

RESEARCH ARTICLE

ASK: Adversarial Soft k-Nearest Neighbor Attack and Defense

REN WANG^{1,2}, (Member, IEEE), TIANQI CHEN³, PHILIP YAO¹, SIJIA LIU⁴,
INDIKA RAJAPAKSE⁵, AND ALFRED O. HERO III¹, (Life Fellow, IEEE)

¹Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA

²Department of Electrical and Computer Engineering, Illinois Institute of Technology, Chicago, IL 60616, USA

³Department of Statistics, University of Michigan, Ann Arbor, MI 48109, USA

⁴Department of Computer Science and Engineering, Michigan State University, East Lansing, MI 48824, USA

⁵Department of Computational Medicine and Bioinformatics, University of Michigan, Ann Arbor, MI 48109, USA

Corresponding authors: Ren Wang (rwang74@iit.edu), Indika Rajapakse (indikar@umich.edu), and Alfred O. Hero III (hero@eecs.umich.edu)

This work was supported in part by the United States Department of Defense, Defense Advanced Research Projects Agency (DARPA), under Grant HR00112020011; in part by the Army Research Office (ARO) under Grant W911NF-15-1-0479; and in part by the Department of Energy National Nuclear Security Administration under Grant DE-NA0003921.

ABSTRACT K-Nearest Neighbor (kNN)-based deep learning methods have been applied to many applications due to their simplicity and geometric interpretability. However, the robustness of kNN-based deep classification models has not been thoroughly explored and kNN attack strategies are underdeveloped. In this paper, we first propose an Adversarial Soft kNN (ASK) loss for developing more effective kNN-based deep neural network attack strategies and designing better defense methods against them. Our ASK loss provides a differentiable surrogate of the expected kNN classification error. It is also interpretable as it preserves the mutual information between the perturbed input and the in-class-reference data. We use the ASK loss to design a novel attack method called the ASK-Attack (ASK-Atk), which shows superior attack efficiency and accuracy degradation relative to previous kNN attacks on hidden layers. We then derive an ASK-Defense (ASK-Def) method that optimizes the worst-case ASK training loss. Experiments on CIFAR-10 (ImageNet) show that (i) ASK-Atk achieves $\geq 13\%$ ($\geq 13\%$) improvement in attack success rate over previous kNN attacks, and (ii) ASK-Def outperforms the conventional adversarial training method by $\geq 6.9\%$ ($\geq 3.5\%$) in terms of robustness improvement. Relevant codes are available at <https://github.com/wangren09/ASK>.

INDEX TERMS Deep learning, K-nearest neighbor, adversarial soft kNN, ASK-attack, ASK-defense.

I. INTRODUCTION

The K-Nearest Neighbor (kNN) classifier is a simple prototype classification algorithm that has been widely used in classification tasks [1], [2], [3]. Combined with deep neural networks (DNNs), kNN has also been successfully applied to a broad range of important learning tasks, e.g., Remote Sensing Image Retrieval [4], machine translation [5], speaker verification [6], and face recognition [7]. Nonetheless, DNN classifiers have been found to be vulnerable against adversarial attacks [8], [9], which generate small perturbations

The associate editor coordinating the review of this manuscript and approving it for publication was Donatella Darsena .

on benign examples to flip class predictions. Although researchers have recently demonstrated that kNN-based deep learning methods are effective in defending against adversarial attacks on DNNs [10], [11], and can be used in more general data quality testing applications [12], as compared to DNN's, the intrinsic adversarial robustness of kNN-based deep classifiers has been less well explored.

Previously proposed strategies for evaluating kNN-based deep learning methods on image classification tasks include: the Deep kNN attack (DkNN-Atk), which is based on a heuristic process to approximate the kNN [13]; and the adversarial kNN attack (AdvKnn) [14], which uses a neural network to approximate the kNN distribution.

As shown in Figure 1, neither work particularly well on the CIFAR-10 DNN model trained by the adversarial training (AT) [15].

In this paper, we develop a more effective optimization-based attack strategy on kNN-type deep classifiers using a novel loss function, called the Adversarial Soft kNN (ASK) loss. ASK loss better approximates the kNN's probability of classification error as compared to the previously introduced loss functions. We show that the ASK loss is bounded by an information-theoretic measure: a mutual information measure between the perturbed input and benign examples, providing intuition for attack power transfer and the ability to defend against hidden-layer attacks. As our major contribution, we develop a more effective attack on kNN-based deep classifiers, called the ASK-Attack (ASK-Atk), that maximizes the ASK loss. When applied to the kNN on DNN structures, our proposed ASK-attack targets its hidden layers. Therefore, ASK-attack can also be leveraged to evaluate the robustness of features extracted in a hidden layer.

We then derive a more robust defense against kNN-based attacks on hidden layers, called the ASK-Defense (ASK-Def), that minimizes the maximum ASK in the robust training process. Previous robust training methods have largely focused on robustifying DNN prediction [15], [16]. To the best of our knowledge, the proposed ASK-defense is the first method to defend against kNN attacks on hidden layers of DNNs via robust training. The left and right panels of Figure 1 demonstrates the effectiveness of both the proposed ASK-Atk and ASK-Def methods, respectively. Importantly, ASK-Def achieves an average 7% robustness with training complexity comparable to conventional adversarial training.

A. CONTRIBUTIONS

Our principal contributions are the following:

- A novel Adversarial Soft kNN (ASK) loss is introduced that approximates the kNN's probability of classification error in the adversarial setting, and establish that it guides the design of effective attack and defense strategies through the perspective of preserving mutual information between the perturbed input and the in-class-reference data.
- A new attack strategy, the ASK-Attack (ASK-Atk), is presented that maximizes the ASK loss. ASK-Attack provides a principled and effective way to evaluate the robustness of kNN-based deep learning methods in both white-box and gray-box settings, and can be leveraged to evaluate the robustness of features extracted in a hidden layer. On CIFAR-10 and Imagenette we show that the proposed ASK-Atk improves on the kNN attacks of [13] and [14]: ASK-Atk achieves $\geq 13\%$ additional degradation in accuracy rates.
- A new defense strategy, the ASK-Defense (ASK-Def), is presented to robustify model weights under various kNN-based attacks targeting hidden layers through minimizing the maximum ASK loss. On CIFAR-10

(Imagenette) the proposed defense improves the robustness of the DkNN over conventional adversarial training [15] by $\geq 6.9\%$ ($\geq 3.5\%$) on the single selected layers and 4.8% (1.3%) on the combined layers.

B. BACKGROUND AND COMPARISONS

Non-parametric methods like kNN have been widely applied to different tasks, e.g., neural language models [5], [6] and the reasoning task of correspondence classification [17]. Recently, kNN has received renewed attention for its potential to improve robustness. In [18] the authors proposed a robust 1-NN classifier that removes a subset corresponding to the oppositely labeled nearby points in the training data. Multi-layer kNN's [10], [11] provide robust classification against adversarial attacks on DNNs. The baseline method we use to evaluate the ASK-Atk and ASK-Def is the Deep k-Nearest Neighbor (DkNN) [10]. In a DkNN the kNN classifier is embedded into selected layers of a DNN, producing an output class decision by majority vote among the kNN classifiers:

$$y_{\text{DkNN}} = \arg \max_c \sum_{l \in \mathcal{L}} p_l^c(\mathbf{x}), \quad c \in [C], \quad (1)$$

where l is the l -th layer of a DNN and \mathcal{L} is the set of the selected layers. Here $p_l^c(\mathbf{x})$ is the confidence score assigned to class c predicted by the kNN in layer l for input \mathbf{x} . $[C]$ denotes the set $\{1, 2, \dots, C\}$, where C is the number of classes. Recent work shows that the DkNN can also be used to remove the mislabeled training data [12].

While robustness of kNN-based deep classifiers against a wide range of adversarial attacks has been established, their robustness against customized attacks is less well studied. Originally designed with cosine similarity, the DkNN attack (DkNN-Atk) attempts to reduce robustness by moving examples towards incorrect classes through some heuristic steps [13]. Another attack, AdvKnn, trains a small neural network, called the deep kNN block, to approximate the kNN classifier and uses it to construct the attack [14]. Our ASK loss framework for designing adversarial attacks establishes a much stronger attack, the ASK-Atk, than these previous attacks. This is illustrated for the DkNN/VGG16 model trained on the CIFAR-10 training set in the left panel of Figure 1. The results also demonstrate that the ASK loss can well approximate the kNN's probability of classification error. One recent work [19] has studied a similar form of loss for representation learning with random minibatch selection. However, the loss is based on pairwise similarity measurement on a single layer and does not consider the adversarial setting. In contrast, the ASK loss is a class-wise loss under the adversarial setting and has a more general extension with a weighted summation of different layers. Moreover, the ASK loss provides interpretability of effective attack in white-box and gray-box settings. The form of ASK loss can be compared with contrastive loss functions, which has often been used in self-supervised representation learning, e.g., for data augmentation by applying operations like cropping, resizing, or rotation [20], [21]. This paper is the first to propose a loss

