

# Learning distinct features helps, provably

Firas Laakom<sup>1</sup>[0000–0001–7436–5692] ✉, Jenni Raitoharju<sup>2,3</sup>[0000–0003–4631–9298],  
Alexandros Iosifidis<sup>4</sup>[0000–0003–4807–1345], and  
Moncef Gabbouj<sup>1</sup>[0000–0002–9788–2323]

<sup>1</sup> Faculty of Information Technology and Communication Sciences  
Tampere University, Finland.

<sup>2</sup> Faculty of Information Technology, University of Jyväskylä, Finland

<sup>3</sup> Programme for Environmental Information, Finnish Environment Institute,  
Jyväskylä, Finland

<sup>4</sup> Department of Electrical and Computer Engineering  
Aarhus University, Denmark

**Abstract.** We study the diversity of the features learned by a two-layer neural network trained with the least squares loss. We measure the diversity by the average  $L_2$ -distance between the hidden-layer features and theoretically investigate how learning non-redundant distinct features affects the performance of the network. To do so, we derive novel generalization bounds depending on feature diversity based on Rademacher complexity for such networks. Our analysis proves that more distinct features at the network’s units within the hidden layer lead to better generalization. We also show how to extend our results to deeper networks and different losses.

**Keywords:** Neural Networks · Generalization Theory · Feature Diversity

## 1 Introduction

Neural networks are a powerful class of non-linear function approximators that have been successfully used to tackle a wide range of problems. They have enabled breakthroughs in many tasks, such as image classification [31], speech recognition [20], and anomaly detection [16]. However, neural networks are often over-parameterized, i.e., have more parameters than the data they are trained on. As a result, they tend to overfit to the training samples and not generalize well on unseen examples [18]. Avoiding overfitting has been extensively studied [14, 15, 43, 45, 47] and various approaches and strategies have been proposed, such as data augmentation [18, 64], regularization [1, 8, 32], and Dropout [21, 38, 39], to close the gap between the empirical loss and the expected loss.

Formally, the output of a neural network consisting of  $P$  layers can be defined as follows:

$$f(\mathbf{x}; \mathbf{W}) = \rho^P(\mathbf{W}^P(\rho^{P-1}(\dots \rho^2(\mathbf{W}^2 \rho^1(\mathbf{W}^1 \mathbf{x}))))), \quad (1)$$

where  $\rho^i(\cdot)$  is the element-wise activation function, e.g., *ReLU* or *Sigmoid*, of the  $i^{th}$  layer and  $\mathbf{W} = \{\mathbf{W}^1, \dots, \mathbf{W}^P\}$  are the weights of the network with the

superscript denoting the layer. By defining  $\Phi(\cdot) = \rho^{P-1}(\dots \rho^2(\mathbf{W}^2 \rho^1(\mathbf{W}^1 \cdot)))$ , the output of neural network becomes

$$f(\mathbf{x}; \mathbf{W}) = \rho^P(\mathbf{W}^P \Phi(\mathbf{x})), \quad (2)$$

where  $\Phi(\mathbf{x}) = [\phi_1(\mathbf{x}), \dots, \phi_M(\mathbf{x})]$  is the  $M$ -dimensional feature representation of the input  $\mathbf{x}$ . This way neural networks can be interpreted as a two-stage process, with the first stage being representation learning, i.e., learning  $\Phi(\cdot)$ , followed by the final prediction layer. Both parts are jointly optimized.


Learning a rich and diverse set of features, i.e., the first stage, is critical for achieving top performance [3, 10, 34]. Studying the different properties of the learned features is an active field of research [11, 13, 29]. For example, [13] showed theoretically that learning a good feature representation can be helpful in few-shot learning. In this paper, we focus on the diversity of the features. This property has been empirically studied in [10, 36, 35] and has been shown to boost performance and reduce overfitting. However, no theoretical guarantees are provided. In this paper, we close this gap and we conduct a theoretical analysis of feature diversity. In particular, we propose to quantify the diversity of the feature set  $\{\phi_1(\cdot), \dots, \phi_M(\cdot)\}$  using the average pairwise  $L_2$ -distance between their outputs. Formally, given a dataset  $\{\mathbf{x}_i\}_{i=1}^N$ , we have

$$diversity = \frac{1}{N} \sum_{k=1}^N \frac{1}{2M(M-1)} \sum_{i \neq j}^M (\phi_i(\mathbf{x}_k) - \phi_j(\mathbf{x}_k))^2. \quad (3)$$

Intuitively, *diversity* measures how distinct the learned features are. If the mappings learned by two different units are redundant, then, given the same input, both units would yield similar output. This yields in low  $L_2$ -distance and as a result a low diversity. In contrast, if the mapping learned by each unit is distinct, the corresponding average distances to the outputs of the other units within the layer are high. Thus, this yields a high global diversity.

To confirm this intuition and further motivate the analysis of this attribute, we conduct empirical simulations. We track the diversity of the representation of the last hidden layer, as defined in equation 3, during the training of three different ResNet [19] models on CIFAR10 [30]. The results are reported in Figure 1. Indeed, diversity consistently increases during the training for all the models. This shows that, in order to solve the task at hand, neural networks learn distinct features.

**Our contributions:** In this paper, we theoretically investigate diversity in the neural network context and study how learning non-redundant features affects the performance of the model. We derive a bound for the generalization gap which is inversely proportional to the proposed diversity measure showing that learning distinct features helps. In our analysis, we focus on the simple neural network model with one-hidden layer trained with mean squared error. This configuration is simple, however, it has been shown to be convenient and insightful for the theoretical analysis [9, 12, 13]. Moreover, we show how to extend our theoretical analysis to different losses and different network architectures.



**Fig. 1.** Preliminary empirical results for additional motivation to theoretically understand feature diversity. The figure shows diversity versus the number of epochs for three different ResNet models trained on CIFAR10 dataset.

Our contributions can be summarized as follows:

- We analyze the effect the feature diversity on the generalization error bound of a neural network. The analysis is presented in Section 3. In Theorem 1, we derive an upper bound for the generalization gap which is inversely proportional to the diversity factor. Thus, we provide theoretical evidence that learning distinct features can help reduce the generalization error.
- We extend our analysis to different losses and general multi-layer networks. These results are presented in Theorems 2, 3, 4, 5, and 6.

**Outline of the paper:** The rest of the paper is organized as follows: Section 2 summarizes the preliminaries for our analysis. Section 3 presents our main theoretical results along with the proofs. Section 4 extends our results for different settings. Section 5 concludes the work with a discussion and several open problems.

