathbf{w}}_i|_{i=1}^n) \quad (35)$$

Then we consider an orthogonal matrix  $\mathbf{R}_p \in \mathbb{R}^{d \times d}$  that is used to transform the  $p$  dimension in the neurons. The equivalent orthogonal transformation for the  $d$ -dimensional neurons  $\tilde{\mathbf{W}}$  is

$$\tilde{\mathbf{R}} = \begin{bmatrix} \mathbf{R}_p & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{n-p} \end{bmatrix} = \begin{bmatrix} \mathbf{R}_p & 0 & \dots & 0 \\ 0 & 1 & \ddots & 0 \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 \end{bmatrix} \quad (36)$$

where  $\mathbf{I}_{n-p}$  is an identity matrix of size  $(n-p) \times (n-p)$ . It is easy to verify that  $\tilde{\mathbf{R}}$  is also an orthogonal matrix:  $\tilde{\mathbf{R}}^\top \tilde{\mathbf{R}} = \mathbf{I}_n$ . Then we permute the order of  $\tilde{\mathbf{W}}$  back to the original neuron set  $\mathbf{W}$  and obtain a new set of neurons  $\mathbf{W}^t = \{\mathbf{w}_1^t, \dots, \mathbf{w}_n^t\}$ .  $\mathbf{W}^t$  is in fact the result of directly performing orthogonal transformation to the  $p$  dimensions in  $\tilde{\mathbf{W}}$ . Because any order permutation of elements in neurons does not change the hyperspherical energy, we have the following equation

$$E(\mathbf{w}_i|_{i=1}^n) = E(\tilde{\mathbf{w}}_i|_{i=1}^n) = E(\tilde{\mathbf{R}}\tilde{\mathbf{w}}_i|_{i=1}^n) = E(\tilde{\mathbf{R}}\mathbf{w}_i^t|_{i=1}^n) \quad (37)$$

which concludes our proof.

## D EXPERIMENTAL SETTINGS

| Layer           | CNN-6 (CIFAR-100)           | CNN-9 (CIFAR-100)           | CNN-10 (ImageNet-2012)   |
|-----------------|-----------------------------|-----------------------------|--|
| Conv1.x         | $[3 \times 3, 64] \times 2$ | $[3 \times 3, 64] \times 3$ | $[7 \times 7, 64]$ , Stride 2<br>$3 \times 3$ , Max Pooling, Stride 2<br>$[3 \times 3, 64] \times 3$ |
| Pool1           | 2x2 Max Pooling, Stride 2   |                             |  |
| Conv2.x         | $[3 \times 3, 64] \times 2$ | $[3 \times 3, 64] \times 3$ | $[3 \times 3, 128] \times 3$   |
| Pool2           | 2x2 Max Pooling, Stride 2   |                             |  |
| Conv3.x         | $[3 \times 3, 64] \times 2$ | $[3 \times 3, 64] \times 3$ | $[3 \times 3, 256] \times 3$   |
| Pool3           | 2x2 Max Pooling, Stride 2   |                             |  |
| Fully Connected | 64                          | 64                          | 256  |

Table 10: Our plain CNN architectures with different convolutional layers. Conv1.x, Conv2.x and Conv3.x denote convolution units that may contain multiple convolution layers. E.g.,  $[3 \times 3, 64] \times 3$  denotes 3 cascaded convolution layers with 64 filters of size  $3 \times 3$ .

| Layer   | ResNet-20 (CIFAR-100)  | ResNet-32 (CIFAR-100)  |
|---------|--|--|
| Conv1.x | $[3 \times 3, 16] \times 1$<br>$\begin{bmatrix} 3 \times 3, 16 \\ 3 \times 3, 16 \end{bmatrix} \times 3$ | $[3 \times 3, 16] \times 1$<br>$\begin{bmatrix} 3 \times 3, 16 \\ 3 \times 3, 16 \end{bmatrix} \times 5$ |
| Conv2.x | $\begin{bmatrix} 3 \times 3, 32 \\ 3 \times 3, 32 \end{bmatrix} \times 3$                                | $\begin{bmatrix} 3 \times 3, 32 \\ 3 \times 3, 32 \end{bmatrix} \times 5$                                |
| Conv3.x | $\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 3$                                | $\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 5$                                |
|         | Average Pooling  |  |

Table 11: Our ResNet architectures with different convolutional layers. Conv0.x, Conv1.x, Conv2.x, Conv3.x and Conv4.x denote convolution units that may contain multiple convolutional layers, and residual units are shown in double-column brackets. Conv1.x, Conv2.x and Conv3.x usually operate on different size feature maps. These networks are essentially the same as [20], but some may have a different number of filters in each layer. The downsampling is performed by convolutions with a stride of 2. E.g.,  $[3 \times 3, 64] \times 4$  denotes 4 cascaded convolution layers with 64 filters of size  $3 \times 3$ , S2 denotes stride 2.

**Reported Results.** For all the experiments on MLPs and CNNs (except CNNs in the few-shot learning), we report testing error rates. For the few-shot learning experiment, we report testing accuracy. For all the experiments on both GCNs and PointNets, we report testing accuracy. All results are averaged over 10 runs of the model.

**Multilayer perceptron.** We conduct digit classification task on MNIST with a three-layer multilayer perceptrons following this repository<sup>1</sup>. The input dimension of each MNIST digit is  $28 \times 28$ , which is 784 dimensions after flattened. Our two hidden layers have 256 output dimensions, *i.e.*, 256 neurons. The output layer will output 10 logits for classification. We use a cross-entropy loss with softmax function. For the optimization, we use a momentum SGD with learning rate 0.01, momentum 0.9 and batch size 100. The training stops in 100 epochs.

**Convolutional neural networks.** The network architectures used in the paper are elaborated in Table 10 and Table 11. For CIFAR-100, we use 128 as the mini-batch size. We use momentum SGD with momentum 0.9 and the learning rate starts with 0.1, divided by 10 when the performance is saturated. For ImageNet-2012, we use batch size 128 and start with learning rate 0.1. The learning rate is divided by 10 when the performance is saturated, and the training is terminated at 700k iterations. For ResNet-20 and ResNet-32 on CIFAR-100, we use exactly the same architecture used on CIFAR-10 as [20]. The rotation matrix is initialized with random normal distribution (mean is 0 and variance is 1). Note that, for all the compared methods, we always use the best possible hyperparameters to make sure that the comparison is fair. The baseline has exactly the same architecture and training settings as the one that OPT uses. If not otherwise specified, standard  $\ell_2$  weight decay ( $5e-4$ ) is applied to all the neural network including baselines and the networks that use OPT training.

**Few-shot learning.** The network architecture (Table 12) we used for few-shot learning experiments is the same as that used in [10]. In our experiments, we show comparison of our OPT training with standard training on ‘baseline’ and ‘baseline++’ settings in [10]. In ‘baseline’ setting, a standard CNN model is pretrained on the whole meta-train dataset (standard non-MAML supervised training) and

<sup>1</sup>[https://github.com/hwalsuklee/tensorflow-mnist-MLP-batch\\_normalization-weight\\_initializers](https://github.com/hwalsuklee/tensorflow-mnist-MLP-batch_normalization-weight_initializers)

| Layer             | CNN-4                       |
|-------------------|-----------------------------|
| Conv1             | 3 × 3, 64                   |
| Pool1             | 2 × 2 Max Pooling, Stride 2 |
| Conv2             | 3 × 3, 64                   |
| Pool2             | 2 × 2 Max Pooling, Stride 2 |
| Conv3             | 3 × 3, 64                   |
| Pool3             | 2 × 2 Max Pooling, Stride 2 |
| Conv4             | 3 × 3, 64                   |
| Pool4             | 2 × 2 Max Pooling, Stride 2 |
| Linear Classifier | number of classes           |

Table 12: Architecture for few-shot learning. The number of classes is different for pretraining and finetuning.

| Layer   | Wide CNN-9 (CIFAR-100)                          | Wide ResNet-18 (ImageNet-2012)                        |
|---------|---|---|
| Conv0.x | N/A   | [7 × 7, 64], Stride 2<br>3 × 3, Max Pooling, Stride 2 |
| Conv1.x | [3 × 3, 64] × 3<br>2 × 2 Max Pooling, Stride 2  | 3 × 3, 64<br>3 × 3, 64 × 2                            |
| Conv2.x | [3 × 3, 128] × 3<br>2 × 2 Max Pooling, Stride 2 | 3 × 3, 128<br>3 × 3, 128 × 2                          |
| Conv3.x | [3 × 3, 256] × 3<br>2 × 2 Max Pooling, Stride 2 | 3 × 3, 256<br>3 × 3, 256 × 2                          |
| Conv4.x | N/A   | 3 × 3, 512<br>3 × 3, 512 × 2                          |
| Final   | 256-dim Fully Connected                         | Average Pooling                                       |

Table 13: Our wide CNN-9 and wide ResNet-18 architectures with different convolutional layers.

later only the classifier layer is finetuned on few-shot dataset. ‘baseline++’ differs from ‘baseline’ on the classifier: in ‘baseline’, each output dimension of the classifier is computed as the inner product between weight  $w$  and input  $x$ , i.e.  $w \cdot x$ ; while in ‘baseline++’ it becomes the scaled cosine distance  $c \frac{w \cdot x}{\|w\| \|x\|}$  where  $c$  is a positive scalar. Following [10], we set  $c = 2$ .