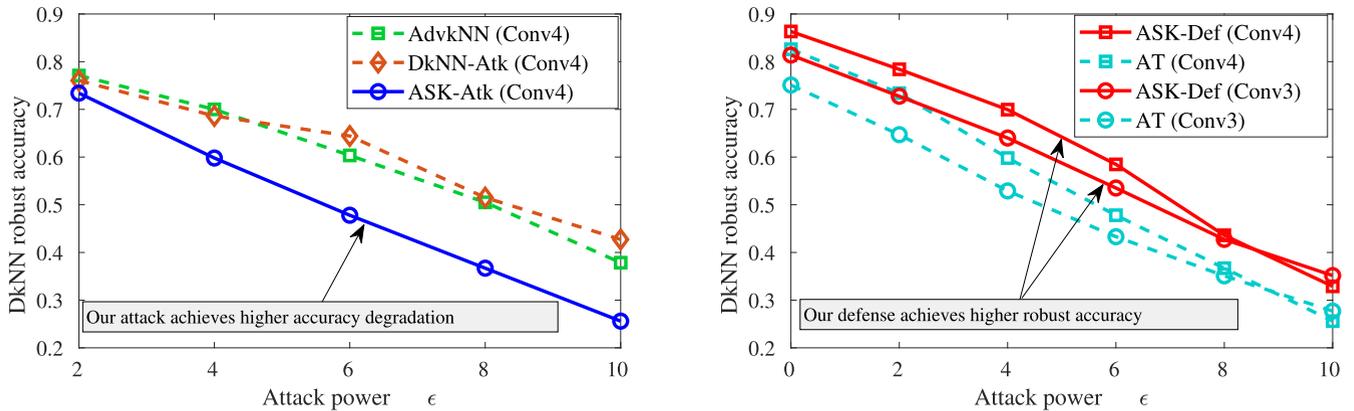


FIGURE 1. Proposed attack (ASK-Atk on left panel) and proposed defense (ASK-Def on right panel) outperform previous methods in terms of impact on accuracy. All accuracy curves correspond to attacks or defenses on the base classifier, a Deep k-Nearest Neighbor (DkNN) [10], with VGG16 CNN architecture, trained on CIFAR-10. Left: ASK-Atk achieves higher accuracy degradation (higher attack success rate) than the existing DkNN-Atk [13] and Advknn [14] attacks on the adversarially trained [15] model. Right: Model trained using the proposed defense (ASK-Def) is more resilient to the ASK-Atk than the adversarially trained baseline model for defending convolutional layer blocks 3 (Conv 3) and 4 (Conv 4) of VGG16. We refer readers to Section VI and Appendix E, F to see more results of ASK-Atk and ASK-Def.

of such form (see (2)) in the supervised and adversarial DNN setting.

Furthermore, as ASK-Atk is based on maximizing the ASK loss function, a stronger defense can be mounted by minimizing the maximum ASK loss during the training process, leading to the proposed ASK-Def. Unlike other robust training methods that defend the output layer of the DNN [15], [16], ASK-Def also defends kNN classifiers on hidden layers leading to improved robustness (right panel of Figure 1). In Figure 2 we illustrate the advantage of ASK-Atk and ASK-Def for one of the CIFAR-10 images contributing to the curves shown in Figure 1. For a clean example of class “truck,” the figure shows that (i) ASK-Atk successfully attacks the example while the Advknn and DkNN-Atk fail, and (ii) the model using standard adversarial training fails to defend against ASK-Atk while the model trained by ASK-Def successfully defends against the attack.

II. PRELIMINARY

A. NOTATION

We consider a DNN with weights θ , and denote θ_l as weights from layer 1 to layer l . $f_{\theta_l}(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^{d_l}$ represents the mapping from the input to the l -th layer feature representation. (\mathbf{x}, y) represents a benign input-label pair. $\mathbf{x}' = \mathbf{x} + \delta$ denotes the perturbed version of \mathbf{x} . In the adversarial context, \mathbf{x}' is called an adversarial example that follows a target adversarial distribution. (X^{l+}, Y^{l+}) denotes a subset containing K data points selected from \mathcal{D}_y (the training data of class y). Similarly, (X_c^{l-}, Y_c^{l-}) is the subset of the training data of class c , where $c \in [C]$, and we denote the $\{X_c^{l-}\}_{c \in [C], c \neq y}$ by X^{l-} . We will call $X^l = \{X^{l+}, X^{l-}\}$ the *reference data* for \mathbf{x} at the l -th layer, X^{l+} the *in-class reference data*, and X^{l-} the *out-of-class reference data*. $A(\cdot, \cdot)$ is a similarity measurement between a data pair, and can be chosen from different similarity functions, e.g., cosine, inner-product, and (negative) ℓ_2 distance. Specifically, the cosine similarity is $A(f_{\theta_l}; \mathbf{x}_1, \mathbf{x}_2) = \frac{1}{\tau_l} \cos(f_{\theta_l}(\mathbf{x}_1), f_{\theta_l}(\mathbf{x}_2))$, where τ_l denotes

a positive (temperature) constant that rescales the similarity of the data pair in the l -th layer. The negative ℓ_2 distance similarity is also used in our experiments and is defined as: $A(f_{\theta_l}; \mathbf{x}_1, \mathbf{x}_2) = -\frac{1}{\tau_l} \|f_{\theta_l}(\mathbf{x}_1) - f_{\theta_l}(\mathbf{x}_2)\|$.

B. ATTACKER’S KNOWLEDGE

We consider both a white-box and a gray-box setting for the adversarial attack. In both settings, the adversary has knowledge of the model parameters, similarity metric, and selected subset of the training set used to perform the kNN search. In the white-box setting, the adversary also knows the layers used for kNN searching and the number of nearest neighbors K . In the gray-box setting, either the layer information or the number K is not available to the attacker. We will focus on the ℓ_∞ attack.

C. DEFENDER’S CAPABILITIES

In our setting, the model trainer is the defender who only has access to the training process while having no control of the model during the inference phase.

III. ADVERSARIAL SOFT kNN (ASK) LOSS

In this section, we will propose a loss that will be used to design effective attacks and defenses on kNN-based deep classifiers. The weights of DNNs are most commonly trained using gradient descent methods. Such methods require a smooth (differentiable or sub-differentiable) network architecture, as a function of the DNN weights, in addition to a smooth loss function. However, neural nets incorporating kNN’s are not differentiable as the kNN mapping is not a smooth function of its inputs. Besides overcoming the non-differentiability challenge, we also hope the attack can work well when the kNN implementation layers and K of the kNN-DNN hybrid system are unknown to attackers, i.e., under our gray-box setting. As for the defense, it needs to be effective in facing kNN-based attacks on deep classifiers. Spurred by these challenges, our design of the loss function



FIGURE 2. Representative example showing superiority of proposed ASK attack (ASK-Atk) and ASK defense (ASK-Def) as compared to previous attack strategies (DkNN-Atk and Adv-Knn). When the DkNN/VGG16 predictor is trained with standard adversarial training [15] on CIFAR-10, for equal perturbation ϵ , ASK-Atk (2nd row) is more successful than other attacks (3rd and 4th row). ASK-Atk maximizes the proposed ASK loss, and its perturbation of the truck image (top left) flips the classes of 4 out of 5 of the image's nearest neighbors. When DkNN/VGG16 is trained with Ask-Def, which minimizes the maximum ASK loss, it successfully defends against all 3 attacks. In particular, the ASK-Atk can only flip the class of a single nearest neighbor (5th row).

is guided by three principles: (1) The loss should provide a differentiable surrogate of the expected kNN classification error, and thus makes it easy to evaluate its robustness. (2) The loss should enable attack power transfer across layers and various choices of K in our gray-box setting (without knowing the exact layers of implementing kNN and exact K). (3) The loss needs to improve model robustness against different attacks on hidden layers.

A. ASK LOSS

We introduce a simple and differentiable Adversarial Soft kNN (ASK) loss, which is a function of the kNN classifier only through the kNN distances. For perturbed input \mathbf{x}' , the ASK loss on layer l is defined as

$$L_{ask}^l(\theta_l; \mathbf{x}', X^l) = -\mathbb{E}_{(X^l, \mathbf{x}')} \left[\log \frac{S(f_{\theta_l}; \mathbf{x}', X^{l+})}{S(f_{\theta_l}; \mathbf{x}', X^{l+}) + \sum_{c=1, c \neq y}^C S(f_{\theta_l}; \mathbf{x}', X_c^{l-})} \right], \quad (2)$$

where $S(f_{\theta_l}; \mathbf{x}', X)$ is a similarity measure between \mathbf{x}' and the data samples X on layer l . f_{θ_l} acts on both \mathbf{x}' and X . We consider two specifications of $S(\cdot, \cdot)$: (i) $S(\cdot, \cdot) = \exp\left(\frac{1}{K} \sum_{k=1}^K A(\cdot, \cdot)\right)$ in our attack (the ASK-Attack Section); and (ii) $S(\cdot, \cdot) = \sum_{k=1}^K \exp(A(\cdot, \cdot))$ in our defense (the ASK-Defense Section). Note that (2) is still well defined when there is no DNN, i.e., f_{θ_l} is an identity map.

B. INFORMATION-THEORETIC INTERPRETATION OF ASK LOSS

The ASK loss (2) can be defined at any layer l and is closely related to the mutual information between X^{l+} and the perturbed input \mathbf{x}' . Specifically, the following proposition shows that the negative ASK loss $-L_{ask}^l$ lower bounds the mutual information $MI(X^{l+}, \mathbf{x}')$ up to an additive constant $\log(C-1)$.

