

# Group Orthogonalization Regularization for Vision Models Adaptation and Robustness

Yoav Kurtz  
yoavkurtz@mail.tau.ac.il

Noga Bar  
nogabar@mail.tau.ac.il

Raja Giryes  
raja@tauex.tau.ac.il

School of Electrical Engineering  
The Iby and Aladar Fleischman Faculty  
of Engineering  
Tel Aviv University  
Tel Aviv, Israel

## Abstract

As neural networks become deeper, the redundancy within their parameters increases. This phenomenon has led to several methods that attempt to reduce the correlation between convolutional filters. We propose a computationally efficient regularization technique that encourages orthonormality between groups of filters within the same layer. Our experiments show that when incorporated into recent adaptation methods for diffusion models and vision transformers (ViTs), this regularization improves performance on downstream tasks. We further show improved robustness when group orthogonality is enforced during adversarial training. Our code is available at <https://github.com/YoavKurtz/GOR>.




Figure 1: Qualitative comparisons on Pokemon-BLIP between a baseline fine-tuned model using LoRA (top) and a model fine-tuned along with GOR (bottom) using the same seed. The green rectangle is zoomed in by a factor of 1.5. Note the improved quality of GOR.

# 1 Introduction

Deep neural networks have achieved remarkable success in various tasks, such as image classification and generation, object detection, and segmentation. However, as networks become deeper and wider, the redundancy within their parameters increases. While increasing the DNN size can lead to improved expressivity and performance, it was shown that there is a high correlation between parameters [45]. This phenomenon has led to several methods that attempt to reduce the correlation between convolutional filters [9, 45].

Additionally, various methods are suggested to regularize the space of filters and features. One of the dominant examples is batch normalization [19], which normalizes the activations of the previous layer to have zero mean and unit variance. For reducing the dependency of batch normalization on the batch size, it was suggested to use group normalization [46], which normalizes for each example the activations of a subset of channels in a given layer to have zero mean and unit variance.

In addition to feature normalization, some methods have been proposed for normalizing the network weights. One example is weight standardization [64], which normalizes the weight vectors to have unit  $\ell_2$  norm. Despite their effectiveness, these methods either normalize each weight independently [52, 59] or all together [9, 21, 52, 47], which may be computationally demanding. To address this limitation, in our work, we address the correlation between filters in the same layer with reduced computational overhead. We propose an efficient regularization method that encourages orthonormality between filter groups within the same layer. Our approach, called Group Orthogonalization Regularization (GOR), is motivated by the observation that orthonormal filters are more diverse, expressive, and less redundant than correlated filters. Yet, instead of enforcing orthogonality between all filters, inspired by group normalization, we enforce orthogonality only between groups of filters.

To assess our proposed method, we first checked it on the well-known classification task of CIFAR-10 showing its improvement in accuracy. Moreover, as our regularization increases the expressivity of the model and imposes diversity on the filter, we use its advantages in extreme regimes where utilizing the expressive power of the model can be crucial for its performance. We evaluate GOR on recent low-rank adaptation methods for diffusion models [56, 58] and vision transformers (ViTs) [17]. Specifically, we propose combining GOR with AdaptMLP [2], which is an adaptation strategy for ViT. We show that incorporating GOR into the AdaptMLP mechanism improves ViT classification accuracy on downstream datasets. We also show how to combine GOR with LoRA [17] to adapt diffusion models. We demonstrate that this can lead to improvement in a variety of datasets (e.g., in Pokemon generation as shown in Figure 1). This further illustrates our regularization’s effective-




Figure 2: Visualization of GOR’s group partition for  $N = 3$ . GOR enforces orthonormal regularization on groups of weights in the network layers. Best viewed in color.

Finally, we show that GOR is useful for improving network robustness. When group orthogonality is enforced during adversarial training, robustness is enhanced. Notably, our regularization method enables increasing layer dimensionality without sacrificing efficiency.

Overall, the main contribution of GOR is addressing the correlation between filters in the same layer by promoting orthonormality between filter groups. Using our approach, layer dimensionality may be increased without compromising efficiency. We empirically demonstrate that our new regularization method complements existing normalization techniques and can be applied to a wide range of deep learning models and tasks.