During pretraining, the model is trained for 200 epochs on the meta-train set of mini-ImageNet with an Adam optimizer (learning rate  $1e - 3$ , weight decay  $5e - 4$ ) and the classifier is discarded after pretraining. The model is later finetuned, with a new classifier, on the few-shot samples (5 way, support size 5) with a momentum SGD optimizer (learning rate  $1e - 2$ , momentum 0.9, dampening 0.9, weight decay  $1e - 3$ , batch size 4) for 100 epochs. We re-initialize the classifier for each few-shot sample.

**Graph neural networks.** We implement the OPT training for GCN on the official repositories<sup>2</sup>. The experimental settings also follow the official repository to ensure a fair comparison. For OPT (CP) method, we use the original hyperparameters and experimental setup except the added rotation matrix. For OPT (OGD) method, we use our own OGD optimizer in Tensorflow to train the rotation matrix in order to maintain orthogonality and use the original optimizer to train the other variables.

Training a GCN with OPT is not that straightforward. Specifically, the forward model of GCN is  $Z = \text{Softmax}(\hat{A} \cdot \text{ReLU}(\hat{A} \cdot X \cdot W_0) \cdot W_1)$  where  $\hat{A} = \tilde{D}^{\frac{1}{2}} \tilde{A} \tilde{D}^{\frac{1}{2}}$ . We note that  $A$  is the adjacency matrix of the graph,  $\tilde{A} = A + I$  ( $I$  is an identity matrix), and  $\tilde{D} = \sum_j \tilde{A}_{ij}$ .  $X \in \mathbb{R}^{n \times d}$  is the feature matrix of  $n$  nodes in the graph (feature dimension is  $d$ ).  $W_1$  is the weights of the classifiers.  $W_0$  is the weights of the classifiers.  $W_0$  is the weight matrix of size  $d \times h$  where  $h$  is the dimension of the hidden space. We treat each column vector of  $W_0$  as a neuron, so there are  $h$  neurons in total. Then we apply OPT to train these  $h$  neurons of dimension  $d$  in GCN. We conduct experiments on Cora and Pubmed datasets [56]. We aim to verify the effectiveness of OPT on GCN instead of achieving state-of-the-art performance on this task.

**Point cloud recognition.** To simplify the comparison and remove all the bells and whistles, we use a vanilla PointNet (without T-Net) as our backbone network. We apply OPT to train the MLPs in PointNet. We follow the same experimental settings as [51] and evaluate on the ModelNet-40 dataset [66]. We exactly follow the same setting in the original paper [51] and the official repositories<sup>3</sup>.

<sup>2</sup><https://github.com/tkipf/gcn><sup>3</sup><https://github.com/charlesq34/pointnet>

Specifically, we multiply the rotation matrix to the original fixed neurons in all the  $1 \times 1$  convolution layers and the fully connected layer except the final classifier. All the rotation matrix is initialized with random normal distribution. For PointNet experiments, we use point number 1024, batch size 32 and Adam optimizer started with learning rate 0.001, the learning rate will decay by 0.7 every 200k iterations, and the training is terminated at 250 epochs.

**Experimental settings for S-OPT.** For the experiment of S-OPT, the architecture of wide CNN-9 and wide ResNet-18 is given in Table 13. CNN-6 is the same as the one in Table 10. We use standard data augmentation for CIFAR-100, following [48]. For ImageNet-2012, we use the same data augmentation in [35, 48]. This data augmentation does not contain as many transformation as the one in [20], so the final performance may be worse than [20]. However, all the compared methods use the same data augmentation in our experiments, so the experiment is still a fair comparison. For CIFAR-100, we use  $N_{\text{out}} = 300$  and  $N_{\text{in}} = 750$ . For ImageNet, we use  $N_{\text{out}} = 700$  and  $N_{\text{in}} = 1000$ . We decrease the learning rate by a factor of 10 when the performance is saturated in the outer iteration.

## E LOSS LANDSCAPE VISUALIZATION (NORMAL DISTRIBUTION PERTURBATION)

### E.1 VISUALIZATION PROCEDURE

We generally follow the visualization procedure in [40]. However, since OPT has a different training process, we use a modified visualization method but still make it comparable to the baseline.

Specifically, if we want to plot the loss landscape of a neural neural with loss  $\mathcal{L}(\theta)$  where  $\theta$  is the learnable model parameters, we need to first choose pretrained model parameters  $\theta^*$  as a center point. Then we choose two random direction vectors  $\delta$  and  $\eta$ . The 2D plot  $f(\alpha, \beta)$  is defined as

$$f(\alpha, \beta) = \mathcal{L}(\theta^* + \alpha\delta + \beta\eta) \quad (38)$$

which can be used as a 2D surface visualization. Note that, after we randomly initialize the direction vectors  $\delta$  and  $\eta$  (with normal distribution), we need to perform the filter normalization [40]. Specifically, we normalize each filter in  $\delta$  and  $\eta$  to have the same norm as the corresponding filter in  $\theta^*$ . The loss landscape of our baseline is plotted using this visualization approach.

In contrast, the learnable parameters in OPT are no longer the weights of neurons. Instead, the learnable parameters are the orthogonal matrices. More precisely, the trainable matrices are used to perform orthogonalization in the neural networks (*i.e.*,  $\mathbf{P}$  in Fig. 2). We denote the combination of all the trainable matrices as  $\tilde{\mathbf{R}}$ , and the corresponding pretrained matrices as  $\tilde{\mathbf{R}}^*$ . Then the 2D visualization of OPT is

$$f(\alpha, \beta) = \mathcal{L}(\tilde{\mathbf{R}}^* + \alpha\gamma + \beta\kappa) \quad (39)$$

where  $\gamma$  and  $\kappa$  are two random direction vectors (which follow the normal distribution) to perturb  $\tilde{\mathbf{R}}^*$ . The visualization procedures of baseline and OPT are essentially the same except that the trainable variables are different. Therefore, their loss landscapes are comparable.

### E.2 EXPERIMENTAL DETAILS

In Fig. 4, we vary  $\alpha$  and  $\beta$  from  $-1.5$  to  $1.5$  for both baseline and OPT, and then plot the surface of 2D function  $f$ . We use the CNN-6 (as specified in Appendix D) on CIFAR-100. We use the same data augmentation as [45]. We train the network with SGD with momentum 0.9 and batch size 128. We start with learning rate 0.1, divide it by 10 at 30k, 50k and 64k iterations, and terminate training at 75k iterations. The training details basically follows [45]. We mostly use CP for OPT due to efficiency. Note that, the other orthogonalization methods in OPT yields similar loss landscapes in general. The pretrained model for standard training yields 37.59% testing error on CIFAR-100, while the pretrained model for OPT yields 33.53% error. This is also reported in Section 6.

### E.3 FULL VISUALIZATION RESULTS FOR THE MAIN PAPER




Figure 6: High-quality rendered loss landscapes of standard training and OPT.

Following the same experimental settings in Appendix E.2, we render the 3D loss landscapes with some color and lighting effects for Fig. 13. The visualization data is exactly the same as Fig. 4, and

we simply use ParaView to plot the figure. The rendered loss landscape better reflects that OPT yields a much more smooth loss geometry.

We also give the large and full version of Fig. 4(b) in the main paper) in the following figure. Fig. 7 is identical to Fig. 4(b) in the main paper except that Fig. 7 has larger size. From Fig. 7, we can better observe the dramatically different loss landscape between standard training and OPT. From the contour plots, we can better see that the red region of standard training is extremely flat and is highly non-convex around the edge. In comparison to standard training, OPT has very smooth and relatively convex contour shape.




Figure 7: Comparison of loss landscapes between standard training and OPT (full results of Fig. 4(b) in the main paper). Top row: loss landscape visualization with Cartesian coordinate system; Bottom row: loss contour visualization.

Then we visualize the testing error landscape. Additionally, we include a high-quality visualization using ParaView. We render the plot with lighting and color effects in order to better demonstrate the difference between OPT and standard training. The visualization results are given in Fig. 8. The comparison of testing error landscape shows that the parameter space of OPT is more robust than standard training, because the testing error of OPT is increased in a slower speed while the model is perturbed away from the pretrained parameters. In other words, the parameter space of OPT is more robust to the random perturbation than standard training. Combining the visualization of loss landscape, it is well justified that OPT can significantly alleviate training difficulty and improve generalization.




Figure 8: Comparison of testing error landscapes between standard training and OPT. Top row: high-quality rendered testing error landscape visualization with lighting effects; Bottom row: testing error landscape visualization with Cartesian coordinate system.



## E.4 LOSS LANDSCAPE VISUALIZATION OF DIFFERENT NEURAL NETWORK

In order to show that the loss landscape of OPT is quite general and consistent across different neural network architectures, we also visualize the loss landscape using a deep network network (CNN-9 as specified in Appendix D). The experiments are conducted on CIFAR-100. The results are given in Fig. 9. We can see that the loss landscape of OPT is much more smooth than standard training, similar to the observation for CNN-6. Therefore, the loss landscape difference between OPT and standard training is consistent across different network architectures, and OPT consistently shows better optimization landscape.




Figure 9: Comparison of loss landscapes between standard training and OPT on CIFAR-100 (CNN-9). Top row: loss landscape visualization with Cartesian coordinate system; Bottom row: loss contour visualization.

Then we show the landscape of testing error for CNN-9 on CIFAR-100. The results are given in Fig. 10. Similar to CNN-6, the testing error landscape of OPT is more smooth and convex than standard training. Moreover, OPT has a more flat local minima of testing error, while standard training has a sharp local minima. The testing error landscape in Fig. 10 generally follows the same pattern as the loss landscape in Fig. 9. The visualization further verifies the superiority of OPT is very consistent across different network architectures.

