

WLD-Reg: A Data-Dependent Within-Layer Diversity Regularizer

Firas Laakom¹, Jenni Raitoharju², Alexandros Iosifidis³, Moncef Gabbouj¹

¹ Faculty of Information Technology and Communication Sciences, Tampere University, Finland

² Faculty of Information Technology, University of Jyväskylä, Finland

³ DIGIT, Department of Electrical and Computer Engineering, Aarhus University, Denmark

{firas.laakom, moncef.gabbouj}@tuni.fi, jenni.k.raitoharju@jyu.fi, ai@ece.au.dk

Abstract

Neural networks are composed of multiple layers arranged in a hierarchical structure jointly trained with a gradient-based optimization, where the errors are back-propagated from the last layer back to the first one. At each optimization step, neurons at a given layer receive feedback from neurons belonging to higher layers of the hierarchy. In this paper, we propose to complement this traditional 'between-layer' feedback with additional 'within-layer' feedback to encourage the diversity of the activations within the same layer. To this end, we measure the pairwise similarity between the outputs of the neurons and use it to model the layer's overall diversity. We present an extensive empirical study confirming that the proposed approach enhances the performance of several state-of-the-art neural network models in multiple tasks. The code is publicly available at <https://github.com/firasl/AAAI-23-WLD-Reg>.

Introduction

Deep learning has been extensively used in the last decade to solve several tasks (Krizhevsky, Sutskever, and Hinton 2012; Golan and El-Yaniv 2018; Hinton et al. 2012a). A deep learning model, i.e., a neural network, is formed of a sequence of layers with parameters optimized during the training process using training data. Formally, an m -layer neural network model can be defined as follows:

$$f(\mathbf{x}; \mathcal{W}) = \phi^m(\mathbf{W}^m(\phi^{m-1}(\dots \phi^2(\mathbf{W}^2\phi^1(\mathbf{W}^1\mathbf{x})))), \quad (1)$$

where $\phi^i(\cdot)$ is the non-linear activation function of the i^{th} layer and $\mathcal{W} = \{\mathbf{W}^1, \dots, \mathbf{W}^m\}$ are the model's weights. Given a training data $\{\mathbf{x}_i, y_i\}_{i=1}^N$, the parameters of $f(\mathbf{x}; \mathcal{W})$ are obtained by minimizing a loss $\hat{L}(\cdot)$:

$$\hat{L}(f) = \frac{1}{N} \sum_{i=1}^N l(f(\mathbf{x}_i; \mathcal{W}), y_i). \quad (2)$$

However, neural networks are often over-parameterized, i.e., have more parameters than data. As a result, they tend to overfit to the training samples and not generalize well on unseen examples (Goodfellow et al. 2016). While research on double descent (Advani, Saxe, and Sompolsky

2020; Belkin et al. 2019; Nakkiran et al. 2020) shows that over-parameterization does not necessarily lead to overfitting, avoiding overfitting has been extensively studied (Dzigaite and Roy 2017; Foret et al. 2020; Nagarajan and Kolter 2019; Neyshabur et al. 2018; Poggio et al. 2017; Grari et al. 2021) and various approaches and strategies, such as data augmentation (Goodfellow et al. 2016; Zhang et al. 2018), regularization (Arora et al. 2019; Bietti et al. 2019; Kukačka, Golkov, and Cremers 2017; Ouali, Hudelot, and Tami 2021; Han and Guo 2021), and Dropout (Hinton et al. 2012b; Lee et al. 2019; Li, Gong, and Yang 2016; Wang et al. 2019), have been proposed to close the gap between the empirical loss and the expected loss.

Diversity of learners is widely known to be important in ensemble learning (Li, Yu, and Zhou 2012; Yu, Li, and Zhou 2011) and, particularly in the deep learning context, diversity of information extracted by the network neurons has been recognized as a viable way to improve generalization (Xie, Liang, and Song 2017; Xie, Deng, and Xing 2015b). In most cases, these efforts have focused on making the set of weights more diverse (Yang, Gkatzelis, and Stoyanovich 2019; Malkin and Bilmes 2009). However, diversity of the activations has not received much attention. Here, we argue that due to the presence of non-linear activations, diverse weights do not guarantee diverse feature representation. Thus, we propose focusing on the diversity on top of feature mapping instead of the weights.

To the best of our knowledge, only (Cogswell et al. 2016; Laakom et al. 2021a) have considered diversity of the activations directly in the neural network context. The work in (Laakom et al. 2021a) studied theoretically how diversity affects generalization showing that it can reduce overfitting. The work in (Cogswell et al. 2016) proposed an additional loss term using cross-covariance of hidden activations, which encourages the neurons to learn diverse or non-redundant representations. The proposed approach, known as DeCov, was empirically proven to alleviate overfitting and to improve the generalization ability of neural networks. However, modeling diversity as the sum of the pairwise cross-covariance, it is not scale-invariant and can lead to trivial solutions. Moreover, it can capture only the pairwise diversity between components and is unable to capture the "higher-order diversity".

In this work, we propose a novel approach to encour-

age activation diversity within the same layer. We propose complementing the 'between-layer' feedback with additional 'within-layer' feedback to penalize similarities between neurons on the same layer. Thus, we encourage each neuron to learn a distinctive representation and to enrich the data representation learned within each layer. We propose three variants for our approach that are based on different global diversity definitions.

Our contributions in this paper are as follows:

- We propose a new approach to encourage the 'diversification' of the layers' output feature maps in neural networks. The proposed approach has three variants. The main intuition is that, by promoting the within-layer activation diversity, neurons within a layer learn distinct patterns and, thus, increase the overall capacity of the model.
- We show empirically that the proposed within-layer activation diversification boosts the performance of neural networks. Experimental results on several tasks show that the proposed approach outperforms competing methods.