## 2 PRELIMINARIES

Generalization theory [50, 28] focuses on the relation between the empirical loss defined as

$$\hat{L}(f) = \frac{1}{N} \sum_{i=1}^N l(f(\mathbf{x}_i; \mathbf{W}), y_i), \quad (4)$$

and the expected risk, for any  $f$  in the hypothesis class  $\mathcal{F}$ , defined as

$$L(f) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{Q}}[l(f(\mathbf{x}), y)], \quad (5)$$

where  $\mathcal{Q}$  is the underlying distribution of the dataset and  $y_i$  the corresponding label of  $x_i$ . Let  $f^* = \arg \min_{f \in \mathcal{F}} L(f)$  be the expected risk minimizer and  $\hat{f} = \arg \min_{f \in \mathcal{F}} \hat{L}(f)$  be the empirical risk minimizer. We are interested in the estimation error, i.e.,  $L(f^*) - L(\hat{f})$ , defined as the gap in the loss between both

minimizers [6]. The estimation error represents how well an algorithm can learn. It usually depends on the complexity of the hypothesis class and the number of training samples [5, 63].

Several techniques have been proposed to quantify the generalization error, such as Probably Approximately Correctly (PAC) learning [50, 53], VC dimension [52], and the Rademacher complexity [50]. The Rademacher complexity has been widely used as it usually leads to a tighter generalization error bound than the other metrics [17, 45, 51]. The formal definition of the empirical Rademacher complexity is given as follows:

**Definition 1.** [7, 50] For a given dataset with  $N$  samples  $\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^N$  generated by a distribution  $\mathcal{Q}$  and for a model space  $\mathcal{F} : \mathcal{X} \rightarrow \mathbb{R}$  with a single dimensional output, the empirical Rademacher complexity  $\mathcal{R}_N(\mathcal{F})$  of the set  $\mathcal{F}$  is defined as follows:

$$\mathcal{R}_N(\mathcal{F}) = \mathbb{E}_\sigma \left[ \sup_{f \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^N \sigma_i f(\mathbf{x}_i) \right], \quad (6)$$

where the variables  $\sigma = \{\sigma_1, \dots, \sigma_N\}$  are independent uniform random variables in  $\{-1, 1\}$ .

In this work, we rely on the Rademacher complexity to study diversity. We recall the following three lemmas related to the Rademacher complexity and the generalization error:

**Lemma 1.** [7] For  $\mathcal{F} \in \mathbb{R}^{\mathcal{X}}$ , assume that  $g : \mathbb{R} \rightarrow \mathbb{R}$  is a  $L_g$ -Lipschitz continuous function and  $\mathcal{A} = \{g \circ f : f \in \mathcal{F}\}$ . Then we have

$$\mathcal{R}_N(\mathcal{A}) \leq L_g \mathcal{R}_N(\mathcal{F}). \quad (7)$$

**Lemma 2.** [58] The Rademacher complexity  $\mathcal{R}_N(\mathcal{F})$  of the hypothesis class  $\mathcal{F} = \{f | f(\mathbf{x}) = \sum_{m=1}^M v_m \phi_m(\mathbf{x}) = \sum_{m=1}^M v_m \phi(\mathbf{w}_m^T \mathbf{x})\}$  can be upper-bounded as follows:

$$\mathcal{R}_N(\mathcal{F}) \leq \frac{2L_p C_{134} M}{\sqrt{N}} + \frac{C_4 |\phi(0)| M}{\sqrt{N}}, \quad (8)$$

where  $C_{134} = C_1 C_3 C_4$  and  $\phi(0)$  is the output of the activation function at the origin.

**Lemma 3.** [7] With a probability of at least  $1 - \delta$ ,

$$L(\hat{f}) - L(f^*) \leq 4\mathcal{R}_N(\mathcal{A}) + B \sqrt{\frac{2 \log(2/\delta)}{N}}, \quad (9)$$

where  $B \geq \sup_{\mathbf{x}, y, f} |l(f(\mathbf{x}), y)|$  and  $\mathcal{R}_N(\mathcal{A})$  is the Rademacher complexity of the loss set  $\mathcal{A}$ .

Lemma 3 upper-bounds the generalization error using the Rademacher complexity defined over the loss set and  $\sup_{x,y,f} |l(f(x), y)|$ . Our analysis aims at expressing this bound in terms of diversity, in order to understand how it affects the generalization.

In order to study the effect of diversity on the generalization, given a layer with  $M$  units  $\{\phi_1(\cdot), \dots, \phi_M(\cdot)\}$ , we make the following assumption:

**Assumption 1** *Given any input  $\mathbf{x}$ , we have*

$$\frac{1}{2M(M-1)} \sum_{i \neq j}^M (\phi_i(\mathbf{x}) - \phi_j(\mathbf{x}))^2 \geq d_{min}^2. \quad (10)$$

$d_{min}$  lower-bounds the average  $L_2$ -distance between the different units' activations within the same representation layer. Intuitively, if several neuron pairs  $i$  and  $j$  have similar outputs, the corresponding  $L_2$  distance is small. Thus, the lower bound  $d_{min}$  is also small and the units within this layer are considered redundant and "not diverse". Otherwise, if the average distance between the different pairs is large, their corresponding  $d_{min}$  is large and they are considered "diverse". By studying how the lower bound  $d_{min}$  affects the generalization of the model, we can analyze how the diversity theoretically affects the performance of neural networks. In the rest of the paper, we derive generalization bounds for neural networks using  $d_{min}$ .

### 3 Learning distinct features helps

In this section, we derive generalization bounds for neural networks depending on their diversity. Here, we consider a simple tow-layer neural network with a hidden layer composed of  $M$  neurons and one-dimensional output trained for a regression task. The full characterization of the setup can be summarized as follows:

- The activation function of the hidden layer,  $\rho(\cdot)$ , is a positive  $L_\rho$ -Lipschitz continuous function.
- The input vector  $\mathbf{x} \in \mathbb{R}^D$  satisfies  $\|\mathbf{x}\|_2 \leq C_1$  and the output scalar  $y \in \mathbb{R}$  satisfies  $|y| \leq C_2$ .
- The weight matrix  $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_M] \in \mathcal{R}^{D \times M}$  connecting the input to the hidden layer satisfies  $\|\mathbf{w}_m\|_2 \leq C_3$ .
- The weight vector  $\mathbf{v} \in \mathbb{R}^M$  connecting the hidden-layer to the output satisfies  $\|\mathbf{v}\|_\infty \leq C_4$ .
- The hypothesis class is  $\mathcal{F} = \left\{ f | f(\mathbf{x}) = \sum_{m=1}^M v_m \phi_m(\mathbf{x}) = \sum_{m=1}^M v_m \rho(\mathbf{w}_m^T \mathbf{x}) \right\}$ .
- Loss function set is  $\mathcal{A} = \left\{ l | l(f(\mathbf{x}), y) = \frac{1}{2} |f(\mathbf{x}) - y|^2 \right\}$ .
- Given an input  $\mathbf{x}$ ,  $\frac{1}{2M(M-1)} \sum_{n \neq m}^M (\phi_n(\mathbf{x}) - \phi_m(\mathbf{x}))^2 \geq d_{min}^2$ .