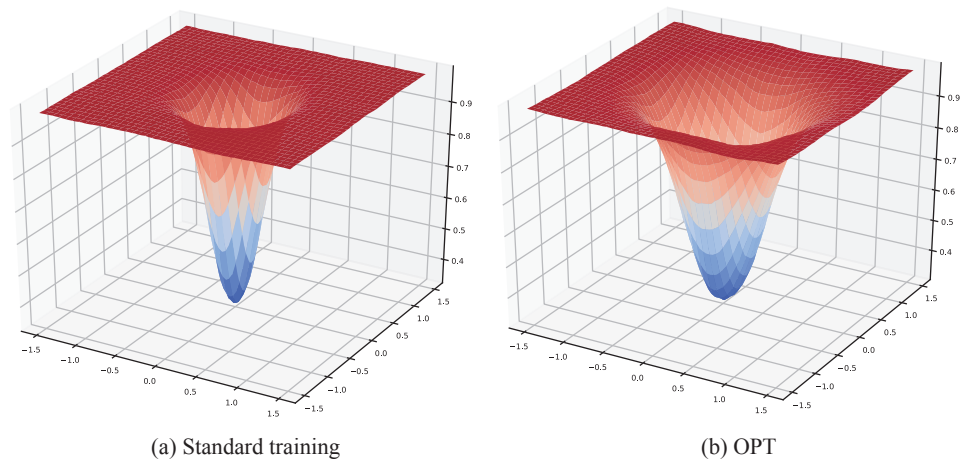


Figure 10: Comparison of testing error landscapes between standard training and OPT on CIFAR-100 (CNN-9).

E.5 LOSS LANDSCAPE VISUALIZATION OF DIFFERENT DATASET

Similar to Appendix F.5, we also visualize the loss landscape of OPT and standard training on a different dataset (CIFAR-10) with CNN-6. The loss landscape visualization is given in Fig. 11. Although the loss landscape on CIFAR-10 is quite different from the one on CIFAR-100, we can still observe that the loss landscape of OPT has a very flat local minima and the loss values are increasing smoothly and slowly. In contrast, the loss landscape of standard training has a sharp local minima and the loss values quickly increase to a large value. The red region of standard training will lead to very small gradient, potentially affecting the training. From the contour plots, the comparison apparently shows that the loss landscape of OPT is much more smooth than standard training.

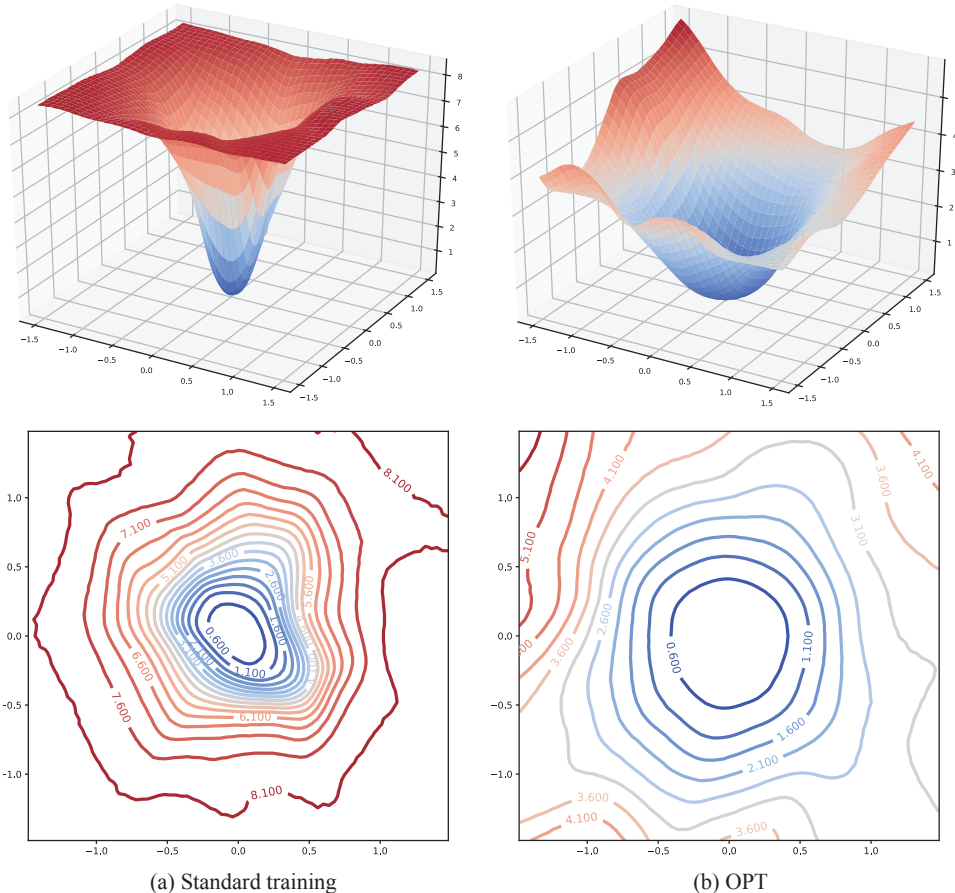


Figure 11: Comparison of loss landscapes between standard training and OPT on CIFAR-10 (CNN-6). Top row: loss landscape visualization with Cartesian coordinate system; Bottom row: loss contour visualization.

We also visualize the landscape of testing error in Fig. 12. The testing error landscape generally follows the pattern in the loss landscape. One can easily observe that the parameter space of standard training is very sensitive to perturbations. A small perturbation can make the model parameters completely fail (*i.e.*, the testing error dramatically increase). Differently, the parameter space of OPT is much more robust to perturbations. The model parameter can still work well with a small perturbation. Both Fig. 11 and Fig. 12 validate that the superiority of OPT is consistent across different training datasets.




Figure 12: Comparison of testing error landscapes between standard training and OPT on CIFAR-10 (CNN-6).

## F LOSS LANDSCAPE VISUALIZATION (UNIFORM DISTRIBUTION PERTURBATION)

### F.1 VISUALIZATION PROCEDURE

Different from Appendix E, we choose two random direction vectors  $\delta$  and  $\eta$  based on  $[0, 1]$  uniform distribution to further justify the effectiveness of OPT. The 2D plot  $f(\alpha, \beta)$  is defined as

$$f(\alpha, \beta) = \mathcal{L}(\theta^* + \alpha\delta + \beta\eta) \quad (40)$$

which can be used as a 2D surface visualization. Note that, after we randomly initialize the direction vectors  $\delta$  and  $\eta$  with  $[0, 1]$  normal distribution, we need to perform the filter normalization [40]. Specifically, we normalize each filter in  $\delta$  and  $\eta$  to have the same norm as the corresponding filter in  $\theta^*$ . The loss landscape of our baseline is plotted using this visualization approach.

In contrast, the learnable parameters in OPT are no longer the weights of neurons. Instead, the learnable parameters are the orthogonal matrices. More precisely, the trainable matrices are used to perform orthogonalization in the neural networks (*i.e.*,  $P$  in Fig. 2). We denote the combination of all the trainable matrices as  $\tilde{R}$ , and the corresponding pretrained matrices as  $\tilde{R}^*$ . Then the 2D visualization of OPT is

$$f(\alpha, \beta) = \mathcal{L}(\tilde{R}^* + \alpha\gamma + \beta\kappa) \quad (41)$$

where  $\gamma$  and  $\kappa$  are two random direction vectors (which follow the  $[0, 1]$  uniform distribution) to perturb  $\tilde{R}^*$ . The visualization procedures of baseline and OPT are essentially the same except that the trainable variables are different. Therefore, their loss landscapes are comparable.

### F.2 EXPERIMENTAL DETAILS

In Fig. 4, we vary  $\alpha$  and  $\beta$  from  $-1$  to  $1$  for both baseline and OPT, and then plot the surface of 2D function  $f$ . We use the CNN-6 (as specified in Appendix D) on CIFAR-100. We use the same data augmentation as [45]. We train the network with SGD with momentum 0.9 and batch size 128. We start with learning rate 0.1, divide it by 10 at 30k, 50k and 64k iterations, and terminate training at 75k iterations. The training details basically follows [45]. We mostly use CP for OPT due to efficiency. Note that, the other orthogonalization methods in OPT yields similar loss landscapes in general. The pretrained model for standard training yields 37.59% testing error on CIFAR-100, while the pretrained model for OPT yields 33.53% error. This is also reported in Section 6.

### F.3 FULL VISUALIZATION RESULTS

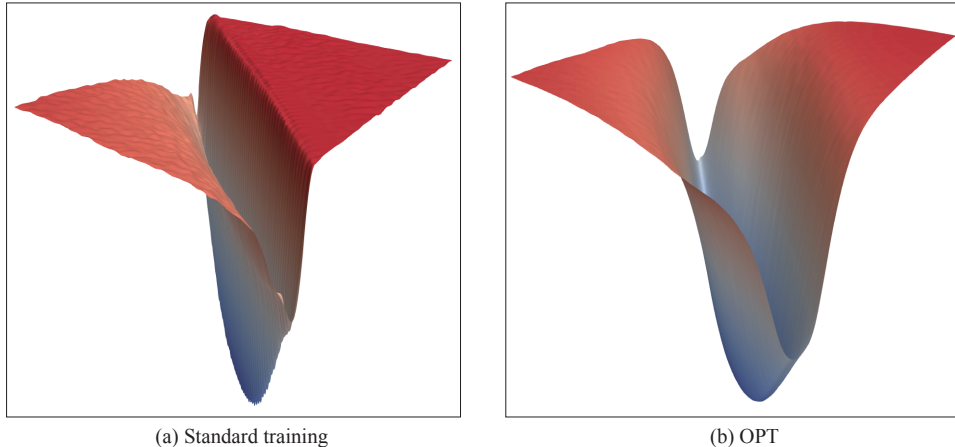


Figure 13: High-quality rendered loss landscapes of standard training and OPT.

