

# Active Subspace of Neural Networks: Structural Analysis and Universal Attacks \*

Chunfeng Cui<sup>†</sup>, Kaiqi Zhang<sup>†</sup>, Talgat Daulbaev<sup>‡</sup>, Julia Gusak<sup>‡</sup>,  
Ivan Oseledets<sup>§</sup>, and Zheng Zhang<sup>†</sup>

**Abstract.** Active subspace is a model reduction method widely used in the uncertainty quantification community. In this paper, we propose analyzing the internal structure and vulnerability of deep neural networks using active subspace. Firstly, we employ the active subspace to measure the number of “active neurons” at each intermediate layer, which indicates that the number of neurons can be reduced from several thousands to several dozens. This motivates us to change the network structure and to develop a new and more compact network, referred to as ASNet, that has significantly fewer model parameters. Secondly, we propose analyzing the vulnerability of a neural network using active subspace by finding an additive universal adversarial attack vector that can misclassify a dataset with a high probability. Our experiments on CIFAR-10 show that ASNet can achieve  $23.98\times$  parameter and  $7.30\times$  flops reduction. The universal active subspace attack vector can achieve around 20% higher attack ratio compared with the existing approaches in our numerical experiments. The PyTorch codes for this paper are available online <sup>1</sup>.

**Key words.** Active Subspace, Deep Neural Network, Network Reduction, Universal Adversarial Perturbation

**AMS subject classifications.** 90C26, 15A18, 62G35

**1. Introduction.** Deep neural networks have achieved impressive performance in many applications, such as computer vision [35], nature language processing [58], and speech recognition [23]. Most neural networks use deep structure (i.e., many layers) and a huge number of neurons to achieve a high accuracy and expressive power [44, 19]. However, it is still unclear how many layers and neurons are necessary. Employing an unnecessarily complicated deep neural network can cause huge extra costs in run-time and hardware resources. Driven by resource-constrained applications such as robotics and internet of things, there is an increasing interest in building smaller neural networks by removing network redundancy. Representative methods include network pruning and sharing [17, 25, 27, 39, 38], low-rank matrix and tensor factorization [49, 26, 18, 36, 43], parameter quantization [12, 15], knowledge distillation [28, 46], and so forth. However, most existing methods delete model parameters directly without changing the network architecture [27, 25, 7, 38].

Another important issue of deep neural networks is the lack of robustness. A deep neural

---

\*Submitted to the editors in October 2019, accepted by SIAM Journal on Mathematics of Data Science in July 2020.

**Funding:** Chunfeng Cui, Kaiqi Zhang, and Zheng Zhang are supported by the UCSB start-up grant. Talgat Daulbaev, Julia Gusak, and Ivan Oseledets are supported by the Ministry of Education and Science of the Russian Federation (grant 14.756.31.0001).

<sup>†</sup>University of California Santa Barbara, Santa Barbara, CA, USA ([chunfengcui@ucsb.edu](mailto:chunfengcui@ucsb.edu), [kzhang07@ucsb.edu](mailto:kzhang07@ucsb.edu), [zhengzhang@ece.ucsb.edu](mailto:zhengzhang@ece.ucsb.edu)).

<sup>‡</sup>Skolkovo Institute of Science and Technology, Moscow, Russia ([talgat.daulbaev@skoltech.ru](mailto:talgat.daulbaev@skoltech.ru), [y.gusak@skoltech.ru](mailto:y.gusak@skoltech.ru)).

<sup>§</sup>Skolkovo Institute of Science and Technology and Institute of Numerical Mathematics of Russian Academy of Sciences, Moscow, Russia ([ivan.oseledets@gamil.com](mailto:ivan.oseledets@gamil.com)).

<sup>1</sup>Codes are available at: <https://github.com/chunfengcui/ASNet>

network is desired to maintain good performance for noisy or corrupted data to be deployed in safety-critical applications such as autonomous driving and medical image analysis. However, recent studies have revealed that many state-of-the-art deep neural networks are vulnerable to small perturbations [54]. A substantial number of methods have been proposed to generate adversarial examples. Representative works can be classified into four classes [52], including optimization methods [8, 41, 40, 54], sensitive features [22, 45], geometric transformations [16, 32], and generative models [4]. However, these methods share a fundamental limitation: each perturbation is designed for a given data point, and one has to implement the algorithm again to generate the perturbation for a new data sample. Recently, several methods have also been proposed to compute a universal adversarial attack to fool a dataset simultaneously (rather than one data sample) in various applications, such as computer vision [40], speech recognition [42], audio [1], and text classifier [5]. However, all the above methods only solve a series of data-dependent sub-problems. In [33], Khrulkov et al. proposed to construct universal perturbation by computing the so-called  $(p, q)$ -singular vectors of the Jacobian matrices of hidden layers of a network.

This paper investigates the above two issues with the active subspace method [48, 9, 10] that was originally developed for uncertainty quantification. The key idea of the active subspace is to identify the low-dimensional subspace constructed by some important directions that can contribute significantly to the variance of the multi-variable function. These directions are corresponding to the principal components of the uncentered covariance matrix of gradients. Afterwards, a response surface can be constructed in this low-dimensional subspace to reduce the number of parameters for partial differential equations [10] and uncertainty quantification [11]. However, the power of active subspace in analyzing and attacking deep neural networks has not been explored.

**1.1. Paper Contributions.** The contribution of this manuscript is twofold.

- Firstly, we apply the active subspace to some intermediate layers of a deep neural network, and try to answer the following question: *how many neurons and layers are important in a deep neural network?* Based on the active subspace, we propose the definition of “active neurons”. Fig. 1 (a) shows that even though there are tens of thousands of neurons, only dozens of them are important from the active subspace point of view. Fig. 1 (b) further shows that most of the neural network parameters are distributed in the last few layers. This motivates us to cut off the tail layers and replace them with a smaller and simpler new framework called ASNet. ASNet contains three parts: the first few layers of a deep neural network, an active-subspace layer that maps the intermediate neurons to a low-dimensional subspace, and a polynomial chaos expansion layer that projects the reduced variables to the outputs. Our numerical experiments show that the proposed ASNet has much fewer model parameters than the original one. ASNet can also be combined with existing structured re-training methods (e.g., pruning and quantization) to get better accuracy while using fewer model parameters.
- Secondly, we use active subspace to develop a new universal attack method to fool deep neural networks on a whole data set. We formulate this problem as a ball-constrained loss maximization problem and propose a heuristic projected gradient descent algorithm to solve it. At each iteration, the ascent direction is the dominant active subspace, and the stepsize

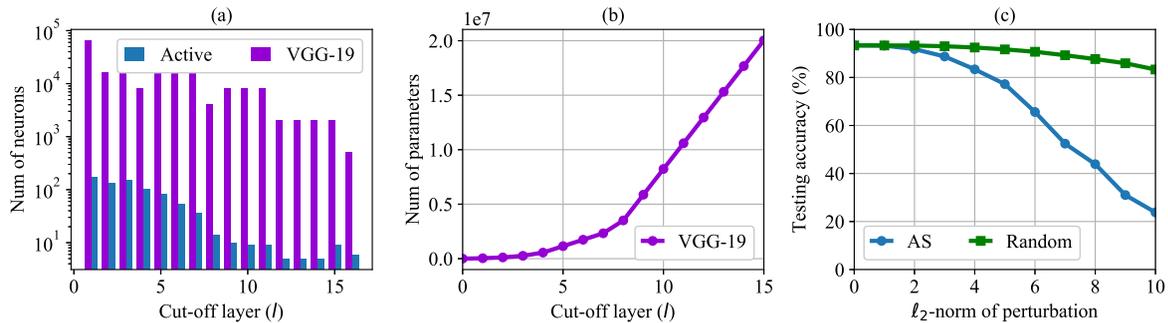


Figure 1: Structural analysis of deep neural networks by the active subspace (AS). All experiments are conducted on CIFAR-10 by VGG-19. (a) The number of neurons can be significantly reduced by the active subspace. Here, the number of active neurons is defined by Definition 3.1 with a threshold  $\epsilon = 0.05$ ; (b) Most of the parameters are distributed in the last few layers; (c) The active subspace direction can perturb the network significantly.

is decided by the backtracking algorithm. Fig. 1 (c) shows that the attack ratio of the active subspace direction is much higher than that of the random vector.

The rest of this manuscript is organized as follows. In Section 2, we review the key idea of active subspace. Based on the active-subspace method, Section 3 shows how to find the number of active neurons in a deep neural network and further proposes a new and compact network, referred to as ASNet. Section 4 develops a new universal adversarial attack method based on active subspace. The numerical experiments for both ASNet and universal adversarial attacks are presented in Section 5. Finally, we conclude this paper in Section 6.

**2. Active Subspace.** Active-subspace is an efficient tool for functional analysis and dimension reduction. Its key idea is to construct a low-dimensional subspace for the input variables in which the function value changes dramatically. Given a continuous function  $c(\mathbf{x})$  with  $\mathbf{x}$  described by the probability density function  $\rho(\mathbf{x})$ , one can construct an uncentered covariance matrix for the gradient:  $\mathbf{C} = \mathbb{E}[\nabla c(\mathbf{x})\nabla c(\mathbf{x})^T]$ . Suppose the matrix  $\mathbf{C}$  admits the following eigenvalue decomposition,

$$(2.1) \quad \mathbf{C} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T,$$

where  $\mathbf{V}$  includes all orthogonal eigenvectors and

$$(2.2) \quad \mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n), \quad \lambda_1 \geq \dots \geq \lambda_n \geq 0$$

are the eigenvalues. All the eigenvalues are nonnegative because  $\mathbf{C}$  is positive semidefinite. One can split the matrix  $\mathbf{V}$  into two parts,

$$(2.3) \quad \mathbf{V} = [\mathbf{V}_1, \mathbf{V}_2], \quad \text{where } \mathbf{V}_1 \in \mathbb{R}^{n \times r} \text{ and } \mathbf{V}_2 \in \mathbb{R}^{n \times (n-r)}.$$

The subspace spanned by matrix  $\mathbf{V}_1 \in \mathbb{R}^{n \times r}$  is called an active subspace [48], because  $c(\mathbf{x})$  is sensitive to perturbation vectors inside this subspace .

*Remark 2.1 (Relationships with the Principal Component Analysis).* Given a set of data  $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^m]$  with each column representing a data sample and each row is zero-mean, the first principal component  $\mathbf{w}_1$  inherits the maximal variance from  $\mathbf{X}$ , namely,

$$(2.4) \quad \mathbf{w}_1 = \operatorname{argmax}_{\|\mathbf{w}\|_2=1} \sum_{i=1}^m (\mathbf{w}_1^T \mathbf{x}^i)^2 = \operatorname{argmax}_{\|\mathbf{w}\|_2=1} \mathbf{w}^T \mathbf{X} \mathbf{X}^T \mathbf{w}.$$

The variance is maximized when  $\mathbf{w}_1$  is the eigenvector associated with the largest eigenvalue of  $\mathbf{X} \mathbf{X}^T$ . The first  $r$  principal components are the  $r$  eigenvectors associated with the  $r$  largest eigenvalues of  $\mathbf{X} \mathbf{X}^T$ . The main difference with the active subspace is that the principal component analysis uses the covariance matrix of input data sets  $\mathbf{X}$ , but the active-subspace method uses the covariance matrix of gradient  $\nabla c(\mathbf{x})$ . Hence, a perturbation along the direction  $\mathbf{w}_1$  from (2.4) only guarantee the variability in the data, and does not necessarily cause a significantly change on the value of  $c(\mathbf{x})$ .

The following lemma quantitatively describes that  $c(\mathbf{x})$  varies more on average along the directions defined by the columns of  $\mathbf{V}_1$  than the directions defined by the columns of  $\mathbf{V}_2$ .

**Lemma 2.2.** [10] *Suppose  $c(\mathbf{x})$  is a continuous function and  $\mathbf{C}$  is obtained from (2.1). For the matrices  $\mathbf{V}_1$  and  $\mathbf{V}_2$  generated by (2.3), and the reduced vector*

$$(2.5) \quad \mathbf{z} = \mathbf{V}_1^T \mathbf{x} \text{ and } \tilde{\mathbf{z}} = \mathbf{V}_2^T \mathbf{x},$$

*it holds that*

$$(2.6) \quad \begin{aligned} \mathbb{E}_{\mathbf{x}}[\nabla_{\mathbf{z}} c(\mathbf{x})^T \nabla_{\mathbf{z}} c(\mathbf{x})] &= \lambda_1 + \dots + \lambda_r, \\ \mathbb{E}_{\mathbf{x}}[\nabla_{\tilde{\mathbf{z}}} c(\mathbf{x})^T \nabla_{\tilde{\mathbf{z}}} c(\mathbf{x})] &= \lambda_{r+1} + \dots + \lambda_n. \end{aligned}$$

*Sketch of proof [10]:*

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}}[\nabla_{\mathbf{z}} c(\mathbf{x})^T \nabla_{\mathbf{z}} c(\mathbf{x})] \\ &= \operatorname{trace}(\mathbb{E}_{\mathbf{x}}[\nabla_{\mathbf{z}} c(\mathbf{x}) \nabla_{\mathbf{z}} c(\mathbf{x})^T]) \\ &= \operatorname{trace}(\mathbb{E}_{\mathbf{x}}[\mathbf{V}_1^T \nabla_{\mathbf{x}} c(\mathbf{x}) \nabla_{\mathbf{x}} c(\mathbf{x})^T \mathbf{V}_1]) \\ &= \operatorname{trace}(\mathbf{V}_1^T \mathbf{C} \mathbf{V}_1) \\ &= \lambda_1 + \dots + \lambda_r. \end{aligned}$$

When  $\lambda_{r+1} = \dots = \lambda_n = 0$ , Lemma 2.2 implies  $\nabla_{\tilde{\mathbf{z}}} c(\mathbf{x})$  is zero everywhere, i.e.,  $c(\mathbf{x})$  is  $\tilde{\mathbf{z}}$ -invariant. In this case, we may reduce  $\mathbf{x} \in \mathbb{R}^n$  to a low-dimensional vector  $\mathbf{z} = \mathbf{V}_1^T \mathbf{x} \in \mathbb{R}^r$  and construct a new response surface  $g(\mathbf{z})$  to represent  $c(\mathbf{x})$ . Otherwise, if  $\lambda_{r+1}$  is small, we may still construct a response surface  $g(\mathbf{z})$  to approximate  $c(\mathbf{x})$  with a bounded error, as shown in the following lemma.