Our main goal is to analyze the generalization error bound of the neural network and to see how its upper-bound is linked to the diversity of the different

units, expressed by  $d_{min}$ . The main result of the paper is presented in Theorem 1. Our proof consists of three steps: At first, we derive a novel bound for the hypothesis class  $\mathcal{F}$  depending on  $d_{min}$ . Then, we use this bound to derive bounds for the loss class  $\mathcal{A}$  and its Rademacher complexity  $\mathcal{R}_N(\mathcal{A})$ . Finally, we plug all the derived bounds in Lemma 3 to complete the proof of Theorem 1.

The first step of our analysis is presented in Lemma 4:

**Lemma 4.** *We have*

$$\sup_{\mathbf{x}, f \in \mathcal{F}} |f(\mathbf{x})| \leq \sqrt{\mathcal{J}}, \quad (11)$$

where  $\mathcal{J} = C_4^2(MC_5^2 + M(M-1)(C_5^2 - d_{min}^2))$  and  $C_5 = L_\rho C_1 C_3 + \phi(0)$ ,

*Proof.*

$$\begin{aligned} f^2(\mathbf{x}) &= \left( \sum_{m=1}^M v_m \phi_m(\mathbf{x}) \right)^2 \leq \left( \sum_{m=1}^M \|v\|_\infty \phi_m(\mathbf{x}) \right)^2 = \|v\|_\infty^2 \left( \sum_{m=1}^M \phi_m(\mathbf{x}) \right)^2 \\ &\leq C_4^2 \left( \sum_{m=1}^M \phi_m(\mathbf{x}) \right)^2 = C_4^2 \left( \sum_{m,n} \phi_m(\mathbf{x}) \phi_n(\mathbf{x}) \right) \\ &= C_4^2 \left( \sum_m \phi_m(\mathbf{x})^2 + \sum_{m \neq n} \phi_m(\mathbf{x}) \phi_n(\mathbf{x}) \right). \end{aligned} \quad (12)$$

We have  $\sup_{w, \mathbf{x}} \phi_m(\mathbf{x}) = \sup_{w, \mathbf{x}} \rho(w^T \mathbf{x}) \leq \sup(L_\rho |w^T \mathbf{x}| + \phi(0))$ , because  $\rho$  is  $L_\rho$ -Lipschitz. Thus,  $\|\phi\|_\infty \leq L_\rho C_1 C_3 + \phi(0) = C_5$ . For the first term in equation 12, we have  $\sum_m \phi_m(\mathbf{x})^2 < M(L_\rho C_1 C_3 + \phi(0))^2 = MC_5^2$ . The second term, using the identity  $\phi_m(\mathbf{x}) \phi_n(\mathbf{x}) = \frac{1}{2} (\phi_m(\mathbf{x})^2 + \phi_n(\mathbf{x})^2 - (\phi_m(\mathbf{x}) - \phi_n(\mathbf{x}))^2)$ , can be rewritten as

$$\sum_{m \neq n} \phi_m(\mathbf{x}) \phi_n(\mathbf{x}) = \frac{1}{2} \left( \sum_{m \neq n} \phi_m(\mathbf{x})^2 + \phi_n(\mathbf{x})^2 - (\phi_m(\mathbf{x}) - \phi_n(\mathbf{x}))^2 \right). \quad (13)$$

In addition, we have  $\frac{1}{2} \sum_{m \neq n} (\phi_m(\mathbf{x}) - \phi_n(\mathbf{x}))^2 \geq M(M-1)d_{min}^2$ . Thus, we have:

$$\sum_{m \neq n} \phi_m(\mathbf{x}) \phi_n(\mathbf{x}) \leq \frac{1}{2} \sum_{m \neq n} (2C_5^2) - M(M-1)d_{min}^2 = M(M-1)(C_5^2 - d_{min}^2). \quad (14)$$

By putting everything back to equation 12, we have:

$$f^2(\mathbf{x}) \leq C_4^2 \left( MC_5^2 + M(M-1)(C_5^2 - d_{min}^2) \right) = \mathcal{J}. \quad (15)$$

Thus,  $\sup_{\mathbf{x}, f} |f(\mathbf{x})| \leq \sqrt{\sup_{\mathbf{x}, f} f(\mathbf{x})^2} \leq \sqrt{\mathcal{J}}$ .

Note that in Lemma 4, we have expressed the upper-bound of  $\sup_{\mathbf{x}, f} |f(\mathbf{x})|$  in terms of  $d_{min}$ . Using this bound, we can now find an upper-bound for  $\sup_{\mathbf{x}, f, y} |l(f(\mathbf{x}), y)|$  in the following lemma:

**Lemma 5.** *We have*

$$\sup_{\mathbf{x}, y, f} |l(f(\mathbf{x}), y)| \leq \frac{1}{2}(\sqrt{\mathcal{J}} + C_2)^2. \quad (16)$$

*Proof.* We have  $\sup_{\mathbf{x}, y, f} |f(\mathbf{x}) - y| \leq \sup_{\mathbf{x}, y, f} (|f(\mathbf{x})| + |y|) = \sqrt{\mathcal{J}} + C_2$ . Thus,  $\sup_{\mathbf{x}, y, f} |l(f(\mathbf{x}), y)| \leq \frac{1}{2}(\sqrt{\mathcal{J}} + C_2)^2$ .

Next, using the result of lemmas 1, 2, and 5, we can derive a bound for the Rademacher complexity of  $\mathcal{A}$ . We have, thus, expressed all the elements of Lemma 3 using the diversity term  $d_{min}$ . By plugging in the derived bounds in Lemmas 4, 5, we obtain Theorem 1.

**Theorem 1.** *With probability at least  $(1 - \delta)$ , we have*

$$L(\hat{f}) - L(f^*) \leq (\sqrt{\mathcal{J}} + C_2) \frac{A}{\sqrt{N}} + \frac{1}{2}(\sqrt{\mathcal{J}} + C_2)^2 \sqrt{\frac{2 \log(2/\delta)}{N}}, \quad (17)$$

where  $C_{134} = C_1 C_3 C_4$ ,  $\mathcal{J} = C_4^2 (M C_5^2 + M(M-1)(C_5^2 - d_{min}^2))$ ,  $A = 4(2L_\rho C_{134} + C_4 |\phi(0)|)M$ , and  $C_5 = L_\rho C_1 C_3 + \phi(0)$ .