Following the same experimental settings in Appendix F.2, we render the 3D loss landscapes with some color and lighting effects for Fig. 13. We first use ParaView to plot a high-quality loss landscape

comparison between standard training and OPT. As expected, the loss landscape of OPT is much more smooth than standard training. Note that, for the flat red region in standard training, we can still observe numerous small local minima, while the red region of OPT is very smooth. Fig. 13 better validates our analysis and discussion in Section 5.1, and also shows the superiority of OPT.

We provide the visualization results in the rest of the subsection. From Fig. 14, we can better observe the dramatically different loss landscape between standard training and OPT.




Figure 14: Comparison of loss landscapes between standard training and OPT (full results of Fig. 4(a) in the main paper). Top row: loss landscape visualization with Cartesian coordinate system; Bottom row: loss contour visualization.

To better understand the difference of the training dynamics between standard training and OPT, we also plot the testing error landscapes in Fig. 15 for both methods. The testing error is computed on the testing set of CIFAR-100 with the perturbed pretrained model ( $\alpha$  and  $\beta$  are the perturbation parameters). From the testing error landscape comparison in Fig. 15, we can see that once the baseline pretrained model is slightly perturbed, the testing error will immediately increase to 99.99% which is random selection-level testing error (because we have 100 balanced classes in total, randomly picking a class leads to 0.01% accuracy). In contrast, the testing error landscape of OPT is much more smooth. Even if we perturb the OPT pretrained model, we still end up with a reasonably low testing error, show that the parameter space of OPT is more smooth and continuous. All these evidences suggest that OPT is a better training framework for neural networks and can significantly alleviate the optimization difficulty. In this following subsections, we aim to show that the loss and testing error landscape difference between standard training and OPT is not a coincidence. We will show that the improvement of OPT on the loss and testing error landscape is both dataset-agnostic and architecture-agnostic.




Figure 15: Comparison of testing error landscapes between standard training and OPT. Top row: high-quality rendered testing error landscape visualization with lighting effects; Bottom row: testing error landscape visualization with Cartesian coordinate system.

F.4 LOSS LANDSCAPE VISUALIZATION OF DIFFERENT NEURAL NETWORK




Figure 16: Comparison of loss landscapes between standard training and OPT on CIFAR-100 (CNN-9). Top row: loss landscape visualization with Cartesian coordinate system; Bottom row: loss contour visualization.

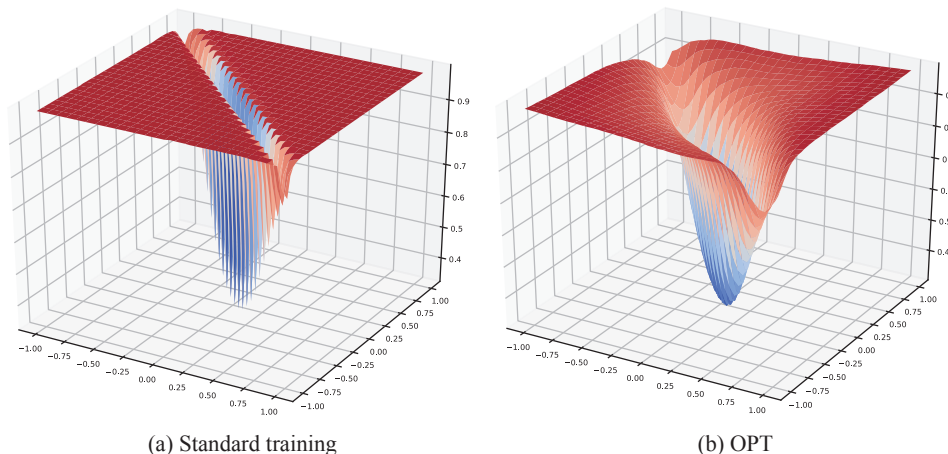


Figure 17: Comparison of testing error landscapes between standard training and OPT on CIFAR-100 (CNN-9).

To show that the difference of the loss landscape between standard training and OPT is consistent across different neural networks. We use a deeper CNN-9 (as specified in Appendix D) to visualize the loss landscape. The experimental settings generally follow Appendix F.2. We use CNN-9 as our backbone architecture and train it on CIFAR-100. The visualization of the loss landscapes is given in



Fig. 16. We observe that the loss landscapes of OPT with CNN-9 and CNN-6 are very similar. In general, the conclusion that OPT yields better loss landscape still holds for a deeper neural network, showing that the effectiveness of OPT is not architecture-dependent.

Moreover, we also visualize the landscape of testing error in Fig. 17. We can observe that the testing error landscapes are somewhat similar to the loss landscape in Fig. 16. The results further validate the superiority of OPT. We can observe in Fig. 17 that OPT has more smooth testing error landscape and can make the training parameter space of the neural network less sensitive to perturbations.

### F.5 LOSS LANDSCAPE VISUALIZATION ON DIFFERENT DATASET

We also perform the same loss and testing error landscape visualization on CIFAR-10. The training details basically follows Appendix F.2. For CIFAR-10, we use the same data augmentation as in Appendix D. The results are given in Fig. 18. From Fig. 18, we can observe even more dramatic difference of the loss landscape between standard training and OPT. In standard training, the loss landscape exhibits highly non-convex and non-smooth behavior. There are countless local minima in the loss landscape. Different from the results in Fig. 16, the loss landscape of standard training on CIFAR-10 has some huge local minima that are hard to escape from. In contrast, the loss landscape of OPT on CIFAR-10 does not show obvious and huge local minima and is far more convex and smooth than standard training. The contour maps show more significant difference between standard training and OPT. The contour map of OPT shows the shape of a single symmetric and convex valley, while the contour map of standard training presents the shape of multiple highly irregular valleys. The visualization further validate that the improvement of OPT on optimization landscape is very consistent across different training datasets.




Figure 18: Comparison of loss landscapes between standard training and OPT on CIFAR-10 (CNN-6). Top row: loss landscape visualization with Cartesian coordinate system; Bottom row: loss contour visualization.

Then we also visualize the testing error landscape of standard training and OPT in Fig. 19. As expected, the testing error landscape of standard training shows that changing the pretrained model parameters with a very small perturbation could lead to a dramatic increase in testing error. It indicates that the parameter space of standard training is very sensitive to even a small perturbation. In comparison, the testing error landscape of OPT shows similar shape to the training loss landscape which is the shape of a single regular, smooth and convex valley. We can conclude that OPT has huge advantages over standard training in terms of the optimization landscape. Although the conclusion

is drawn from a simple visualization method, it can still shed some light on why OPT yields better training dynamics and generalization ability.

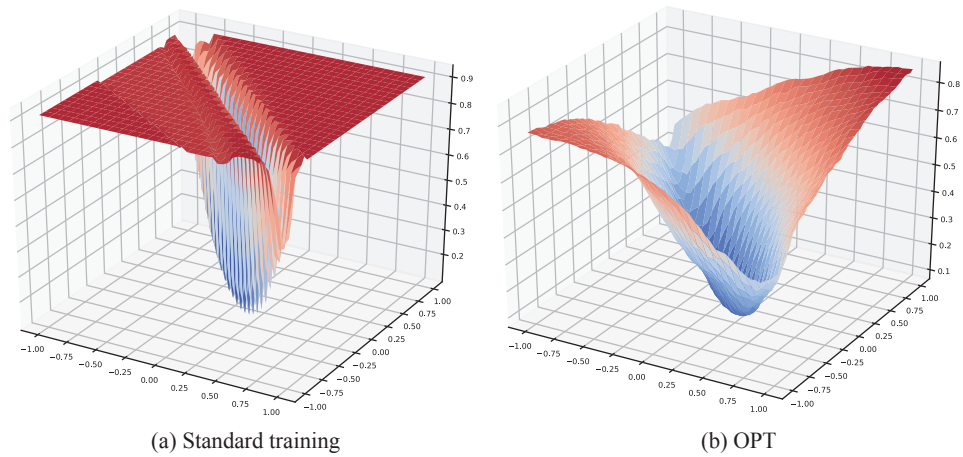


Figure 19: Comparison of testing error landscapes between standard training and OPT on CIFAR-10 (CNN-6).

## G THEORETICAL DISCUSSION ON OPTIMIZATION AND GENERALIZATION

The key problem we discuss in this section is why OPT may lead to easier optimization and better generalization. We have already shown that OPT can guarantee the minimum hyperspherical energy (MHE) in a probabilistic sense. Although empirical evidences [45] have shown significant and consistent performance gain by minimizing hyperspherical energy, why lower hyperspherical energy will lead to better generalization is still unclear. We argue that OPT leads to better generalization from two aspects: how OPT may affect the training and generalization, and why minimum hyperspherical energy serves as a good inductive bias. We note that rigorously proving that OPT generalizes better is out of the scope of this paper and remains our future work. The section serves as a very preliminary discussion for this topic, and hopefully the discussion can inspire more theoretical studies about OPT.

Our goal here is to leverage and apply existing theoretical results [1, 12, 31, 37, 59, 67] to explain the role that MHE plays rather than proving sharp and novel generalization bounds. We emphasize that our paper is *NOT* targeted as a theoretical one that proves novel generalization bounds.