**2.1. Response Surface.** For a fixed  $\mathbf{z}$ , the best guess for  $g$  is the conditional expectation of  $c$  given  $\mathbf{z}$ , i.e.,

$$(2.7) \quad g(\mathbf{z}) = \mathbb{E}_{\tilde{\mathbf{z}}}[c(\mathbf{x})|\mathbf{z}] = \int c(\mathbf{V}_1 \mathbf{z} + \mathbf{V}_2 \tilde{\mathbf{z}}) \rho(\tilde{\mathbf{z}}|\mathbf{z}) d\tilde{\mathbf{z}}.$$

Based on the Poincaré inequality, the following approximation error bound is obtained [10].

**Lemma 2.3.** *Assume that  $c(\mathbf{x})$  is absolutely continuous and square integrable with respect to the probability density function  $\rho(\mathbf{x})$ , then the approximation function  $g(\mathbf{z})$  in (2.7) satisfies:*

$$(2.8) \quad \mathbb{E}[(c(\mathbf{x}) - g(\mathbf{z}))^2] \leq O(\lambda_{r+1} + \dots + \lambda_n).$$

*Sketch of proof [10]:*

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}}[(c(\mathbf{x}) - g(\mathbf{z}))^2] \\ &= \mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\mathbf{z}}[(c(\mathbf{x}) - g(\mathbf{z}))^2 | \mathbf{z}]] \\ &\leq \text{const} \times \mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\mathbf{z}}[\nabla_{\mathbf{z}} c(\mathbf{x})^T \nabla_{\mathbf{z}} c(\mathbf{x}) | \mathbf{z}]] \quad (\text{Poincaré inequality}) \\ &= \text{const} \times \mathbb{E}_{\mathbf{x}}[\nabla_{\mathbf{z}} c(\mathbf{x})^T \nabla_{\mathbf{z}} c(\mathbf{x})] \\ &= \text{const} \times (\lambda_{r+1} + \dots + \lambda_n) \quad (\text{Lemma 2.2}) \\ &= O(\lambda_{r+1} + \dots + \lambda_n). \end{aligned}$$

In other words, the active-subspace approximation error will be small if  $\lambda_{r+1}, \dots, \lambda_n$  are negligible.

### 3. Active Subspace for Structural Analysis and Compression of Deep Neural Networks.

This section applies the active subspace to analyze the internal layers of a deep neural network to reveal the number of important neurons at each layer. Afterward, a new network called ASNet is built to reduce the storage and computational complexity.

**3.1. Deep Neural Networks.** A deep neural network can be described as

$$(3.1) \quad f(\mathbf{x}_0) = f_L(f_{L-1} \dots (f_1(\mathbf{x}_0))),$$

where  $\mathbf{x}_0 \in \mathbb{R}^{n_0}$  is an input,  $L$  is the total number of layers, and  $f_l : \mathbb{R}^{n_{l-1}} \rightarrow \mathbb{R}^{n_l}$  is a function representing the  $l$ -th layer (e.g., combinations of convolution or fully connected, batch normalization, ReLU, or pooling layers). For any  $1 \leq l \leq L$ , we rewrite the above feed-forward model as a superposition of functions, i.e.,

$$(3.2) \quad f(\mathbf{x}_0) = f_{\text{post}}^l(f_{\text{pre}}^l(\mathbf{x}_0)),$$

where the **pre-model**  $f_{\text{pre}}^l(\cdot) = f_l \dots (f_1(\cdot))$  denotes all operations before the  $l$ -th layer and the **post-model**  $f_{\text{post}}^l(\cdot) = f_L \dots (f_{l+1}(\cdot))$  denotes all succeeding operations. The intermediate neuron  $\mathbf{x}_l = f_{\text{pre}}^l(\mathbf{x}_0) \in \mathbb{R}^{n_l}$  usually lies in a high dimension. We aim to study whether such a high dimensionality is necessary. If not, how can we reduce it?

**3.2. The Number of Active Neurons.** Denote  $\text{loss}(\cdot)$  as the loss function, and

$$(3.3) \quad c_l(\mathbf{x}) = \text{loss}(f_{\text{post}}^l(\mathbf{x})).$$

The covariance matrix  $\mathbf{C} = \mathbb{E}[\nabla c_l(\mathbf{x}) \nabla c_l(\mathbf{x})^T]$  admits the eigenvalue decomposition  $\mathbf{C} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T$  with  $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_{n_l})$ . We try to extract the active subspace of  $c_l(\mathbf{x})$  and reduce the intermediate vector  $\mathbf{x}$  to a low dimension. Here the intermediate neuron  $\mathbf{x}$ , the covariance matrix  $\mathbf{C}$ , eigenvalues  $\mathbf{\Lambda}$ , and eigenvectors  $\mathbf{V}$  are also related to the layer index  $l$ , but we ignore the index for simplicity.

**Definition 3.1.** Suppose  $\mathbf{\Lambda}$  is computed by (2.2). For any layer index  $1 \leq l \leq L$ , we define the number of active neurons  $n_{l,AS}$  as follows:

$$(3.4) \quad n_{l,AS} = \arg \min \left\{ i : \frac{\lambda_1 + \dots + \lambda_i}{\lambda_1 + \dots + \lambda_{n_l}} \geq 1 - \epsilon \right\},$$

where  $\epsilon > 0$  is a user-defined threshold.

Based on Definition 3.1, the post-model can be approximated by an  $n_{l,AS}$ -dimensional function with a high accuracy, i.e.,

$$(3.5) \quad g_l(\mathbf{z}) = \mathbb{E}_{\tilde{\mathbf{z}}} [c_l(\mathbf{x}) | \mathbf{z}].$$

Here  $\mathbf{z} = \mathbf{V}_1^T \mathbf{x} \in \mathbb{R}^{n_{l,AS}}$  plays the role of active neurons,  $\tilde{\mathbf{z}} = \mathbf{V}_2^T \mathbf{x} \in \mathbb{R}^{n-n_{l,AS}}$ , and  $\mathbf{V} = [\mathbf{V}_1, \mathbf{V}_2]$ .

**Lemma 3.1.** Suppose the input  $\mathbf{x}_0$  is bounded. Consider a deep neural network with the following operations: convolution, fully connected, ReLU, batch normalization, max-pooling, and equipped with the cross entropy loss function. Then for any  $l \in \{1, \dots, L\}$ ,  $\mathbf{x} = f_{pre}^l(\mathbf{x}_0)$ , and  $c_l(\mathbf{x}) = \text{loss}(f_{post}^l(\mathbf{x}))$ , the  $n_{l,AS}$ -dimensional function  $g_l(\mathbf{z})$  defined in (3.5) satisfies

$$(3.6) \quad \mathbb{E}_{\mathbf{z}} \left[ (g_l(\mathbf{z}))^2 \right] \leq 2\mathbb{E}_{\mathbf{x}_0} \left[ (c_0(\mathbf{x}_0))^2 \right] + O(\epsilon).$$

*Proof.* Denote  $c_l(\mathbf{x}) = \text{loss}(f_L(\dots(f_{l+1}(\mathbf{x})))$ , where  $\text{loss}(\mathbf{y}) = -\log \frac{\exp(y_b)}{\sum_{i=1}^{n_L} \exp(y_i)}$  is the cross entropy loss function,  $b$  is the true label, and  $n_L$  is the total number of classes. We first show  $c_l(\mathbf{x})$  is absolutely continuous and square integrable, and then apply Lemma 2.3 to derive (3.6).

Firstly, all components of  $c_l(\mathbf{x})$  are Lipschitz continuous because (1) the convolution, fully connected, and batch normalization operations are all linear; (2) the max pooling and ReLU functions are non-expansive. Here, a mapping  $m$  is non-expansive if  $\|m(\mathbf{x}) - m(\mathbf{y})\| \leq \|\mathbf{x} - \mathbf{y}\|$ ; (3) the cross entropy loss function is smooth with an upper bounded gradient, i.e.,  $\|\nabla \text{loss}(\mathbf{y})\| = \|\mathbf{e}_b - \exp(\mathbf{y}) / \sum_{i=1}^{n_L} \exp(y_i)\| \leq \sqrt{n_L}$ . The composition of two Lipschitz continuous functions is also Lipschitz continuous: suppose the Lipschitz constants for  $f_1$  and  $f_2$  are  $\alpha_1$  and  $\alpha_2$ , respectively, it holds that  $\|f_1(f_2(\bar{\mathbf{x}})) - f_1(f_2(\underline{\mathbf{x}}))\| \leq \alpha_1 \|f_2(\bar{\mathbf{x}}) - f_2(\underline{\mathbf{x}})\| \leq \alpha_1 \alpha_2 \|\bar{\mathbf{x}} - \underline{\mathbf{x}}\|$  for any vectors  $\bar{\mathbf{x}}$  and  $\underline{\mathbf{x}}$ . By recursively applying the above rule,  $c_l(\mathbf{x})$  is Lipschitz continuous:

$$\begin{aligned} \|c_l(\bar{\mathbf{x}}) - c_l(\underline{\mathbf{x}})\|_2 &= \|\text{loss}(f_L(\dots(f_{l+1}(\bar{\mathbf{x}})))) - \text{loss}(f_L(\dots(f_{l+1}(\underline{\mathbf{x}}))))\|_2 \\ &\leq \sqrt{n_L} \alpha_L \dots \alpha_{l+1} \|\bar{\mathbf{x}} - \underline{\mathbf{x}}\|_2. \end{aligned}$$

The intermediate neuron  $\mathbf{x}$  is in a bounded domain because the input  $\mathbf{x}_0$  is bounded and all functions  $f_i(\cdot)$  are either continuous or non-expansive. Based on the fact that any Lipschitz-continuous function is also absolutely continuous on a compact domain [47], we conclude that  $c_l(\mathbf{x})$  is absolutely continuous.

Secondly, because  $\mathbf{x}$  is bounded and  $c_l(\mathbf{x})$  is continuous, both  $c_l(\mathbf{x})$  and its square integral will be bounded, i.e.,  $\int (c_l(\mathbf{x}))^2 \rho(\mathbf{x}) d\mathbf{x} < \infty$ .

Finally, by Lemma 2.3, it holds that

$$\mathbb{E}_{\mathbf{x}}[(c_l(\mathbf{x}) - g_l(\mathbf{z}))^2] \leq O(\lambda_{n_{l,AS}+1} + \dots + \lambda_n).$$

From Definition 3.1, we have

$$\lambda_{n_{l,AS}+1} + \dots + \lambda_n \leq (\lambda_1 + \dots + \lambda_n)\epsilon = \|\mathbf{C}^{1/2}\|_F^2 \epsilon = O(\epsilon).$$

In the last equality, we used that  $\|\mathbf{C}^{1/2}\|_F$  is upper bounded because  $c_l(\mathbf{x})$  is Lipschitz continuous with a bounded gradient. Consequently, we have

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}}[(g_l(\mathbf{z}))^2] \\ &= \mathbb{E}_{\mathbf{x}}[(g_l(\mathbf{z}) - c_l(\mathbf{x}) + c_l(\mathbf{x}))^2] \\ &\leq 2\mathbb{E}_{\mathbf{x}}[(c_l(\mathbf{x}))^2] + 2\mathbb{E}_{\mathbf{x}}[(c_l(\mathbf{x}) - g_l(\mathbf{z}))^2] \\ &= 2\mathbb{E}_{\mathbf{x}_0}[(c_0(\mathbf{x}_0))^2] + 2\mathbb{E}_{\mathbf{x}}[(c_l(\mathbf{x}) - g_l(\mathbf{z}))^2] \\ &\leq 2\mathbb{E}_{\mathbf{x}_0}[(c_0(\mathbf{x}_0))^2] + O(\epsilon). \end{aligned}$$

The proof is completed. ■

The above lemma shows that the active subspace method can reduce the number of neurons of the  $l$ -th layer from  $n_l$  to  $n_{l,AS}$ . The loss for the low-dimensional function  $g_l(\mathbf{z})$  is bounded by two terms: the loss  $c_0(\mathbf{x}_0)$  of the original network, and the threshold  $\epsilon$  related to  $n_{l,AS}$ . This loss function is the cross entropy loss, not the classification error. However, it is believed that a small loss will result in a small classification error. Further, the result in Lemma 3.1 is valid for the fixed parameters in the pre-model. In practice, we can fine-tune the pre-model to achieve better accuracy.

Further, a small active neurons  $n_{l,AS}$  is critical to get a high compress ratio. From Definition 3.1,  $n_{l,AS}$  depends on the eigenvalue distribution of the covariance matrix  $\mathbf{C}$ . For a proper network structure and a good choice of the layer index  $l$ , if the eigenvalues of  $\mathbf{C}$  are dominated by the first few eigenvalues, then  $n_{l,AS}$  will be small. For instance, in Fig. 5(a), the eigenvalues for layers  $4 \leq l \leq 7$  of VGG-19 are nearly exponential decreasing to zero.