Proposition 1: For perturbed input \mathbf{x}' that follows a target distribution,

$$MI(X^{l+}, \mathbf{x}') \geq -L_{ask}^l + \log(C-1) \quad (3)$$

The proof is given in Appendix A, which follows a similar line of argument to that used in infoNCE [22]. One can also check that the ASK losses under two forms of $S(\cdot, \cdot)$ (used for design of the attack and the defense) lower bound the mutual information. In practice, we would hope the encoding network θ_l is deep enough in order to provide good feature extraction.

Using kNN search, X^{l+} and X^{l-} are the K closest data points to the example \mathbf{x} on layer l from each respective class. The kNN classifier works well when the data points within the same class are consistent and close to each other, while data points from different classes are far apart. Proposition 1 makes it clear that the attacker and defender should take different strategies. From the defender's point of view, minimizing L_{ask}^l can enhance the model robustness by increasing the mutual information between the perturbed input and the in-class reference data, which helps defend against different attacks on hidden layers without knowing the exact attack

forms (principle (3)). In contrast, an attacker would want to minimize the mutual information to create maximum ambiguity between the perturbed and unperturbed inputs which tends to make ASK loss large when there are a large number C of classes. In particular, if the attacker can force the mutual information to zero, then L_{ask}^l can't be smaller than $\log(C-1)$. Proposition 1 also indicates that in the sense of minimizing the mutual information, the attack can be effective without knowing the exact layers of implementing kNN and exact K (principle (2)). Intuitively, minimizing (2) forces (\mathbf{x}', X^{l+}) to be mapped close together, while forcing pairs (\mathbf{x}', X^{l-}) further away in the embedding space induced by the l -th layer. The ASK loss has the interpretation as a temperature-scaled cross-entropy [23], which is a differentiable surrogate of the expected kNN classification error (principle (1)). We will consider two types of ASK loss in the following Sections: ASK loss obtained through locally searched reference data in the attack and ASK loss obtained through non-locally selected reference data with class constraints in our defense.

IV. ASK-ATTACK

The attacker's goal is to find a small perturbation δ such that $\mathbf{x}' = \mathbf{x} + \delta$ flips the kNN classifier prediction (at the input layer or at a deep layer). The strength of the attack δ is constrained to the ℓ_p ball: $\mathcal{C} = \{\delta \mid \|\delta\|_p \leq \epsilon\}$.

A. ATTACK STRATEGY

The attacker's goal is to push the perturbed input \mathbf{x}' away from the y classification region, i.e., the region in which the kNN classifier would correctly predict y , and move \mathbf{x}' towards some data points from other classes. Motivated by the properties of the ASK loss discussed in the previous Section, we adopt a attack strategy that maximizes the weighted summation of ASK losses, i.e., it solves the optimization problem $\arg \max_{\delta \in \mathcal{C}} \sum_{l \in \mathcal{L}} \omega_l L_{ask}^l$, or more explicitly

$$\arg \max_{\delta \in \mathcal{C}} \sum_{l \in \mathcal{L}} -\omega_l \cdot \log \frac{S(f_{\theta_l}; \mathbf{x} + \delta, X^{l+})}{S(f_{\theta_l}; \mathbf{x} + \delta, X^{l+}) + \sum_{c=1, c \neq y}^C S(f_{\theta_l}; \mathbf{x} + \delta, X_c^{l-})}, \quad (4)$$

where $S(f_{\theta_l}; \mathbf{x} + \delta, X) = \exp \frac{1}{K} \sum_{k=1}^K A(f_{\theta_l}; \mathbf{x} + \delta, X)$. ω_l is a weighting coefficient applied to layer l , and $\sum_{l \in \mathcal{L}} \omega_l = 1$. Note that (4) corresponds to a non-targeted attack. For a targeted attack the strategy is to move an example $\mathbf{x} + \delta$ to a target class c_t that is different from class y . For the targeted attack we replace $\sum_{c=1, c \neq y}^C S(f_{\theta_l}; \mathbf{x} + \delta, X_c^{l-})$ with $S(f_{\theta_l}; \mathbf{x} + \delta, X_{c_t}^{l-})$. The target class can be randomly selected or pre-selected by calculating the average distance to \mathbf{x} . In the sequel ASK-Atk will always denote the non-targeted attack, unless otherwise specified. Results for the targeted attack are presented in Appendix E. (4) can be efficiently solved using gradient ascent with projection, with iterations defined as

$$\mathbf{x}' \leftarrow \prod_{\mathcal{B}(\mathbf{x}', \epsilon)} [\mathbf{x}' + \kappa \text{sign}(\nabla_{\mathbf{x}'} \sum_{l \in \mathcal{L}} \omega_l L_{ask}^l(\theta_l, \mathbf{x}', X^l))] \quad (5)$$

where \prod is a projection operator, and $\mathcal{B}(\mathbf{x}', \epsilon)$ represents the projection set that satisfies the prior constraints on the data domain (e.g., $\mathbf{x}' \in [0, 255]^d$ when \mathbf{x}' is an image vector) and the ℓ_p ball constraint. κ is the step size. Note that implementing the ASK-Atk on any layer and any K can affect the mutual information between the perturbed input and the in-class reference data. This enables the attack to work in the gray-box setting. More details are included in Section VI and Appendix B. Next we introduce important parameter selection strategies and computational cost.

B. SELECTION OF X^{l+} AND X^{l-}

The in-class reference points X^{l+} for \mathbf{x} are selected by kNN search for the K data points closest to \mathbf{x} in class y . The out-of-class reference points X^{l-} for \mathbf{x} are also selected using kNN search, but possibly with a different value of K . We consider K to be the same by default. The exact kNN selection improves upon an alternative random selection method to constitute X^{l-} presented in Appendix E.

C. SELECTION OF τ_l AND ω_l

Attacker selection of the hyperparameter τ_l that scales the similarity between the data pair is selected via cross-validation, which only requires a one-time kNN search on a small batch. We also use cross-validation to select ω_l , which controls the power of attacking layer l under the multi-layer attack. This is discussed in more detail in the Section VI.

D. COMPUTATIONAL COST

During training the bottleneck for ASK-Atk is the kNN search, which must be performed over all target samples selected from the training set and requires on the order of $O(kn \log n)$ computations. This is a fixed setup cost that can be handled off-line by parallel or distributed computing, fast approximate kNN approximation [24], [25]. After training, the cost for on-line prediction is cheap, e.g., using a ball-tree to perform kNN search is only of order $O(k \log n)$. τ_l and ω_l only requires a one-time kNN search on a small batch, therefore has limited impact on the computational time.

V. ASK-DEFENSE

The proposed defense ASK-Def solves a minimax optimization problem to robustify the DkNN against attacks and, in particular, kNN-based attacks. ASK-Def is constructed in two stages corresponding to adversarial examples generation and model update,

$$\begin{aligned} & \arg \min_{\theta} \sum_{(\mathbf{x}, y) \in \mathcal{D}} L_{ce}(\theta, \mathbf{x}^{\text{adv}1}, y) \\ & \quad + \lambda \sum_{l \in \mathcal{L}} L_{ask}^l(\theta_l, \mathbf{x}^{\text{adv}2}, X^l) \\ & \text{subject to } \mathbf{x}^{\text{adv}1}, \mathbf{x}^{\text{adv}2} := G1(\mathbf{x}), G2(\mathbf{x}), \end{aligned} \quad (6)$$

where L_{ce} denotes the DNN cross-entropy loss. G1 and G2 denote adversarial attack generators discussed below. The $\arg \min_{\theta}$ operation in (6) performs the update that integrates

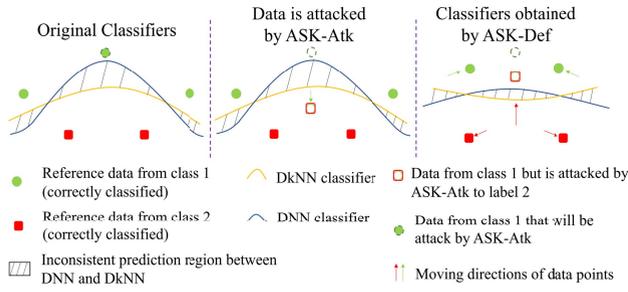


FIGURE 3. Conceptual diagram illustrating how proposed ASK-Def defense strategy emulates ASK-Atk and improves robustness of both the DNN and the DkNN. **Left:** A DNN binary classifier has equivalent classification performance as a deeper kNN classifier but is locally more complex (higher curvature decision boundary). Benign data points (red class and green class) are well separated in the feature space. **Middle:** The adversarial generator G2 generates an adversarial example that crosses both decision boundaries leading to a classification error. **Right:** Minimizing the ASK loss forces the adversarial example and the reference data from class 1 to be closer to each other, resulting in locally smoothed decision boundaries of the DkNN and DNN that increase attack resiliency.

the DNN cross-entropy with the ASK loss, using a penalty parameter λ to control the balance between the two loss terms. In ASK-Def, we use $S(f_{\theta_l}; \mathbf{x}', X_c^l) = \sum_{k=1}^K e^{A(f_{\theta_l}; \mathbf{x}', X_c^l(k))}$ as the class similarity measure. Following the same practical way of dealing with the computational burden as in the nearest neighbor approximation work [19], we find $\{X^{l+}, X^{l-}\}$ from a reference batch, which is obtained by non-local sampling with input label-aware class constraints.