*Proof.* Given that  $l(\cdot)$  is  $K$ -Lipschitz with a constant  $K = \sup_{\mathbf{x}, y, f} |f(\mathbf{x}) - y| \leq \sqrt{\mathcal{J}} + C_2$ , and using Lemma 1, we can show that  $\mathcal{R}_N(\mathcal{A}) \leq K \mathcal{R}_N(\mathcal{F}) \leq (\sqrt{\mathcal{J}} + C_2) \mathcal{R}_N(\mathcal{F})$ . For  $\mathcal{R}_N(\mathcal{F})$ , we use the bound found in Lemma 2. Using Lemmas 3 and 5, we have

$$L(\hat{f}) - L(f^*) \leq 4(\sqrt{\mathcal{J}} + C_2) \left( 2L_\rho C_{134} + C_4 |\phi(0)| \right) \frac{M}{\sqrt{N}} + \frac{1}{2}(\sqrt{\mathcal{J}} + C_2)^2 \sqrt{\frac{2 \log(2/\delta)}{N}}, \quad (18)$$

where  $C_{134} = C_1 C_3 C_4$ ,  $\mathcal{J} = C_4^2 (M C_5^2 + M(M-1)(C_5^2 - d_{min}^2))$ , and  $C_5 = L_\rho C_1 C_3 + \phi(0)$ . Thus, setting  $A = 4(2L_\rho C_{134} + C_4 |\phi(0)|)M$  completes the proof.

Theorem 1 provides an upper-bound for the generalization gap. We note that it is a decreasing function of  $d_{min}$ . Thus, this suggests that higher  $d_{min}$ , i.e., more diverse activations, yields a lower generalization error bound. This shows that learning distinct features helps in neural network context.

We note that the bound in Theorem 1 is non-vacuous in the sense that it converges to zero when the number of training samples  $N$  goes to infinity. Moreover, we note that in this paper we do not claim to reach a tighter generalization bound for neural networks in general [14, 24, 44, 48]. Our main claim is that we derive a generalization bound which depends on the diversity of learned features, as measured by  $d_{min}$ . To the best of our knowledge, this is the first work that performs such theoretical analysis based on the average  $L_2$ -distance between the units within the hidden layer.

### Connection to prior studies

Theoretical analysis of the properties of the features learned by neural network models is an active field of research. Feature representation has been theoretically studied in the context of few-shot learning in [13], where the advantage of learning a good representation in the case of scarce data was demonstrated. [2] showed the same in the context of imitation learning, demonstrating that it has sample complexity benefits for imitation learning. [55] developed similar findings for the self-supervised learning task. [42] derived novel bounds showing the statistical benefits of multitask representation learning in linear Markov Decision Processes. Opposite to the aforementioned works, the main focus of this paper is not on the large sample complexity problems. Instead, we focused on feature diversity in the learned representation and showed that learning distinct features leads to better generalization.

Another line of research related to our work is weight-diversity in neural networks [4, 33, 57, 58, 61]. Diversity in this context is defined based on dissimilarity between the weight component using, e.g., cosine distance and weight matrix covariance [59]. In [58], theoretical benefits of weight-diversity have been demonstrated. We note that, in our work, diversity is defined in a fundamentally different way. We do not consider dissimilarity between the parameters of the neural network. Our main scope is the feature representation and, to this end, diversity is defined based on the  $L_2$  distance between the feature maps directly and not the weights. Empirical analysis of the deep representation of neural networks has drawn attention lately [10, 11, 29, 36]. For example, [10, 36] showed empirically that learning decorrelated features reduces overfitting. However, theoretical understanding of the phenomena is lacking. Here, we close this gap by studying how feature diversity affects generalization.

## 4 Extensions

In this section, we show how to extend our theoretical analysis for classification, for general multi-layer networks, and for different losses.

### 4.1 Binary classification

Here, we extend our analysis of the effect of learning a diverse feature representation on the generalization error to the case of a binary classification task, i.e.,  $y \in \{-1, 1\}$ . Here, we consider the special cases of a hinge loss and a logistic loss. To derive diversity-dependent generalization bounds for these cases, similar to the proofs of Lemmas 7 and 8 in [58], we can show the following two lemmas:

**Lemma 6.** *Using the hinge loss, we have with probability at least  $(1 - \delta)$*

$$L(\hat{f}) - L(f^*) \leq 4 \left( 2L_\rho C_{134} + C_4 |\phi(0)| \right) \frac{M}{\sqrt{N}} + (1 + \sqrt{\mathcal{J}}) \sqrt{\frac{2 \log(2/\delta)}{N}}, \quad (19)$$



where  $C_{134} = C_1 C_3 C_4$ ,  $\mathcal{J} = C_4^2(MC_5^2 + M(M-1)(C_5^2 - d_{min}^2))$ , and  $C_5 = L_\rho C_1 C_3 + \phi(0)$ .

**Lemma 7.** *Using the logistic loss  $l(f(x), y) = \log(1 + e^{-yf(x)})$ , we have with probability at least  $(1 - \delta)$*

$$L(\hat{f}) - L(f^*) \leq \frac{4}{1 + e^{\sqrt{-\mathcal{J}}}} \left( 2L_\rho C_{134} + C_4 |\phi(0)| \right) \frac{M}{\sqrt{N}} + \log(1 + e^{\sqrt{\mathcal{J}}}) \sqrt{\frac{2 \log(2/\delta)}{N}}, \quad (20)$$

where  $C_{134} = C_1 C_3 C_4$ ,  $\mathcal{J} = C_4^2(MC_5^2 + M(M-1)(C_5^2 - d_{min}^2))$ , and  $C_5 = L_\rho C_1 C_3 + \phi(0)$ .

Using the above lemmas, we can now derive a diversity-dependant bound for the binary classification case. The extensions of Theorem 1 in the cases of a hinge loss and a logistic loss are presented in Theorems 2 and 3, respectively.

**Theorem 2.** *Using the hinge loss, with probability at least  $(1 - \delta)$ , we have*

$$L(\hat{f}) - L(f^*) \leq A/\sqrt{N} + (1 + \sqrt{\mathcal{J}}) \sqrt{\frac{2 \log(2/\delta)}{N}}, \quad (21)$$

where  $\mathcal{J} = C_4^2(MC_5^2 + M(M-1)(C_5^2 - d_{min}^2))$ ,  $A = 4(2L_\rho C_{134} + C_4 |\phi(0)|)M$ , and  $C_5 = L_\rho C_1 C_3 + \phi(0)$ .