**3.3. Active Subspace Network (ASNet).** This subsection proposes a new network called ASNet that can reduce both the storage and computational cost. Given a deep neural network, we first choose a proper layer  $l$  and project the high-dimensional intermediate neurons to a low-dimensional vector in the active subspace. Afterward, the post-model is deleted completely and replaced with a nonlinear model that maps the low-dimensional active feature vector to the output directly. This new network, called ASNet, has three parts:

- (1) **Pre-model:** the pre-model includes the first  $l$  layers of a deep neural network.
- (2) **Active subspace layer:** a linear projection from the intermediate neurons to the low-dimensional active subspace.
- (3) **Polynomial chaos expansion layer:** the polynomial chaos expansion [20, 56] maps the active-subspace variables to the output.

The initialization for the active subspace layer and polynomial chaos expansion layer are presented in Sections 3.4 and 3.5, respectively. We can also retrain all the parameters to increase the accuracy. The whole procedure is illustrated in Fig. 2 (b) and Algorithm 3.1.

---

**Algorithm 3.1** The training procedure of the active subspace network (ASNet)

---

**Input:** A pretrained deep neural network, the layer index  $l$ , and the number of active neurons  $r$ .

Step 1 **Initialize the active subspace layer.** The active subspace layer is a linear projection where the projection matrix  $\mathbf{V}_1 \in \mathbb{R}^{n \times r}$  is computed by Algorithm 3.2. If  $r$  is not given, we use  $r = n_{\text{AS}}$  defined in (3.4) by default.

Step 2 **Initialize the polynomial chaos expansion layer.** The polynomial chaos expansion layer is a nonlinear mapping from the reduced active subspace to the outputs, as shown in (3.10). The weights  $\mathbf{c}_\alpha$  is computed by (3.12).

Step 3 **Construct the ASNet.** Combine the pre-model (the first  $l$  layers of the deep neural network) with the active subspace and polynomial chaos expansion layers as a new network, referred to as ASNet.

Step 4 **Fine-tuning.** Retrain all the parameters in pre-model, active subspace layer and polynomial chaos expansion layer in ASNet for several epochs by stochastic gradient descent.

**Output:** A new network ASNet

---

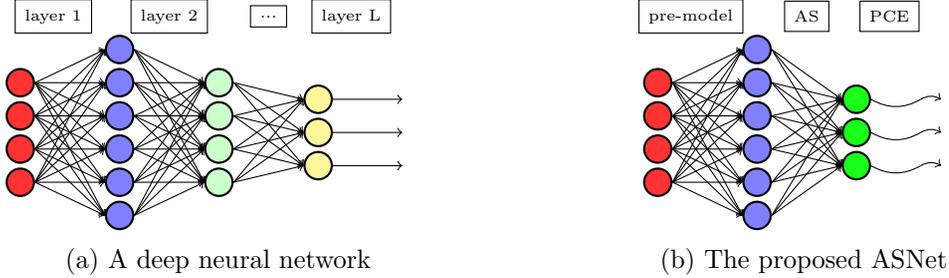


Figure 2: (a) The original deep neural network; (b) The proposed ASNet with three parts: a pre-model, an active subspace (AS) layer, and a polynomial chaos expansion (PCE) layer.

**3.4. The Active Subspace Layer.** This subsection presents an efficient method to project the high dimensional neurons to the active subspace. Given a dataset  $\mathcal{D} = \{\mathbf{x}^1, \dots, \mathbf{x}^m\}$ , the empirical covariance matrix is computed by  $\hat{\mathbf{C}} = \frac{1}{m} \sum_{i=1}^m \nabla c_l(\mathbf{x}^i) \nabla c_l(\mathbf{x}^i)^T$ . When ReLU is applied as an activation,  $c_l(\mathbf{x})$  is not differentiable. In this case,  $\nabla$  denotes the sub-gradient with a little abuse of notation.

Instead of calculating the eigenvalue decomposition of  $\hat{\mathbf{C}}$ , we compute the singular value decomposition of  $\hat{\mathbf{G}}$  to save the computation cost:

$$(3.7) \quad \hat{\mathbf{G}} = [\nabla c_l(\mathbf{x}^1), \dots, \nabla c_l(\mathbf{x}^m)] = \hat{\mathbf{V}} \hat{\mathbf{\Sigma}} \hat{\mathbf{U}}^T \in \mathbb{R}^{n_l \times m} \text{ with } \hat{\mathbf{\Sigma}} = \text{diag}(\hat{\sigma}_1, \dots, \hat{\sigma}_{n_l}).$$

The eigenvectors of  $\mathbf{C}$  are approximated by the left singular vectors  $\hat{\mathbf{V}}$  and the eigenvalues of  $\mathbf{C}$  are approximated by the singular values of  $\hat{\mathbf{G}}$ , i.e.,  $\mathbf{\Lambda} \approx \hat{\mathbf{\Sigma}}^2$ .

We use the memory-saving frequent direction method [21] to compute the  $r$  dominant singular value components, i.e.,  $\hat{\mathbf{G}} \approx \hat{\mathbf{V}}_r \hat{\mathbf{\Sigma}}_r \hat{\mathbf{U}}_r^T$ . Here  $r$  is smaller than the total number of

---

**Algorithm 3.2** The frequent direction algorithm for computing the active subspace

---

**Input:** A dataset with  $m_{AS}$  input samples  $\{\mathbf{x}_0^j\}_{j=1}^{m_{AS}}$ , a pre-model  $f_{\text{pre}}^l(\cdot)$ , a subroutine for computing  $\nabla c_l(\mathbf{x})$ , and the dimension of truncated singular value decomposition  $r$ .

- 1: Select  $r$  samples  $\mathbf{x}_0^i$ , compute  $\mathbf{x}^i = f_{\text{pre}}^l(\mathbf{x}_0^i)$ , and construct an initial matrix  $\mathbf{S} \leftarrow [\nabla c_l(\mathbf{x}^1), \dots, \nabla c_l(\mathbf{x}^r)]$ .
- 2: **for**  $t=1, 2, \dots$ , **do**
- 3: Compute the singular value decomposition  $\mathbf{V}\mathbf{\Sigma}\mathbf{U}^T \leftarrow \text{svd}(\mathbf{S})$ , where  $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_r)$ .
- 4: If the maximal number of samples  $m_{AS}$  is reached, stop.
- 5: Update  $\mathbf{S}$  by the soft-thresholding (3.8).
- 6: Get a new sample  $\mathbf{x}_0^{\text{new}}$ , compute  $\mathbf{x}^{\text{new}} = f_{\text{pre}}^l(\mathbf{x}_0^{\text{new}})$ , and replace the last column of  $\mathbf{S}$  (now all zeros) by the gradient vector  $\mathbf{S}(:, r) \leftarrow \nabla c_l(\mathbf{x}^{\text{new}})$ .
- 7: **end for**

**Output:** The projection matrix  $\mathbf{V} \in \mathbb{R}^{n_l \times r}$  and the singular values  $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$ .

---

samples. The frequent direction approach only stores an  $n \times r$  matrix  $\mathbf{S}$ . At the beginning, each column of  $\mathbf{S} \in \mathbb{R}^{n \times r}$  is initialized by a gradient vector. Then the randomized singular value decomposition [24] is used to generate  $\mathbf{S} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ . Afterwards,  $\mathbf{S}$  is updated in the following way,

$$(3.8) \quad \mathbf{S} \leftarrow \mathbf{V} \sqrt{\mathbf{\Sigma}^2 - \sigma_r^2}.$$

Now the last column of  $\mathbf{S}$  is zero and we replace it with the gradient vector of a new sample. By repeating this process,  $\mathbf{S}\mathbf{S}^T$  will approximate  $\hat{\mathbf{G}}\hat{\mathbf{G}}^T$  with a high accuracy and  $\mathbf{V}$  will approximate the left singular vectors of  $\hat{\mathbf{G}}$ . The algorithm framework is presented in Algorithm 3.2.

After obtaining  $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_r)$ , we can approximate the number of active neurons as

$$(3.9) \quad \hat{n}_{l,AS} = \arg \min \left\{ i : \frac{\sqrt{\sigma_1^2 + \dots + \sigma_i^2}}{\sqrt{\sigma_1^2 + \dots + \sigma_r^2}} \geq 1 - \epsilon \right\}.$$

Under the condition that  $\sigma_i^2 \rightarrow \lambda_i$  for  $i = 1, \dots, r$  and  $\lambda_i \rightarrow 0$  for  $i = r + 1, \dots, n_l$ , (3.9) can approximate  $n_{l,AS}$  in (3.4) with a high accuracy. Further, the projection matrix  $\hat{\mathbf{V}}_1$  is chosen as the first  $\hat{n}_{l,AS}$  columns of  $\mathbf{V}$ . The storage cost is reduced from  $O(n_l^2)$  to  $O(n_l r)$  and the computational cost is reduced from  $O(n_l^2 r)$  to  $O(n_l r^2)$ .

**3.5. Polynomial Chaos Expansion Layer.** We continue to construct a new surrogate model to approximate the post-model of a deep neural network. This problem can be regarded as an uncertainty quantification problem if we set  $\mathbf{z}$  as a random vector. We choose the nonlinear polynomial because it has higher expressive power than linear functions.

By the polynomial chaos expansion [55], the network output  $\mathbf{y} \in \mathbb{R}^{n_L}$  is approximated by

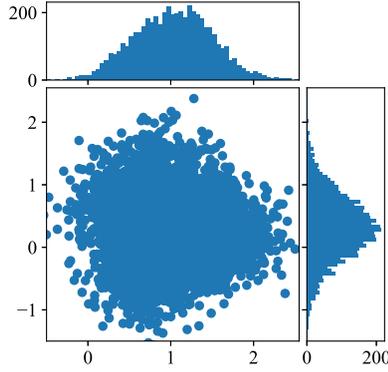


Figure 3: Distribution of the first two active subspace variables at the 6-th layer of VGG-19 for CIFAR-10.

a linear combination of the orthogonal polynomial basis functions:

$$(3.10) \quad \hat{\mathbf{y}} \approx \sum_{|\boldsymbol{\alpha}|=0}^p \mathbf{c}_{\boldsymbol{\alpha}} \phi_{\boldsymbol{\alpha}}(\mathbf{z}), \text{ where } |\boldsymbol{\alpha}| = \alpha_1 + \dots + \alpha_d.$$

Here  $\phi_{\boldsymbol{\alpha}}(\mathbf{z})$  is a multivariate polynomial basis function chosen based on the probability density function of  $\mathbf{z}$ . When the parameters  $\mathbf{z} = [z_1, \dots, z_r]^T$  are independent, both the joint density function and the multi-variable basis function can be decomposed into products of one-dimensional functions, i.e.,  $\rho(\mathbf{z}) = \rho_1(z_1) \dots \rho_r(z_r)$ ,  $\phi_{\boldsymbol{\alpha}}(\mathbf{z}) = \phi_{\alpha_1}(z_1) \phi_{\alpha_2}(z_2) \dots \phi_{\alpha_r}(z_r)$ . The marginal basis function  $\phi_{\alpha_j}(z_j)$  is uniquely determined by the marginal density function  $\rho_i(z_i)$ . The scatter plot in Fig. 3 shows that the marginal probability density of  $ez_i$  is close to a Gaussian distribution.

Suppose  $\rho_i(z_i)$  follows a Gaussian distribution, then  $\phi_{\alpha_j}(z_j)$  will be a Hermite polynomial [37], i.e.,

$$(3.11) \quad \phi_0(z) = 1, \phi_1(z) = z, \phi_2(z) = 4z^2 - 2, \phi_{p+1}(z) = 2z\phi_p(z) - 2p\phi_{p-1}(z).$$

In general, the elements in  $\mathbf{z}$  can be non-Gaussian correlated. In this case, the basis functions  $\{\phi_{\boldsymbol{\alpha}}(\mathbf{z})\}$  can be built via the Gram-Schmidt approach described in [13].

The coefficient  $\mathbf{c}_{\boldsymbol{\alpha}}$  can be computed by a linear least-square optimization. Denote  $\mathbf{z}^j = \hat{\mathbf{V}}_1^T f_{\text{pre}}^l(\mathbf{x}_0^j)$  as the random samples and  $\mathbf{y}^j$  as the network output for  $j = 1, \dots, m_{\text{PCE}}$ . The coefficient vector  $\mathbf{c}_{\boldsymbol{\alpha}}$  can be computed by

$$(3.12) \quad \min_{\{\mathbf{c}_{\boldsymbol{\alpha}}\}} \frac{1}{m_{\text{PCE}}} \sum_{j=1}^{m_{\text{PCE}}} \left\| \mathbf{y}^j - \sum_{|\boldsymbol{\alpha}|=0}^p \mathbf{c}_{\boldsymbol{\alpha}} \phi_{\boldsymbol{\alpha}}(\mathbf{z}^j) \right\|^2.$$

Based on the Nyquist-Shannon sampling theorem, the number of samples to train  $\mathbf{c}_{\boldsymbol{\alpha}}$  needs to satisfy  $m_{\text{PCE}} \geq 2n_{\text{basis}} = 2 \binom{r+p}{p}$ . However, this number can be reduced to a smaller set of “important” samples by the D-optimal design [59] or the sparse regularization approach [14].

The polynomial chaos expansion builds a surrogate model to approximate the deep neural network output  $\mathbf{y}$ . This idea is similar to the knowledge distillation [28], where a pre-trained teacher network teaches a smaller student network to learn the output feature. However, our polynomial-chaos layer uses one nonlinear projection whereas the knowledge distillation uses a series of layers. Therefore, the polynomial chaos expansion is more efficient in terms of computational and storage cost. The polynomial chaos expansion layer is different from the polynomial activation because the dimension of  $\mathbf{z}$  may be different from that of output  $\mathbf{y}$ .