Within-Layer Diversity Regularizer

In this section, we propose a novel diversification strategy, where we encourage neurons within a layer to activate in a mutually different manner, i.e., to capture different patterns. In this paper, we define as "feature layer" the last intermediate layer in a neural network. In the rest of the paper, we focus on this layer and propose a data-dependent regularizer which forces each unit within this layer to learn a distinct pattern and penalizes the similarities between the units. Intuitively, the proposed approach reduces the reliance of the model on a single pattern and, thus, can improve generalization.

We start by modeling the global similarity between two units. Let $\phi_n(\mathbf{x}_j)$ and $\phi_m(\mathbf{x}_j)$ be the outputs of the n^{th} and m^{th} unit in the feature layer for the same input sample \mathbf{x}_j . The similarity s_{nm} between the n^{th} and m^{th} neurons can be obtained as the average similarity measure of their outputs for N input samples. We use the radial basis function to express the similarity:

$$s_{nm} := \frac{1}{N} \sum_{j=1}^N \exp(-\gamma \|\phi_n(\mathbf{x}_j) - \phi_m(\mathbf{x}_j)\|^2), \quad (3)$$

where γ is a hyper-parameter. The similarity s_{nm} can be computed over the whole dataset or batch-wise. Intuitively, if two neurons n and m have similar outputs for many samples, their corresponding similarity s_{nm} will be high. Otherwise, their similarity s_{nm} is small and they are considered "diverse".

Next, based on these pairwise similarities, we propose three variants for obtaining the overall similarity J of all the units within the feature layer:

- **Direct:** $J := \sum_{n \neq m} s_{nm}$. In this variant, we model the global layer similarity directly as the sum of the pairwise similarities between the neurons. By minimizing their sum, we encourage the neurons to learn different representations.

• **Det:** $J := -\det(\mathbf{S})$, where \mathbf{S} is a similarity matrix defined as $S_{nm} = s_{nm}$. This variant is inspired by the Determinantal Point Process (DPP) (Kulesza and Taskar 2010, 2012), as the determinant of \mathbf{S} measures the global diversity of the set. Geometrically, $\det(\mathbf{S})$ is the volume of the parallelepiped formed by vectors in the feature space associated with s . Vectors that result in a larger volume are considered more "diverse". Thus, maximizing $\det(\cdot)$ (minimizing $-\det(\cdot)$) encourages the diversity of the learned features.

- **Logdet:** $J := -\log \det(\mathbf{S})$ ¹. This variant has the same motivation as the second one. We use Logdet instead of Det as Logdet is a convex function over the positive definite matrix space.

It should be noted here that the first proposed variant, i.e., direct, similar to DeCov (Cogswell et al. 2016), captures only the pairwise similarity between components and is unable to capture the higher-order "diversity", whereas the other two variants consider the global similarity and are able to measure diversity in a more global manner. Promoting diversity of activations within a layer can lead to tighter generalization bound and can theoretically decrease the gap between the empirical and the true risks (Laakom et al. 2021a).

The proposed global similarity measures J can be minimized by using them as an additional loss term. However, we note that the pair-wise similarity measure s_{nm} , expressed in (3), is not scale-invariant. In fact, it can be trivially minimized by making all activations of the feature layer high, i.e., by multiplying by a high scaling factor, which has no effect on the performance, since the model can rescale high activations to normal values simply by learning small weights on the next layer. To alleviate this problem, we propose an additional term, which penalizes high activation values. The total proposed additional loss is defined as follows:

$$\hat{L}_{WLD-Reg} := \lambda_1 J + \lambda_2 \sum_{i=1}^N \|\Phi(\mathbf{x}_i)\|_2^2, \quad (4)$$

where $\Phi(\mathbf{x}) = [\phi_1(\mathbf{x}), \dots, \phi_C(\mathbf{x})]$ is the feature vector, C is the number of units within the feature layer, and λ_1 and λ_2 are two hyper-parameters controlling the contribution of each term to the diversity loss. Intuitively, the first term of (4) penalizes the similarity between the units and promotes diversity, whereas the second term ensures the scale-invariance of the proposed regularizer.

The total loss function $\hat{L}(f)$ defined in (2) is augmented as follows:

$$\begin{aligned} \hat{L}_{aug}(f) &:= \hat{L}(f) + \hat{L}_{WLD-Reg} \\ &= \hat{L}(f) + \lambda_1 J + \lambda_2 \sum_{i=1}^N \|\Phi(\mathbf{x}_i)\|_2^2. \end{aligned} \quad (5)$$

The proposed approach is summarized in Algorithm 1. We note that our approach can be incorporated in a plug-and-

¹This is defined only if \mathbf{S} is positive definite. It can be shown that in our case \mathbf{S} is positive semi-definite. Thus, in practice, we use a regularized version $(\mathbf{S} + \epsilon \mathbf{I})$ to ensure the positive definiteness.

Algorithm 1: One epoch of training with WLD-Reg

Model: Given a neural network $f(\cdot)$ with a feature representation $\phi(\cdot)$, i.e., last intermediate layer.
Input: Training Data: $\{\mathbf{x}_i, y_i\}_{i=1}^N$
Parameters: λ_1 and λ_2 in (4)

```

1: for every mini-batch:  $\{\mathbf{x}_i, y_i\}_{i=1}^m \in \{\mathbf{x}_i, y_i\}_{i=1}^N$  do
2:   Forward pass the inputs  $\{\mathbf{x}_i\}_{i=1}^m$  into the model to
   obtain the outputs  $\{f(\mathbf{x}_i)\}_{i=1}^m$  and the feature repre-
   sentations  $\{\Phi(\mathbf{x}_i)\}_{i=1}^m$ 
3:   Compute the standard loss  $\hat{L}(f)$  ((2)).
4:   Compute the extra loss  $\hat{L}_{WLD-Reg}$  ((4)).
5:   Compute the total loss  $\hat{L}_{aug}(f)$  ((5))
6:   Compute the gradient of the total loss and use it to
   update the weights of  $f$ .
7: end for
8: return Return  $f$ .