## 2 Related Work

**Feature regularization** techniques aim to reduce the correlation between features in the same layer of a neural network. Batch normalization (BN) [19] and group normalization (GN) [46] are two widely used feature normalization techniques that have been shown to improve the performance of deep neural networks. BN normalizes the layer’s activations to have zero mean and unit variance, while GN normalizes, for each example, the activations of a subset of channels within a given layer to have zero mean and unit variance. Alternatives to BN and GN include layer normalization [2] and instance normalization [43].

The DeCov approach [8] encourages diverse or non-redundant representations in deep neural networks by minimizing the cross-covariance of hidden activations. This method regularizes the feature maps of convolutional layers by encouraging them to be uncorrelated. In [2], a relationship has been drawn between dropout and having an Extreme Value Theory Factorization (ETF) structure in auto-encoders. It has been suggested that imposing such a structure on neural network learning can improve learning.

**Weight regularization** techniques aim to scale the weights of neural networks to have unit norm or zero mean and unit variance. Weight normalization [39] is a technique that normalizes the weight vectors to have unit  $\ell_2$  norm. Weight standardization (WS) [34] is a method that attempts to smoothen the loss landscape by standardizing the weights (i.e., shifting them to have zero mean and unit variance) in convolutional layers. WS has been shown to improve the generalization performance of deep neural networks.

Unlike the reparametrization techniques listed above, which operate on each weight independently, the methods introduced in [4, 47] perform regularization according to relations between filters within the layer. Specifically, they explore orthogonal regularization by promoting the Gram matrix of each weight matrix to be close to identity under the Frobenius norm. The approach presented in [27] proposes to impose orthogonality on specific network components, inspired by methods from sparse dictionary learning. Additional works suggested to enforce orthogonality on the filters [21, 32] which require computing their singular value decomposition. In order to reduce the number of computations, it was further suggested to compute the singular value decomposition of groups of filters [18].

Recently, a framework presented by [41] enforces strict orthogonality by using orthogonal filter bank parametrization. The paper also demonstrates the equivalence between different orthogonal convolutional layers in the spatial domain and the para-unitary systems in the spectral domain.

Our proposed method, Group Orthogonalization Regularization (GOR), is a weight regularization technique that promotes orthonormality between groups of filters within the same layer. GOR complements existing normalization techniques, such as BN and GN, and can be applied to a wide range of deep-learning layers and models. As enforcing orthogonality

on all the weights in a given layer is quadratic in complexity, GOR enables increasing the layer dimensionality without sacrificing efficiency as it operates on groups and thus its complexity scales linearly with respect to the number of groups, and it scales quadratically only with respect to the group size that can be kept constant. Thus, it improves orthogonality between the weights while keeping the computational complexity reasonable.

### 3 Group Orthogonalization

**Preliminaries.** We start by describing some useful notation. To train a DNN, the common practice is to use a loss function for learning a specific task. It is calculated according to the model parameters  $W_{(1)}, \dots, W_{(L)}$ , and is denoted as  $\mathcal{L}_{task}(W_{(1)}, \dots, W_{(L)})$ . Note that  $L$  is the number of layers in the model and  $W_{(l)}$  the parameters of a specific layer ( $l = 1, \dots, L$ ). The regularization we present in this work supports convolutional and fully connected layers. Convolution layers are tensors of parameters, and we use their reshaped form. Let  $\mathbf{W}_{(l)} \in \mathbb{R}^{w \times h \times c \times C_{out}}$  be the convolutional tensor, where  $w, h, c, C_{out}$  are width, height, input channel number and output channel number, respectively. We reshape  $\mathbf{W}_{(l)}$  to be  $W_{(l)} \in \mathbb{R}^{whc \times C_{out}}$ . For simplicity, we define:  $C_{in} = whc$ . Denote by  $\|\cdot\|_F$  the Frobenius norm of a matrix.

#### 3.1 Group Orthogonalization Regularization (GOR)

Similarly to [4], we define *orthogonalization* as the process of enforcing the Gram matrix of a set of filters to be close to identity according to a defined metric.