A. ADVERSARIAL GENERATORS G1, G2

ASK-Def uses the adversarial generators G1, G2 that generate attacks that seek to maximize two different losses

$$\begin{aligned}
 \text{(G1)} \quad \mathbf{x}^{\text{adv1}} &= \mathbf{x} + \delta^*, \\
 \delta^* &= \arg \max_{\delta \in \mathcal{C}} L_{ce}(\theta, \mathbf{x} + \delta, y), \\
 \text{(G2)} \quad \mathbf{x}^{\text{adv2}} &= \mathbf{x} + \delta_{(l)}^*, \\
 \delta_{(l)}^* &= \arg \max_{\delta \in \mathcal{C}} L_{ask}^l(\theta_l, \mathbf{x} + \delta, X^l \cup X^{\text{adv2}}), \quad (7)
 \end{aligned}$$

where X^{adv2} is the adversarial counterpart of X^l . G1 generates attacks that maximize the DNN cross-entropy loss, while G2 generates attacks that maximize ASK loss. Thus ASK-Def reduces to conventional adversarial training (AT) [15] when $\lambda = 0$ in (6), while it reduces to hardening against the customized ASK-Atk attack when $\lambda \gg 0$. We could also set G1 to generate attacks by a variant of AT, e.g., TRADES [16]. The pseudocode for the full ASK-Def procedure is given in Appendix C.

B. UNIVERSAL ROBUSTNESS ENHANCEMENT

Based on Proposition 1, minimizing the ASK loss improves the model robustness by increasing the mutual information between the perturbed input and the in-class reference data. This indicates that ASK-Def is able to increase the model resilience against different kNN-based attacks without

knowing the exact attack forms. ASK-Def is also conceptually illustrated in Figure 3, in which we show how using G2 followed by minimizing the ASK loss term help smooth local decision boundaries of DNN and (deep) kNN, increasing resiliency to attack.

VI. EXPERIMENTAL RESULTS

We focus on image classification tasks using CIFAR-10 and Imagenette (10 sub-classes of ImageNet) databases. VGG16 and ResNet18 are trained on CIFAR-10 and Imagenette, respectively. More details on the use of these datasets and models are given in Appendix D. The DkNN [10] is used for illustration of the proposed ASK-Atk and ASK-Def methods. We name the i -th convolutional layer block as Conv i . For VGG16 (ResNet18), Conv 4 denotes the 10-th layer (13-th layer). Conv 3 and Conv 4 are designated as targets. If not otherwise specified, we use $\epsilon = 8$ as the attack power in evaluations, and use $K = 5$ nearest neighbors in the kNN. We show results using 95% confidence intervals over 20 (10) trials in ASK-Atk (ASK-Def) studies.

A. EMPIRICAL STUDY OF ASK-ATK

1) COMPARISONS WITH OTHER ATTACKS

We compare ASK-Atk with AdvKnn [14] and DkNN-Atk [13] on the DkNN classifier [10]. AdvKnn and DkNN-Atk are two state-of-art attacks on kNN-based deep classifiers. We first apply different attacks on a CIFAR-10 model adversarially trained with $\epsilon = 4$. We refer readers to Appendix D for additional details about the attack parameter settings. Performance comparisons are shown in Table 1 demonstrating that ASK-Atk causes higher accuracy degradation than other attacks, for both cosine similarity and ℓ_2 similarity measures used in the ASK loss function (2). The left figure in Figure 1 demonstrates that ASK-Atk outperforms other attacks for various attack power ϵ . It shows almost a linear response on the ASK-Attack's accuracy degradation. The results are consistent with the observations in Figure 2 in [15], demonstrating that increasing the attack power can cause significant accuracy degradation before pushing the robust accuracy to a close-to-zero region. A similar comparison is conducted on an Imagenette model adversarially trained with $\epsilon = 2$. We compare ASK-Atk with other attacks using $\epsilon = 4$ and $\epsilon = 8$, and present the results in Table 2. The table shows that ASK-Atk is stronger than the other attacks on Imagenette. The results on both datasets show that the attack performance of ASK-Atk does not decrease for the lower layers, unlike the other attacks. The ASK-Atk under the targeted attack and random sampling strategy are shown in Tables 9 and 10 in the Appendix and demonstrate that ASK-Atk also has superior performance in these settings.

2) EVALUATIONS ON DIFFERENT LAYER COMBINATIONS

We also show in Table 3 that ASK-Atk can evaluate feature robustness on different layer combinations. ASK-Atk

TABLE 1. ASK-Attack (ASK-Atk) outperforms AdvKnn [14] and DkNN-Atk [13] on CIFAR-10. All the methods are applied on the same DNN model. kNN attacks and DkNN are implemented on the same layers and use either the cosine similarity or ℓ_2 similarity in the ASK loss (2).

	cosine similarity			ℓ_2 similarity		
	ASK-Atk	AdvKnn	DkNN-Atk	ASK-Atk	AdvKnn	DkNN-Atk
Conv 3	35.1%±0.2%	58.9%±0.3%	60.7%±0.44%	32.21%±0.16%	56.55%±0.24%	54.45%±0.52%
Conv 4	36.7%±0.18%	50.5%±0.34%	50.8%±0.39%	34.46%±0.16%	53.9%±0.26%	49.15%±0.49%
Conv 3,4	41%±0.19%	54.1%±0.25%	54.85%±0.41%	39.26%±0.17%	57.15%±0.28%	54.65%±0.5%

TABLE 2. ASK-Attack (ASK-Atk) outperforms AdvKnn [14] and DkNN-Atk [13] on Imagenette under $\epsilon = 8$ and $\epsilon = 4$. All indicated methods are applied to the same DNN model. kNN attacks and DkNN attacks are implemented on the same layers.

	$\epsilon = 8$			$\epsilon = 4$		
	ASK-Atk	AdvKnn	DkNN-Atk	ASK-Atk	AdvKnn	DkNN-Atk
Conv 3	34.4%±0.24%	58.8%±0.27%	51.4%±0.41%	55.2%±0.2%	66.2%±0.16%	58.4%±0.38%
Conv 4	26.7%±0.17%	44.9%±0.3%	56.5%±0.35%	56.2%±0.22%	68.6%±0.25%	67.6%±0.44%
Conv 3,4	34.5%±0.21%	50.5%±0.26%	58.4%±0.4%	60.2%±0.19%	70.2%±0.23%	68.1%±0.4%

retains good performance when number of layers increases to five.

3) SENSITIVITY TO τ_l

The temperature τ_l is selected using cross-validation on a small batch, as discussed in the ASK-Attack Section. We study the sensitivity of ASK-Atk to τ_l by fixing the number of nearest neighbors K to be the same in ASK-Atk and in DkNN, and study how τ_l affects the ASK-Atk performance under different values of K . $\tau_4 = 0.03$ is the optimal value on convolutional layer 4. According to the results shown in the left figure of Figure 4, varying K has only small impact on the performance of ASK-Atk when $\tau_4 = 0.03$. Increasing the temperature to τ_4 , the performance of ASK-Atk remains stable even as τ_4 is increased by a factor of 10, 100, 1000 times the optimal τ_4 . Note that the performance improves as K increases when τ_4 is sufficiently large. Though the performances of ASK-Atk under non-optimal τ_4 are worse than the performance under optimal τ_4 , they are still better than AdvKnn and DkNN-Atk.

4) SENSITIVITY TO THE NEAREST NEIGHBOR K

We next fix τ_4 and study the gray box threat model where the attacker has no information on the value of K used by DkNN. The right figure of Figure 4 shows the performance of ASK-Atk using different values of K . ASK-Atk performs well when the classifier's K is less than or equal to the attacker's K , and ASK-Atk performance can remain relatively good for larger values of the classifier's K when the attacker's value of K exceeds 2. Experimental results of larger variation of K (≥ 15) can be found in Table 11 in the Appendix.

5) SENSITIVITY TO THE ATTACK IMPLEMENTATION LAYERS

Here we conduct experiments on another gray-box setting where ASK-Atk is applied on different layers from the DkNN layers. We apply ASK-Atk on one of Conv 3, Conv 4, and

Conv 3, 4 and also apply DkNN on the one of Conv 3, Conv 4, and Conv 3, 4. Results of these nine combinations are shown in Table 4. One can see that the ASK-Atk still works well in these scenarios, although the optimal attack on a single layer is still achieved when ASK-Atk is applied on the same layer as DkNN. A good choice in the gray-box setting is to attack multi-layer, e.g., Conv 3,4 in the example shown in Table 4.

6) FAILURE CASES ANALYSIS

We compare all the attack failure cases and the success cases under the ASK-Atk from the quantitative perspective. We first calculate the prediction confidence scores of all the benign examples using 50-Nearest Neighbor. We find that the proportion of data points with a confidence score equalling one is 86.89% for failure cases and 22.91% for success cases. The proportion of data points with a confidence score larger than 0.95 is 93.57% for failure cases and 33.88% for success cases. These observations indicate that the failure cases are generally more compact with their neighbors belonging to ground truth classes and thus hard to attack. The analysis also demonstrates that ASK-Atk can be leveraged to evaluate the compactness and robustness of features extracted in hidden layers.