The problem (3.12) is convex and any first order method can get a global optimal solution. Denote the optimal coefficients as  $\mathbf{c}_\alpha^*$  and the final objective value as  $\delta^*$ , i.e.,

$$(3.13) \quad \delta^* = \frac{1}{m_{\text{PCE}}} \sum_{j=1}^{m_{\text{PCE}}} \|\mathbf{y}^j - \psi^*(\mathbf{z}^j)\|^2, \quad \text{where } \psi^*(\mathbf{z}^j) = \sum_{|\alpha|=0}^p \mathbf{c}_\alpha^* \phi_\alpha(\mathbf{z}^j).$$

If  $\delta^* = 0$ , the polynomial chaos expansion is a good approximation to the original deep neural network on the training dataset. However, the approximation loss of the testing dataset may be large because of the overfitting phenomena.

The objective function in (3.12) is an empirical approximation to the expected error

$$(3.14) \quad \mathbb{E}_{(\mathbf{z}, \mathbf{y})} [\|\mathbf{y} - \psi(\mathbf{z})\|^2], \quad \text{where } \psi(\mathbf{z}) = \sum_{|\alpha|=0}^p \mathbf{c}_\alpha \phi_\alpha(\mathbf{z}).$$

According to the Hoeffding's inequality [29], the expected error (3.14) is close to the empirical error (3.12) with a high probability. Consequently, the loss for ASNet with polynomial chaos expansion layer is bounded as follows.

**Lemma 3.2.** *Suppose that the optimal solution for solving problem (3.12) is  $\mathbf{c}_\alpha^*$ , the optimal polynomial chaos expansion is  $\psi^*(\mathbf{z})$ , and the optimal residue is  $\delta^*$ . Assume that there exist constants  $a, b$  such that for all  $j$ ,  $\|\mathbf{y}^j - \psi^*(\mathbf{z}^j)\|^2 \in [a, b]$ . Then the loss of ASNet will be upper bounded*

$$(3.15) \quad \mathbb{E}_{\mathbf{z}} [(\text{loss}(\psi^*(\mathbf{z})))^2] \leq 2\mathbb{E}_{\mathbf{x}_0} [(c_0(\mathbf{x}_0))^2] + 2n_L(\delta^* + t) \quad \text{w.p. } 1 - \gamma^*,$$

where  $t$  is a user-defined threshold, and  $\gamma^* = \exp(-\frac{2t^2 m_{\text{PCE}}}{(b-a)^2})$ .

*Proof.* Since the cross entropy loss function is  $\sqrt{n_L}$ -Lipschitz continuous, we have

$$(3.16) \quad \mathbb{E}_{(\mathbf{y}, \mathbf{z})} [(\text{loss}(\mathbf{y}) - \text{loss}(\psi^*(\mathbf{z})))^2] \leq n_L \mathbb{E}_{(\mathbf{y}, \mathbf{z})} [\|\mathbf{y} - \psi^*(\mathbf{z})\|^2],$$

Denote  $\mathcal{T}^j = \|\mathbf{y}^j - \psi^*(\mathbf{z}^j)\|^2$  for  $i = 1, \dots, n_L$ .  $\{\mathcal{T}^j\}$  are independent under the assumption that the data samples are independent. By the Hoeffding's inequality, for any constant  $t$ , it holds that

$$(3.17) \quad \mathbb{E}[\mathcal{T}] \leq \frac{1}{m_{\text{PCE}}} \sum_j \mathcal{T}^j + t \quad \text{w.p. } 1 - \gamma^*,$$

with  $\gamma^* = \exp(-\frac{2t^2 m_{\text{PCE}}}{(b-a)^2})$ . Equivalently,

$$(3.18) \quad \mathbb{E}_{(\mathbf{y}, \mathbf{z})} [\|\mathbf{y} - \psi^*(\mathbf{z})\|^2] \leq \delta^* + t \quad \text{w.p. } 1 - \gamma^*,$$

Consequently, there is

$$\begin{aligned} & \mathbb{E}_{\mathbf{z}}[(\text{loss}(\psi^*(\mathbf{z})))^2] \\ & \leq 2\mathbb{E}_{\mathbf{y}}[(\text{loss}(\mathbf{y}))^2] + 2\mathbb{E}_{(\mathbf{y}, \mathbf{z})}[(\text{loss}(\psi^*(\mathbf{z})) - \text{loss}(\mathbf{y}))^2] \\ & \leq 2\mathbb{E}_{\mathbf{x}_0}[(c_0(\mathbf{x}_0))^2] + 2n_L(\delta^* + t) \quad \text{w.p. } 1 - \gamma^*. \end{aligned}$$

The last inequality follows from  $c_0(\mathbf{x}_0) = c_l(\mathbf{x}_l) = \text{loss}(\mathbf{y})$ , equations (3.16) and (3.18). This completes the proof.  $\blacksquare$

Lemma 3.2 shows with a high probability  $1 - \gamma^*$ , the expected error of ASNet without fine-tuning is bounded by the pre-trained error of the original network, the accuracy loss in solving the polynomial chaos subproblem (3.13), and the number of classes  $n_L$ . The probability  $\gamma^*$  is controlled by the threshold  $t$  as well as the number of training samples  $m_{\text{PCE}}$ .

In practice, we always re-train ASNet for several epochs and the accuracy of ASNet is beyond the scope of Lemma 3.2.

**3.6. Structured Re-training of ASNet.** The pre-model can be further compressed by various techniques such as network pruning and sharing [25], low-rank factorization [43, 36, 18], or data quantization [15, 12]. Denote  $\boldsymbol{\theta}$  as the weights in ASNet and  $\{\mathbf{x}_0^1, \dots, \mathbf{x}_0^m\}$  as the training dataset. Here,  $\boldsymbol{\theta}$  denotes all the parameters in the pre-model, active subspace layer, and the polynomial chaos expansion layer. We re-train the network by solving the following regularized optimization problem:

$$(3.19) \quad \boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \frac{1}{m} \sum_{i=1}^m \text{loss}(f(\boldsymbol{\theta}; \mathbf{x}_0^i)) + \lambda R(\boldsymbol{\theta}).$$

Here  $(\mathbf{x}_0^i, \mathbf{y}^i)$  is a training sample,  $m$  is the total number of training samples,  $\text{loss}(\cdot)$  is the cross-entropy loss function,  $R(\boldsymbol{\theta})$  is a regularization function, and  $\lambda$  is a regularization parameter. Different regularization functions can result in different model structures. For instance, an  $\ell_1$  regularizer  $R(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|_1$  [2, 50, 57] will return a sparse weight, an  $\ell_{1,2}$ -norm regularizer will result in a column-wise sparse weights, a nuclear norm regularizer will result in low-rank weights. At each iteration, we solve (3.19) by a stochastic proximal gradient decent algorithm [53]

$$(3.20) \quad \boldsymbol{\theta}^{k+1} = \arg \max_{\boldsymbol{\theta}} (\boldsymbol{\theta} - \boldsymbol{\theta}^k)^T \mathbf{g}^k + \frac{1}{2\alpha_k} \|\boldsymbol{\theta} - \boldsymbol{\theta}^k\|_2^2 + \lambda R(\boldsymbol{\theta}).$$

Here  $\mathbf{g}^k = \frac{1}{|\mathcal{B}_k|} \sum_{i \in \mathcal{B}_k} \nabla_{\boldsymbol{\theta}} \text{loss}(f(\boldsymbol{\theta}; \mathbf{x}_0^i), \mathbf{y}^i)$  is the stochastic gradient,  $\mathcal{B}_k$  is a batch at the  $k$ -th step, and  $\alpha_k$  is the stepsize.

In this work, we chose the  $\ell_1$  regularization to get sparse weight matrices. In this case, problem (3.20) has a closed-form solution:

$$(3.21) \quad \boldsymbol{\theta}^{k+1} = \mathcal{S}_{\alpha_k \lambda}(\boldsymbol{\theta}^k - \alpha_k \mathbf{g}^k),$$

where  $\mathcal{S}_{\lambda}(\mathbf{x}) = \mathbf{x} \odot \max(0, 1 - \lambda/|\mathbf{x}|)$  is a soft-thresholding operator.

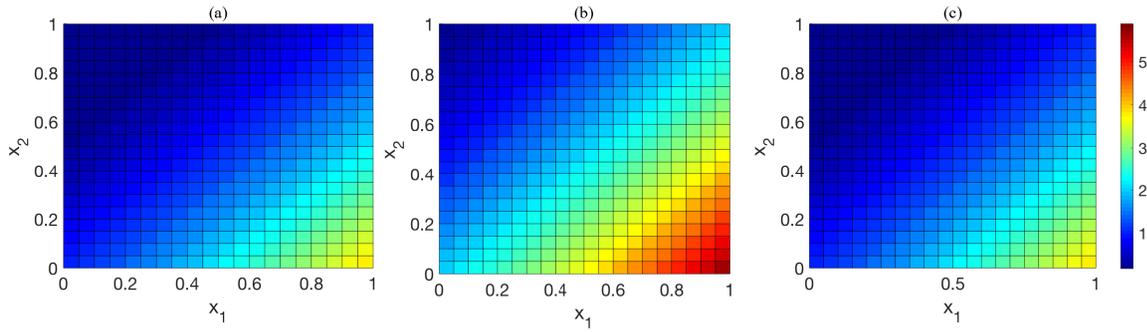


Figure 4: Perturbations along the directions of an active-subspace direction and of principal component, respectively. (a) The function  $f(\mathbf{x}) = \mathbf{a}^T \mathbf{x} - b$ . (b) The perturbed function along the active-subspace direction. (c) The perturbed function along the principal component analysis direction.

**4. Active-Subspace for Universal Adversarial Attacks.** This section investigates how to generate a universal adversarial attack by the active-subspace method. Given a function  $f(\mathbf{x})$ , the maximal perturbation direction is defined by

$$(4.1) \quad \mathbf{v}_\delta^* = \operatorname{argmax}_{\|\mathbf{v}\|_2 \leq \delta} \mathbb{E}_{\mathbf{x}}[(f(\mathbf{x} + \mathbf{v}) - f(\mathbf{x}))^2].$$

Here,  $\delta$  is a user-defined perturbation upper bound. By the first order Taylor expansion, we have  $f(\mathbf{x} + \mathbf{v}) \approx f(\mathbf{x}) + \nabla f(\mathbf{x})^T \mathbf{v}$ , and problem (4.1) can be reduced to

$$(4.2) \quad \mathbf{v}_{AS} = \operatorname{argmax}_{\|\mathbf{v}\|_2=1} \mathbb{E}_{\mathbf{x}}[(\nabla f(\mathbf{x})^T \mathbf{v})^2] = \operatorname{argmax}_{\|\mathbf{v}\|_2=1} \mathbf{v}^T \mathbb{E}_{\mathbf{x}}[\nabla f(\mathbf{x}) \nabla f(\mathbf{x})^T] \mathbf{v}.$$

The vector  $\mathbf{v}_{AS}$  is exactly the dominant eigenvector of the covariance matrix of  $\nabla f(\mathbf{x})$ . The solution for (4.1) can be approximated by  $+\delta \mathbf{v}_{AS}$  or  $-\delta \mathbf{v}_{AS}$ . Here, both  $\mathbf{v}_{AS}$  and  $-\mathbf{v}_{AS}$  are solutions of (4.2) but their effect on (4.1) are different.

*Example 4.1.* Consider a two-dimensional function  $f(\mathbf{x}) = \mathbf{a}^T \mathbf{x} - b$  with  $\mathbf{a} = [1, -1]^T$  and  $b = 1$ , and  $\mathbf{x}$  follows a uniform distribution in a two-dimensional square domain  $[0, 1]^2$ , as shown in Fig. 4 (a). It follows from direct computations that  $\nabla f(\mathbf{x}) = \mathbf{a}$  and the covariance matrix  $\mathbf{C} = \mathbf{a} \mathbf{a}^T$ . The dominant eigenvector of  $\mathbf{C}$  or the active-subspace direction is  $\mathbf{v}_{AS} = \mathbf{a} / \|\mathbf{a}\|_2 = [1/\sqrt{2}, -1/\sqrt{2}]$ . We apply  $\mathbf{v}_{AS}$  to perturb  $f(\mathbf{x})$  and plot  $f(\mathbf{x} + \delta \mathbf{v}_{AS})$  in Fig. 4 (b), which shows a significant difference even for a small permutation  $\delta = 0.3$ . Furthermore, we plot the perturbed function along the first principal component direction  $\mathbf{w}_1 = [1/\sqrt{2}, 1/\sqrt{2}]^T$  in Fig. 4 (c). Here,  $\mathbf{w}_1$  is the eigenvector of the covariance matrix  $\mathbb{E}_{\mathbf{x}}[\mathbf{x} \mathbf{x}^T] = \begin{bmatrix} 1/3 & 1/4 \\ 1/4 & 1/3 \end{bmatrix}$ . However,  $\mathbf{w}_1$  does not result in any perturbation because  $\mathbf{a}^T \mathbf{w}_1 = 0$ . This example indicates the difference between the active-subspace and principal component analysis: the active-subspace direction can capture the sensitivity information of  $f(\mathbf{x})$  whereas the principal component is independent of  $f(\mathbf{x})$ .