Previous works have proposed adding regularization on top of the task loss so that the Gram matrix of the parameters is approximately identity,

$$\mathcal{L}_{total}(W_{(1)}, \dots, W_{(L)}) = \mathcal{L}_{task}(W_{(1)}, \dots, W_{(L)}) + \lambda \sum_{l=1}^L \|W_{(l)}^T W_{(l)} - I\|_F^2, \quad (1)$$

Yet, those methods may require an excessive amount of additional computation, especially for wide layers. A detailed explanation can be found in below.

In this work, we suggest reducing the amount of computations by grouping the filters in each layer and orthogonalizing and normalizing only the filters within each group.

$$\mathcal{L}_{total}(W_{(1)}, \dots, W_{(L)}) = \mathcal{L}_{task}(W_{(1)}, \dots, W_{(L)}) + \lambda \sum_{l=1}^L \sum_{i=1}^N \|W_{(i,l)}^T W_{(i,l)} - I\|_F^2. \quad (2)$$

Note that we have  $N$  groups, each with  $\frac{C_{out}}{N}$  filters in it. We flatten and stack the filters in each group to form a matrix  $W_{(i,l)}$  of the  $i$ th group. Then, we enforce orthogonality within each of these groups. The partition into groups is visually demonstrated in Figure 2.

In this way, we do not promote all the filters in a layer to be orthogonal to each other, which might be too restrictive, but rather orthogonal only to a subset of the filters. This can be beneficial, especially when  $C_{in} \ll C_{out}$ . In this case,  $W_{(l)}^T W_{(l)} \in C_{out} \times C_{out}$  is necessarily not orthogonal and the rank of the Gram matrix is bounded by  $C_{in}$ . In this case, using Equation (1) leads to a large magnitude of regularization with limited ability to orthogonalize the parameters. With the partition into groups, this problem is smaller because we now penalize matrices with smaller output dimensions, which can be close to rank  $C_{in}$  more easily.

**Computational Efficiency.** A prominent advantage of GOR is its reduction in computation compared to whole-layer regularization. The complexity of using orthogonalization between

all weights as in Equation (1) is  $C_{out}^2 C_{in}$  per layer. Yet, computing the regularization with groups as in Equation (2) requires only  $O(\frac{C_{out}^2 C_{in}}{N})$ . This can be further improved by parallelizing the matrix multiplication of each group within the layer. The bounds are calculated according to the naive school book algorithm. Runtime graphs appear in the sup. mat.

### 3.2 Group Partition

Given the number of groups,  $N$ , we divide the kernel along the number of output channels,  $C_{out}$ , to  $N$  consecutive groups - each contains  $G = \frac{C_{out}}{N}$  filters. The partition is demonstrated in Figure 2. Given a partition, we distinguish between two cases:

- (a) Inter-Group regularization: The orthonormality is enforced on filters within the same group. Thus the regularization term (for a single layer) is  $\sum_{i=1}^N \|W_i^T W_i - I\|_F^2$ , where

$$W_i^T = (f_{(i-1)G+1} | f_{(i-1)G+2} | \dots | f_{iG}).$$

- (b) Intra-Group regularization: The orthonormality is enforced on filters between different groups. In this case, the per-layer regularization term is  $\sum_{i=1}^G \|W_i^T W_i - I\|_F^2$ , where

$$W_i^T = (f_i | f_{G+i} | f_{2G+i} | \dots | f_{(N-1)G+i}).$$

We focus on Inter-Group regularization.

### 3.3 Relation to Group Normalization

In the general case, where no other partition is applied to the intermediate features and/or filters, the groups can be chosen arbitrarily without affecting the learned function. This is because the order of the filters is not significant to the trained model.

One prominent method used for reducing the computational overhead of training DNNs is GN [46], which replaces BN operations. For GN, the intermediate output features of the network are partitioned into groups, and each group is centered and normalized. In this case, it is no longer correct to state that the partition of the filters in each layer is arbitrary.

In this work, we distinguish between two possible options for partitioning the filters. The first, which we call “inter-group” regularization where we use filters that match the normalization group and apply the regularization within this group. The second is “intra-group” regularization, where we choose to apply our regularization to filters that match features from different groups of normalization. In particular, each filter in a regularization group corresponds to a different normalization group.

The “intra-group” regularization may implicitly amplify the orthogonality between the normalization groups, while the “inter-group” regularization enhances orthogonality within normalization groups. Both settings are illustrated in the supplementary material and further explained in Section 3.2.