**Theorem 3.** *Using the logistic loss  $l(f(x), y) = \log(1 + e^{-yf(x)})$ , with probability at least  $(1 - \delta)$ , we have*

$$L(\hat{f}) - L(f^*) \leq \frac{A}{(1 + e^{\sqrt{-\mathcal{J}}})\sqrt{N}} + \log(1 + e^{\sqrt{\mathcal{J}}}) \sqrt{\frac{2 \log(2/\delta)}{N}}, \quad (22)$$

where  $\mathcal{J} = C_4^2(MC_5^2 + M(M-1)(C_5^2 - d_{min}^2))$ ,  $A = 4(2L_\rho C_{134} + C_4 |\phi(0)|)M$ , and  $C_5 = L_\rho C_1 C_3 + \phi(0)$ .

As we can see, also for the binary classification task, the generalization bounds for the hinge and logistic losses are decreasing with respect to  $d_{min}$ . Thus, this shows that learning distinct features helps and can improve the generalization also in binary classification.

## 4.2 Multi-layer networks

Here, we extend our result for networks with  $P (> 1)$  hidden layers. We assume that the pair-wise distances between the activations within layer  $p$  are lower-bounded by  $d_{min}^{(p)}$ . In this case, the hypothesis class can be defined recursively. In addition, we assume that:  $\|\mathbf{W}^{(p)}\|_\infty \leq C_3^{(p)}$  for every  $\mathbf{W}^{(p)}$ , i.e., the weight matrix of the  $p$ -th layer. In this case, the main theorem is extended as follows:

**Theorem 4.** *With probability of at least  $(1 - \delta)$ , we have*

$$L(\hat{f}) - L(f^*) \leq (\sqrt{\mathcal{J}^P} + C_2) \frac{A}{\sqrt{N}} + \frac{1}{2} \left( \sqrt{\mathcal{J}^P} + C_2 \right)^2 \sqrt{\frac{2 \log(2/\delta)}{N}}, \quad (23)$$

where  $A = 4((2L_\rho)^P C_1 C_3^0 \prod_{p=0}^{P-1} \sqrt{M^{(p)}} C_3^{(p)} + |\phi(0)| \sum_{p=0}^{P-1} (2L_\rho)^{P-1-p} \prod_{j=p}^{P-1} \sqrt{M^j} C_3^j)$ , and  $\mathcal{J}^P$  is defined recursively using the following identities:  $\mathcal{J}^0 = C_3^0 C_1$  and  $\mathcal{J}^{(p)} = M^{(p)} C^{p2} (M^{p2} (L_\rho \mathcal{J}^{p-1} + \phi(0))^2 - M(M-1) d_{min}^{(p)2})$ , for  $p = 1, \dots, P$ .

*Proof.* Lemma 5 in [58] provides an upper-bound for the hypothesis class. We denote by  $\mathbf{v}^{(p)}$  the outputs of the  $p^{th}$  hidden layer before applying the activation function:

$$\mathbf{v}^0 = [\mathbf{w}_1^{0T} \mathbf{x}, \dots, \mathbf{w}_{M^0}^{0T} \mathbf{x}], \quad (24)$$

$$\mathbf{v}^{(p)} = \left[ \sum_{j=1}^{M^{p-1}} w_{j,1}^{(p)} \phi(\mathbf{v}_j^{p-1}), \dots, \sum_{j=1}^{M^{p-1}} w_{j,M^{(p)}}^{(p)} \phi(\mathbf{v}_j^{p-1}) \right], \quad (25)$$

$$\mathbf{v}^{(p)} = \left[ \mathbf{w}_1^{(p)T} \boldsymbol{\phi}^{(p)}, \dots, \mathbf{w}_{M^{(p)}}^{(p)T} \boldsymbol{\phi}^{(p)} \right], \quad (26)$$

where  $\boldsymbol{\phi}^{(p)} = [\phi(v_1^{p-1}), \dots, \phi(v_{M^{p-1}}^{p-1})]$ . We have  $\|\mathbf{v}^{(p)}\|_2^2 = \sum_{m=1}^{M^{(p)}} (\mathbf{w}_m^{(p)T} \boldsymbol{\phi}^{(p)})^2$  and  $\mathbf{w}_m^{(p)T} \boldsymbol{\phi}^{(p)} \leq C_3^{(p)} \sum_n \phi_n^{(p)}$ . Thus,

$$\|\mathbf{v}^{(p)}\|_2^2 \leq \sum_{m=1}^{M^{(p)}} \left( C_3^{(p)} \sum_n \phi_n^{(p)} \right)^2 = M^{(p)} C_3^{p2} \left( \sum_n \phi_n^{(p)} \right)^2 = M^{(p)} C_3^{p2} \sum_{mn} \phi_m^{(p)} \phi_n^{(p)}. \quad (27)$$

We use the same decomposition trick of  $\phi_m^{(p)} \phi_n^{(p)}$  as in the proof of Lemma 2. We need to bound  $\sup_x \phi^{(p)}$ :

$$\sup_x \phi^{(p)} < \sup_x (L_\rho |\mathbf{v}^{p-1}| + \phi(0)) < L_\rho \|\mathbf{v}^{p-1}\|_2^2 + \phi(0). \quad (28)$$

Thus, we have

$$\|\mathbf{v}^{(p)}\|_2^2 \leq M^{(p)} C_3^{p2} (M^2 (L_\rho \|\mathbf{v}^{p-1}\|_2^2 + \phi(0))^2 - M(M-1) d_{min}^2) = \mathcal{J}^P. \quad (29)$$

We found a recursive bound for  $\|\mathbf{v}^{(p)}\|_2^2$  and we note that for  $p = 0$  we have  $\|\mathbf{v}^0\|_2^2 \leq \|W^0\|_\infty C_1 \leq C_3^0 C_1 = \mathcal{J}^0$ . Thus,

$$\sup_{\mathbf{x}, f^P \in \mathcal{F}^P} |f(\mathbf{x})| = \sup_{\mathbf{x}, f^P \in \mathcal{F}^P} |\mathbf{v}^P| \leq \sqrt{\mathcal{J}^P}. \quad (30)$$

By replacing the variables in Lemma 3, we have

$$\begin{aligned} L(\hat{f}) - L(f^*) &\leq 4(\sqrt{\mathcal{J}^P} + C_2) \left( \frac{(2L_\rho)^P C_1 C_3^0}{\sqrt{N}} \prod_{p=0}^{P-1} \sqrt{M^{(p)}} C_3^{(p)} \right. \\ &\quad \left. + \frac{|\phi(0)|}{\sqrt{N}} \sum_{p=0}^{P-1} (2L_\rho)^{P-1-p} \prod_{j=p}^{P-1} \sqrt{M^j} C_3^j \right) + \frac{1}{2} \left( \sqrt{\mathcal{J}^P} + C_2 \right)^2 \sqrt{\frac{2 \log(2/\delta)}{N}}, \end{aligned}$$

Taking  $A = 4((2L_\rho)^P C_1 C_3^0 \prod_{p=0}^{P-1} \sqrt{M^{(p)}} C_3^{(p)} + |\phi(0)| \sum_{p=0}^{P-1} (2L_\rho)^{P-1-p} \prod_{j=p}^{P-1} \sqrt{M^j} C_3^j)$  completes the proof.