**4.1. Universal Perturbation of Deep Neural Networks.** Given a dataset  $\mathcal{D}$  and a classification function  $j(\mathbf{x})$  that maps an input sample to an output label. The universal perturbation seeks for a vector  $\mathbf{v}^*$  whose norm is upper bounded by  $\delta$ , such that the class label can be perturbed with a high probability, i.e.,

$$(4.3) \quad \mathbf{v}^* = \underset{\|\mathbf{v}\| \leq \delta}{\operatorname{argmax}} \operatorname{prob}_{\mathbf{x} \in \mathcal{D}} [j(\mathbf{x} + \mathbf{v}) \neq j(\mathbf{x})] = \underset{\|\mathbf{v}\| \leq \delta}{\operatorname{argmax}} \mathbb{E}_{\mathbf{x}} [1_{j(\mathbf{x}+\mathbf{v}) \neq j(\mathbf{x})}],$$

where  $1_d$  equals one if the condition  $d$  is satisfied and zero otherwise. Solving problem (4.3) directly is challenging because both  $1_d$  and  $j(\mathbf{x})$  are discontinuous. By replacing  $j(\mathbf{x})$  with the loss function  $c(\mathbf{x}) = \operatorname{loss}(f(\mathbf{x}))$  and the indicator function  $1_d$  with a quadratic function, we reformulate problem (4.3) as

$$(4.4) \quad \max_{\mathbf{v}} \mathbb{E}_{\mathbf{x}} [(c(\mathbf{x} + \mathbf{v}) - c(\mathbf{x}))^2] \quad \text{s.t.} \quad \|\mathbf{v}\|_2 \leq \delta.$$

The ball-constrained optimization problem (4.4) can be solved by various numerical techniques such as the spectral gradient descent method [6] and the limited-memory projected quasi-Newton [51]. However, these methods can only guarantee convergence to a local stationary point. Instead, we are interested in computing a direction that can achieve a better objective value by a heuristic algorithm.

**4.2. Recursive Projection Method.** Using the first order Taylor expansion  $c(\mathbf{x} + \mathbf{v}) \approx c(\mathbf{x}) + \mathbf{v}^T \nabla c(\mathbf{x})$ , we reformulate problem (4.4) as a ball constrained quadratic problem

$$(4.5) \quad \max_{\mathbf{v}} \mathbf{v}^T \mathbb{E}_{\mathbf{x}} [\nabla c(\mathbf{x}) \nabla c(\mathbf{x})^T] \mathbf{v} \quad \text{s.t.} \quad \|\mathbf{v}\|_2 \leq \delta.$$

Problem (4.5) is easy to solve because its closed-form solution is exactly the dominant eigenvector of the covariance matrix  $\mathbf{C} = \mathbb{E}_{\mathbf{x}} [\nabla c(\mathbf{x}) \nabla c(\mathbf{x})^T]$  or the first active-subspace direction. However, the dominant eigenvector in (4.5) may not be efficient because  $c(\mathbf{x})$  is nonlinear. Therefore, we compute  $\mathbf{v}$  recursively by

$$(4.6) \quad \mathbf{v}^{k+1} = \operatorname{proj}(\mathbf{v}^k + s^k d_{\mathbf{v}}^k),$$

where  $\operatorname{proj}(\mathbf{v}) = \mathbf{v} \times \min(1, \delta/\|\mathbf{v}\|_2)$ ,  $s^k$  is the stepsize, and  $d_{\mathbf{v}}^k$  is approximated by

$$(4.7) \quad d_{\mathbf{v}}^k = \underset{d_{\mathbf{v}}}{\operatorname{argmax}} \quad d_{\mathbf{v}}^T \mathbb{E}_{\mathbf{x}} \left[ \nabla c(\mathbf{x} + \mathbf{v}^k) \nabla c(\mathbf{x} + \mathbf{v}^k)^T \right] d_{\mathbf{v}}, \quad \text{s.t.} \quad \|d_{\mathbf{v}}\|_2 \leq 1.$$

Namely,  $d_{\mathbf{v}}^k$  is the dominant eigenvector of  $\mathbf{C}^k = \mathbb{E}_{\mathbf{x}} \left[ \nabla c(\mathbf{x} + \mathbf{v}^k) \nabla c(\mathbf{x} + \mathbf{v}^k)^T \right]$ . Because  $d_{\mathbf{v}}^k$  maximizes the changes in  $\mathbb{E}_{\mathbf{x}} [(c(\mathbf{x} + \mathbf{v} + d_{\mathbf{v}}) - c(\mathbf{x} + \mathbf{v}))^2]$ , we expect that the attack ratio keeps increasing, i.e.,  $r(\mathbf{v}^{k+1}; \mathcal{D}) \geq r(\mathbf{v}^k; \mathcal{D})$ , where

$$(4.8) \quad r(\mathbf{v}; \mathcal{D}) = \frac{1}{|\mathcal{D}|} \sum_{\mathbf{x}^i \in \mathcal{D}} 1_{j(\mathbf{x}^i + \mathbf{v}) \neq j(\mathbf{x}^i)}.$$

The backtracking line search approach [3] is employed to choose  $s^k$  such that the attack ratio of  $\mathbf{v}^k + s^k d_{\mathbf{v}}^k$  is higher than the attack ratio of both  $\mathbf{v}^k$  and  $\mathbf{v}^k - s^k d_{\mathbf{v}}^k$ , i.e.,

$$(4.9) \quad s^k = \min_i \{s_{i,t}^k : r(\mathbf{v}_{i,t}^{k+1}; \mathcal{D}) > \max(r(\mathbf{v}_{i,-t}^{k+1}; \mathcal{D}), r(\mathbf{v}^k; \mathcal{D}))\},$$

**Algorithm 4.1** Recursive Active Subspace Universal Attack

**Input:** A pre-trained deep neural network denoted as  $c(\mathbf{x})$ , a classification oracle  $j(\mathbf{x})$ , a training dataset  $\mathcal{D}^0$ , an upper bound for the attack vector  $\delta$ , an initial stepsize  $s_0$ , a decrease ratio  $\gamma < 1$ , and the parameter in the stopping criterion  $\alpha$ .

- 1: Initialize the attack vector as  $\mathbf{v}^0 = 0$ .
- 2: **for**  $k = 0, 1, \dots$  **do**
- 3:   Select the training dataset as  $\mathcal{D} = \{\mathbf{x}^i + \mathbf{v}^k : \mathbf{x}^i \in \mathcal{D}^0 \text{ and } j(\mathbf{x}^i + \mathbf{v}^k) = j(\mathbf{x}^i)\}$ , then compute the dominate active subspace direction  $d\mathbf{v}$  by Algorithm 3.2.
- 4:   **for**  $i = 0, 1, \dots, I$  **do**
- 5:     Let  $s_{i,\pm}^k = (-1)^\pm s_0 \gamma^i$  and  $\mathbf{v}_{i,\pm}^{k+1} = \text{proj}(\mathbf{v}^k + s_{i,\pm}^{k+1} d_{\mathbf{v}}^k)$ . Compute the attack ratios  $r(\mathbf{v}_{i,1}^{k+1})$  and  $r(\mathbf{v}_{i,-1}^{k+1})$  by (4.8).
- 6:     If either  $r(\mathbf{v}_{i,1}^{k+1})$  or  $r(\mathbf{v}_{i,-1}^{k+1})$  is greater than  $r(\mathbf{v}^k)$ , stop the process. Return  $s^k = (-1)^t s_{i,1}^k$ , where  $t = 1$  if  $r(\mathbf{v}_{i,1}^{k+1}) \geq r(\mathbf{v}_{i,-1}^{k+1})$  and  $t = -1$  otherwise.
- 7:   **end for**
- If no stepsize  $s^k$  is returned, let  $s^k = s_0 r^I$  and record this step as a failure. Compute the next iteration  $\mathbf{v}^{k+1}$  by the projection (4.6).
- 8:   If the number of failure is greater the threshold  $\alpha$ , stop.
- 9: **end for**

**Output:** The universal active adversarial attack vector  $\mathbf{v}_{AS}$ .

where  $s_{i,t}^k = (-1)^t s_0 \gamma^i$ ,  $t \in \{1, -1\}$ ,  $s_0$  is the initial stepsize,  $\gamma < 1$  is the decrease ratio, and  $\mathbf{v}_{i,t}^{k+1} = \text{proj}(\mathbf{v}^k + s_{i,t}^{k+1} d_{\mathbf{v}}^k)$ . If such a stepsize  $s^k$  exists, we update  $\mathbf{v}^{k+1}$  by (4.6) and repeat the process. Otherwise, we record the number of failures and stop the algorithm when the number of failure is greater than a threshold.

The overall flow is summarized in Algorithm 4.1. In practice, instead of using the whole dataset to train this attack vector, we use a subset  $\mathcal{D}^0$ . The impact for different number of samples is discussed in section 5.2.2.

**5. Numerical Experiments.** In this section, we show the power of active-subspace in revealing the number of active neurons, compressing neural networks, and computing the universal adversarial perturbation. All codes are implemented in PyTorch and are available online<sup>2</sup>.

**5.1. Structural Analysis and Compression.** We test the ASNet constructed by Algorithm 3.1, and set the polynomial order as  $p = 2$ , the number of active neurons as  $r = 50$ , and the threshold in Equation (3.4) as  $\epsilon = 0.05$  on default. Inspired by the knowledge distillation [28], we retrain all the parameters in the ASNet by minimizing the following loss function

$$\min_{\theta} \sum_{i=1}^m \beta H(\text{ASNet}_{\theta}(\mathbf{x}_0^i), f(\mathbf{x}_0^i)) + (1 - \beta) H(\text{ASNet}_{\theta}(\mathbf{x}_0^i), \mathbf{y}^i).$$

<sup>2</sup><https://github.com/chunfenge/ASNet>

Table 1: Comparison of number of neurons  $r$  of VGG-19 on CIFAR-10. For the storage speedup, the higher is better. For the accuracy reduction before or after finetuning, the lower is better.

	$r = 25$				$r = 50$				$r = 75$			
	$\epsilon$	Storage	Accu. Reduce		$\epsilon$	Storage	Accu. Reduce		$\epsilon$	Storage	Accu. Reduce	
			Before	After			Before	After			Before	After
ASNet(5)	0.34	<b>20.7</b> $\times$	7.06	2.82	0.18	14.4 $\times$	4.40	1.82	0.11	11.0 $\times$	<b>3.64</b>	<b>1.66</b>
ASNet(6)	0.24	<b>12.8</b> $\times$	2.14	0.59	0.11	10.1 $\times$	1.62	0.27	0.05	8.3 $\times$	<b>1.40</b>	<b>0.21</b>
ASNet(7)	0.15	<b>9.3</b> $\times$	0.79	0.11	0.06	7.8 $\times$	<b>0.63</b>	<b>-0.10</b>	0.03	6.7 $\times$	0.77	0.00

Here, the cross entropy  $H(\mathbf{p}, \mathbf{q}) = \sum_j s(\mathbf{p})_j \log s(\mathbf{q})_j$ , the softmax function  $s(\mathbf{x})_j = \frac{\exp(x_j)}{\sum_j \exp(x_j)}$ , and the parameter  $\beta = 0.1$  on default. We retrain ASNet for 50 epochs by ADAM [34]. The stepsizes for the pre-model are set as  $10^{-4}$  and  $10^{-3}$  for VGG-19 and ResNet, and the stepsize for the active subspace layer and the polynomial chaos expansion layer is set as  $10^{-5}$ , respectively,

We also seek for sparser weights in ASNet by the proximal stochastic gradient descent method in Section 3.6. On default, we set the stepsize as  $10^{-4}$  for the pre-model and  $10^{-5}$  for the active subspace layer and the polynomial chaos expansion layer. The maximal epoch is set as 100. The obtained sparse model is denoted as ASNet-s.

In all figures and tables, the numbers in the bracket of ASNet( $\cdot$ ) or ASNet-s( $\cdot$ ) indicate the index of a cut-off layer. We report the performance for different cut-off layers in terms of *accuracy, storage, and computational complexities*.

**5.1.1. Choices of Parameters.** We first show the influence of number of reduced neurons  $r$ , tolerance  $\epsilon$ , and cutting-off layer index  $l$  of VGG-19 on CIFAR-10 in Table 1. The VGG-19 can achieve 93.28% testing accuracy with 76.45 Mb storage consumption. Here,  $\epsilon = \frac{\lambda_{r+1} + \dots + \lambda_n}{\lambda_1 + \dots + \lambda_n}$ . For different choices of  $r$ , we display the corresponding tolerance  $\epsilon$ , the storage speedup compared with the original teacher network, and the testing accuracy reduction for ASNet before and after fine-tuning compared with the original teacher network.

Table 1 shows that when the cutting-off layer is fixed, a larger  $r$  usually results in a smaller tolerance  $\epsilon$  and a smaller accuracy reduction but also a smaller storage speedup. This is corresponding to Lemma 3.1 that the error of ASNet before fine-tuning is upper bounded by  $O(\epsilon)$ . Comparing  $r = 50$  with  $r = 75$ , we find that  $r = 50$  can achieve almost the same accuracy with  $r = 75$  with a higher storage speedup.  $r = 50$  can even achieve better accuracy than  $r = 75$  in layer 7 probably because of overfitting. This guides us to chose  $r = 50$  in the following numerical experiments. For different layers, we see a later cutting-off layer index can produce a lower accuracy reduction but a smaller storage speedup. In other words, the choice of layer index is a trade-off between accuracy reduction with storage speedup.