```

play manner into any neural network-based approach to augment the original loss and to ensure learning diverse features. We also note that although in this paper, we focus only on applying diversity regularizer to a single layer, i.e., the feature layer, our proposed diversity loss, as in (Cogswell et al. 2016), can be applied to multiple layers within the model.

Our newly proposed loss function defined in (5) has two terms. The first term is the classic loss function. It computes the loss with respect to the ground-truth. In the back-propagation, this feedback is back-propagated from the last layer to the first layer of the network. Thus, it can be considered as a between-layer feedback, whereas the second term is computed within a layer. From (5), we can see that our proposed approach can be interpreted as a regularization scheme. However, regularization in deep learning is usually applied directly on the parameters, i.e., weights (Goodfellow et al. 2016; Kukačka, Golkov, and Cremers 2017), while in our approach a data-dependent additional term is defined over the output maps of the layers. For a feature layer with C units and a batch size of m , the additional computational cost is $O(C^2(m + 1))$ for Direct variant and $O(C^3 + C^2m)$ for both Det and Logdet variants.

Related Work

Diversity promoting strategies have been widely used in ensemble learning (Li, Yu, and Zhou 2012; Yu, Li, and Zhou 2011), sampling (Biyik et al. 2019; Derezhinski, Calandriello, and Valko 2019; Gartrell et al. 2019), energy-based models (Laakom et al. 2021b; Zhao, Mathieu, and Le-Cun 2017), ranking (Gan et al. 2020; Yang, Gkatzelis, and Stoyanovich 2019), pruning by reducing redundancy (He et al. 2019; Kondo and Yamauchi 2014; Lee et al. 2020; Singh et al. 2020), and semi-supervised learning (Zboncar et al. 2021). In the deep learning context, various approaches have used diversity as a direct regularizer on top of the weight parameters. Here, we present a brief overview of these regularizers. Based on the way diversity is de-

fined, we can group these approaches into two categories. The first group considers the regularizers that are based on the pairwise dissimilarity of the components, i.e., the overall set of weights is diverse if every pair of weights is dissimilar. Given the weight vectors $\{\mathbf{w}_m\}_{m=1}^M$, (Yu, Li, and Zhou 2011) defines the regularizer as $\sum_{mn}(1 - \theta_{mn})$, where θ_{mn} represents the cosine similarity between \mathbf{w}_m and \mathbf{w}_n . In (Bao et al. 2013), an incoherence score defined as $-\log\left(\frac{1}{M(M-1)}\sum_{mn}\beta|\theta_{mn}|^{\frac{1}{\beta}}\right)$, where β is a positive hyperparameter, is proposed. In (Xie, Deng, and Xing 2015a; Xie, Zhu, and Xing 2016), $\text{mean}(\theta_{mn}) - \text{var}(\theta_{mn})$ is used to regularize Boltzmann machines. The authors theoretically analyzed its effect on the generalization error bounds in (Xie, Deng, and Xing 2015b) and extend it to kernel space in (Xie, Liang, and Song 2017). The second group of regularizers considers a more global view of diversity. For example, in (Malkin and Bilmes 2008, 2009; Xie, Singh, and Xing 2017), a weight regularization based on the determinant of the weights' covariance is proposed based on determinantal point process (Kulesza and Taskar 2012; Kwok and Adams 2012).

Unlike the aforementioned methods which promote diversity on the weight level and similar to our method, (Cogswell et al. 2016; Laakom et al. 2022) proposed to enforce dissimilarity on the feature map outputs, i.e., on the activations. To this end, they proposed an additional loss based on the pairwise covariance of the activation outputs. Their additional loss, L_{Decov} , is defined as the squared sum of the non-diagonal elements of the global covariance matrix \mathbf{C} of the activations:

$$L_{Decov} = \frac{1}{2}(\|\mathbf{C}\|_F^2 - \|\text{diag}(\mathbf{C})\|_2^2), \quad (6)$$

where $\|\cdot\|_F$ is the Frobenius norm. Their approach, $Decov$, yielded superior empirical performance. However, correlation is highly sensitive to noise (Kim, Kim, and Ergün 2015), as opposite to the RBF-based distance used in our approach (Savas and Dovis 2019; Haykin 2010). Moreover, the $Decov$ approach only captures the pairwise diversity between the components, whereas we propose variants of our approach which consider a global view of diversity. Moreover, based on the cross-covariance, their approach is not scale-invariant. In fact, it can be trivially minimized by making all activations in the latent representation small, which has no effect on the generalization since the model can rescale tiny activations to normal values simply by learning large weights on the next layer.

Experimental Results

CIFAR10 & CIFAR100

We start by evaluating our proposed diversity approach on two image datasets: CIFAR10 and CIFAR100 (Krizhevsky, Hinton et al. 2009). They contain 60,000 (50,000 train/10,000 test) 32×32 images grouped into 10 and 100 distinct categories, respectively. We split the original training set (50,000) into two sets: we use the first 40,000 images as the main training set and the last 10,000 as a validation set for hyperparameters optimization. We use our approach on two state-of-the-art CNNs:

- **ResNext-29-08-16:** we consider the standard ResNext Model (Xie et al. 2017) with a 29-layer architecture, a cardinality of 8, and a width of 16.
- **ResNet50:** we consider the standard ResNet model (He et al. 2016) with 50 layers.