In Theorem 4, we see that  $\mathcal{J}^P$  is decreasing with respect to  $d_{min}^{(p)}$ . This extends our results to the multi-layer neural network case.

### 4.3 Multiple outputs

Finally, we consider the case of a neural network with a multi-dimensional output, i.e.,  $\mathbf{y} \in R^D$ . In this case, we can extend Theorem 1 with the following two theorems:

**Theorem 5.** *For a multivariate regression trained with the squared error, there exists a constant  $A$  such that, with probability at least  $(1 - \delta)$ , we have*

$$L(\hat{f}) - L(f^*) \leq (\sqrt{\mathcal{J}} + C_2) \frac{A}{\sqrt{N}} + \frac{D}{2} (\sqrt{\mathcal{J}} + C_2)^2 \sqrt{\frac{2 \log(2/\delta)}{N}} \quad (31)$$

where  $\mathcal{J} = C_4^2(MC_5^2 + M(M-1)(C_5^2 - d_{min}^2))$ ,  $C_5 = L_\rho C_1 C_3 + \phi(0)$ , and  $A = 4D(2L_\rho C_{134} + C_4|\phi(0)|)M$ .

*Proof.* The squared loss  $\frac{1}{2} \|f(\mathbf{x}) - \mathbf{y}\|_2^2$  can be decomposed into  $D$  terms  $\frac{1}{2} (f(\mathbf{x})_k - y_k)^2$ . Using Theorem 1, we can derive the bound for each term and, thus, we have:

$$L(\hat{f}) - L(f^*) \leq 4D(\sqrt{\mathcal{J}} + C_2) \left( 2L_\rho C_{134} + C_4|\phi(0)| \right) \frac{M}{\sqrt{N}} + \frac{D}{2} (\sqrt{\mathcal{J}} + C_2)^2 \sqrt{\frac{2 \log(2/\delta)}{N}}, \quad (32)$$

where  $C_{134} = C_1 C_3 C_4$ ,  $\mathcal{J} = C_4^2(MC_5^2 + M(M-1)(C_5^2 - d_{min}^2))$ , and  $C_5 = L_\rho C_1 C_3 + \phi(0)$ . Taking  $A = 4D(2L_\rho C_{134} + C_4|\phi(0)|)M$  completes the proof.

**Theorem 6.** *For a multi-class classification task using the cross-entropy loss, there exists a constant  $A$  such that, with probability at least  $(1 - \delta)$ , we have*


$$L(\hat{f}) - L(f^*) \leq \frac{A}{(D-1 + e^{-2\sqrt{\mathcal{J}}})\sqrt{N}} + \log \left( 1 + (D-1)e^{2\sqrt{\mathcal{J}}} \right) \sqrt{\frac{2 \log(2/\delta)}{N}} \quad (33)$$

where  $\mathcal{J} = C_4^2(MC_5^2 + M(M-1)(C_5^2 - d_{min}^2))$  and  $C_5 = L_\rho C_1 C_3 + \phi(0)$ , and  $A = 4D(D-1)(2L_\rho C_{134} + C_4|\phi(0)|)M$ .

*Proof.* Using Lemma 9 in [58], we have  $\sup_{f, \mathbf{x}, \mathbf{y}} l = \log(1 + (D-1)e^{2\sqrt{\mathcal{J}}})$  and  $l$  is  $\frac{D-1}{D-1+e^{-2\sqrt{\mathcal{J}}}}$ -Lipschitz. Thus, using the decomposition property of the Rademacher complexity, we have

$$\mathcal{R}_n(\mathcal{A}) \leq \frac{4D(D-1)}{D-1 + e^{-2\sqrt{\mathcal{J}}}} (2L_\rho C_{134} + C_4|\phi(0)|) \frac{M}{\sqrt{N}}. \quad (34)$$

Taking  $A = 4D(D-1)(2L_\rho C_{134} + C_4|\phi(0)|)M$  completes the proof.




**Fig. 2.** Generalization gap, i.e., train error - test error, and the theoretical bound, i.e.,  $(C_5^2 - d_{min}^2)/\sqrt{N}$ , as a function of the number of training samples on MNIST dataset for neural networks with intermediate layer sizes from left to right: 128 (correlation=0.9948), 256 (correlation=0.9939), and 512 (correlation=0.9953). The theoretical term has been scaled in the same range as the generalization gap. All results are averaged over 5 random seeds.

Theorems 5 and 6 extend our result for the multi-dimensional regression and classification tasks, respectively. Both bounds are inversely proportional to the diversity factor  $d_{min}$ . We note that for the classification task the upper-bound is exponentially decreasing with respect to  $d_{min}$ . This shows that learning a diverse and rich feature representation yields a tighter generalization gap and, thus, theoretically guarantees a stronger generalization performance.

## 5 Discussion and open problems

In this paper, we showed how the diversity of the features learned by a two-layer neural network trained with the least-squares loss affects generalization. We quantified the diversity by the average  $L_2$ -distance between the hidden-layer features and we derived novel diversity-dependant generalization bounds based on Rademacher complexity for such models. The derived bounds are inversely-proportional to the diversity term, thus demonstrating that more distinct features within the hidden layer can lead to better generalization. We also showed how to extend our results to deeper networks and different losses.

The bound found in Theorem 1 suggests that the generalization gap, with respect to diversity, is inversely proportional to  $d_{min}$  and scales as  $\sim (C_5^2 - d_{min}^2)/\sqrt{N}$ . We validate this finding empirically in Figure 2. We train a two-layer neural network on the MNIST dataset for 100 epochs using SGD with a learning rate of 0.1 and batch size of 256. We show the generalization gap, i.e., test error - train error, and the theoretical bound, i.e.,  $(C_5^2 - d_{min}^2)/\sqrt{N}$ , for different training set sizes.  $d_{min}$  is the lower bound of diversity. Empirically, it can be estimated as the minimum feature diversity over the training data  $S$ :  $d_{min} = \min_{x \in S} \frac{1}{2M(M-1)} \sum_{n \neq m}^M (\phi_n(x) - \phi_m(x))^2$ . We experiment with different sizes of the hidden layer, namely 128, 256, and 512. The average results using 5 random seeds are reported for different training sizes in Figure 2 showing that the theoretical bound correlates consistently well (correlation  $> 0.9939$ ) with the generalization error.