### 3.4 GOR with Adapters

A common paradigm in the field of deep learning is the use of foundation models [6, 24, 35, 36, 38, 42]. It relies on large-scale pretraining on general domain data which is followed by adaptation to particular tasks or domains. As models become larger, full fine-tuning becomes

Table 1: CIFAR-10 Top-1 accuracy. ResNet110 GN is a 110-layer ResNet where all BN layers are replaced with GN. Both ResNet110 models are comprised of the basic block as the building block.

Model	CE	SO	GOR inter	GOR intra
ResNet110	93.95 $\pm$ 0.04	<b>94.44 <math>\pm</math> 0.05</b>	94.23 $\pm$ 0.07	-
ResNet110 GN	92.02 $\pm$ 0.04	92.33 $\pm$ 0.03	<b>92.73 <math>\pm</math> 0.03</b>	92.24 $\pm$ 0.02

less feasible due to the heavy computation and the large storage needed. Some works that were introduced for large language models [16, 17, 29], propose using lightweight adapters as an alternative to full model fine-tuning. One of the most prominent works for this task is LoRA [18]. Such adaptation approaches have also been proposed for the computation vision domain. The AdaptFormer [19] is a recent strategy that has shown promising results in the computer vision domain, especially when combined with vision transformers [20].

The motivation for using GOR for model adaptation is that models trained with it are more likely to explore unknown directions in the parameter space, which may be beneficial when transforming between domains.

In our work, we focus on low-rank adapter modules as presented in [4, 19]. Note that the AdaptMLP block follows the same decomposition as in LoRA but with a larger rank (64 vs. 4) and a non-linear layer in between. Both of these techniques fine-tune the foundation model weights by learning the residual that should be added to some selected layers. This residual is parameterized by a low-rank matrix that is composed of two matrices. The first reduces the size (‘down sampling block’), and the second increases it (‘up sampling block’).

GOR is applied to the adapter’s matrices to encourage weights’ expressibility. Utilizing parameter space is more crucial for model performance in low-rank adapters. Regularizing the downsampling block, i.e., the matrix that transforms from a high dimension to a small one, has little effect as there are few filters in each group. Additionally, due to the extremely small output dimension of this block, it is most likely that features span the whole space. Thus, we only focus on applying GOR to the up-sampling block of the adapter.

## 4 Experiments

We evaluate GOR by conducting extensive experiments across several architectures and tasks. For each section, we describe the experimental setting and training details.

### 4.1 Image Classification on CIFAR10

To evaluate our method, we first perform image classification experiments.

**Experimental setting.** Using the 110-layer ResNet [13] as a backbone, we benchmark our approach on the CIFAR-10 [25] dataset. To assess how well our technique complements GN, we also train the aforementioned architecture in its GN variant - meaning we replace all BN layers with GN layers. The rest of the model stays untouched. GOR is added to the optimization of all convolutional layers. Accuracy is compared to two baselines: non-regularized training (CE) and Soft-Orthogonalization (SO) [8], which enforces the entire ( $N = 1$  groups) set of filters to be orthonormal. For the GN variant, we tested both “inter-group” and “intra-group” GOR. Only “inter-group” is reported for the BN variant since, for this type of normalization, the two GOR variants are equivalent. As discussed in Section 3.3.

Table 2: Fine-tuning self-supervised ViT-B via AdaptMLP. Full-tuning and AdaptMLP results are taken from the original work. Pre-trained weights are from MAE [42]. We report mean and std across 3 runs.

Method	CIFAR-100	SVHN	Food-101
Full-tuning	85.9	97.67	90.09
AdaptFormer	85.9	<b>96.89</b>	87.61
AdaptFormer + GOR	<b>86.16</b> $\pm$ 0.07	96.87 $\pm$ 0.09	<b>87.73</b> $\pm$ 0.1

**Training details.** We train on the 50,000 training images set and evaluate on the 10,000 images test set. Top-1 accuracy is reported. With regard to batch size, number of epochs, weight decay, etc., training follows the exact same protocol as in [43], both for the GN model and the original model. Let  $G$  be the number of groups for the GN layer. Following [54], the value of  $G$  is determined to be the minimum value between 32 and  $(\# \text{ of channels})/4$ . As for the regularization’s hyperparameters, the value of  $\lambda$  is set to  $10^{-2}$ , and  $N$  (number of groups for GOR) follows the same logic as  $G$ . For each table entry, we report the mean and std across three different random seeds.