We compare against the standard networks², as well as networks trained with the DeCov diversity strategy (Cogswell et al. 2016). All the models are trained using stochastic gradient descent (SGD) with a momentum of 0.9, weight decay of 0.0001, and a batch size of 128 for 200 epochs. The initial learning rate is set to 0.1 and is then decreased by a factor of 5 after 60, 120, and 160 epochs, respectively. We also adopt a standard data augmentation scheme that is widely used for these two datasets (He et al. 2016; Huang et al. 2017). For all models, the additional diversity term is applied on top the last intermediate layer. The penalty coefficients λ_1 and λ_2 , in (4), for our approach and the penalty coefficient of DeCov are chosen from $\{0.0001, 0.001, 0.01, 0.1\}$, and γ in the radial basis function is chosen from $\{1, 10\}$. For each approach, the model with the best validation performance is used in the test phase. We report the average performance over three random seeds.

Table 1 reports the average top-1 errors of the different approaches with the two basis networks. We note that, compared to the standard approach, employing a diversity strategy consistently boosts the results for all the two models and that our approach consistency outperforms both competing methods (standard and DeCov) in all the experiments. With ResNet50, the three variants of our proposed approach significantly reduce the test errors compared to the standard approach over both datasets: 0.51% – 0.63% improvement on CIFAR10 and 1.25% – 1.44% on CIFAR100.

For CIFAR10, the best performance is achieved by the direct variant and the Logdet variant for ResNext and ResNet models, respectively. For example, with ResNext, our direct variant yields 0.65 boost compared to the standard approach and 0.54 boost compared to DeCov. For CIFAR100, the best performance is achieved by our Logdet variant for both models. This variant leads to 1.4% and 0.85% boost for ResNet and ResNext, respectively. Overall, our three variants consistently outperform DeCov and standard approach in all testing configurations.

ImageNet

To further demonstrate the effectiveness of our approach and its ability to boost the performance of state-of-the-art neural networks, we conduct additional image classification experiments on the ImageNet-2012 classification dataset (Russakovsky et al. 2015) using four different models: ResNet50 (He et al. 2016), Wide-ResNet50 (Zagoruyko and Komodakis 2016), ResNeXt50 (Xie et al. 2017), and ResNet101 (He et al. 2016). The diversity term is applied on the last intermediate layer, i.e., the global average pooling layer for both DeCov and our method.

²For the standard approach, the only difference is not using an additional diversity loss. The remaining regularizers, data augmentation, weight decay etc., are all applied as specified per experiment.

For the hyperparameters, we fix $\lambda_1 = \lambda_2 = 0.001$ and $\gamma = 10$ for all the different approaches. The scope of this paper is feature diversity. However, in this experiment, we also report results with weight diversity approaches. In particular, we compare with the methods in (Yu, Li, and Zhou 2011), (Xie, Deng, and Xing 2015b), (Rodríguez et al. 2016), and (Ayinde, Inanc, and Zurada 2019).

We use the standard augmentation practice for this dataset as in (Zhang et al. 2018; Huang et al. 2017; Cogswell et al. 2016). All the models are trained with a batch size of 256 for 100 epoch using SGD with Nesterov Momentum of 0.9. The learning rate is initially set to 0.1 and decreases at epochs 30, 60, 90 by a factor of 10.

Table 2 reports the test errors of the different approaches on ImageNet dataset. As can be seen, feature diversity (our approach and DeCov) reduces the test error of the model and yields better performance compared to the standard approach. We note that, as opposed to feature diversity, weight diversity does not always yield performance improvement and it can sometimes hurt generalization. Compared to decov, our three variants consistently reach better performance.

For ResNet50 and ResNeXt50, the best performance is achieved by our direct variant, yielding more than 0.5% improvement compared to the standard approach for both models. For Wide-ResNet50 and ResNet101, our Det variant yields the top performance with over 0.6% boost for Wide-ResNet50. We note that our approach has a small additional time cost. For example for ResNet50, our direct, Det and Logdet variants take only 0.29%, 0.39%, and 0.49% extra training time, respectively.

Sensitivity Analysis

To further investigate the effect of the proposed diversity strategy, we conduct a sensitivity analysis using ImageNet on the hyperparameters of our methods: λ_1 and λ_2 which controls the contribution of the global diversity term to the global loss. We analyze the effect of the two parameters on the final performance of ResNet50 on ImageNet dataset. The analysis is presented in Figure 1.

As shown in Figure 1, using a diversity strategy, i.e., three variants of our method, consistently outperform the standard approach and are robust to the hyperparameters. For the Direct variant, the best performance is reached with $\lambda_1 = 0.005$ and $\lambda_2 = 0.001$. With this configuration, the model achieves 0.71% improvement compared to the standard approach. For the Det and the Logdet variants, using $\lambda_1 = 0.001$ and $\lambda_2 = 0.0005$, the model reaches the lowest error rate (23.09%) corresponding to 0.75% accuracy boost. Emphasizing diversity and using high weights (λ_1 and λ_2) still lead to better results compared to the standard approach but can make the total loss dominated by the diversity term. In general, we recommend using $\lambda_1 = \lambda_2 = 0.001$. However, this depends on the problem at hand.