7) ASK-ATK ON THE DATA SPACE

Although the ASK-Atk is designed to operate on inner representations of a deep neural network, we demonstrate that ASK-Atk performs well on the data space. We remark that both CIFAR-10 and Imagenette have poor kNN classification accuracy in the data space. Therefore, we tested ASK-Atk on the data space of MNIST [26] and Fashion-MNIST [27]. The results are shown in Table 5. The standard accuracy for MNIST dataset (Fashion-MNIST dataset) is 95.5% (82.3%) when $K = 5$ and 94% (85%) when $K = 10$. On MNIST (Fashion-MNIST), ASK-Atk can achieve 25% (40%) accuracy degradation under the ℓ_∞ norm constraint with $\epsilon = 60$.

TABLE 3. ASK-Attack (ASK-Atk) retains good performance on different layer combinations. All the methods are applied on the same DNN model. Experiments are conducted on CIFAR-10.

	Conv 4	Conv 3,4	Conv 2,3,4	Conv 1,2,3,4	Conv 1,2,3,4,5
$K = 5$	36.7%±0.18%	41%±0.19%	41.5%±0.23%	41.5%±0.21%	40.1%±0.27%
$K = 10$	38.6%±0.14%	42.2%±0.19%	41.7%±0.24%	41.3%±0.2%	40.7%±0.2%

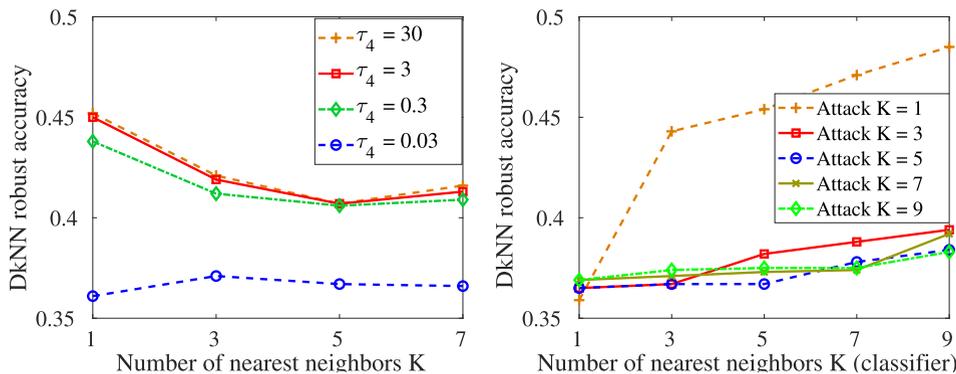


FIGURE 4. The ASK-Atk is robust to variations of hyperparameters τ_4 and K . Left: Varying K has small impact on the performance of ASK-Atk when $\tau_4 = 0.03$. The performance of ASK-Atk remains stable even as τ_4 is increased by a factor of 10, 100, 1000. Right: ASK-Atk performs well when the classifier’s K is less or equal to the attacker’s K , and retains relatively good performance for larger classifier K when the attacker’s K exceeds 2. More results can be found in Tables 11 and 12 in the Appendix.

TABLE 4. ASK-Atk is applied on different layers from the DkNN layers but can still cause high accuracy degradation (attack success rate). All the methods are applied on the same DNN model. Experiments are conducted on CIFAR-10.

	ASK-Atk (Conv 3)	ASK-Atk (Conv 4)	ASK-Atk (Conv 3,4)
Conv 3 (DkNN)	35.1%±0.2%	49%±0.17%	38.6%±0.23%
Conv 4 (DkNN)	55.2%±0.25%	36.7%±0.18%	37.1%±0.16%
Conv 3,4 (DkNN)	48.2%±0.2%	41.9%±0.17%	41%±0.19%

TABLE 5. Performance of ASK-Attack (ASK-Atk) on the data space using MNIST and Fashion-MNIST. The standard accuracy for MNIST dataset (Fashion-MNIST dataset) is 95.5% (82.3%) when $K = 5$ and 94% (85%) when $K = 10$.

	$\epsilon = 40$	$\epsilon = 60$
MNIST ($K = 5$)	76.42%±0.05%	70.5%±0.08%
Fashion-MNIST ($K = 5$)	49.31%±0.1%	41.2%±0.08%
MNIST ($K = 10$)	77.2%±0.1%	70.3%±0.06%
Fashion-MNIST ($K = 10$)	48.7%±0.05%	41.8%±0.11%

B. EMPIRICAL STUDY OF ASK-DEF

1) IMPROVING DkNN ROBUSTNESS BY ASK-DEF

We first implement the ASK-Atk as the evaluation method to test DkNN robustness on models trained by different robust training methods. To demonstrate the effectiveness of ASK-Def, we compare it to (i) using Adversarial Training (AT) [15]; and (ii) AT [15]+ Soft Nearest Neighbor (SNN) [19], where SNN was proposed as a regularizer to improve feature representation and model generalization. The results are shown in Tables 6, which demonstrates that there is additional robustness brought on by the auxiliary ASK loss term. We empirically find that applying ASK loss on the Conv 4 is enough to robustify both Conv 3 and Conv

4. One can see that the model trained by ASK-Def attains the highest robust accuracy against ASK-Atk on the two datasets across on different layers. Specifically, on CIFAR-10/Imagenette, ASK-Def improves the robustness of DkNN over AT (AT+SNN) by $\geq 6.9\%/3.5\%$ ($\geq 1.9\%/0.3\%$) on single selected layer, and $4.8\%/1.3\%$ ($4\%/1.7\%$) on combined layers. The right panel in Figure 1 demonstrates that the model trained by ASK-Def has improved DkNN robustness over that of the model trained by AT over a range of attack power ϵ . We also demonstrate that the model trained by ASK-Def improves robustness over a range of values of K , see Figure 5 in the Appendix. Our theoretical analysis indicates that ASK-Def can improve universal

TABLE 6. ASK-Defense (ASK-Def) outperforms Adversarial Training (AT) [15] and AT [15]+Soft Nearest Neighbor [19] (AT+SNN). The ASK-Def does a better robustification of different layers under the cosine similarity using CIFAR-10 and Imagenette. We add ASK loss to convolutional layer 4.

	CIFAR-10			Imagenette		
	ASK-Def (ASK loss on Conv 4)	AT	AT+SNN (SNN on Conv 4)	ASK-Def (ASK loss on Conv 4)	AT	AT+SNN (SNN on Conv 4)
Conv 3	42.7%±0.16%	35.1%±0.2%	38.3%±0.21%	37.9%±0.19%	34.4%±0.24%	32.1%±0.24%
Conv 4	43.6%±0.19%	36.7%±0.18%	41.7%±0.16%	30.6%±0.17%	26.7%±0.17%	30.3%±0.2%
Conv 3,4	45.9%±0.14%	41%±0.19%	41.9%±0.18%	35.8%±0.2%	34.5%±0.21%	34.1%±0.21%

TABLE 7. ASK-Defense (ASK-Def) outperforms Adversarial Training (AT) [15] and AT [15]+Soft Nearest Neighbor [19] (AT+SNN) in defending AdvkNN [14] and DkNN-Atk [13]. We use cosine similarity and CIFAR-10.

	AdvkNN			DkNN-Atk		
	ASK-Def (ASK loss on Conv 4)	AT	AT+SNN (SNN on Conv 4)	ASK-Def (ASK loss on Conv 4)	AT	AT+SNN (SNN on Conv 4)
Conv 3	61.4%±0.24%	58.9%±0.3%	55.85%±0.31%	62.3%±0.26%	60.7%±0.44%	62%±0.34%
Conv 4	53.4%±0.27%	50.5%±0.34%	41.15%±0.27%	68.1%±0.35%	50.8%±0.39%	65.3%±0.31%
Conv 3,4	56.45%±0.31%	54.1%±0.25%	45.45%±0.17%	72%±0.24%	54.85%±0.41%	71%±0.28%

TABLE 8. ASK-Defense (ASK-Def) improves DNN prediction robustness over Adversarial Training (AT) [15] and AT [15]+Soft Nearest Neighbor [19] (AT+SNN) under different PGD attack power. Both CIFAR-10 and Imagenette are included.

CIFAR-10	$\epsilon = 4$	$\epsilon = 6$	$\epsilon = 8$
ASK-Def (Conv 4)	65.94%±0.23%	52.72%±0.18%	40.36%±0.2%
AT	58.34%±0.21%	43.53%±0.16%	31.65%±0.16%
AT+SNN (Conv 4)	59.2%±0.18%	43.9%±0.21%	30.9%±0.14%
Imagenette	$\epsilon = 4$	$\epsilon = 6$	$\epsilon = 8$
ASK-Def (Conv 4)	55.2%±0.25%	40.2%±0.15%	29.9%±0.13%
AT	49.1%±0.28%	32.3%±0.22%	19.6%±0.1%
AT+SNN (Conv 4)	50.5%±0.19%	34.9%±0.15%	23.1%±0.07%

robustness beyond ASK-Atk. We next conduct experiments on AdvKnn and DkNN-Atk. Table 7 shows that ASK-Def can defend against various attacks and outperforms AT and AT+SNN. Specifically, ASK-Def outperforms the AT+SNN (AT) by 12% (17%) under the AdvkNN (DkNN-Atk) on the convolutional layer 4. These results are comparable to the improvements shown in Table 6, and we believe the results in Table 7 are strong enough to show the additional robustness brought on by the auxiliary ASK loss term. ASK-Def can also be combined with variants of adversarial training, e.g., TRADES [16]. Additional experimental results on comparing with TRADES are shown in Table 13 in the Appendix. We choose TRADES because AT and TRADES are the two most widely used robust training methods. Although ASK-Def can be combined with other variants of AT, we believe that current results are enough to show the effectiveness of our defense method.