We simply consider one-hidden-layer networks as the hypothesis class:

$$\mathcal{F} = \{f(x) = \sum_{j=1}^n v_j \sigma(\mathbf{w}_j^\top \mathbf{x}) : v_j \in \{\pm 1\}, \sum_{j=1}^n \|\mathbf{w}_j\| \leq C_w\} \quad (42)$$

where  $\sigma(\cdot) = \max(0, \cdot)$  is ReLU. Since the magnitude of  $v_j$  can be scaled into  $\mathbf{w}_j$ , we can restrict  $v_j$  to be  $\pm 1$ . Given a set of *i.i.d.* training sample  $\{\mathbf{x}_i, y_i\}_{i=1}^m$  where  $\mathbf{x} \in \mathbb{R}^d$  is drawn uniformly from the unit hypersphere, we minimize the least square loss  $\mathcal{L} = \frac{1}{2m} \sum_{i=1}^m (y_i - f(\mathbf{x}_i))^2$ . The gradient *w.r.t.*  $\mathbf{w}_i$  is

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}_j} = \frac{1}{m} \sum_{i=1}^m (f(\mathbf{x}_i) - y_i) v_j \sigma'(\mathbf{w}_j^\top \mathbf{x}_i) \mathbf{x}_i. \quad (43)$$

Let  $\mathbf{W} := \{\mathbf{w}_1^\top, \dots, \mathbf{w}_n^\top\}^\top$  be the column concatenation of neuron weights. We aim to identify the conditions under which there are no spurious local minima. We rewrite that

$$\frac{\partial \mathcal{L}}{\partial \mathbf{W}} = \left( \frac{\partial \mathcal{L}}{\partial \mathbf{w}_1}, \dots, \frac{\partial \mathcal{L}}{\partial \mathbf{w}_n} \right)^\top = \mathbf{D} \mathbf{r} \quad (44)$$

where  $\mathbf{r} \in \mathbb{R}^m$   $\mathbf{r}_i = \frac{1}{m} f(\mathbf{x}_i) - y_i$ ,  $\mathbf{D} \in \mathbb{R}^{n \times m}$ , and  $\mathbf{D}_{ij} = v_i \sigma'(\mathbf{w}_i^\top \mathbf{x}_j) \mathbf{x}_j$ . Therefore, we can obtain the following inequality:

$$\|\mathbf{r}\| \leq \frac{1}{s_m(\mathbf{D})} \left\| \frac{\partial \mathcal{L}}{\partial \mathbf{W}} \right\| \quad (45)$$

where  $\|\mathbf{r}\|$  is the training error and  $s_m(\mathbf{D})$  is the minimum singular value of  $\mathbf{D}$ . If we need the training error to be small, then we have to lower bound  $s_m(\mathbf{D})$  away from zero. Therefore, the essential problem now becomes the relationship between MHE and the lower bound of  $s_m(\mathbf{D})$ . We have the following result from [67]:

**Lemma 4.** *With probability larger than  $1 - m \exp(-m\gamma_m/8) - 2m^2 \exp(-4 \log^2 d) - \delta$ , we will have that  $s_m(\mathbf{D})^2 \geq \frac{1}{2} nm\gamma_m - cn\rho(\mathbf{W})$  where*

$$\begin{aligned} \rho(\mathbf{W}) \leq & \frac{\log d}{\sqrt{d}} \sqrt{2L_2(\mathbf{W})} m \left( \frac{4}{m} \log \frac{1}{\delta} \right)^{1/4} \\ & + \frac{2 \log d}{\sqrt{d}} m \sqrt{\frac{4}{3m} \log \frac{1}{\delta}} + \frac{\log d}{\sqrt{d}} mL_2(\mathbf{W}) + 2, \end{aligned} \quad (46)$$

and  $L_2(\mathbf{W}) = \frac{1}{n^2} \sum_{i,j=1}^n k(\mathbf{w}_i, \mathbf{w}_j)^2 - \mathbb{E}_{\mathbf{u}, \mathbf{v}}[k(\mathbf{u}, \mathbf{v})^2]$ . The kernel function  $k(\mathbf{u}, \mathbf{v})$  is  $\frac{1}{2} - \frac{1}{2\pi} \arccos\left(\frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\|\mathbf{u}\| \|\mathbf{v}\|}\right)$ .

Once MHE is achieved, the neurons will be uniformly distributed on the unit hypersphere. From Lemma 4, we can see that if the neurons are uniformly distributed on the unit hypersphere,  $L_2(\mathbf{W})$  will be very small and close to zero. Then  $\rho(\mathbf{W})$  will also be small, leading to large lower bound for

$s_m(\mathcal{D})$ . Therefore, MHE can result in small training error once the gradient norm  $\|\frac{\partial \mathcal{L}}{\partial \mathbf{W}}\|$  is small. The result implies no spurious local minima if we use OPT for training.

Furthermore, suppose that  $\|\frac{\partial \mathcal{L}}{\partial \mathbf{W}}\|^2 \leq \epsilon$ , [67] also proves a training error bound  $\tilde{\mathcal{O}}(\epsilon)$  and a generalization bound  $\tilde{\mathcal{O}}(\epsilon + \frac{1}{\sqrt{m}})$  based on the assumption that  $\mathbf{W}$  belongs to a specific set  $\mathcal{G}_{\mathbf{W}}$  (for the definition of  $\mathcal{G}_{\mathbf{W}}$ , please refer to [67]). Therefore, MHE is also connected to the training and generalization error. Note that, the analysis here is highly simplified and the purpose here is to give some justifications rather than rigorously proving any bound.

We further argue that MHE induced by OPT serves as an important inductive bias for neural networks. As the standard regularizer for neural networks, weight decay controls the norm of the neuron weights, regularizing essentially one dimension of the weight. In contrast, MHE completes an important missing piece by regularizing the remaining dimensions of the weight. MHE encourages minimum hyperspherical redundancy between neurons. In the linear classifier case, MHE impose a prior of maximal inter-class separability.

## H MORE DISCUSSIONS

**Semi-randomness.** OPT fixes the randomly initialized neuron weight vectors and simply learns layer-shared orthogonal matrices, so OPT naturally imposes strong randomness to the neurons. OPT well combines the good generalizability from randomness and the strong approximation power from neural networks. Such randomness suggests that the specific configuration of relative position among neurons does not matter that much, and the coordinate system is more crucial for generalization. [32, 53, 60] also show that randomness can be beneficial to generalization.

**Flexible training.** First, OPT can used in multi-task training [49] where each set of orthogonal matrices represent one task. OPT can learn different set of orthogonal matrices for different tasks with the neuron weights remain the same. Second, we can perform progressive training with OPT. For example, after learning a set of orthogonal matrices on a large coarse-grained dataset (*i.e.*, pretraining), we can multiple the orthogonal matrices back to the neuron weights and construct a new set of neuron weights. Then we can use the new neuron weights as a starting point and apply OPT to train on a small fine-grained dataset (*i.e.*, finetuning).

**Limitations and open problems** The limitations of OPT include more GPU memory consumption and heavy computation during training, more numerical issues when ensuring orthogonality and weak scalability for ultra wide neural networks. Therefore, there will be plenty of open problems in OPT, such as scalable and efficient training. Most significantly, OPT opens up a new possibility for studying theoretical generalization of deep networks. With the decomposition to hyperspherical energy and coordinate system, OPT provides a new perspective for future research.

## I ON PARAMETER-EFFICIENT OPT

### I.1 FORMULATION

Since OPT over-parameterizes the neurons, it will consume more GPU memory in training (note that, the number of parameters will not increase in testing). For a  $d$ -dimensional neuron, OPT will learn an orthogonal matrix of size  $d \times d$  that applies to the neuron. Therefore, we will need  $d^2$  extra parameters for one layer of neurons, making the training more expensive in terms of the GPU memory. Although the extra training overhead in OPT will not affect the inference speed of the trained neural networks, we still desire to achieve better parameter efficiency in OPT. To this end, we discuss some design possibilities for the *parameter-efficient OPT* (PE-OPT) in this section.

Original OPT over-parameterize a neuron  $\mathbf{v} \in \mathbb{R}^{n \times n}$  with  $\mathbf{R}\mathbf{v}$  where  $\mathbf{R}$  is a layer-shared orthogonal matrix of size  $d \times d$ . We aim to reduce the effective parameters of this  $d \times d$  orthogonal matrix. We incorporate a block-diagonal structure to the orthogonal matrix  $\mathbf{R}$ . Specifically, we formulate  $\mathbf{R}$  as  $\text{Diag}(\mathbf{R}^{(1)}, \mathbf{R}^{(2)}, \dots, \mathbf{R}^{(k)})$  where  $\mathbf{R}^{(i)}$  is an orthogonal matrix with size  $d_i \times d_i$  (it is easy to see that we need  $d = \sum_i d_i$ ). As an example, we only consider the case where all  $\mathbf{R}^{(i)}$  are of the same size (*i.e.*,  $d_1 = d_2 = \dots = d_k = \frac{d}{k}$ ). It is also obvious that as long as each block is an orthogonal matrix, then the overall matrix  $\mathbf{R}$  remains an orthogonal matrix.




Figure 20: Comparison between the block-shared matrix  $\mathbf{R}_s$  and the unconstrained block matrix  $\mathbf{R}_u$ .