Feature Diversity Reduces Overfitting

In (Laakom et al. 2021a; Cogswell et al. 2016), it has been observed that feature diversity can reduce overfitting. To study the effect of feature diversity on the generalization

method	ResNext-29-08-16		ResNet50	
	CIFAR10	CIFAR100	CIFAR10	CIFAR100
Standard	6.93 ± 0.10	26.73 ± 0.10	8.28 ± 0.41	33.39 ± 0.42
DeCov	6.82 ± 0.15	26.70 ± 0.10	8.03 ± 0.11	32.26 ± 0.22
Ours(Direct)	6.28 ± 0.11	26.20 ± 0.18	7.77 ± 0.09	32.09 ± 0.11
Ours(Det)	6.51 ± 0.16	26.35 ± 0.23	7.75 ± 0.12	32.14 ± 0.28
Ours(Logdet)	6.38 ± 0.08	25.88 ± 0.21	7.65 ± 0.10	31.99 ± 0.05

Table 1: Classification errors of the different approaches on CIFAR10 and CIFAR100 with three different models. Results are averaged over three random seeds.

	ResNet50	Wide-ResNet50	ResNeXt50	ResNet101
Standard	23.84	22.42	22.70	22.33
(Yu, Li, and Zhou 2011)	23.87	22.48	22.57	22.23
(Ayinde, Inanc, and Zurada 2019)	23.95	22.41	22.67	22.36
(Rodríguez et al. 2016)	24.23	22.70	22.80	23.10
(Xie, Deng, and Xing 2015b)	23.79	22.66	22.64	22.71
DeCov	23.62	22.68	22.57	22.31
Ours(Direct)	23.24	21.95	22.25	22.14
Ours(Det)	23.34	21.75	22.44	21.87
Ours(Logdet)	23.32	21.96	22.40	22.04

Table 2: Performance of different models with different diversity strategies on ImageNet dataset

	ERM	DeCov	direct*	det*	Ldet*
ResNet50	2.87	2.70	1.15	1.23	1.21
WResNet50	6.33	6.34	4.44	4.34	4.58
ResNeXt50	5.99	5.85	4.41	4.59	4.48
ResNet101	4.64	4.61	3.68	3.38	3.71

Table 3: Generalization Gap, i.e., training error - test error, of different models with different diversity strategies on ImageNet dataset. * denotes our approach. Ldet refers to our Logdet variant

gap, in Table 3, we report the final training errors and the generalization gap, i.e., training accuracy - test accuracy for the different feature diversity approaches on ImageNet dataset.

As shown in Table 3, we note that using diversity indeed can reduce overfitting and decreases the empirical generalization gap of neural networks. The three variants of our approach significantly reduce overfitting for all four models by more than 1% compared to standard and DeCov for all the models. For example, our Det variant reduces the empirical generalization gap, compared to the standard approach and DeCov, by 2% for Wide-ResNet model and over 1.2% for the ResNet101 model.

MLP-Based Models

Beyond CNN models, we also evaluate the performance of our diversity strategy on modern attention-free, multi-layer perceptron (MLP) based models for image classifica-

tion (Tolstikhin et al. 2021; Liu et al. 2021; Lee-Thorp et al. 2021). Such models are known to exhibit high overfitting and require regularization. We evaluate how diversity affects the accuracy of such models on CIFAR10. In particular, we conduct a simple experiment using two models: MLP-Mixer (Tolstikhin et al. 2021), gMLP (Liu et al. 2021) with four blocks each.

For the diversity strategies, i.e., ours and Decov, similar to our other experiments, the additional loss has been added on top of the last intermediate layer. The input images are resized to 72×72 . We use a patch size of 8×8 and an embedding dimension of 256. All models are trained for 100 epochs using Adam with a learning rate of 0.002, weight decay with a rate of 0.0001, and batch size of 256. Standard data augmentation, i.e., random horizontal flip and random zoom with a factor of 20%, is used. We use 10% of the training data for validation. We also reduce the learning rate by a factor of 2 if the validation loss does not improve for 5 epochs and use early stopping when the validation loss does not improve for 10 epochs. All experiments are repeated over 10 random seeds and the average results are reported.

The results in Table 4 show that employing a diversity strategy can indeed improve the performance of these models, thanks to its ability to help learn rich and robust representations of the input. Our proposed approach consistently outperforms the competing methods for both the MLP-Mixer and gMLP. For example, our direct variant leads to 1.15% and 0.3% boost for MLP-Mixer and gMLP, respectively.