2) IMPROVING DNN ROBUSTNESS

As illustrated in Section V, ASK-Def is able to improve the robustness of both DkNN and DNN at the same time. For the

PGD attack, Table 8 shows that ASK-Def can improve DNN robustness under different attack power levels, without appreciable loss of standard accuracy. Compared the $\epsilon = 8$ column in Table 8 with the results shown in Table 6 on Conv 4, one can see that the differences of robust accuracy values on CIFAR-10/Imagenette for ASK-Def are 3.24%/0.7%, while the differences of robust accuracy values for AT are 5.05%/7.1%, indicating that the robust accuracy gap between DkNN and DNN becomes smaller in the model trained by ASK-Def.

VII. CONCLUSION

We presented a novel attack and defense strategy for kNN-based deep classifiers. These strategies both include our proposed Adversarial Soft kNN (ASK) loss term, which we demonstrated is a lower bound on the mutual information between perturbed inputs and the in-class reference data. Maximization of the ASK loss led to our proposed ASK-Attack (ASK-Atk). Experiments show that ASK-Atk outperforms previous kNN-based attacks. A minimax optimization strategy applied to the ASK loss was introduced leading to our robust training method called ASK-Defense (ASK-Def).

Experiments demonstrate that ASK-Def provides additional robustness for kNN-based deep classifiers compared to conventional adversarial training. ASK-Atk and ASK-Def are applicable to a wide array of non-parametric metric learning applications beyond the computer vision applications and kNN-DNN hybrid system developed here.

**APPENDIX A
PROOF OF PROPOSITION 1**

Our proof follows the similar proof technique in [22]. We will show that (i) the optimal similarity measurement $[S(f_{\theta_l}; \mathbf{x}', X_c^l)]_{opt}$ is proportional to density ratio between the conditional distribution $p(X_c^l | \mathbf{x}')$ and the marginal distribution $p(X_c^l)$, where X_c^l denotes the reference data from class $c \in [C]$ (ii) Minimizing the ASK loss maximizes a lower bound on the mutual information between the perturbed input and the in-class-reference data. Given a reference data set $X^l = \{X^{l+}, X_c^{l-}\}$, $c \in [C]/y$ containing one in-class reference sample from the distribution $p(\cdot | \mathbf{x}')$ and $C - 1$ out-of-class reference samples from the distribution $p(\cdot)$, we optimize the ASK loss to correctly select the in-class reference sample out of the given set. In other words, the optimal probability for the ASK loss should imply the fact that X^{l+} comes from $p(\cdot | \mathbf{x}')$ and X_c^{l-} , $c \in [C]/y$ come from $p(\cdot)$. We use $P(X_c^l = l + | X^l, \mathbf{x}')$ to represent such probability. We then have Lemma 1 for the optimal choice of the similarity measurement $S(f_{\theta_l}; \mathbf{x}', X_c^l)$ under the optimal probability by minimizing the ASK loss.

Lemma 1: $[S(f_{\theta_l}; \mathbf{x}', X_c^l)]_{opt} \propto \frac{P(X_c^l | \mathbf{x}')}{P(X_c^l)}$.

The proof details are shown as follows.

Proof: $\frac{S(f_{\theta_l}; \mathbf{x}', X_c^l)}{S(f_{\theta_l}; \mathbf{x}', X_c^l) + \sum_{i=1, i \neq c}^C S(f_{\theta_l}; \mathbf{x}', X_i^l)}$ can be viewed as the prediction score of predicting X_c^l to belonging to the distribution $p(\cdot | \mathbf{x}')$. First, we calculate the optimal probability of $P(X_c^l = l + | X^l, \mathbf{x}')$. We have

$$\begin{aligned} P(X_c^l = l + | X^l, \mathbf{x}') &= \frac{p(X_c^l | \mathbf{x}') \prod_{i \neq c} p(X_i^l)}{p(X_c^l | \mathbf{x}') \prod_{i \neq c} p(X_i^l) + \sum_{i \neq c} p(X_i^l | \mathbf{x}') \prod_{j \neq i} p(X_j^l)} \\ &= \frac{\frac{p(X_c^l | \mathbf{x}')}{p(X_c^l)}}{\frac{p(X_c^l | \mathbf{x}')}{p(X_c^l)} + \sum_{i \neq c} \frac{p(X_i^l | \mathbf{x}')}{p(X_i^l)}}, \end{aligned} \tag{8}$$

Compare the result of (8) with $\frac{S(f_{\theta_l}; \mathbf{x}', X_c^l)}{S(f_{\theta_l}; \mathbf{x}', X_c^l) + \sum_{i \neq c} S(f_{\theta_l}; \mathbf{x}', X_i^l)}$, the proof is done. \square

Leveraging Lemma 1, we provide the proof of Proposition 1 below.

Proof:

$$\begin{aligned} L_{ask}^l &\geq L_{ask}^l(OPT) \\ &= -\mathbb{E}_{(X^l, \mathbf{x}')} \log \left[\frac{S(f_{\theta_l}; \mathbf{x}', X^{l+})}{S(f_{\theta_l}; \mathbf{x}', X^{l+}) + \sum_{c=1, c \neq y}^C S(f_{\theta_l}; \mathbf{x}', X_c^{l-})} \right]_{opt} \end{aligned} \tag{9}$$

$$\tag{10}$$

$$= -\mathbb{E}_{(X^l, \mathbf{x}')} \log \frac{\frac{p(X^{l+} | \mathbf{x}')}{p(X^{l+})}}{\frac{p(X^{l+} | \mathbf{x}')}{p(X^{l+})} + \sum_{c \neq y} \frac{p(X_c^{l-} | \mathbf{x}')}{p(X_c^{l-})}} \tag{11}$$

$$= \mathbb{E}_{(X^l, \mathbf{x}')} \log \left[1 + \frac{p(X^{l+})}{p(X^{l+} | \mathbf{x}')} \sum_{c \neq y} \frac{p(X_c^{l-} | \mathbf{x}')}{p(X_c^{l-})} \right] \tag{12}$$

$$\approx \mathbb{E}_{(X^l, \mathbf{x}')} \log \left[1 + \frac{p(X^{l+})}{p(X^{l+} | \mathbf{x}')} (C - 1) \mathbb{E}_{X^{l-}} \frac{p(X_c^{l-} | \mathbf{x}')}{p(X_c^{l-})} \right] \tag{13}$$

$$= \mathbb{E}_{(X^{l+}, \mathbf{x}')} \log \left[1 + \frac{p(X^{l+})}{p(X^{l+} | \mathbf{x}')} (C - 1) \right] \tag{14}$$

$$\geq \mathbb{E}_{(X^{l+}, \mathbf{x}')} \log(C - 1) \frac{p(X^{l+})}{p(X^{l+} | \mathbf{x}')} \tag{15}$$

$$= \mathbb{E}_{(X^{l+}, \mathbf{x}')} \log \frac{p(X^{l+}) p(\mathbf{x}')}{p(X^{l+}, \mathbf{x}')} + \log(C - 1) \tag{16}$$

$$= -\text{MI}(X^{l+}, \mathbf{x}') + \log(C - 1), \tag{17}$$

\square

where (9) holds because any loss value is larger or equal to the optimal value. (11) comes from Lemma 1. For simplicity, we only consider one data point from each class. $(C - 1)$ will be replaced by $K(C - 1)$ if each class has K reference data. (13) becomes more accurate when the number of classes or number of reference data in each class increases. In practice, the number of classes in most classification problems is large enough for the approximation to work decently well. One can also check that the ASK losses under two forms of $S(\cdot, \cdot)$ (used for design of the attack and the defense) lower bound the mutual information. They can be treated equivalent under Proposition 1.

Algorithm 1 ASK-Atk

Require: Model weights before the target layer l , θ_l ; attack steps s ; step size κ ; attack power ϵ ; reference data points X^l .

Uniformly initialize δ from $[-\epsilon, \epsilon]$, and $\mathbf{x}'[0] = \mathbf{x} + \delta$.

for $j = 1, \dots, s$ **do**

$\mathbf{x}'[j] \leftarrow \mathbf{x}'[j - 1] + \kappa \text{sign}(\nabla_{\mathbf{x}'} \sum_{l \in \mathcal{L}} \omega_l L_{ask}^l(\theta_l, \mathbf{x}'[j - 1], X^l))$.

$\mathbf{x}'[j] \leftarrow \text{Clip}_{[0, 1]}(\mathbf{x}'[j])$.

$\mathbf{x}'[j] \leftarrow \mathbf{x}'[j] + \text{Clip}_{[-\epsilon, \epsilon]}(\delta)$.

end for

Output: \mathbf{x} .