**5.1.2. Efficiency of Active-subspace.** We show the effectiveness of ASNet constructed by Steps 1-3 of Algorithm 3.1 without fine-tuning. We investigate the following three properties. (1) **Redundancy of neurons.** The distributions of the first 200 singular values of the matrix  $\hat{\mathbf{G}}$  (defined in (3.7)) are plotted in Fig. 5 (a). The singular values decrease almost

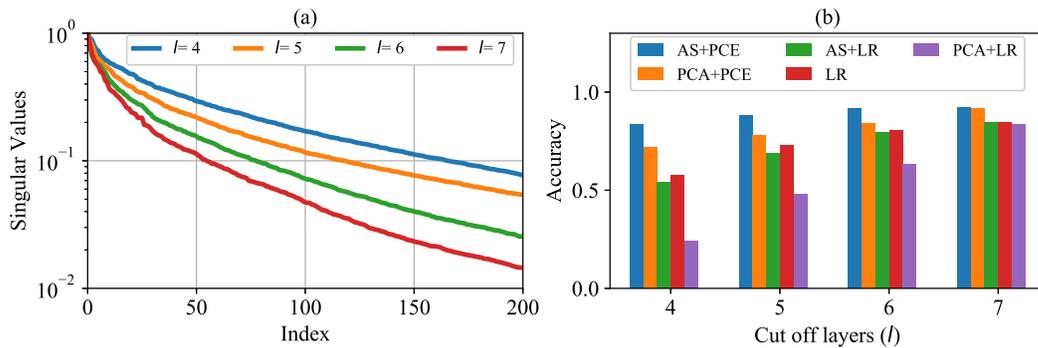


Figure 5: Structural analysis of VGG-19 on the CIFAR-10 dataset. (a) The first 200 singular values for layers  $4 \leq l \leq 7$ ; (b) The accuracy (without any fine-tuning) obtained by active-subspace (AS) and polynomial chaos expansions (PCE) compared with principal component analysis (PCA) and logistic regression (LR).

exponentially for layers  $l \in \{4, 5, 6, 7\}$ . Although the total numbers of neurons are 8192, 16384, 16384, and 16384, the numbers of active neurons are only 105, 84, 54, and 36, respectively. (2) **Redundancy of the layers.** We cut off the deep neural network at an intermediate layer and replace the subsequent layers with one simple logistic regression [30]. As shown by the red bar in Fig. 5 (b), the logistic regression can achieve relatively high accuracy. This verifies that the features trained from the first few layers already have a high expression power since replacing all subsequent layers with a simple expression loses little accuracy. (3) **Efficiency of the active-subspace and polynomial chaos expansion.** We compare the proposed active-subspace layer with the principal component analysis [31] in projecting the high-dimensional neuron to a low-dimensional space, and also compare the polynomial chaos expansion layer with logistic regression in terms of their efficiency to extract class labels from the low-dimensional variables. Fig. 5 (b) shows that the combination of active-subspace and polynomial chaos expansion can achieve the best accuracy.

**5.1.3. CIFAR-10.** We continue to present the results of ASNet and ASNet-s on CIFAR-10 by two widely used networks: VGG-19 and ResNet-110 in Tables 2 and 3, respectively. The second column shows the testing accuracy for the corresponding network. We report the storage and computational costs for the pre-model, post-model (i.e., active-subspace plus polynomial chaos expansion for ASNet and ASNet-s), and overall results, respectively. For both examples, ASNet and ASNet-s can achieve a similar accuracy with the teacher network yet with much smaller storage and computational cost. For VGG-19, ASNet achieves  $14.43\times$  storage savings and  $3.44\times$  computational reduction; ASNet-s achieves  $23.98\times$  storage savings and  $7.30\times$  computational reduction. For most ASNet and ASNet-s networks, the storage and computational costs of the post-models achieve significant performance boosts by our proposed network structure changes. It is not surprising to see that increasing the layer index (i.e., cutting off the deep neural network at a later layer) can produce a higher accuracy. However, increasing the layer index also results in a smaller compression ratio. In other words, the choice of layer index is a trade-off between the accuracy reduction with the compression

Table 2: Accuracy and storage on VGG-19 for CIFAR-10. Here, “Pre-M” denotes the pre-model, i.e., layers 1 to  $l$  of the original deep neural networks, “AS” and “PCE” denote the active subspace and polynomial chaos expansion layer, respectively.

Network	Accuracy	Storage (MB)			Flops ( $10^6$ )		
		Pre-M	AS+PCE	Overall	Pre-M	AS+PCE	Overall
VGG-19	93.28%	76.45			398.14		
ASNet(5)	91.46%	2.12	3.18	5.30	115.02	0.83	115.85
			(23.41×)	(14.43×)		(340.11×)	(3.44×)
ASNet-s(5)	90.40%	1.14	2.05	3.19	54.03	0.54	54.56
		(1.86×)	(36.33×)	(23.98×)	(2.13×)	(527.91×)	(7.30×)
ASNet(6)	93.01%	4.38	3.18	7.55	152.76	0.83	153.60
			(22.70×)	(10.12×)		(294.76×)	(2.59×)
ASNet-s(6)	91.08%	1.96	1.81	3.77	67.37	0.48	67.85
		(2.24×)	(39.73×)	(20.27×)	(2.27×)	(515.98×)	(5.87×)
ASNet(7)	93.38%	6.63	3.18	9.80	190.51	0.83	191.35
			(21.99×)	(7.80×)		(249.41×)	(2.08×)
ASNet-s(7)	90.87%	2.61	1.91	4.52	80.23	0.50	80.73
		(2.54×)	(36.64×)	(16.92×)	(2.37×)	(415.68×)	(4.93×)

Table 3: Accuracy and storage on ResNet-110 for CIFAR-10. Here, “Pre-M” denotes the pre-model, i.e., layers 1 to  $l$  of the original deep neural networks, “AS” and “PCE” denote the active subspace and polynomial chaos expansion layer, respectively.

Network	Accuracy	Storage (MB)			Flops ( $10^6$ )		
		Pre-M	AS+PCE	Overall	Pre-M	AS+PCE	Overall
ResNet-110	93.78%	6.59			252.89		
ASNet(61)	89.56%	1.15	1.61	2.77	140.82	0.42	141.24
			(3.37×)	(2.38×)		(265.03×)	(1.79×)
ASNet-s(61)	89.26%	0.83	1.23	2.06	104.05	0.32	104.37
		(1.39×)	(4.41×)	(3.19×)	(1.35×)	(346.82×)	(2.42×)
ASNet(67)	90.16%	1.37	1.61	2.98	154.98	0.42	155.40
			(3.24×)	(2.21×)		(231.55×)	(1.63×)
ASNet-s(67)	89.69%	1.00	1.22	2.22	116.38	0.32	116.70
		(1.36×)	(4.29×)	(2.97×)	(1.33×)	(306.72×)	(2.17×)
ASNet(73)	90.48%	1.58	1.61	3.19	169.13	0.42	169.55
			(3.11×)	(2.06×)		(198.07×)	(1.49×)
ASNet-s(73)	90.02%	1.18	1.16	2.34	128.65	0.30	128.96
		(1.34×)	(4.32×)	(2.82×)	(1.31×)	(275.74×)	(1.96×)

ratio.

For Resnet-110, our results are not as good as those on VGG-19. We find that the eigenvalues for its covariance matrix are not exponentially decreasing as that of VGG-19, which results in a large number of active neurons or a large error  $\epsilon$  when fixing  $r = 50$ . A possible reason is that ResNet updates as  $\mathbf{x}_{l+1} = \mathbf{x}_l + f_l(\mathbf{x}_l)$ . Hence, the partial gradient  $\partial \mathbf{x}_{l+1} / \partial \mathbf{x}_l = I + \nabla f_l(\mathbf{x}_l)$  is less likely to be low-rank.

Table 4: Accuracy and storage on VGG-19 for CIFAR-100. Here, “Pre-M” denotes the pre-model, i.e., layers 1 to  $l$  of the original deep neural networks, “AS” and “PCE” denote the active subspace and polynomial chaos expansion layer, respectively.

Network	Top-1	Top-5	Storage (MB)			Flops ( $10^6$ )		
			Pre-M	AS+PCE	Overall	Pre-M	AS+PCE	Overall
VGG-19	71.90%	89.57%	76.62			398.18		
ASNet(7)	70.77%	91.05%	6.63	3.63 (19.23×)	10.26 (7.45×)	190.51	0.83 (249.41×)	191.35 (2.08×)
ASNet-s(7)	70.20%	90.90%	5.20 (1.27×)	3.24 (21.56×)	8.44 (9.06×)	144.81 (1.32×)	0.85 (244.57×)	145.66 (2.73×)
ASNet(8)	69.50%	90.15%	8.88	1.29 (52.50×)	10.17 (7.52×)	228.26	0.22 (779.04×)	228.48 (1.74×)
ASNet-s(8)	69.17%	89.73%	6.87 (1.29×)	1.22 (55.36×)	8.09 (9.45×)	172.69 (1.32×)	0.32 (530.92×)	173.01 (2.30×)
ASNet(9)	<b>72.00%</b>	<b>90.61%</b>	13.39	2.07 (30.49×)	15.46 (4.95×)	247.14	0.42 (357.10×)	247.56 (1.61×)
ASNet-s(9)	71.38%	90.28%	9.38 (1.43×)	1.94 (32.49×)	11.32 (6.75×)	183.27 (1.35×)	0.51 (296.74×)	183.78 (2.17×)

Table 5: Accuracy and storage on ResNet-110 for CIFAR-100. Here, “Pre-M” denotes the pre-model, i.e., layers 1 to  $l$  of the original deep neural networks, “AS” and “PCE” denote the active subspace and polynomial chaos expansion layer, respectively.

Network	Top-1	Top-5	Storage (MB)			Flops ( $10^6$ )		
			Pre-M	AS+PCE	Overall	Pre-M	AS+PCE	Overall
ResNet-110	71.94%	91.71 %	6.61			252.89		
ASNet(75)	63.01%	88.55%	1.79	1.29 (3.73×)	3.08 (2.14×)	172.67	0.22 (367.88×)	172.89 (1.46×)
ASNet-s(75)	63.16%	88.65%	1.47 (1.22×)	1.20 (3.99×)	2.67 (2.46×)	143.11 (1.21×)	0.31 (254.69×)	143.42 (1.76×)
ASNet(81)	65.82%	90.02%	2.64	1.29 (3.07×)	3.93 (1.68×)	186.83	0.22 (302.96×)	187.04 (1.35×)
ASNet-s(81)	65.73%	89.95%	2.20 (1.20×)	1.21 (3.27×)	3.41 (1.93×)	155.61 (1.20×)	0.32 (208.38×)	155.93 (1.62×)
ASNet(87)	<b>67.71%</b>	<b>90.17%</b>	3.48	1.29 (2.41×)	4.77 (1.38×)	200.98	0.22 (238.04×)	201.20 (1.26×)
ASNet-s(87)	67.65%	90.10%	2.91 (1.20×)	1.21 (2.56×)	4.12 (1.60×)	166.50 (1.21×)	0.32 (163.50×)	166.81 (1.52×)

**5.1.4. CIFAR-100.** Next, we present the results of VGG-19 and ResNet-110 on CIFAR-100 in Tables 4 and 5, respectively. On VGG-19, ASNet can achieve  $7.45\times$  storage savings and  $2.08\times$  computational reduction, and ASNet-s can achieve  $9.06\times$  storage savings and  $2.73\times$  computational reduction. The accuracy loss is negligible for VGG-19 but larger for ResNet-110. The performance boost of ASNet is obtained by just changing the network structures and without any model compression (e.g., pruning, quantization, or low-rank factorization).

**5.2. Universal Adversarial Attacks.** This subsection demonstrates the effectiveness of active-subspace in identifying a universal adversarial attack vector. We denote the result generated by Algorithm 4.1 as “AS” and compare it with the “UAP” method in [40] and with “random” Gaussian distribution vector. The parameters in Algorithm 4.1 are set as  $\alpha = 10$  and  $\delta = 5, \dots, 10$ . The default parameters of UAP are applied except for the maximal iteration. In the implementation of [40], the maximal iteration is set as infinity, which is time-consuming when the training dataset or the number of classes is large. In our experiments, we set the maximal iteration as 10. In all figures and tables, we report the average attack ratio and CPU time in training out of ten repeated experiments with different training datasets. A higher attack ratio means the corresponding algorithm is better in fooling the given deep neural network. The datasets are chosen in two ways. We firstly test data points from one class (e.g., trousers in Fashion-MNIST) because these data points share lots of common features and have a higher probability to be attacked by a universal perturbation vector. We then conduct experiments on the whole dataset to show our proposed algorithm can also provide better performance compared with the baseline even if the dataset has diverse features.

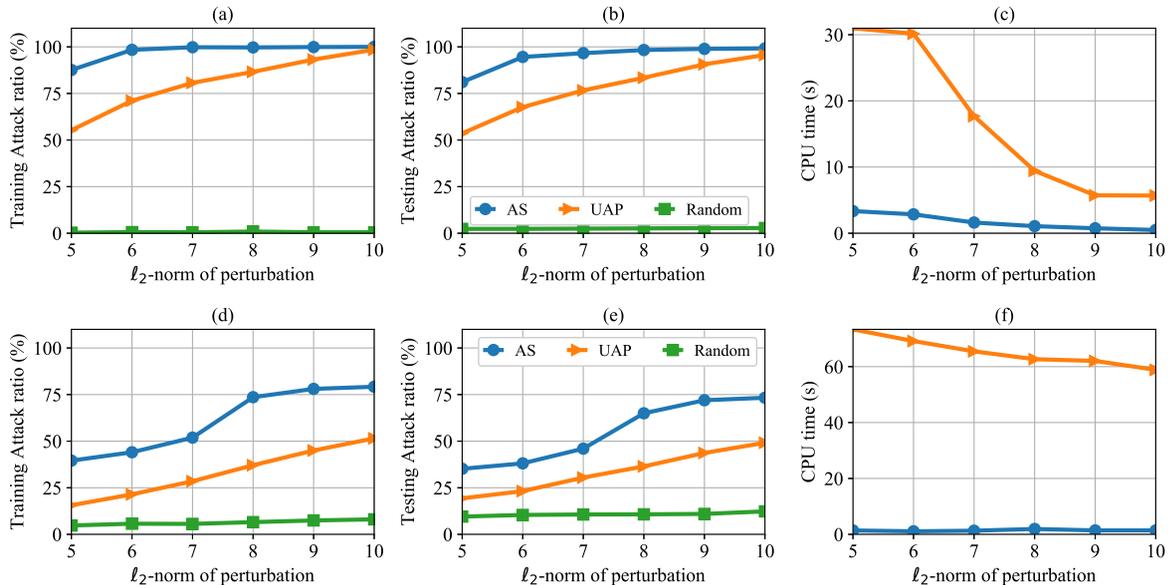


Figure 6: Universal adversarial attacks for the Fashion-MNIST with respect to different  $\ell_2$ -norms. (a)-(c): the results for attacking one class dataset. (d)-(f): the results for attacking the whole dataset.