For the MLP-mixer, the top performance is achieved by

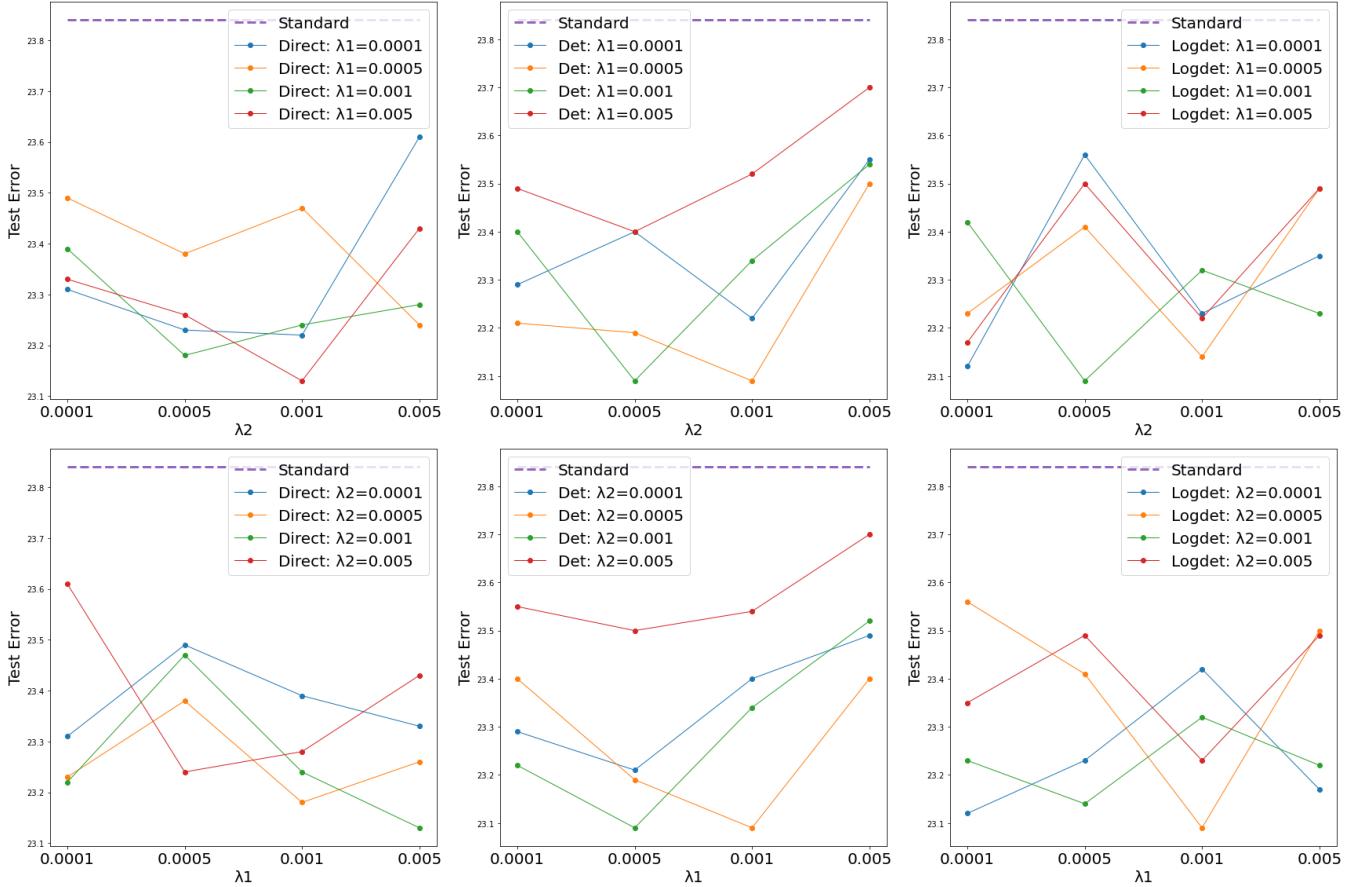


Figure 1: Sensitivity analysis of λ_1 and λ_2 on the test error using ResNet50 trained on ImageNet. The first row contains experiments with fixed λ_1 and the second row contains experiments with fixed λ_2 . From left to right: our Direct variant, our Det variant, and our Logdet variant. γ is fixed to 10 in all experiments.

	MLP-Mixer	gMLP
Standard	23.93	22.26
DeCov	24.10	22.00
Ours(Direct)	22.78	21.95
Ours(Det)	22.66	21.62
Ours(Logdet)	22.84	21.56

Table 4: Classification errors of modern MLP-based approaches on CIFAR10. Results are averaged over ten random seeds.

the Det variant of our approach reducing the error rates by 1.27% and 1.44% compared to the standard approach and DeCov, respectively. For the gMLP model, the top performance is achieved by the Logdet variant of our approach boosting the results by 0.7% and 0.44% compared to the standard approach and DeCov, respectively.

Learning in the Presence of Label Noise

To further demonstrate the usefulness of promoting diversity, we test the robustness of our approach in the presence of

label noise. In such situations, standard neural network tend to overfit to the noisy samples and not generalize well to the test set. Enforcing diversity can lead to better and richer representations attenuating the effect of noise. To show this, we performed additional experiments with label noise (20% and 40%) on CIFAR10 and CIFAR100 using ResNet50. We use the same training protocol used for the original CIFAR10 and CIFAR100: all models are trained using SGD with a momentum of 0.9, weight decay of 0.0001, and a batch size of 128 for 200 epochs. The initial learning rate is set to 0.1 and is then decreased by a factor of 5 after 60, 120, and 160 epochs, respectively. We also adopt a standard data augmentation scheme that is widely used for these two datasets (He et al. 2016; Huang et al. 2017). For all models, the additional diversity term is applied on top of the last intermediate layer. For the hyperparameters: The loss weights is chosen from $\{0.0001, 0.001, 0.01, 0.1\}$ for both our approach (λ_1 and λ_2) and Decov and γ in the radial basis function is chosen from $\{1, 10\}$. For each approach, the model with the best validation performance is used in the test phase. The average errors over three random seeds are reported.

The results are reported in Table 5. As it can be seen, in

Method	20% label noise		40% label noise	
	CIFAR10	CIFAR100	CIFAR10	CIFAR100
Standard	14.38 ± 0.29	45.11 ± 0.52	19.40 ± 0.80	48.81 ± 0.57
DeCov	13.75 ± 0.19	41.93 ± 0.40	17.60 ± 0.66	48.23 ± 0.48
Ours(Direct)	13.31 ± 0.40	40.10 ± 0.31	16.96 ± 0.32	46.73 ± 0.23
Ours(Det)	13.21 ± 0.21	40.35 ± 0.31	17.49 ± 0.04	46.93 ± 0.62
Ours(Logdet)	13.01 ± 0.40	39.97 ± 0.19	17.24 ± 0.31	46.52 ± 0.22

Table 5: Classification errors of ResNet50 using different diversity strategies on CIFAR10 and CIFAR100 datasets with different label noise ratios. Results are averaged over three random seeds.

the presence of noise, the gap between the standard approach and diversity (Decov and ours) increases. For example, our Logdet variant boosts the results by 1.91% and 2.29% on CIFAR10 and CIFAR100 with 40% noise, respectively.