First, we consider that all the block matrices on the diagonal of the orthogonal matrix  $\mathbf{R}$  are shared, meaning that  $\mathbf{R} = \text{Diag}(\mathbf{R}^{(1)}, \mathbf{R}^{(1)}, \dots, \mathbf{R}^{(1)})$  (*i.e.*,  $\mathbf{R}^{(1)} = \mathbf{R}^{(2)} = \dots = \mathbf{R}^{(k)}$ ). Therefore, we have a block-diagonal matrix  $\mathbf{R}_s$  with shared block  $\mathbf{R}^{(1)}$  as the final orthogonal matrix for the neuron  $\mathbf{v}$ :

$$\mathbf{R}_s = \begin{bmatrix} \mathbf{R}^{(1)} & 0 & \dots & 0 \\ 0 & \mathbf{R}^{(1)} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \mathbf{R}^{(1)} \end{bmatrix} \quad (47)$$

where  $\mathbf{R}^{(1)} \in \mathbb{R}^{\frac{d}{k} \times \frac{d}{k}}$ . The effective number of parameters for the orthogonal matrix  $\mathbf{R}_s$  immediately reduces to  $\frac{d^2}{k^2}$ . The left figure in Fig. 20 gives an intuitive illustration for the block-shared matrix  $\mathbf{R}_s$ . Therefore, PE-OPT only needs to learn  $\mathbf{R}^{(1)}$  in order to construct the orthogonal matrix of size  $d \times d$ .

Second, we consider that all the diagonal block matrices are independent, indicating that  $\mathbf{R} = \text{Diag}(\mathbf{R}^{(1)}, \mathbf{R}^{(2)}, \dots, \mathbf{R}^{(k)})$  where  $\mathbf{R}^{(i)}, \forall i$  are different orthogonal matrices in general. We term such matrix  $\mathbf{R}$  as unconstrained block matrix. Therefore, we have the unconstrained block diagonal matrix  $\mathbf{R}_u$  as

$$\mathbf{R}_u = \begin{bmatrix} \mathbf{R}^{(1)} & 0 & \dots & 0 \\ 0 & \mathbf{R}^{(2)} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \mathbf{R}^{(k)} \end{bmatrix} \quad (48)$$

where the orthogonal matrices  $\mathbf{R}^{(i)}, \forall i$  will be learned independently. The effective number of parameters for the orthogonal matrix  $\mathbf{R}_u$  is  $\frac{d^2}{k}$ , making it more flexible than the block-shared matrix  $\mathbf{R}_s$ .

Let’s consider a convolution neuron (*i.e.*, convolution filter)  $\mathbf{v} \in \mathbb{R}^{c_1 \times c_2 \times c_3}$  (*e.g.*, a typical convolution neuron is of size  $3 \times 3 \times 64$ ) as an example. The orthogonal matrix  $\mathbf{R}$  for the convolution neuron is of size  $(c_1 c_2 c_3) \times (c_1 c_2 c_3)$ . Typically, we will divide the neuron into  $k$  sub-neuron along the  $c_3$ -axis, each with size  $c_1 \times c_2 \times \frac{c_3}{k}$ . Then in order to learn a block-shared orthogonal matrix  $\mathbf{R}_s$ , we will essentially learn a shared orthogonal matrix of size  $(\frac{1}{k} c_1 c_2 c_3) \times (\frac{1}{k} c_1 c_2 c_3)$  that applies to each sub-neuron (there are  $k$  sub-neurons of size  $c_1 \times c_2 \times \frac{c_3}{k}$  in total). For the case of learning a unconstrained block-diagonal orthogonal matrix  $\mathbf{R}_u$ , we simply learn different orthogonal matrices for different sub-neurons.

## I.2 EXPERIMENTS AND RESULTS

We conduct the image recognition experiments on CIFAR-100 with CNN-6 described in Table 10. The setting is exactly the same as Section 6.2. For the convolution filter, we use the size of  $3 \times 3 \times 64$ , *i.e.*,  $c_1 = 3, c_2 = 3, c_3 = 64$ . The results are given in Table 14 and Table 15. “# Parameters” in both tables denote the number of effective parameters for the orthogonal matrix  $\mathbf{R}$  in a single layer. The baseline with fixed neurons is only to train the final classifiers with the randomly initialized neuron weights staying fixed. It means that this baseline basically removes the learnable orthogonal matrices but still fixes the neuron weights, so it only achieves 73.81% testing error. As expected, as the number of effective parameters goes down, the performance of PE-OPT generally decreases. One can also observe that using separate orthogonal matrices generally yields better performance than shared orthogonal matrices.  $k = 2$  and  $k = 4$  seems to be a reasonable trade-off between better accuracy and less parameters.

When  $k$  becomes larger (*i.e.*, the number of parameters become less) in the case of block-shared orthogonal matrices, we find that PE-OPT (LS) performs the best among all the variants. When  $k$  becomes larger (*i.e.*, the number of parameters become less) in the case of unconstrained block orthogonal matrices, we can see that both PE-OPT (GS) and PE-OPT (LS) performs better than the other variants.

| Method   | # Parameters | PE-OPT (CP) | PE-OPT (GS)  | PE-OPT (HR) | PE-OPT (LS)  | PE-OPT (OGD) |
|--|--------------|-------------|--------------|-------------|--------------|--------------|
| $c_3/k = 64$ ( $k = 1$ ) ( <i>i.e.</i> , Original OPT) | 331.7K       | 33.53       | <b>33.02</b> | 35.67       | 34.48        | 33.33        |
| $c_3/k = 32$ ( $k = 2$ )                               | 82.9K        | 34.93       | <b>34.39</b> | 35.83       | 34.50        | 35.06        |
| $c_3/k = 16$ ( $k = 4$ )                               | 20.7K        | 39.40       | 39.13        | 39.67       | <b>37.58</b> | 39.80        |
| $c_3/k = 8$ ( $k = 8$ )                                | 5.2K         | 47.77       | 46.65        | 46.69       | <b>45.62</b> | 47.43        |
| $c_3/k = 4$ ( $k = 16$ )                               | 1.3K         | 56.65       | 55.91        | 55.69       | <b>54.75</b> | 57.15        |
| $c_3/k = 2$ ( $k = 32$ )                               | 0.3K         | 63.46       | 62.65        | 62.38       | <b>61.60</b> | 62.46        |
| $c_3/k = 1$ ( $k = 64$ )                               | 0.1K         | 67.36       | 67.11        | 67.05       | <b>66.61</b> | 67.23        |
| Baseline   | -            | -           | -            | 37.59       | -            | -            |
| Baseline with fixed random neurons                     | -            | -           | -            | 73.81       | -            | -            |

Table 14: Testing error (%) on CIFAR-100 with different settings of PE-OPT (with block-shared orthogonal matrix  $\mathbf{R}_s$ ).

| Method   | # Parameters | PE-OPT (CP) | PE-OPT (GS)  | PE-OPT (HR) | PE-OPT (LS)  | PE-OPT (OGD) |
|--|--------------|-------------|--------------|-------------|--------------|--------------|
| $c_3/k = 64$ ( $k = 1$ ) ( <i>i.e.</i> , Original OPT) | 331.7K       | 33.53       | <b>33.02</b> | 35.67       | 34.48        | 33.33        |
| $c_3/k = 32$ ( $k = 2$ )                               | 165.9K       | 33.54       | <b>33.15</b> | 35.65       | 34.09        | 34.27        |
| $c_3/k = 16$ ( $k = 4$ )                               | 82.9K        | 34.77       | <b>34.50</b> | 35.71       | 34.96        | 35.97        |
| $c_3/k = 8$ ( $k = 8$ )                                | 41.5K        | 37.25       | 36.43        | 36.40       | <b>36.17</b> | 39.75        |
| $c_3/k = 4$ ( $k = 16$ )                               | 20.7K        | 40.74       | <b>39.89</b> | 39.98       | 39.93        | 43.43        |
| $c_3/k = 2$ ( $k = 32$ )                               | 10.4K        | 45.36       | 44.77        | 44.83       | <b>44.61</b> | 48.98        |
| $c_3/k = 1$ ( $k = 64$ )                               | 5.2K         | 50.94       | <b>49.16</b> | 49.57       | 49.23        | 54.93        |
| Baseline   | -            | -           | -            | 37.59       | -            | -            |
| Baseline with fixed random neurons                     | -            | -           | -            | 73.81       | -            | -            |

Table 15: Testing error (%) on CIFAR-100 with different settings of PE-OPT (with unconstrained block orthogonal matrix  $\mathbf{R}_u$ ).

## J ON GENERALIZING OPT: OVER-PARAMETERIZED TRAINING WITH CONSTRAINT

OPT opens many new possibilities in training neural networks. We consider a simple generalization to OPT in this section to showcase the great potential of OPT. Instead of constraining the over-parameterization matrix  $\mathbf{R} \in \mathbb{R}^{d \times d}$  in Eq. 1 to be orthogonal, we can use any meaningful structural constraints for this matrix, and even regularize it in a task-driven way. Furthermore, instead of a linear over-parameterization (*i.e.*, multiplying a matrix  $\mathbf{R}$ ) to the neuron, we can also consider nonlinear mapping. We come up with the following straightforward generalization to OPT (the settings and notations exactly follow Eq. 1):

$$\begin{aligned}
 \text{Standard:} \quad & \min_{\mathbf{v}_i, u_i, \forall i} \sum_{j=1}^m \mathcal{L}(y, \sum_{i=1}^n u_i \mathbf{v}_i^\top \mathbf{x}_j) \\
 \text{Original OPT:} \quad & \min_{\mathbf{R}, u_i, \forall i} \sum_{j=1}^m \mathcal{L}(y, \sum_{i=1}^n u_i (\mathbf{R} \mathbf{v}_i)^\top \mathbf{x}_j) \\
 & \text{s.t. } \mathbf{R}^\top \mathbf{R} = \mathbf{R} \mathbf{R}^\top = \mathbf{I} \\
 \text{Generalized OPT:} \quad & \min_{\mathbf{R}, u_i, \forall i} \sum_{j=1}^m \mathcal{L}(y, \sum_{i=1}^n u_i (\mathcal{T}(\mathbf{v}_i))^\top \mathbf{x}_j) \\
 & \text{s.t. Some constraints on } \mathcal{T}(\cdot)
 \end{aligned} \tag{49}$$

where  $\mathcal{T}(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d$  denotes some transformation (including both linear and nonlinear). Notice that the generalized OPT (G-OPT) no longer requires orthogonality. Such formulation of G-OPT can immediately inspire a number of instances. We will discuss some obvious ones here.