**Comparing regularization methods.** We compare the model’s accuracy when trained with the two different regularization methods in Table 1. Both regularization methods improve generalization ability as both are superior to training without orthogonal regularization. For the original BN model, GOR showed competitive results to SO while being more efficient. The results show the benefit of using GOR along with GN, where the regularization enhances orthogonality within normalization groups.

## 4.2 Adapting ViT with GOR for image classification

In this section, we report accuracy on several image datasets when fine-tuning a pre-trained ViT via the AdaptMLP [4] bottleneck model.

**Experimental setting.** Following experiments done by [4], we fine-tune two pre-trained Vision Transformers (both supervised and self-supervised) on multiple downstream datasets. The same ViT backbone model is used. We compare our results with a full fine-tuning setting and non-regularized AdaptFormer. In each training sequence, the pre-trained weights are frozen, and only the newly added modules are optimized. GOR regularization is applied to the AdaptMLP modules, specifically only to the up-projection layers.

**Datasets.** We use three common datasets: CIFAR-100 [45] contains 50,000 training images and 10,000 test images of resolution  $32 \times 32$  with 100 labels. Street View House Numbers (SVHN) [30] is a digit classification benchmark dataset. In total, the dataset comprises over 600,000 labeled images containing 73,257 training samples, 26,032 testing samples, and 531,131 extra training data. The Food-101 [43] dataset consists of 101 food categories with a total of 101k images, including 750 training and 250 testing samples per category.

**Training details.** We follow the protocol described in the original work and use the same weights for the supervised and self-supervised baselines. The bottleneck dimension is set to 64 as recommended in the original paper [4]. As for GOR configuration, We set  $N$  to 16 and  $\lambda$  to  $10^{-4}$ . Training is performed on  $4 \times$  NVIDIA GeForce RTX 2080 Ti GPUs together with gradient accumulation to match the original works’ batch size.

**Results.** Table 2 and Table 3 show that indeed our regularization improves the expressiveness of the bottleneck module and thus allows it to generalize better and achieve higher accuracy. GOR improves fine-tuning performance across almost all datasets and pre-trained weights.

Table 3: Fine-tuning supervised ViT-B via AdaptMLP. Full-tuning and AdaptMLP results are taken from the original work. Pre-trained weights are from the ImageNet-21k [10] supervised pre-trained model. We report mean and std across 3 runs.

Method	CIFAR-100	SVHN	Food-101
Full-tuning	89.12	95.41	90.96
AdaptFormer	91.86	97.29	90.89
AdaptFormer + GOR	<b>92.49</b> $\pm$ 0.11	<b>97.36</b> $\pm$ 0.04	<b>91.17</b> $\pm$ 0.02

Table 4: FID score comparisons on downstream datasets. Lower is better.

Method	Oxford102	Pokemon-BLIP	FS-COCO
LoRA	11.01	13.6	30.75
LoRA + GOR	<b>10.57</b>	<b>13.14</b>	<b>30.29</b>

### 4.3 Adapting Diffusion models with GOR for text-to-image generation

We further evaluate the effectiveness of applying GOR to adapters by regularizing the fine-tuning process of image-generation models adapted with LoRA [10]. Efficient fine-tuning of diffusion models by injecting LoRA bottleneck modules into the U-Net’s cross-attention layers was first proposed by [10]. While originally used to fine-tune large language models (LLMs) on downstream tasks, the writer suggested applying the rank-decomposition matrices to image generation models. To our best knowledge, our work is the first time diffusion models adapted with LoRA are quantitatively evaluated on the task of text-to-image generation.

**Experimental setting.** We measure how well the diffusion model adapts to the target data distribution when using GOR to regularize the LoRA matrices. The generated image fidelity of the adapted text-to-image model is evaluated using the Fréchet inception distance (FID) [10] metric. For each dataset, the distance is computed between the generated images and all available real images. To reduce the impact of random variation, we compute FID three times in each experiment and report the minimum. LoRA [10] matrices are added to all four ( $W_q, W_k, W_v, W_o$ ) of the U-Net’s self-attention modules. As previously mentioned, during the model’s optimization, we only enforce orthogonality via GOR on the "up" (denoted as  $B$  in the original paper) matrices. Furthermore, we only regularize the adapter added to the upsampling blocks of the U-Net.