**APPENDIX B
ALGORITHM FOR ASK-ATK**

We formulate ASK-Atk as an optimization problem that maximizes the ASK loss under the given constraints on perturbations and adversarial examples. Algorithm 1 shows the details of implementing ASK-Atk, which includes steps of gradient ascent and projections.

Algorithm 2 ASK-Defense

Require: Initial model θ ; model weights before the target layer l , θ_l ; number of total training epochs E ; number of batches in one epoch T ; attack power ϵ ; penalty parameter λ ; number of reference data points in each class K ; reference data S_{rl} in each batch.

```

for  $e = 1, \dots, E$  do
  for  $t = 1, \dots, T$  do
     $S_t(\epsilon) \leftarrow \text{Attack}^{\text{CE}}(\theta, S_t)$ .  $S_t^l(\epsilon) \leftarrow$ 
     $\text{Attack}^{\text{ASK}}(\theta_l, S_t)$ .
     $\theta \leftarrow \text{Update}^{\text{CE}}(\theta, S_t(\epsilon)) +$ 
     $\lambda \text{Update}^{\text{ASK}}(\theta_l, S_t^l(\epsilon), S_{rl}(\epsilon))$ .
  end for
end for
Output:  $\theta$ .

```

APPENDIX C ALGORITHM FOR ASK-DEF

ASK-Def is constructed in two stages corresponding to adversarial examples generation and model update. The inner maximization can be efficiently solved using gradient ascent with projection, and the model weights are updated using gradient descent. The pseudocode for the full ASK-Def procedure is given in Algorithm 2. $\text{Attack}^{\text{CE}}$ and $\text{Attack}^{\text{ASK}}$ refer to our generators on DNN cross-entropy and on ASK loss. $\text{Update}^{\text{CE}}$ and $\text{Update}^{\text{ASK}}$ denote the gradient updates of first and second terms in model update stage.

APPENDIX D DATASETS, MODELS, AND SETTINGS

A. DATASETS

We mainly utilize two datasets, CIFAR-10 and Imagenette (10 sub-classes of ImageNet), to demonstrate the effectiveness of our methods. CIFAR-10 contains 50000 training samples and 10000 test samples. Each image is in the size of $32 \times 32 \times 3$. Imagenette contains 10 classes from ImageNet and includes 9469 training samples and 3925 test samples. We rescale each image to $128 \times 128 \times 3$. These two datasets are also used in many state-of-the-art empirical attacks and defenses. As a comparison, DkNN and DkNN-Atk only consider simpler datasets like MNIST and SVHN. In the experiments on data space, we utilize MNIST [26] and Fashion-MNIST [27]. The reason is that both CIFAR-10 and Imagenette have poor kNN classification accuracy in the data space. Both MNIST and Fashion-MNIST consist of 60000 training samples and 10000 test samples.

B. MODELS

VGG16 and ResNet18 architectures are used for CIFAR-10 and Imagenette, respectively. The CIFAR-10 model used for attack evaluation (the same baseline model used in ASK-Def comparison) is adversarially trained with $\epsilon = 4$. The Imagenette model used for attack evaluation (the same baseline model used in ASK-Def comparison) is adversarially trained

with $\epsilon = 2$. We apply Deep k-Nearest Neighbor (DkNN) [10] on hidden layers of VGG16 and ResNet18.

C. DkNN LAYER SELECTION

Note that both VGG16 and ResNet18 contain **five** convolutional layer blocks. We name the i th block Conv i . Conv 5 is close to the output layer, thus the robustness is similar to the output, which is lower than Conv 4 and Conv 3. In contrast, Conv 1 and Conv 2 are shallow layers that have not learned good features and have low standard/robust accuracy. We consider the third and fourth convolutional layer blocks of VGG16 and ResNet18 since these two layers have the highest robust accuracy compared to others. We also study different combinations of layers and show results in Table 3.

D. TOTAL AMOUNT OF COMPUTE AND TYPE OF RESOURCES

We use 1 GPU (Tesla V100) with 64GB memory and 2 cores for all the experiments.

E. ATTACK SETTINGS

We compare ASK-Atk with AdvKnn [14] and DkNN-Atk [13] on DkNN classifiers. We use 20 iteration steps for all the attacks. For AdvKnn, we train the deep kNN block with 100 epochs. For DkNN-Atk, we select the optimal hyperparameter in the sigmoid function on each attack. We consider the third and fourth convolutional layers of VGG16 and ResNet18 since these two layers have the highest robust accuracy compared to others. For all the tests on CIFAR-10 (Imagenette), we randomly select 20000 (8500) training samples that are equally distributed across different classes. All the test data examples are used for evaluation. We show results using 95% confidence intervals over 20 trials. If not otherwise specified, we use $\epsilon = 8$ for evaluating both CIFAR-10 and Imagenette and use the nearest neighbor number $K = 5$ by default. We find the optimal τ_l (and ω_l for combined layers) in the first batch and then fix these parameters in the following batches. For both the CIFAR-10 and Imagenette baseline models, we use $\tau_3 = \tau_4 = 0.03$. For the combined layers of the CIFAR-10 baseline model, we use $\omega_3 = \omega_4 = 0.5$. For the combined layers of the Imagenette baseline model, we use $\omega_3 = 0.1$ and $\omega_4 = 0.9$.

F. DEFENSE SETTINGS

Our main goal is to harden both Conv 3 and Conv 4 via ASK-Def. Note that Conv 4 will not be robustified if applying ASK loss to Conv 3. We find that applying ASK loss to Conv 4 can reach similar robustness compare to applying ASK loss to both Conv 3 and Conv 4, and in the meantime, enjoys less computational cost. In ASK-Def, we set $\lambda = 1$, and $\tau_4 = 0.1$. We show results using 95% confidence intervals over 10 trials. Throughout the experiments on defense, we implement ASK-Atk, DkNN-Atk, AdvKnn as the evaluation methods to test DkNN robustness on models trained by Adversarial Training (AT) [15], AT [15]+ Soft Nearest Neighbor (SNN) [19], and the models trained by

TABLE 9. The targeted attack of ASK-Attack (ASK-Atk) can achieve the similarly good performance as non-targeted attack of ASK-Atk on both CIFAR-10 and Imagenette All the methods are applied on the same DNN model. ASK-Atk and DkNN are implemented on the same layers and use the ℓ_2 similarity in the ASK loss (2).

	CIFAR-10		Imagenette	
	ASK-Atk (non-targeted)	ASK-Atk (targeted)	ASK-Atk (non-targeted)	ASK-Atk (targeted)
Conv 3	32.21%±0.16%	32.4%±0.13%	35.9%±0.18%	29.4%±0.2%
Conv 4	34.46%±0.16%	35.9%±0.14%	26.6%±0.13%	27.4%±0.11%
Conv 3,4	39.26%±0.17%	39.9%±0.14%	35.2%±0.2%	33.9%±0.22%

TABLE 10. The random sampling strategy of ASK-Attack (ASK-Atk) performs not as good as using the exact kNN search, but still outperforms AdvKnn [14] and DkNN-Atk [13] on both CIFAR-10 and Imagenette All the methods are applied on the same DNN model. ASK-Atk and DkNN are implemented on the same layers and use the cosine similarity in the ASK loss (2).

	CIFAR-10		Imagenette	
	ASK-Atk	ASK-Atk (random sampling)	ASK-Atk	ASK-Atk (random sampling)
Conv 3	35.1%±0.2%	46.4%±0.35%	34.4%±0.24%	47.5%±0.38%
Conv 4	36.7%±0.18%	39.2%±0.24%	26.7%±0.17%	30.8%±0.21%
Conv 3,4	41%±0.19%	49.4%±0.3%	34.5%±0.21%	49.1%±0.34%

the ASK-Def. Similarly, we apply SSN to Conv 4. When evaluating the models trained by different robust training methods, we find the optimal τ_l and ω_l in the same way as discussed in the attack setting.

APPENDIX E ADDITIONAL EXPERIMENTAL RESULTS OF ASK-ATK A. TARGETED ATTACK AND RANDOM SAMPLING STRATEGY

For the targeted attack we simply replace $\sum_{c \neq y} S(f_{\theta_l}; \mathbf{x} + \delta, X_c^{l-})$ with $S(f_{\theta_l}; \mathbf{x} + \delta, X_c^{l-})$ in the ASK loss. Here we select the target class by calculating the average distance to \mathbf{x} . We show that the targeted attack of ASK-Atk can achieve the similarly good performance as non-targeted attack on both CIFAR-10 and Imagenette, as shown in Table 9. In Table 10, we show how the attack perform when we select X_y^{l-} randomly. The results indicate that selecting X_y^{l-} randomly can still provide a stronger attack than AdvKnn [14] and DkNN-Atk [13] (results for these attacks are shown in Tables 1, 2). To achieve an optimal attack, X_y^{l-} should be selected using kNN search.

B. EXPERIMENTS WHEN VARYING K (ATTACK) AND K (CLASSIFIER) FROM 15 TO 75

Here we focus on larger K ($K \geq 9$). The following Table 11 shows that the ASK-Atk can achieve a good performance when K increases from 15 to 75.