If we consider  $\mathcal{T}(\cdot)$  to be a linear mapping, we may constrain  $\mathbf{R}$  to be symmetric positive definite other than orthogonal. A simple way to achieve that is to use Cholesky factorization  $\mathbf{L} \mathbf{L}^\top$  where  $\mathbf{L}$  is a lower triangular matrix to parameterize the matrix  $\mathbf{R}$ . Essentially, we learn a lower triangular matrix  $\mathbf{L}$  and use  $\mathbf{L} \mathbf{L}^\top$  to replace  $\mathbf{R}$  in OPT. The positive definiteness provides the transformation  $\mathbf{R}$  with some geometric constraint. Specifically, a positive definite  $\mathbf{R}$  only transforms the neuron weight  $\mathbf{v}$  to the direction that has the angle less than  $\frac{\pi}{2}$  to  $\mathbf{v}$ , because  $\mathbf{v}^\top \mathbf{R} \mathbf{v} > 0$ . Moreover, we can also require the transformation to have structural constraints on  $\mathbf{R}$ . For example,  $\mathbf{R}$  can be upper (lower) triangular, banded, symmetric, skew-symmetric, upper (lower) Hessenberg, etc.

We can also consider  $\mathcal{T}(\cdot)$  to be a nonlinear mapping. A obvious example is to use a neural network (*e.g.*, MLP, CNN) as  $\mathcal{T}(\cdot)$ . Then the nonlinear G-OPT will share some similarities with HyperNetworks [17] and Network-in-Network [43]. If we further consider  $\mathcal{T}(\cdot)$  to be dependent on the input, then the nonlinear G-OPT will have close connections to dynamic neural networks [29, 46].

To summarize, OPT provides a novel and effective framework to train neural networks and may inspire many different threads of future research.



## K HYPERSPHERICAL ENERGY TRAINING DYNAMICS OF INDIVIDUAL LAYERS

We also plot the hyperspherical energy ( $\mathbf{E}(\hat{\mathbf{v}}_i|_{i=1}^n) = \sum_{i=1}^n \sum_{j=1, j \neq i}^n \|\hat{\mathbf{v}}_i - \hat{\mathbf{v}}_j\|^{-1}$  in which  $\hat{\mathbf{v}}_i = \frac{\mathbf{v}_i}{\|\mathbf{v}_i\|}$  is the  $i$ -th neuron weight projected onto the unit hypersphere.) in every layer of CNN-6 during training to show how these hyperspherical energies are being minimized. From Fig. 21, we can observe that OPT can always maintain the minimum hyperspherical energy during the entire training process, while the MHE regularization cannot. Moreover, the hyperspherical energy of the baseline will also decrease as the training proceeds, but it is still much higher than the OPT training.

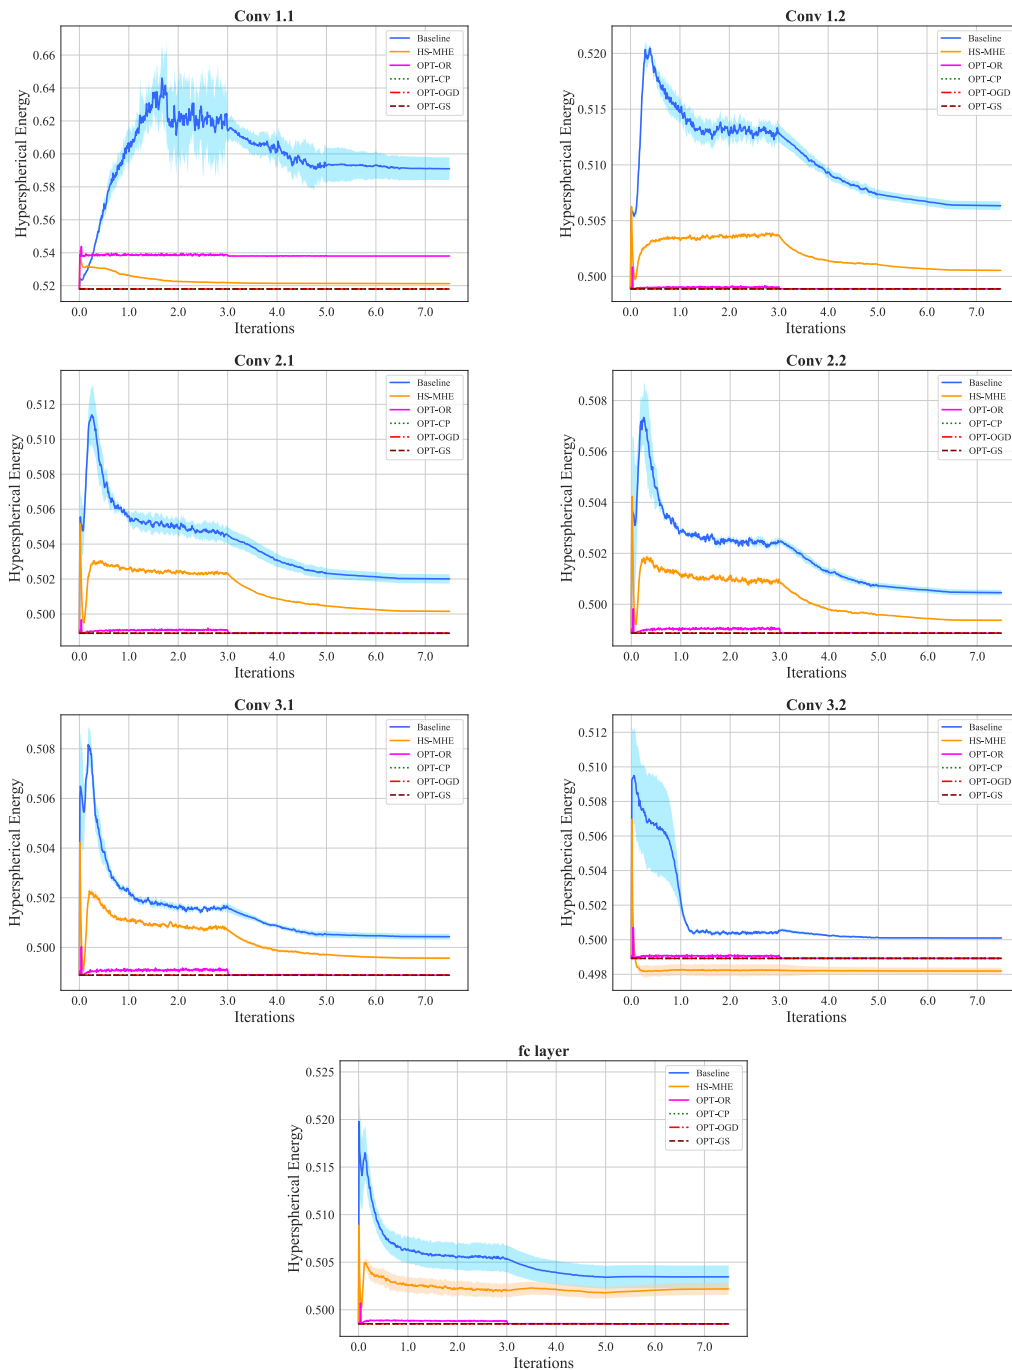


Figure 21: Training dynamics of hyperspherical energy in each layer of CNN-6. We average results with 10 runs.

## L GEOMETRIC PROPERTIES OF RANDOMLY INITIALIZED NEURONS

There are many interesting geometric properties [5, 6] of random points distributed independently and uniformly on the unit hypersphere. We summarize a few of them that make randomly initialized neurons distinct from any deterministic neuron configuration. Note that, there exist many deterministic neuron configurations that can also achieve very low hyperspherical energy, and this section aims to describe a few unique geometric properties of randomly initialized neurons.

There are two widely used geometric properties corresponding to a neuron configuration (*i.e.*, a set of neurons)  $\hat{\mathbf{W}}_N = \{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_N \in \mathbb{S}^d\}$ . In the main paper, we define neurons on  $\mathbb{S}^{d-1}$ , but without loss of generality we define neurons on  $\mathbb{S}^d$  here for convenience. The first one is the *covering radius*:

$$\alpha(\hat{\mathbf{W}}_N) := \alpha(\hat{\mathbf{W}}_N; \mathbb{S}^{d-1}) := \max_{\mathbf{u} \in \mathbb{S}^d} \min_{1 \leq i \leq N} \arccos(\mathbf{u}, \hat{\mathbf{w}}_i) \quad (50)$$

which is the biggest geodesic distance from a neuron in  $\mathbb{S}^d$  to the nearest point in  $\hat{\mathbf{W}}_N$ . The second one is the *separation distance*:

$$\psi(\hat{\mathbf{W}}_N) := \min_{1 \leq i, j, \leq N, i \neq j} \arccos(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j) \quad (51)$$

which gives the least geodesic distance between arbitrary two points in  $\hat{\mathbf{W}}_N$ . Random points (*i.e.*, randomly initialized neurons) typically have poor separation properties, since the separation is sensitive to the specific placement of points. [5] shows an example on  $\mathbb{S}^1$  to illustrate this observation.