**Datasets.** The Oxford 102 category flower dataset [10] consists of 102 flower categories. Each class consists of between 40 and 258 images. Images have large scale, pose, and light variations. The second dataset used is Pokemon-BLIP [10]. This image dataset consists of 833 text-image pairs. The pairs include a Pokemon image and a BLIP-generated [10] caption. FS-COCO comprises 10,000 freehand scene vector sketches drawn by 100 non-expert individuals, offering both object and scene-level abstraction. Each sketch is augmented with its text description. Representative examples can be seen in the supplementary material.

**Note on FID calculation.** When measuring FID on the Oxford102 and FS-COCO datasets, we use 10,000 generated images to ensure that representative statistics are calculated. Due to the small size of the Pokemon-BLIP dataset, we only use 583 generated images. For all datasets, we randomly sample prompts and use them to condition the diffusion model during generation. More information on measuring the FID is listed the supplementary material.

**Training details.** Pre-trained latent diffusion model weights taken from the HuggingFace Diffusers library [10]. Specifically, the *stable-Diffusion-v1-5* checkpoint is used. Following



the original work, we set the LoRA rank ( $r$  from paper) to 4 and only train the adapter’s matrices while keeping the rest of the weights frozen. We set  $N = 32$  to match the groups of normalized features in the GN layer and train the model for 15K iterations. The value of  $\lambda$  is set to  $10^{-5}$  for the Oxford102 dataset and  $10^{-6}$  for the other two datasets. All experiments are conducted on  $2 \times$  NVIDIA GeForce RTX 2080 Ti. The rest of the training protocol is listed in the supplementary material.

**Effect of regularization on FID score.** Table 4 shows that enforcing group-orthogonality of the LoRA layers throughout the model’s fine-tuning improves its ability to learn the underlying distribution. A qualitative comparison between for and the baseline is presented in the supplementary material.

## 4.4 Adversarial Robustness

A major challenge faced by deep neural networks is their vulnerability to small changes in input data, which can result in incorrect predictions. This presents a significant difficulty, particularly for applications that require high safety standards. To tackle this issue, the development of adversarial defenses has emerged as a critical research area across various fields, including machine learning, computer vision, natural language processing, and others.

One of the approaches utilized to defend against such attacks is adversarial training [22, 28, 40, 50, 51]. It involves generating adversarial examples during the training process and utilizing them to update the model parameters. The adversarial training scheme can be viewed as a form of regularization that prevents the model from overfitting to the distributions of clean training data.

Recent works [20, 48] have shown that model orthogonality improves robustness. Motivated by these results, we investigate how adversarial training methods can benefit from our efficient group orthogonalization.

**Experimental setting.** We evaluate the robustness of two adversarial training methods when combined with GOR on CIFAR-10 in Table 5. The Trained model is evaluated under several white-box and black-box attacks. The two well-known training techniques we consider are TRADES [50] and FAT [51]. We train both GN and BN models.

**Adversarial Attacks** The adversarial test samples are bounded by  $L_\infty$  perturbations with  $\varepsilon = 8/255$ , which are generated by FGSM, PGD<sub>20</sub>, PGD<sub>100</sub> and CW<sub>∞</sub>. Where the subscript numbers indicate the number of iterations used for calculating the attack and CW<sub>∞</sub> is the  $L_\infty$  version of C&W loss [5] optimized by PGD<sub>20</sub>. Trained models are also attacked by the AutoAttack [10] method, which consists of APGD-CE, APGD-DLR, FAB [9] and Square [11].

**Training and evaluation details.** We follow the original settings for each of the training methods, with Wide ResNet [49] as the chosen architecture. We train a WideResNet-34-10 for TRADES and a WideResNet-32-10 for FAT. The GN model is created by replacing all BN layers with GN, with  $G = 32$ . We report the test accuracy of the deep model at the last training epoch (76 for TRADES and 120 for FAT). As for GOR configuration, we keep  $N = G = 32$  and show results for  $\lambda$  at  $10^{-4}$  and  $10^{-5}$  for all models. Inter-group regularization is used.