**5.2.1. Fashion-MNIST.** Firstly, we present the adversarial attack result on Fashion-MNIST by a 4-layer neural network. There are two convolutional layers with kernel size equals  $5 \times 5$ . The size of output channels for each convolutional layer is 20 and 50, respectively. Each convolutional layer is followed by a ReLU activation layer and a max-pooling layer with a kernel size of  $2 \times 2$ . There are two fully connected layers. The first fully connected layer has an input feature 800 and an output feature 500.

Fig. 6 presents the attack ratio of our active-subspace method compared with the baselines

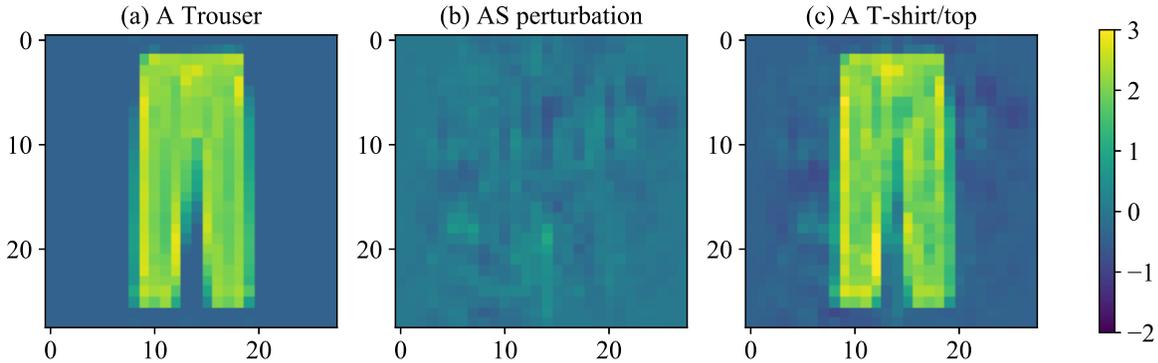


Figure 7: The effect of our attack method on one data sample in the Fashion-MNIST dataset. (a) A trouser from the original dataset. (b) An active-subspace perturbation vector with the  $\ell_2$  norm equals to 5. (c) The perturbed sample is misclassified as a t-shirt/top by the deep neural network.

UAP method [40] and Gaussian random vectors. The top figures show the results for just one class (i.e., trouser), and the bottom figures show the results for all ten classes. For all perturbation norms, the active-subspace method can achieve around 30% higher attack ratio than UAP while more than 10 times faster. This verifies that the active-subspace method has better universal representation ability compared with UAP because the active-subspace can find a universal direction while UAP solves data-dependent subproblems independently. By the active-subspace approach, the attack ratio for the first class and the whole dataset are around 100% and 75%, respectively. This coincides with our intuition that the data points in one class have higher similarity than data points from different classes.

In Fig. 7, we plot one image from Fashion-MNIST and its perturbation by the active-subspace attack vector. The attacked image in Fig. 7 (c) still looks like a trouser for a human. However, the deep neural network misclassifies it as a t-shirt/top.

**5.2.2. CIFAR-10.** Next, we show the numerical results of attacking VGG-19 on CIFAR-10. Fig. 8 compares the active-subspace method compared with the baseline UAP and Gaussian random vectors. The top figures show the results by the dataset in the first class (i.e., automobile), and the bottom figures show the results for all ten classes. For both two cases, the proposed active-subspace attack can achieve 20% higher attack ratios while three times faster than UAP. This is similar to the results in Fashion-MNIST because the active-subspace has a better ability to capture the global information.

We further show the effects of *different number of training samples* in Fig. 9. When the number of samples is increased, the testing attack ratio is getting better. In our numerical experiments, we set the number of samples as 100 for one-class experiments and 200 for all-classes experiments.

We continue to show the *cross-model* performance on four different ResNet networks and one VGG network. We test the performance of the attack vector trained from one model on all other models. Each row in Table 6 shows the results on the same deep neural network and each column shows the results of the same attack vector. It shows that ResNet-20 is easier

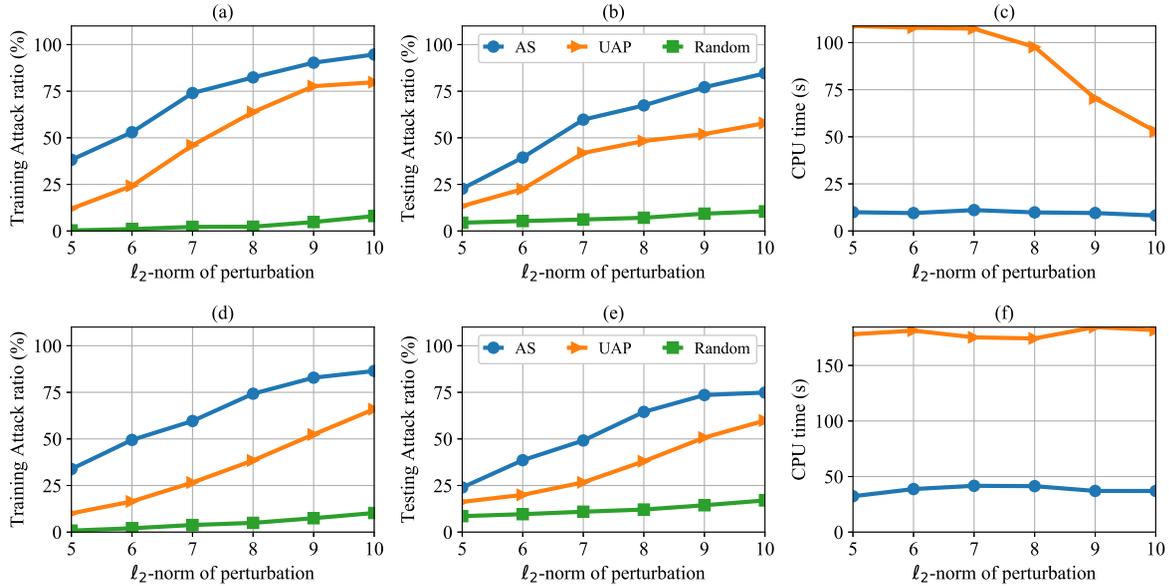


Figure 8: Universal adversarial attacks of VGG-19 on CIFAR-10 with respect to different  $\ell_2$ -norm perturbations. (a)-(c): The training attack ratio, the testing attack ratio, and the CPU time in seconds for attacking one class dataset. (d)-(f): The results for attacking ten classes dataset together.

Table 6: Cross-model performance for CIFAR-10

	ResNet-20	ResNet-44	ResNet-56	ResNet-110	VGG-19
ResNet-20	<b>91.35%</b>	87.74%	86.28%	87.38%	81.16%
ResNet-44	84.75%	<b>92.28%</b>	87.03%	85.44%	83.44%
ResNet-56	83.63%	86.67%	<b>90.15%</b>	87.39%	84.38%
ResNet-110	71.02%	77.58%	74.19%	<b>92.77%</b>	77.32%
VGG-19	53.61%	59.74%	61.49%	66.29%	<b>80.02%</b>

to be attacked compared with other models. This agrees with our intuition that a simple network structure such as ResNet-20 is less robust. On the contrary, VGG-19 is the most robust. The success of cross-model attacks indicates that these neural networks could find a similar feature.

**5.2.3. CIFAR-100.** Finally, we show the results on CIFAR-100 for both the first class (i.e., dolphin) and all classes. Similar to Fashion-MNIST and CIFAR-10, Fig. 10 shows that active-subspace can achieve higher attack ratios than both UAP and Gaussian random vectors. Further, compared with CIFAR-10, CIFAR-100 is easier to be attacked partially because it has more classes.

We summarize the results for different datasets in Table 7. The second column shows the number of classes in the dataset. In terms of testing attack ratio for the whole dataset, active-

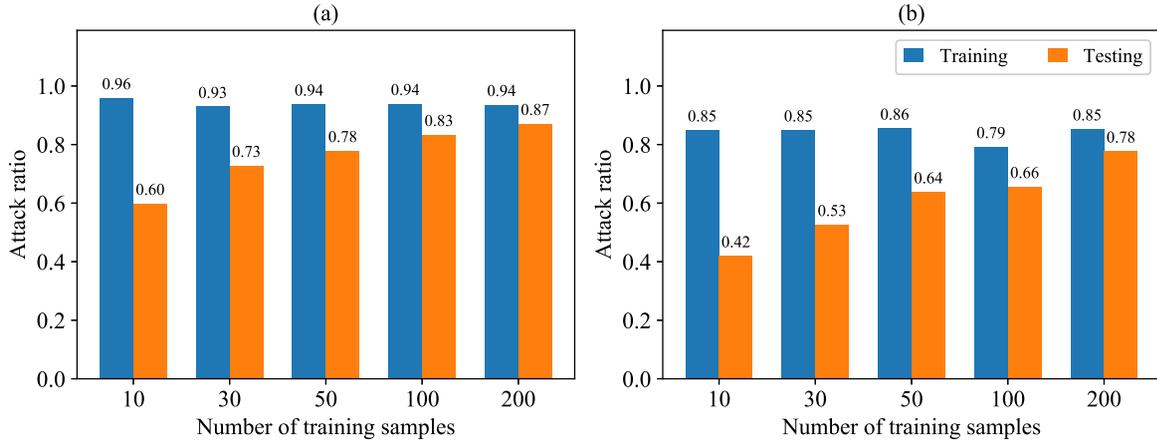


Figure 9: Adversarial attack of VGG-19 on CIFAR-10 with different number of training samples. The  $\ell_2$ -norm perturbation is fixed as 10. (a) The results of attacking the dataset from the first class; (b) The results of attacking the whole dataset with 10 classes.

Table 7: Summary of the universal attack for different datasets by the active-subspace compared with UAP and the random vector. The norm of perturbation is equal to 10.

	# Class	Training Attack ratio			Testing Attack ratio			CPU time (s)	
		AS	UAP	Rand	AS	UAP	Rand	AS	UAP
Fashion-MNIST	1	100.0%	93.6%	1.8%	<b>98.0%</b>	91.3%	3.0%	<b>0.15</b>	5.49
	10	79.2%	51.5%	8.0%	<b>73.3%</b>	49.1%	12.3%	<b>1.40</b>	58.85
CIFAR-10	1	94.7%	79.8%	8.0%	<b>84.5%</b>	57.9%	10.6%	<b>8.18</b>	52.83
	10	86.5%	65.9%	10.2%	<b>74.9%</b>	59.9%	17.0%	<b>37.01</b>	181.72
CIFAR-100	1	97.2%	87.9%	19.7%	<b>92.1%</b>	84.3%	37.9%	<b>13.32</b>	248.78
	100	93.7%	86.5%	38.7%	<b>83.5%</b>	77.4%	52.0%	<b>14.32</b>	204.50

subspace achieves 24.2%, 15%, and 6.1% higher attack ratios than UAP for Fashion-MNIST, CIFAR-10, and CIFAR-100, respectively. In terms of the CPU time, active-subspace achieves  $42\times$ ,  $5\times$ , and  $14\times$  speedup than UAP on the Fashion-MNIST, CIFAR-10, and CIFAR-100, respectively.

**6. Conclusions and Discussions.** This paper has analyzed deep neural networks by the active subspace method originally developed for dimensionality reduction of uncertainty quantification. We have investigated two problems: how many neurons and layers are necessary (or important) in a deep neural network, and how to generate a universal adversarial attack vector that can be applied to a set of testing data? Firstly, we have presented a definition of “the number of active neurons” and have shown its theoretical error bounds for model reduction. Our numerical study has shown that many neurons and layers are not needed. Based on this observation, we have proposed a new network called ASNet by cutting off the whole neural

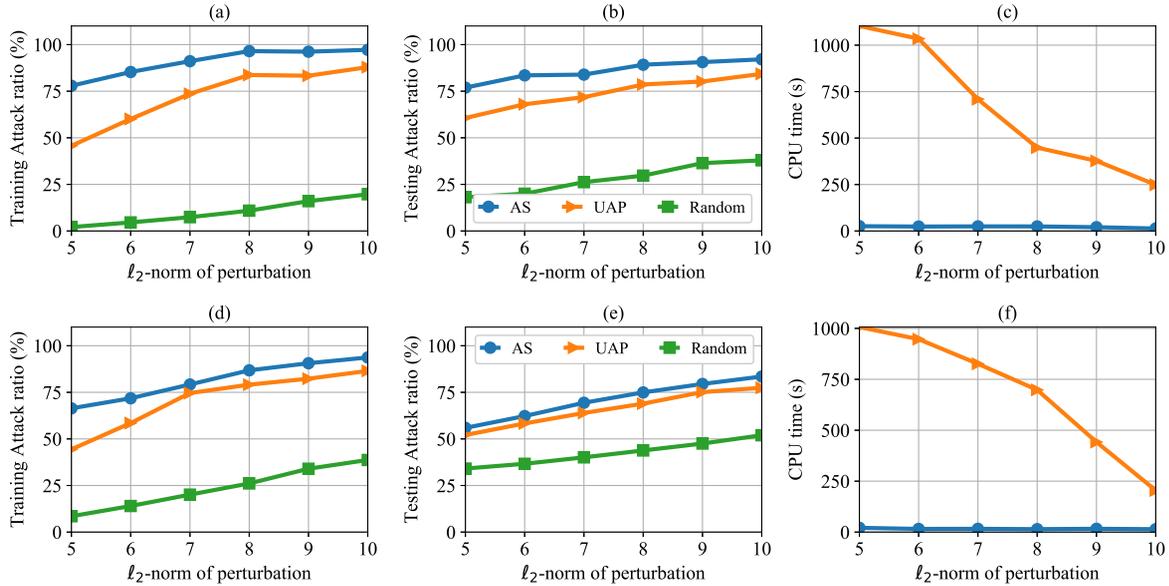


Figure 10: Results for universal adversarial attack for CIFAR-100 with respect to different  $\ell_2$ -norm perturbations. (a)-(c): The results for attacking the dataset from the first class. (d)-(f): The results for attacking ten classes dataset together.

network at a proper layer and replacing all subsequent layers with an active subspace layer and a polynomial chaos expansion layer. The numerical experiments show that the proposed deep neural network structural analysis method can produce a new network with significant storage savings and computational speedup yet with little accuracy loss. Our methods can be combined with existing model compression techniques (e.g., pruning, quantization and low-rank factorization) to develop compact deep neural network models that are more suitable for the deployment on resource-constrained platforms. Secondly, we have applied the active subspace to generate a universal attack vector that is independent of a specific data sample and can be applied to a whole dataset. Our proposed method can achieve a much higher attack ratio than the existing work [40] and enjoys a lower computational cost.