[5] considers a different but related quantity, *i.e.*, the sum of powers of the ‘‘hole radii’’. A set of neurons  $\hat{\mathbf{W}}_N$  on  $\mathbb{S}^d$  uniquely defines a convex polytope, which can be viewed as the convex hull of the neuron configuration. Each facet of the polytope defines a ‘‘hole’’. Such a hole denotes the maximal spherical cap for a facet that contains neurons of  $\hat{\mathbf{W}}_N$  only on the boundary. It is easy to see that the geodesic radius of the largest hole is the covering radius  $\alpha(\hat{\mathbf{W}}_N)$ . We assume that for the set of neurons  $\hat{\mathbf{W}}_N$ , there are  $f_d$  holes (*i.e.*, facets) in total. Therefore, the  $i$ -th hole radius is defined as  $\rho_i = \rho_i(\hat{\mathbf{W}}_N)$  which is the Euclidean distance in  $\mathbb{R}^{d+1}$  between the center of the  $i$ -th spherical cap and the boundary. The  $i$ -th spherical cap is located on the sphere corresponding to the  $i$ -th facet. We have that  $\rho_i = 2 \sin(\frac{\alpha_i}{2})$  where  $\alpha_i$  is the geodesic radius of the  $i$ -th spherical cap. We are interested in the sums of the  $p$ -th powers of the hole radii, *i.e.*,

$$\mathcal{P} = \sum_{i=1}^{f_d} (\rho_i)^p \quad (52)$$

where  $p$  is larger than zero. For large  $p$ , the largest hole dominates:

$$\lim_{p \rightarrow \infty} (\mathcal{P})^{\frac{1}{p}} = \lim_{p \rightarrow \infty} \left( \sum_{i=1}^{f_d} (\rho_i)^p \right)^{\frac{1}{p}} = \max_{1 \leq i \leq f_d} \rho_i = 2 \sin\left(\frac{\alpha(\hat{\mathbf{W}}_N)}{2}\right) \quad (53)$$

where  $\rho(\hat{\mathbf{W}}_N) := \max_{1 \leq i \leq f_d} \rho_i$ . Then we introduce some useful notations to state the geometric properties. Let  $\psi_d$  be the surface area of  $\mathbb{S}^d$ , and we have that

$$\psi_d = \frac{2\pi^{\frac{d+1}{2}}}{\Gamma(\frac{d+1}{2})}, \quad (54)$$

and we also define the following quantities (with  $\psi_0 = 2$ ):

$$\begin{aligned} \kappa_d &:= \frac{1}{d} \frac{\psi_{d-1}}{\psi_d} = \frac{1}{d} \frac{\Gamma(\frac{d+1}{2})}{\sqrt{\pi} \Gamma(\frac{d}{2})} \\ B_d &:= \frac{2}{d+1} \frac{\kappa_d^2}{(\kappa_d)^d} \end{aligned} \quad (55)$$

where  $\kappa_d$  can be alternatively defined with the recursion:  $\kappa_1 = \frac{1}{\pi}$  and  $\kappa_d = \frac{1}{2\pi d \kappa_{d-1}}$ . [55] gives the expected number of facets constructed from  $N$  random neurons that are independently and uniformly distributed on the unit hypersphere  $\mathbb{S}^d$ :

$$\mathbb{E}[f_d] = B_d N (1 + o(1)) \quad (56)$$

where  $N \rightarrow \infty$ . Then we introduce the main results of [5] (asymptotics for the expected moments of the hole radii) in the following lemma:

**Lemma 5.** *If  $p \geq 0$  and  $\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_N$  are  $N$  neurons on  $\mathbb{S}^d$  that are independently and randomly distributed with respect to the normalized surface area measure  $\sigma_d$  on  $\mathbb{S}^d$ , then we have that*

$$\begin{aligned} \mathbb{E}[\mathcal{P}] &= B_d(\kappa_d)^{-\frac{p}{d}} \frac{\Gamma(d + \frac{p}{d})\Gamma(N + 1)}{\Gamma(d)\Gamma(N + \frac{p}{d})} (1 + \mathcal{O}(N^{-\frac{2}{d}})) \\ &= c_{d,p} N^{1-\frac{p}{d}} (1 + \mathcal{O}(N^{-\frac{2}{d}})) \end{aligned} \tag{57}$$

as  $N \rightarrow \infty$ , where  $\rho_i = \rho_{i,N}$  is the Euclidean hole radius associated with the  $i$ -th facet of the convex hull of  $\hat{\mathbf{W}}_N$ ,  $c_{d,p} := B_d B_{d,p}$ , and  $B_{d,p} := \frac{\Gamma(d+\frac{p}{d})}{\Gamma(d)} (\kappa_d)^{-\frac{p}{d}}$ . The  $\mathcal{O}$ -terms above depend on  $d$  and  $p$ .

As we mentioned, there are many deterministic point (*i.e.*, neuron) configurations such as minimizing hyperspherical energy (*i.e.*, Riesz  $s$ -energy) [45] (as  $s \rightarrow \infty$ , the minimal  $s$ -energy points approach the best separation), maximizing the determinant for polynomial interpolation [57], Fibonacci points, spherical  $t$ -designs, minimizing covering radius (*i.e.*, best covering problem), maximizing the separation (*i.e.*, best packing problem) and maximizing the  $s$ -polarization, etc. We note that randomly initialized neurons are quite different from these deterministic neuron configurations and have unique geometric properties.

## M ADDITIONAL EXPERIMENTAL RESULTS

### M.1 HIGH VS. LOW HYPERSPHERICAL ENERGY

We empirically verify that high hyperspherical energy corresponds to inferior generalization performance. To initialize neurons with high hyperspherical energy, we use a random initialization with mean equal to  $1e-3$ ,  $1e-2$ ,  $2e-2$ ,  $3e-2$  and  $5e-2$ . We use CNN-6 to conduct experiments on CIFAR-100. The results in Table 16 (“N/C” denotes not converged) show that the network with higher hyperspherical energy is more difficult to converge. Moreover, we find that if the hyperspherical energy is larger than a certain value, then the network cannot converge at all. Note that, when the hyperspherical energy is small (near the minima), a little change in hyperspherical energy (*e.g.*, from 3.5109 to 3.5160) can lead to dramatic generalization gap (*e.g.*, from 32.49% error rate to 39.51%). As expected, one can also observe that higher hyperspherical energy leads to worse generalization.

| Mean | Energy        | Error (%)    |
|------|---------------|--------------|
| 0    | <b>3.5109</b> | <b>32.49</b> |
| 1e-3 | 3.5117        | 33.11        |
| 1e-2 | 3.5160        | 39.51        |
| 2e-2 | 3.5531        | 53.89        |
| 3e-2 | 3.6761        | N/C          |
| 5e-2 | 4.2776        | N/C          |

Table 16: Testing error for different initial energy.

### M.2 TRAINING WITHOUT BATCH NORMALIZATION

We further evaluate how OPT performs without batch normalization (BN). In specific, we use CNN-6 as our backbone network and test on CIFAR-100. From Table 17, one can observe that all OPT variants again outperform both the baseline and HS-MHE [45], validating that OPT is robust and can work reasonably well even without batch normalization. Among all the OPT variants, Cayley parameterization achieves the best testing error with approximately 4% lower than the standard training.

| Method    | Error (%)    |
|-----------|--------------|
| Baseline  | 38.95        |
| HS-MHE    | 36.90        |
| OPT (GS)  | 35.61        |
| OPT (HR)  | 37.51        |
| OPT (LS)  | 35.83        |
| OPT (CP)  | <b>34.88</b> |
| OPT (OGD) | 35.38        |

Table 17: Training without BN on CIFAR-100.

### M.3 STANDARD DEVIATION

Due to the space limit, we have not provided the standard deviation for the results of OPT. In general, OPT performs quite stably and consistently outperforms all the compared methods. We show the standard deviation for the results in Table 4. The results are given in Table 18. We can observe that the standard deviations are quite small for OPT, showing that OPT’s performance gain is very consistent and significant.

| Method    | CNN-6             | CNN-9             |
|-----------|-------------------|-------------------|
| Baseline  | 37.59±0.24        | 33.55±0.26        |
| HS-MHE    | 34.97±0.15        | 32.87±0.17        |
| OPT (GS)  | <b>33.02±0.11</b> | <b>31.03±0.14</b> |
| OPT (HR)  | 35.67±0.14        | 32.75±0.15        |
| OPT (LS)  | 34.48±0.09        | 31.22±0.11        |
| OPT (CP)  | 33.53±0.13        | 31.28±0.12        |
| OPT (OGD) | 33.33±0.21        | 31.47±0.23        |
| OPT (OR)  | 34.70±0.18        | 32.63±0.17        |

Table 18: Error (%) with standard deviation on CIFAR-100.

### M.4 TRAINING RUNTIME

One of the weaknesses of OPT is its training runtime. As expected, OPT is more costly than standard training, but fortunately, OPT does not affect the inference speed of a trained model. We use CNN-6

and perform image classification on CIFAR-100. The batch size is 128. We use a single RTX 6000 for the experiment. The results are given in Table 19. We can see that the slowest OPT is around 35X slower than the baseline. However, OPT with OGD, OR and CP performs just slightly worse than the baseline. In general, the training runtime does not matter that much in practice, and we usually care more about the inference speed during testing. We emphasize again that OPT does not affect the inference speed of neural networks.

| Method    | Runtime |
|-----------|---------|
| Baseline  | 3.28    |
| OPT (GS)  | 14.01   |
| OPT (HR)  | 71.32   |
| OPT (LS)  | 107.47  |
| OPT (CP)  | 7.06    |
| OPT (OGD) | 4.30    |
| OPT (OR)  | 3.43    |

Table 19: Training runtime (s) per 100 iterations.