**Effect of regularization on robustness.** Table 5 justifies the assumption that GOR allows the model to learn more diverse and robust features that are less correlated and more informative. Models regulated by GOR show better natural accuracy and, in most cases, are more robust to attacks. The gains are more significant for the GN models. This affirms our claim that GOR complements GN well, especially when regularizing according to the normalized groups.

Table 5: Test accuracy under different training methods on CIFAR10 dataset under attacks bounded by  $L_\infty$ . 34-layer wide ResNet is used for TRADES [50] training, and a 32-layer wide ResNet is used for FAT training [51].

Method	$\lambda$	Natural	FGSM	PGD <sub>20</sub>	PGD <sub>100</sub>	CW <sub><math>\infty</math></sub>	AutoAttack
TRADES + BN	0	83.73	66.12	<b>55.85</b>	55.52	53.55	52.45
TRADES + BN	$10^{-4}$	<b>84.63</b>	<b>66.65</b>	55.16	55.00	53.66	52.33
TRADES + BN	$10^{-5}$	84.01	66.5	55.13	<b>55.93</b>	<b>54.33</b>	<b>53.14</b>
TRADES + GN	0	81.88	63.72	53.27	53.11	50.87	49.72
TRADES + GN	$10^{-4}$	<b>83.86</b>	<b>65.67</b>	<b>54.76</b>	<b>54.58</b>	<b>52.87</b>	<b>51.56</b>
TRADES + GN	$10^{-5}$	82.72	64.42	54.04	53.77	51.77	50.66
FAT + BN	0	88.6	<b>65.05</b>	46.34	45.5	46.87	43.62
FAT + BN	$10^{-4}$	88.83	64.49	45.82	45.1	46.58	43.39
FAT + BN	$10^{-5}$	<b>88.85</b>	65	<b>46.72</b>	<b>45.83</b>	<b>47.16</b>	<b>43.71</b>
FAT + GN	0	82.83	54.13	36.34	35.68	38.47	34.19
FAT + GN	$10^{-4}$	87.87	64.84	46.31	45.77	47.34	44.52
FAT + GN	$10^{-5}$	<b>88.17</b>	<b>66.73</b>	<b>47.07</b>	<b>46.4</b>	<b>48.37</b>	<b>45.04</b>

## 4.5 Limitations

While GOR shows promising results on multiple tasks, some limitations need to be addressed in future work. First, it introduces a new hyper-parameter,  $\lambda$ , that might need to be re-tuned depending on the architecture and task. For non-GN models,  $N$  might need to be tuned as well. Second, and similarly to other orthogonalization regularizations [9, 47], it introduces some computational overhead that is proportional to the number of regularized layers.

## 5 Conclusion

In this study, we propose a novel regularization technique that encourages orthonormality between groups of filters within the same layer. This technique is computationally efficient and can significantly reduce the redundancy within the parameters of deep neural networks.

Our experiments show that incorporating our regularization technique into recent adaptation methods for diffusion models and vision transformers (ViTs) leads to improved performance on downstream tasks. Moreover, enforcing group orthogonality during adversarial training results in better model robustness. Overall, our proposed regularization technique effectively enhances the performance and robustness of deep neural networks. By reducing the redundancy within the network’s parameters, we can create more efficient and accurate models that perform better on a variety of tasks.

**Acknowledgements** This work was partially supported by the European research council under Grant ERC-StG 757497.

## References

- [1] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII*, pages 484–501. Springer, 2020.
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey Hinton. Layer normalization. In *ICLR*, 2017.
- [3] Dor Bank and Raja Giryes. An etf view of dropout regularization. In *BMVC*, 2020.
- [4] Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. Can we gain more from orthogonality regularizations in training deep networks? *Advances in Neural Information Processing Systems*, 31, 2018.
- [5] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. Ieee, 2017.
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.
- [7] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition, 2022.
- [8] Michael Cogswell, Faruk Ahmed, Ross B. Girshick, Larry Zitnick, and Dhruv Batra. Reducing overfitting in deep networks by decorrelating representations. In *ICLR*, 2016.
- [9] Francesco Croce and Matthias Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *International Conference on Machine Learning*, pages 2196–2205. PMLR, 2020.
- [10] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International conference on machine learning*, pages 2206–2216. PMLR, 2020.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.