ASNet has two main goals: to detect the necessary neurons and layers, and to compress the existing network. To fulfill the first goal, we require a pre-trained model because from Lemmas 3.1, and 3.2, the accuracy of the reduced model will approach that of the original one. For the second task, the pre-trained model helps us to get a good estimation for the number of active neurons, a proper layer to cut off, and a good initialization for the active subspace layer and polynomial chaos expansion layer. However, a pre-trained model is not required because we can construct ASNet in a heuristic way (as done in most DNN): a reasonable guess for the number of active neurons and cut-off layer, and a random parameter initialization for the pre-model, the active subspace layer and the polynomial chaos expansion layer.

**Acknowledgement.** We thank the associate editor and referees for their valuable comments and suggestions.

## REFERENCES

- [1] S. ABDOLI, L. G. HAFEMANN, J. RONY, I. B. AYED, P. CARDINAL, AND A. L. KOERICH, *Universal adversarial audio perturbations*, arXiv preprint arXiv:1908.03173, (2019).
- [2] A. AGHASI, A. ABDI, N. NGUYEN, AND J. ROMBERG, *Net-trim: Convex pruning of deep neural networks with performance guarantee*, in Advances in Neural Information Processing Systems, 2017, pp. 3177–3186.
- [3] L. ARMIJO, *Minimization of functions having lipschitz continuous first partial derivatives*, Pacific Journal of mathematics, 16 (1966), pp. 1–3.
- [4] S. BALUJA AND I. FISCHER, *Adversarial transformation networks: Learning to generate adversarial examples*, arXiv preprint arXiv:1703.09387, (2017).
- [5] M. BEHJATI, S.-M. MOOSAVI-DEZFOOLI, M. S. BAGHSHAH, AND P. FROSSARD, *Universal adversarial attacks on text classifiers*, in ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2019, pp. 7345–7349.
- [6] E. G. BIRGIN, J. M. MARTÍNEZ, AND M. RAYDAN, *Nonmonotone spectral projected gradient methods on convex sets*, SIAM Journal on Optimization, 10 (2000), pp. 1196–1211.
- [7] H. CAI, L. ZHU, AND S. HAN, *ProxyllessNAS: Direct neural architecture search on target task and hardware*, arXiv preprint arXiv:1812.00332, (2018).
- [8] N. CARLINI AND D. WAGNER, *Towards evaluating the robustness of neural networks*, in 2017 IEEE Symposium on Security and Privacy (SP), IEEE, 2017, pp. 39–57.
- [9] P. G. CONSTANTINE, *Active subspaces: Emerging ideas for dimension reduction in parameter studies*, vol. 2, SIAM, 2015.
- [10] P. G. CONSTANTINE, E. DOW, AND Q. WANG, *Active subspace methods in theory and practice: applications to kriging surfaces*, SIAM Journal on Scientific Computing, 36 (2014), pp. A1500–A1524.
- [11] P. G. CONSTANTINE, M. EMORY, J. LARSSON, AND G. IACCARINO, *Exploiting active subspaces to quantify uncertainty in the numerical simulation of the hyshot ii scramjet*, Journal of Computational Physics, 302 (2015), pp. 1–20.
- [12] M. COURBARIAUX, I. HUBARA, D. SOUDRY, R. EL-YANIV, AND Y. BENGIO, *Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1*, arXiv preprint arXiv:1602.02830, (2016).
- [13] C. CUI AND Z. ZHANG, *Stochastic collocation with non-gaussian correlated process variations: Theory, algorithms and applications*, IEEE Transactions on Components, Packaging and Manufacturing Technology, (2018).
- [14] C. CUI AND Z. ZHANG, *High-dimensional uncertainty quantification of electronic and photonic ic with non-gaussian correlated process variations*, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, (2019).
- [15] L. DENG, P. JIAO, J. PEI, Z. WU, AND G. LI, *Gznr-net: Training deep neural networks with ternary weights and activations without full-precision memory under a unified discretization framework*, Neural Networks, 100 (2018), pp. 49–58.
- [16] G. K. DZIUGAITE, Z. GHAHRAMANI, AND D. M. ROY, *A study of the effect of jpg compression on adversarial images*, arXiv preprint arXiv:1608.00853, (2016).
- [17] J. FRANKLE AND M. CARBIN, *The lottery ticket hypothesis: Finding sparse, trainable neural networks*, arXiv preprint arXiv:1803.03635, (2018).
- [18] T. GARISOV, D. PODOPRIKHIN, A. NOVIKOV, AND D. VETROV, *Ultimate tensorization: compressing convolutional and FC layers alike*, arXiv preprint arXiv:1611.03214, (2016).
- [19] R. GE, R. WANG, AND H. ZHAO, *Mildly overparametrized neural nets can memorize training data efficiently*, arXiv preprint arXiv:1909.11837, (2019).
- [20] R. G. GHANEM AND P. D. SPANOS, *Stochastic finite element method: Response statistics*, in Stochastic Finite Elements: A Spectral Approach, Springer, 1991, pp. 101–119.
- [21] M. GHASHAMI, E. LIBERTY, J. M. PHILLIPS, AND D. P. WOODRUFF, *Frequent directions: Simple and deterministic matrix sketching*, SIAM Journal on Computing, 45 (2016), pp. 1762–1792.
- [22] I. J. GOODFELLOW, J. SHLENS, AND C. SZEGEDY, *Explaining and harnessing adversarial examples*, arXiv preprint arXiv:1412.6572, (2014).
- [23] A. GRAVES, S. FERNÁNDEZ, F. GOMEZ, AND J. SCHMIDHUBER, *Connectionist temporal classification:*

- labelling unsegmented sequence data with recurrent neural networks, in Proceedings of the 23rd international conference on Machine learning, ACM, 2006, pp. 369–376.
- [24] N. HALKO, P.-G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM review, 53 (2011), pp. 217–288.
- [25] S. HAN, H. MAO, AND W. J. DALLY, *Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding*, arXiv preprint arXiv:1510.00149, (2015).
- [26] C. HAWKINS AND Z. ZHANG, *Bayesian tensorized neural networks with automatic rank selection*, arXiv preprint arXiv:1905.10478, (2019).
- [27] Y. HE, J. LIN, Z. LIU, H. WANG, L.-J. LI, AND S. HAN, *Amc: Automl for model compression and acceleration on mobile devices*, in Proceedings of the European Conference on Computer Vision (ECCV), 2018, pp. 784–800.
- [28] G. HINTON, O. VINYALS, AND J. DEAN, *Distilling the knowledge in a neural network*, stat, 1050 (2015), p. 9.
- [29] W. HOEFFDING, *Probability inequalities for sums of bounded random variables*, in The Collected Works of Wassily Hoeffding, Springer, 1994, pp. 409–426.
- [30] D. W. HOSMER JR, S. LEMESHOW, AND R. X. STURDIVANT, *Applied logistic regression*, vol. 398, John Wiley & Sons, 2013.
- [31] I. JOLLIFFE, *Principal component analysis*, in International encyclopedia of statistical science, Springer, 2011, pp. 1094–1096.
- [32] C. KANBAK, S.-M. MOOSAVI-DEZFOOLI, AND P. FROSSARD, *Geometric robustness of deep networks: analysis and improvement*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 4441–4449.
- [33] V. KHRULKOV AND I. OSELEDETS, *Art of singular vectors and universal adversarial perturbations*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 8562–8570.
- [34] D. P. KINGMA AND J. BA, *Adam: A method for stochastic optimization*, arXiv preprint arXiv:1412.6980, (2014).
- [35] A. KRIZHEVSKY, I. SUTSKEVER, AND G. E. HINTON, *Imagenet classification with deep convolutional neural networks*, in Advances in neural information processing systems, 2012, pp. 1097–1105.
- [36] V. LEBEDEV, Y. GANIN, M. RAKHUBA, I. OSELEDETS, AND V. LEMPITSKY, *Speeding-up convolutional neural networks using fine-tuned cp-decomposition*, arXiv preprint arXiv:1412.6553, (2014).
- [37] D. R. LIDE, *Handbook of mathematical functions*, in A Century of Excellence in Measurements, Standards, and Technology, CRC Press, 2018, pp. 135–139.
- [38] L. LIU, L. DENG, X. HU, M. ZHU, G. LI, Y. DING, AND Y. XIE, *Dynamic sparse graph for efficient deep learning*, arXiv preprint arXiv:1810.00859, (2018).
- [39] Z. LIU, M. SUN, T. ZHOU, G. HUANG, AND T. DARRELL, *Rethinking the value of network pruning*, arXiv preprint arXiv:1810.05270, (2018).
- [40] S.-M. MOOSAVI-DEZFOOLI, A. FAWZI, O. FAWZI, AND P. FROSSARD, *Universal adversarial perturbations*, in Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 1765–1773.
- [41] S.-M. MOOSAVI-DEZFOOLI, A. FAWZI, AND P. FROSSARD, *Deepfool: a simple and accurate method to fool deep neural networks*, in Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 2574–2582.
- [42] P. NEEKHARA, S. HUSSAIN, P. PANDEY, S. DUBNOV, J. MCAULEY, AND F. KOUSHANFAR, *Universal adversarial perturbations for speech recognition systems*, arXiv preprint arXiv:1905.03828, (2019).
- [43] A. NOVIKOV, D. PODOPRIKHIN, A. OSOKIN, AND D. P. VETROV, *Tensorizing neural networks*, in Advances in Neural Information Processing Systems, 2015, pp. 442–450.
- [44] S. OYMAK AND M. SOLTANOLKOTABI, *Towards moderate overparameterization: global convergence guarantees for training shallow neural networks*, arXiv preprint arXiv:1902.04674, (2019).
- [45] N. PAPERNOT, P. MCDANIEL, S. JHA, M. FREDRIKSON, Z. B. CELIK, AND A. SWAMI, *The limitations of deep learning in adversarial settings*, in 2016 IEEE European Symposium on Security and Privacy (EuroS&P), IEEE, 2016, pp. 372–387.
- [46] A. ROMERO, N. BALLAS, S. E. KAHOU, A. CHASSANG, C. GATTA, AND Y. BENGIO, *Fitnets: Hints for*

- thin deep nets*, arXiv preprint arXiv:1412.6550, (2014).
- [47] H. L. ROYDEN, *Real Analysis*, Macmillan, 2010.
  - [48] T. M. RUSSI, *Uncertainty quantification with experimental data and complex system models*, PhD thesis, UC Berkeley, 2010.
  - [49] T. N. SAINATH, B. KINGSBURY, V. SINDHWANI, E. ARISOY, AND B. RAMABHADRAN, *Low-rank matrix factorization for deep neural network training with high-dimensional output targets*, in IEEE international conference on acoustics, speech and signal processing, 2013, pp. 6655–6659.
  - [50] S. SCARDAPANE, D. COMMINELO, A. HUSSAIN, AND A. UNCINI, *Group sparse regularization for deep neural networks*, Neurocomputing, 241 (2017), pp. 81–89.
  - [51] M. SCHMIDT, E. BERG, M. FRIEDLANDER, AND K. MURPHY, *Optimizing costly functions with simple constraints: A limited-memory projected quasi-newton algorithm*, in Artificial Intelligence and Statistics, 2009, pp. 456–463.
  - [52] A. C. SERBAN AND E. POLL, *Adversarial examples—a complete characterisation of the phenomenon*, arXiv preprint arXiv:1810.01185, (2018).
  - [53] S. SHALEV-SHWARTZ AND T. ZHANG, *Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization*, in International Conference on Machine Learning, 2014, pp. 64–72.
  - [54] C. SZEGEDY, W. ZAREMBA, I. SUTSKEVER, J. BRUNA, D. ERHAN, I. GOODFELLOW, AND R. FERGUS, *Intriguing properties of neural networks*, arXiv preprint arXiv:1312.6199, (2013).
  - [55] D. XIU AND G. E. KARNIADAKIS, *Modeling uncertainty in steady state diffusion problems via generalized polynomial chaos*, Computer methods in applied mechanics and engineering, 191 (2002), pp. 4927–4948.
  - [56] D. XIU AND G. E. KARNIADAKIS, *The wiener–askey polynomial chaos for stochastic differential equations*, SIAM journal on scientific computing, 24 (2002), pp. 619–644.
  - [57] S. YE, X. FENG, T. ZHANG, X. MA, S. LIN, Z. LI, K. XU, W. WEN, S. LIU, J. TANG, ET AL., *Progressive DNN compression: A key to achieve ultra-high weight pruning and quantization rates using ADMM*, arXiv preprint arXiv:1903.09769, (2019).
  - [58] T. YOUNG, D. HAZARIKA, S. PORIA, AND E. CAMBRIA, *Recent trends in deep learning based natural language processing*, IEEE Computational intelligence magazine, 13 (2018), pp. 55–75.
  - [59] V. P. ZANKIN, G. V. RYZHAKOV, AND I. OSELEDETS, *Gradient descent-based D-optimal design for the least-squares polynomial approximation*, arXiv preprint arXiv:1806.06631, (2018).