Conclusions

In this paper, we proposed a new approach to encourage the ‘diversification’ of the layer-wise feature map outputs in neural networks. The main motivation is that by promoting within-layer activation diversity, units within the same layer learn to capture mutually distinct patterns. We proposed an additional loss term that can be added on top of any fully-connected layer. This term complements the traditional ‘between-layer’ feedback with an additional ‘within-layer’ feedback encouraging diversity of the activations. Extensive experimental results show that such a strategy can indeed improve the performance of different state-of-the-art networks across different datasets and different tasks, i.e., image classification, and label noise. We are confident that these results will spark further research in diversity-based approaches to improve the performance of neural networks.

Acknowledgments

This work was supported by NSF-Business Finland Center for Visual and Decision Informatics (CVDI) project AMALIA.

References

Advani, M. S.; Saxe, A. M.; and Sompolinsky, H. 2020. High-dimensional dynamics of generalization error in neural networks. *Neural Networks*, 132: 428–446.

Arora, S.; Cohen, N.; Hu, W.; and Luo, Y. 2019. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems*, 7413–7424.

Ayinde, B. O.; Inanc, T.; and Zurada, J. M. 2019. Regularizing deep neural networks by enhancing diversity in feature extraction. *IEEE transactions on neural networks and learning systems*, 30(9): 2650–2661.

Bao, Y.; Jiang, H.; Dai, L.; and Liu, C. 2013. Incoherent training of deep neural networks to de-correlate bottleneck features for speech recognition. In *International Conference on Acoustics, Speech and Signal Processing*, 6980–6984.

Belkin, M.; Hsu, D.; Ma, S.; and Mandal, S. 2019. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32): 15849–15854.

Bietti, A.; Mialon, G.; Chen, D.; and Mairal, J. 2019. A kernel perspective for regularizing deep neural networks. In *International Conference on Machine Learning*, 664–674.

Biyik, E.; Wang, K.; Anari, N.; and Sadigh, D. 2019. Batch active learning using determinantal point processes. *arXiv preprint arXiv:1906.07975*.

Cogswell, M.; Ahmed, F.; Girshick, R. B.; Zitnick, L.; and Batra, D. 2016. Reducing Overfitting in Deep Networks by Decorrelating Representations. In *International Conference on Learning Representations*.

Derezhinski, M.; Calandriello, D.; and Valko, M. 2019. Exact sampling of determinantal point processes with sublinear time preprocessing. In *Advances in Neural Information Processing Systems*, 11546–11558.

Dziugaite, G. K.; and Roy, D. M. 2017. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv preprint arXiv:1703.11008*.

Foret, P.; Kleiner, A.; Mobahi, H.; and Neyshabur, B. 2020. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*.

Gan, L.; Nurbakova, D.; Laporte, L.; and Calabretto, S. 2020. Enhancing Recommendation Diversity using Determinantal Point Processes on Knowledge Graphs. In *Conference on Research and Development in Information Retrieval*, 2001–2004.

Gartrell, M.; Brunel, V.-E.; Dohmatob, E.; and Krichene, S. 2019. Learning nonsymmetric determinantal point processes. In *Advances in Neural Information Processing Systems*, 6718–6728.

Golan, I.; and El-Yaniv, R. 2018. Deep anomaly detection using geometric transformations. In *Advances in Neural Information Processing Systems*, 9758–9769.

Goodfellow, I.; Bengio, Y.; Courville, A.; and Bengio, Y. 2016. *Deep learning*. MIT Press.

Grari, V.; Hajouji, O. E.; Lamprier, S.; and Detyniecki, M. 2021. Learning Unbiased Representations via Rényi Minimization. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 749–764. Springer.

Han, X.; and Guo, Y. 2021. Continual Learning with Dual Regularizations. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 619–634. Springer.

Haykin, S. 2010. *Neural networks and learning machines*, 3/E. Pearson Education India.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

He, Y.; Liu, P.; Wang, Z.; Hu, Z.; and Yang, Y. 2019. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4340–4349.

Hinton, G.; Deng, L.; Yu, D.; Dahl, G. E.; Mohamed, A. r.; Jaitly, N.; Senior, A.; Vanhoucke, V.; Nguyen, P.; Sainath, T. N.; et al. 2012a. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *Signal processing magazine*, 29(6): 82–97.

Hinton, G. E.; Srivastava, N.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. R. 2012b. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*.

Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4700–4708.

Kim, Y.; Kim, T.-H.; and Ergün, T. 2015. The instability of the Pearson correlation coefficient in the presence of coincidental outliers. *Finance Research Letters*, 13: 243–257.

Kondo, Y.; and Yamauchi, K. 2014. A dynamic pruning strategy for incremental learning on a budget. In *International Conference on Neural Information Processing*, 295–303. Springer.

Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. *Technical report, University of Toronto*.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 1097–1105.

Kukačka, J.; Golkov, V.; and Cremers, D. 2017. Regularization for deep learning: A taxonomy. *arXiv preprint arXiv:1710.10686*.

Kulesza, A.; and Taskar, B. 2010. Structured determinantal point processes. In *Advances in Neural Information Processing Systems*, 1171–1179.

Kulesza, A.; and Taskar, B. 2012. Determinantal point processes for machine learning. *arXiv preprint arXiv:1207.6083*.

Kwok, J. T.; and Adams, R. P. 2012. Priors for diversity in generative latent variable models. In *Advances in Neural Information Processing Systems*, 2996–3004.

Laakom, F.; Raitoharju, J.; Iosifidis, A.; and Gabbouj, M. 2021a. Learning distinct features helps, provably. *arXiv preprint arXiv:2106.06012*.

Laakom, F.; Raitoharju, J.; Iosifidis, A.; and Gabbouj, M. 2021b. On Feature Diversity in Energy-based models. In *Energy Based Models Workshop-ICLR 2021*.