**Fig. 3.** From left to right: (a)-(b) Tracking the diversity during the training for different models on ImageNet. (c) Final diversity as a function of depth for different models on MNIST.

As shown in Figure 1, diversity increases for neural networks along the training phase. To further investigate this observation, we conduct additional experiments on ImageNet [49] dataset using 4 different state-of-the-art models: **ResNet50** and **ResNet101**, i.e., the standard ResNet model [19] with 50 layers and 101 layers, **ResNext50** [60], and **WideResNet50** [62] with 50 layers. All models are trained with SGD using standard training protocol [10, 22, 64]. We track the diversity, as defined in equation 3, of the features of the last intermediate layer. The results are shown in Figure 3 (a) and (b). As it can be seen, SGD without any explicit regularization implicitly optimizes diversity and converges toward regions with high features’ distinctness. These observations suggest the following conjecture:

*Conjecture 1.* Standard training with SGD implicitly optimizes the diversity of intermediate features.

Studying the fundamental properties of SGD is extremely important to understand generalization in deep learning [23, 25, 27, 54, 65]. Conjecture 1 suggests a new implicit bias for SGD, showing that it favors regions with high feature diversity.

Another research question related to diversity that is worth investigating is: *How does the network depth affect diversity?* In order to answer this question, we conduct an empirical experiment using MNIST dataset [37]. We use fully connected networks (FCNs) with ReLU activation and different depths (1 to 12). We experiment with three models with different widths, namely FCN-256, FCN-512, and FCN-1024, with 256, 512, and 1024 units per layer, respectively. We measure the final diversity of the last hidden layer for the different depths. The average results using 5 random seeds are reported in Figure 3 (c). Interestingly, in this experiment, increasing the depth consistently leads to learning more distinct features and higher diversity for the different models. However, by looking at Figure 1, we can see that having more parameters does not always lead to higher diversity. This suggests the following open question:

**Open Problem 1** *When does having more parameters/depth lead to higher diversity?*

Understanding the difference between shallow and deep models and why deeper models generalize better is one of the puzzles of deep learning [26, 40, 47]. The insights gained by studying Open Problem 1 can lead to a novel key advantage of depth: deeper models are able to learn a richer and more diverse set of features.

Another interesting line of research is adversarial robustness [40, 41, 46, 56]. Intuitively, learning distinct features can lead to a richer representation and, thus, more robust networks. However, the theoretical link is missing. This leads to the following open problem:

**Open Problem 2** *Can the theoretical tools proposed in this paper be used to prove the benefits of feature diversity for adversarial robustness?*

## Ethical consideration

This is a theoretical work and does not present any foreseeable societal consequences. The data used in this work comes from publicly accessible channels.

**Acknowledgements** This work has been supported by the NSF-Business Finland Center for Visual and Decision Informatics (CVDI) project AMALIA. The work of Jenni Raitoharju was funded by the Academy of Finland (project 324475). Alexandros Iosifidis acknowledges funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 957337.

## References

1. Arora, S., Cohen, N., Hu, W., Luo, Y.: Implicit regularization in deep matrix factorization. In: Advances in Neural Information Processing Systems. pp. 7413–7424 (2019)
2. Arora, S., Du, S., Kakade, S., Luo, Y., Saunshi, N.: Provable representation learning for imitation learning via bi-level optimization. In: International Conference on Machine Learning. PMLR (2020)
3. Arpit, D., Jastrzebski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M.S., Maharaj, T., Fischer, A., Courville, A., Bengio, Y., et al.: A closer look at memorization in deep networks. In: International Conference on Machine Learning. pp. 233–242. PMLR (2017)
4. Bao, Y., Jiang, H., Dai, L., Liu, C.: Incoherent training of deep neural networks to de-correlate bottleneck features for speech recognition. In: International Conference on Acoustics, Speech and Signal Processing. pp. 6980–6984 (2013)
5. Barron, A.R.: Universal approximation bounds for superpositions of a sigmoidal function. IEEE Transactions on Information theory pp. 930–945 (1993)
6. Barron, A.R.: Approximation and estimation bounds for artificial neural networks. Machine Learning pp. 115–133 (1994)
7. Bartlett, P.L., Mendelson, S.: Rademacher and gaussian complexities: Risk bounds and structural results. Journal of Machine Learning Research pp. 463–482 (2002)

8. Bietti, A., Mialon, G., Chen, D., Mairal, J.: A kernel perspective for regularizing deep neural networks. In: International Conference on Machine Learning. pp. 664–674 (2019)
9. Bubeck, S., Sellke, M.: A universal law of robustness via isoperimetry. Neural Information Processing Systems (Neurips) (2021)
10. Cogswell, M., Ahmed, F., Girshick, R.B., Zitnick, L., Batra, D.: Reducing overfitting in deep networks by decorrelating representations. In: International Conference on Learning Representations (2016)
11. Deng, H., Ren, Q., Chen, X., Zhang, H., Ren, J., Zhang, Q.: Discovering and explaining the representation bottleneck of dnns. arXiv preprint arXiv:2111.06236 (2021)
12. Deng, Z., Zhang, L., Vodrahalli, K., Kawaguchi, K., Zou, J.: Adversarial training helps transfer learning via better representations. Neural Information Processing Systems (Neurips) (2021)
13. Du, S.S., Hu, W., Kakade, S.M., Lee, J.D., Lei, Q.: Few-shot learning via learning the representation, provably. International Conference on Learning Representations (2021)
14. Dziugaite, G.K., Roy, D.M.: Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. arXiv preprint arXiv:1703.11008 (2017)
15. Foret, P., Kleiner, A., Mobahi, H., Neyshabur, B.: Sharpness-aware minimization for efficiently improving generalization. arXiv preprint arXiv:2010.01412 (2020)
16. Golan, I., El-Yaniv, R.: Deep anomaly detection using geometric transformations. In: Advances in Neural Information Processing Systems. pp. 9758–9769 (2018)
17. Golowich, N., Rakhlin, A., Shamir, O.: Size-independent sample complexity of neural networks. In: Conference On Learning Theory. pp. 297–299 (2018)
18. Goodfellow, I., Bengio, Y., Courville, A., Bengio, Y.: Deep learning. MIT Press (2016)
19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
20. 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.: Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. Signal processing magazine **29**(6), 82–97 (2012)
21. Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.R.: Improving neural networks by preventing co-adaptation of feature detectors. arXiv preprint arXiv:1207.0580 (2012)
22. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4700–4708 (2017)
23. Ji, Z., Telgarsky, M.: The implicit bias of gradient descent on nonseparable data. In: Proceedings of the Thirty-Second Conference on Learning Theory. pp. 1772–1798 (2019)
24. Jiang, Y., Neyshabur, B., Mobahi, H., Krishnan, D., Bengio, S.: Fantastic generalization measures and where to find them. International Conference on Learning Representations (2019)
25. Kalimeris, D., Kaplun, G., Nakkiran, P., Edelman, B., Yang, T., Barak, B., Zhang, H.: Sgd on neural networks learns functions of increasing complexity. Neural Information Processing Systems **32**, 3496–3506 (2019)