- [14] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022.
- [15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [16] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Larous-silhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp, 2019.
- [17] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021.
- [18] Lei Huang, Xianglong Liu, Bo Lang, Adams Yu, Yongliang Wang, and Bo Li. Orthogonal weight normalization: Solution to optimization over multiple dependent stiefel manifolds in deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [19] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456, 2015.
- [20] Mohammad AAK Jalwana, Naveed Akhtar, Mohammed Bennamoun, and Ajmal Mian. Orthogonal deep models as defense against black-box attacks. *IEEE Access*, 8:119744–119757, 2020.
- [21] Kui Jia, Dacheng Tao, Shenghua Gao, and Xiangmin Xu. Improving training of deep neural networks via singular value bounding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4344–4352, 2017.
- [22] Xiaojun Jia, Yong Zhang, Baoyuan Wu, Ke Ma, Jue Wang, and Xiaochun Cao. Las-at: adversarial training with learnable attack strategy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13398–13408, 2022.
- [23] Parneet Kaur, Karan Sikka, and Ajay Divakaran. Combining weakly and weakly supervised learning for classifying food images. *arXiv preprint arXiv:1712.08730*, 2017.
- [24] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023.
- [25] A Krizhevsky. Learning multiple layers of features from tiny images. *Master’s thesis, University of Tront*, 2009.
- [26] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022.
- [27] Yishi Li, Kunran Xu, Rui Lai, and Lin Gu. Towards an effective orthogonal dictionary convolution strategy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 1473–1481, 2022.

- [28] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [29] Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter layers, 2021.
- [30] Yuval Netzer, Tao Wang, Adam Coates, A. Bissacco, Bo Wu, and A. Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- [31] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008.
- [32] Mete Ozay and Takayuki Okatani. Optimization on submanifolds of convolution kernels in cnns. *arXiv preprint arXiv:1610.07008*, 2016.
- [33] Justin N. M. Pinkney. Pokemon blip captions. <https://huggingface.co/datasets/lambdalabs/pokemon-blip-captions/>, 2022.
- [34] Siyuan Qiao, Huiyu Wang, Chenxi Liu, Wei Shen, and Alan Yuille. Micro-batch training with batch-channel normalization and weight standardization. *arXiv preprint arXiv:1903.10520*, 2019.
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [36] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831, 2021.
- [37] Simo Ryu. lora. <https://github.com/cloneofsimon/lora>, 2022.
- [38] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, volume 35, pages 36479–36494, 2022.
- [39] Tim Salimans and Diederik P. Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks, 2016.
- [40] Ali Shafahi, Mahyar Najibi, Mohammad Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! *Advances in Neural Information Processing Systems*, 32, 2019.
- [41] Jiahao Su, Wonmin Byeon, and Furong Huang. Scaling-up diverse orthogonal convolutional networks by a paraunitary framework. In *International Conference on Machine Learning*, pages 20546–20579. PMLR, 2022.

- [42] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [43] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. In *ICCV*, pages 3500–3508, 2017.
- [44] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>, 2022.
- [45] Jiayun Wang, Yubei Chen, Rudrasis Chakraborty, and Stella X Yu. Orthogonal convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11505–11515, 2020.
- [46] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, pages 3–19, 2018.
- [47] Di Xie, Jiang Xiong, and Shiliang Pu. All you need is beyond a good init: Exploring better solution for training extremely deep convolutional neural networks with orthonormality and modulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6176–6185, 2017.
- [48] Cong Xu, Xiang Li, and Min Yang. An orthogonal classifier for improving the adversarial robustness of neural networks. *Information Sciences*, 591:251–262, 2022.
- [49] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [50] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International conference on machine learning*, pages 7472–7482. PMLR, 2019.
- [51] Jingfeng Zhang, Xilie Xu, Bo Han, Gang Niu, Lizhen Cui, Masashi Sugiyama, and Mohan Kankanhalli. Attacks which do not kill training make adversarial learning stronger. In *International conference on machine learning*, pages 11278–11287, 2020.