Laakom, F.; Raitoharju, J.; Iosifidis, A.; and Gabbouj, M. 2022. Reducing Redundancy in the Bottleneck Representation of the Autoencoders. *arXiv preprint arXiv:2202.04629*.

Lee, H. B.; Nam, T.; Yang, E.; and Hwang, S. J. 2019. Meta Dropout: Learning to Perturb Latent Features for Generalization. In *International Conference on Learning Representations*.

Lee, S.; Heo, B.; Ha, J.-W.; and Song, B. C. 2020. Filter Pruning and Re-Initialization via Latent Space Clustering. *IEEE Access*, 8: 189587–189597.

Lee-Thorp, J.; Ainslie, J.; Eckstein, I.; and Ontanon, S. 2021. FNet: Mixing Tokens with Fourier Transforms. *arXiv preprint arXiv:2105.03824*.

Li, N.; Yu, Y.; and Zhou, Z.-H. 2012. Diversity regularized ensemble pruning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 330–345.

Li, Z.; Gong, B.; and Yang, T. 2016. Improved dropout for shallow and deep learning. In *Advances in Neural Information Processing Systems*, 2523–2531.

Liu, H.; Dai, Z.; So, D. R.; and Le, Q. V. 2021. Pay Attention to MLPs. *arXiv preprint arXiv:2105.08050*.

Malkin, J.; and Bilmes, J. 2008. Ratio semi-definite classifiers. In *International Conference on Acoustics, Speech and Signal Processing*, 4113–4116.

Malkin, J.; and Bilmes, J. 2009. Multi-layer ratio semi-definite classifiers. In *International Conference on Acoustics, Speech and Signal Processing*, 4465–4468.

Nagarajan, V.; and Kolter, J. Z. 2019. Uniform convergence may be unable to explain generalization in deep learning. In *Advances in Neural Information Processing Systems*, 11615–11626.

Nakkiran, P.; Kaplun, G.; Bansal, Y.; Yang, T.; Barak, B.; and Sutskever, I. 2020. Deep Double Descent: Where Bigger Models and More Data Hurt. In *International Conference on Learning Representations*.

Neyshabur, B.; Li, Z.; Bhojanapalli, S.; LeCun, Y.; and Srebro, N. 2018. The role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations*.

Ouali, Y.; Hudelot, C.; and Tami, M. 2021. Spatial contrastive learning for few-shot classification. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 671–686. Springer.

Poggio, T.; Kawaguchi, K.; Liao, Q.; Miranda, B.; Rosasco, L.; Boix, X.; Hidary, J.; and Mhaskar, H. 2017. Theory of deep learning III: explaining the non-overfitting puzzle. *arXiv preprint arXiv:1801.00173*.

Rodríguez, P.; Gonzalez, J.; Cucurull, G.; Gonfaus, J. M.; and Roca, X. 2016. Regularizing cnns with locally constrained decorrelations. *arXiv preprint arXiv:1611.01967*.

Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252.

Savas, C.; and Dovis, F. 2019. The impact of different kernel functions on the performance of scintillation detection based on support vector machines. *Sensors*, 19(23): 5219.

Singh, P.; Verma, V. K.; Rai, P.; and Namboodiri, V. 2020. Leveraging filter correlations for deep model compression. In *The IEEE Winter Conference on Applications of Computer Vision*, 835–844.

Tolstikhin, I.; Houlsby, N.; Kolesnikov, A.; Beyer, L.; Zhai, X.; Unterthiner, T.; Yung, J.; Keysers, D.; Uszkoreit, J.; Lucic, M.; et al. 2021. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*.

Wang, H.; Yang, W.; Zhao, Z.; Luo, T.; Wang, J.; and Tang, Y. 2019. Rademacher dropout: An adaptive dropout for deep neural network via optimizing generalization gap. *Neurocomputing*, 177–187.

Xie, B.; Liang, Y.; and Song, L. 2017. Diverse neural network learns true target functions. In *Artificial Intelligence and Statistics*, 1216–1224.

Xie, P.; Deng, Y.; and Xing, E. 2015a. Diversifying restricted boltzmann machine for document modeling. In *International Conference on Knowledge Discovery and Data Mining*, 1315–1324.

Xie, P.; Deng, Y.; and Xing, E. 2015b. On the generalization error bounds of neural networks under diversity-inducing mutual angular regularization. *arXiv preprint arXiv:1511.07110*.

Xie, P.; Singh, A.; and Xing, E. P. 2017. Uncorrelation and evenness: a new diversity-promoting regularizer. In *International Conference on Machine Learning*, 3811–3820.

Xie, P.; Zhu, J.; and Xing, E. 2016. Diversity-promoting bayesian learning of latent variable models. In *International Conference on Machine Learning*, 59–68.

Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; and He, K. 2017. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1492–1500.

Yang, K.; Gkatzelis, V.; and Stoyanovich, J. 2019. Balanced Ranking with Diversity Constraints. In *International Joint Conference on Artificial Intelligence*, 6035–6042.

Yu, Y.; Li, Y.-F.; and Zhou, Z.-H. 2011. Diversity regularized machine. In *International Joint Conference on Artificial Intelligence*.

Zagoruyko, S.; and Komodakis, N. 2016. Wide Residual Networks. In Richard C. Wilson, E. R. H.; and Smith, W. A. P., eds., *Proceedings of the British Machine Vision Conference (BMVC)*, 87.1–87.12. BMVA Press. ISBN 1-901725-59-6.

Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; and Deny, S. 2021. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. *arXiv preprint arXiv:2103.03230*.

Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond empirical risk minimization. *International Conference on Learning Representations*.

Zhao, J.; Mathieu, M.; and LeCun, Y. 2017. Energy-based generative adversarial network. *International Conference on Learning Representations*.