26. Kawaguchi, K., Bengio, Y.: Depth with nonlinearity creates no bad local minima in resnets. *Neural Networks* **118**, 167–174 (2019)
27. Kawaguchi, K., Huang, J.: Gradient descent finds global minima for generalizable deep neural networks of practical sizes. In: 2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton). pp. 92–99. IEEE (2019)
28. Kawaguchi, K., Kaelbling, L.P., Bengio, Y.: Generalization in deep learning. arXiv preprint arXiv:1710.05468 (2017)
29. Kornblith, S., Chen, T., Lee, H., Norouzi, M.: Why do better loss functions lead to less transferable features? *Advances in Neural Information Processing Systems* **34** (2021)
30. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
31. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems* (2012)
32. Kukačka, J., Golkov, V., Cremers, D.: Regularization for deep learning: A taxonomy. arXiv preprint arXiv:1710.10686 (2017)
33. Kwok, J.T., Adams, R.P.: Priors for diversity in generative latent variable models. In: *Advances in Neural Information Processing Systems*. pp. 2996–3004 (2012)
34. Laakom, F., Raitoharju, J., Iosifidis, A., Gabbouj, M.: Efficient cnn with uncorrelated bag of features pooling. In: 2022 IEEE Symposium Series on Computational Intelligence (SSCI) (2022)
35. Laakom, F., Raitoharju, J., Iosifidis, A., Gabbouj, M.: Reducing redundancy in the bottleneck representation of the autoencoders. arXiv preprint arXiv:2202.04629 (2022)
36. Laakom, F., Raitoharju, J., Iosifidis, A., Gabbouj, M.: Wld-reg: A data-dependent within-layer diversity regularizer. the 37th AAAI Conference on Artificial Intelligence (2023)
37. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
38. Lee, H.B., Nam, T., Yang, E., Hwang, S.J.: Meta dropout: Learning to perturb latent features for generalization. In: *International Conference on Learning Representations* (2019)
39. Li, Z., Gong, B., Yang, T.: Improved dropout for shallow and deep learning. In: *Advances in Neural Information Processing Systems*. pp. 2523–2531 (2016)
40. Liao, Q., Miranda, B., Rosasco, L., Banburski, A., Liang, R., Hidary, J., Poggio, T.: Generalization puzzles in deep networks. *International Conference on Learning Representations* (2020)
41. Mao, C., Gupta, A., Nitin, V., Ray, B., Song, S., Yang, J., Vondrick, C.: Multitask learning strengthens adversarial robustness. In: *European Conference on Computer Vision*. pp. 158–174. Springer (2020)
42. Maurer, A., Pontil, M., Romera-Paredes, B.: The benefit of multitask representation learning. *Journal of Machine Learning Research* (2016)
43. Nagarajan, V., Kolter, J.Z.: Uniform convergence may be unable to explain generalization in deep learning. In: *Advances in Neural Information Processing Systems* (2019)
44. Neyshabur, B., Bhojanapalli, S., McAllester, D., Srebro, N.: Exploring generalization in deep learning. *NIPS* (2017)
45. Neyshabur, B., Li, Z., Bhojanapalli, S., LeCun, Y., Srebro, N.: The role of over-parametrization in generalization of neural networks. In: *International Conference on Learning Representations* (2018)



46. Pinot, R., Meunier, L., Araujo, A., Kashima, H., Yger, F., Gouy-Pailler, C., Atif, J.: Theoretical evidence for adversarial robustness through randomization. In: *Advances in Neural Information Processing Systems (Neurips)* (2019)
47. Poggio, T., Kawaguchi, K., Liao, Q., Miranda, B., Rosasco, L., Boix, X., Hidary, J., Mhaskar, H.: Theory of deep learning III: explaining the non-overfitting puzzle. *arXiv preprint arXiv:1801.00173* (2017)
48. Rodriguez-Galvez, B., Bassi, G., Thobaben, R., Skoglund, M.: Tighter expected generalization error bounds via wasserstein distance. *Advances in Neural Information Processing Systems* (2021)
49. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* (2015)
50. Shalev-Shwartz, S., Ben-David, S.: *Understanding machine learning: From theory to algorithms*. Cambridge university press (2014)
51. Sokolic, J., Giryas, R., Sapiro, G., Rodrigues, M.R.: *Lessons from the rademacher complexity for deep learning* (2016)
52. Sontag, E.D.: VC dimension of neural networks. *NATO ASI Series F Computer and Systems Sciences* pp. 69–96 (1998)
53. Valiant, L.: A theory of the learnable. *Commun. of the ACM* **27**(1), 134–1 (1984)
54. Volhejn, V., Lampert, C.: Does sgd implicitly optimize for smoothness? In: *DAGM German Conference on Pattern Recognition*. pp. 246–259. Springer (2020)
55. Wang, X., Chen, X., Du, S.S., Tian, Y.: Towards demystifying representation learning with non-contrastive self-supervision. *arXiv preprint arXiv:2110.04947* (2021)
56. Wu, B., Chen, J., Cai, D., He, X., Gu, Q.: Do wider neural networks really help adversarial robustness? *Advances in Neural Information Processing Systems* **34** (2021)
57. Xie, B., Liang, Y., Song, L.: Diverse neural network learns true target functions. In: *Artificial Intelligence and Statistics*. pp. 1216–1224 (2017)
58. Xie, P., Deng, Y., Xing, E.: On the generalization error bounds of neural networks under diversity-inducing mutual angular regularization. *arXiv preprint arXiv:1511.07110* (2015)
59. Xie, P., Singh, A., Xing, E.P.: Uncorrelation and evenness: a new diversity-promoting regularizer. In: *International Conference on Machine Learning*. pp. 3811–3820 (2017)
60. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1492–1500 (2017)
61. Yu, Y., Li, Y.F., Zhou, Z.H.: Diversity regularized machine. In: *International Joint Conference on Artificial Intelligence* (2011)
62. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: *Proceedings of the British Machine Vision Conference (BMVC)* (2016)
63. Zhai, K., Wang, H.: Adaptive dropout with rademacher complexity regularization. In: *International Conference on Learning Representations* (2018)
64. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. *International Conference on Learning Representations 2018* (2018)
65. Zou, D., Wu, J., Gu, Q., Foster, D.P., Kakade, S., et al.: The benefits of implicit regularization from sgd in least squares problems. *Neural Information Processing Systems* **34** (2021)