C. WHEN K IS DIFFERENT FOR X_y^{l+} AND X_y^{l-}

Here we study the performance of ASK-Atk when K is different for X_y^{l+} and X_y^{l-} . We use K^+ and K^- to represent

the K for X_y^{l+} and X_y^{l-} , respectively. We vary K^+ from 3 to 7, and vary K^- from 1 to 9. The number of nearest neighbors in DkNN is always set as K^+ . We conduct all experiments on CIFAR-10 to Conv 4. We use the default $\epsilon = 8$. The results are shown in Figure 12. One can see that ASK-Atk can still achieve good performances when X_y^{l+} and X_y^{l-} have different K . When $K^+ \geq 3$, ASK-Atk becomes stronger for $K^- \geq 3$.

D. ADDITIONAL NOTES ON OUR RESULTS

Despite the higher success rate of ASK-Atk, we also find that the ASK-Atk is much faster than DkNN-Atk and AdvKnn. Under our setting, ASK-Atk is around 20 times faster.

APPENDIX F ADDITIONAL EXPERIMENTAL RESULTS OF ASK-DEF A. ROBUSTNESS UNDER DIFFERENT K

We also check how the model (trained by ASK-Def) robustness changes when K in ASK-Atk and DkNN varies. One can see from Figure 5 that the model trained by ASK-Def achieves higher robustness under different K compared to the model trained by conventional adversarial training [15].

B. VARIANTS OF ASK-DEFENSE

ASK-Defense can be combined with variants of adversarial training, e.g., TRADES [16] or other state-of-the-art robust training methods. Here we replace the first loss term (together with the adversarial generator) with TRADES loss term. We name the ASK-Defense with TRADES loss term to ASK-Defense (TRADES). We conduct additional experiments on comparing with TRADES. Table 13 shows that ASK-Defense (TRADES) outperforms TRADES.

TABLE 11. ASK-Attack (ASK-Atk) retains good performance when varying K (Attack) and K (Classifier) from 15 to 75. Each entry shows the classification accuracy of the classifier. ASK-Atk and DkNN are applied on Conv 3, 4.

	$K = 15$ (Classifier)	$K = 35$ (Classifier)	$K = 55$ (Classifier)	$K = 75$ (Classifier)
$K = 15$ (Attack)	40.72%±0.16%	43.32%±0.2%	45.7%±0.23%	46.8%±0.21%
$K = 35$ (Attack)	40.6%±0.19%	40.2%±0.18%	41.63%±0.16%	43.2%±0.11%
$K = 55$ (Attack)	40.2%±0.22%	40.4%±0.15%	40.2%±0.21%	41.33%±0.13%
$K = 75$ (Attack)	41.57%±0.14%	40.9%±0.16%	41.11%±0.14%	40.8%±0.25%

TABLE 12. ASK-Atk performs well when K is different for $x_i^{!+}$ and $x_i^{!-}$. Here we use K^+ and K^- to represent the K for $x_i^{!+}$ and $x_i^{!-}$, respectively. All experiments are conducted on CIFAR-10 to convolutional layer 4.

	$K^- = 1$	$K^- = 3$	$K^- = 5$	$K^- = 7$	$K^- = 9$
$K^+ = 3$	40.8%±0.2%	36%±0.14%	36.1%±0.19%	36.3%±0.17%	36.4%±0.22%
$K^+ = 5$	42.3%±0.23%	36.9%±0.18%	36.7%±0.18%	35.9%±0.12%	35.9%±0.13%
$K^+ = 7$	42.8%±0.18%	38%±0.15%	36.6%±0.2%	36.3%±0.13%	36%±0.17%

TABLE 13. ASK-Defense (TRADES) outperforms TRADES [16]. We combine ASK-Defense with the TRADES loss term. The ASK-Defense (TRADES) has a better performance on the robustification of different layers under the cosine similarity using CIFAR-10. We evaluate the robustness of different defenses using ASK-Atk with $\epsilon = 4, 8$.

	ASK-Defense (TRADES) ($\epsilon = 4$)	TRADES ($\epsilon = 4$)	ASK-Defense (TRADES) ($\epsilon = 8$)	TRADES ($\epsilon = 8$)
Conv 3	64.63%±0.25%	64.11%±0.2%	38.3%±0.12%	37.56%±0.1%
Conv 4	67.58%±0.22%	64.74%±0.16%	36.92%±0.1%	30.62%±0.07%
Conv 3,4	67.79%±0.19%	65.26%±0.28%	40.05%±0.17%	33.08%±0.07%

TABLE 14. ASK-Defense (ASK-Def) outperforms Adversarial Training (AT) [15] and AT [15]+Soft Nearest Neighbor [19] (AT+SNN) under the ASK-Attack with an increasing number of attack iteration steps. We fix the attack power $\epsilon = 8$ and use the cosine similarity. All experiments are conducted on CIFAR-10 and are applied on convolutional layer 4. Other settings follow the default settings in the paper.

	$iter = 20$	$iter = 40$	$iter = 60$	$iter = 80$	$iter = 100$
ASK-Def	43.6%±0.19%	43.1%±0.19%	43.02%±0.12%	42.95%±0.18%	42.87%±0.16%
AT	36.7%±0.18%	36.21%±0.15%	36.15%±0.11%	35.98%±0.13%	35.96%±0.1%
AT+SNN	41.7%±0.16%	41.04%±0.15%	40.93%±0.2%	40.9%±0.18%	40.9%±0.15%

C. DEFENSE AGAINST ATTACKS WITH LARGE ITERATION NUMBER

We conduct experiments on increasing the iteration steps of ASK-Attack. Experimental results in Table 14 show that the ASK-Defense (ASK-Def) outperforms Adversarial Training (AT) [15] and AT [15]+Soft Nearest Neighbor [19] (AT+SNN) under the ASK-Attack with an increasing number of attack iteration steps. Another observation is that increasing the number of iteration steps has little impact on the performance when $iter \geq 20$.

APPENDIX G FAILURE CASES ANALYSIS AND SOCIAL IMPACT

A. FAILURE CASES ANALYSIS

Here we provide more analysis on those failure cases. A consistent observation with the analysis in the main paper is

that the average similarity gap between the largest class and the runner-up class for failure cases is 0.1555, while it is only 0.0519 for success cases (here we use cosine similarity). Again, these observations indicate that the failure cases are generally more compact with their neighbors belonging to ground truth classes and thus hard to attack. We also analyze the failure of ASK-Attack from the visualization perspective. We found that the failure cases usually contain more distinctive features belonging to ground truth classes. For example, the truck looks more stereoscopic.

B. SOCIETAL IMPACT

The broad motivation of our work is to explore the robustness of kNN-based deep learning models, which has not been thoroughly studied. We believe this goal is highly relevant to the machine learning/artificial intelligence community, and

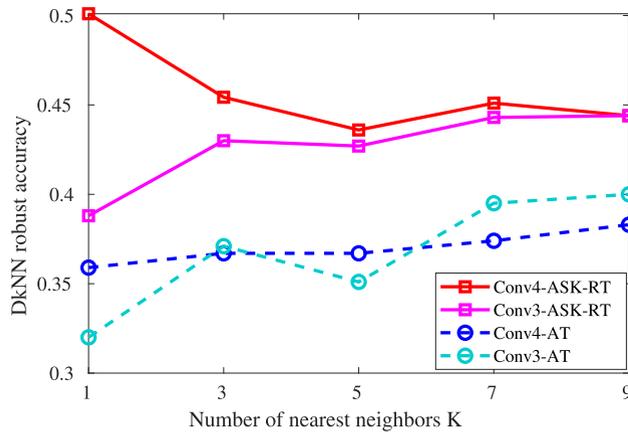


FIGURE 5. ASK-Defense (ASK-Def) outperforms Adversarial Training (AT) [15] when attacks and DkNN vary the number of nearest neighbors K . The ASK-Def does a better robustification of different layers using CIFAR-10. We add ASK loss to convolutional layer 4 and test on convolutional layers 3 and 4.

the methods that our paper introduces can be brought to bear on other deep learning/non-parametric learning problems of interest.

REFERENCES

- [1] T. Cover and P. Hart, "Nearest neighbor pattern classification," *IEEE Trans. Inf. Theory*, vol. IT-13, no. 1, pp. 21–27, Jan. 1967.
- [2] H. Tang, Y. Xu, A. Lin, A. A. Heidari, M. Wang, H. Chen, Y. Luo, and C. Li, "Predicting green consumption behaviors of students using efficient firefly grey wolf-assisted K-nearest neighbor classifiers," *IEEE Access*, vol. 8, pp. 35546–35562, 2020.
- [3] A. Balsubramani, S. Dasgupta, and S. Moran, "An adaptive nearest neighbor rule for classification," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 1–10.
- [4] F. Ye, M. Dong, W. Luo, X. Chen, and W. Min, "A new re-ranking method based on convolutional neural network and two image-to-class distances for remote sensing image retrieval," *IEEE Access*, vol. 7, pp. 141498–141507, 2019.
- [5] U. Khandelwal, A. Fan, D. Jurafsky, L. Zettlemoyer, and M. Lewis, "Nearest neighbor machine translation," in *Proc. Int. Conf. Learn. Represent.*, 2021, pp. 1–14.
- [6] U. Khan, M. India, and J. Hernando, "I-Vector transformation using K-nearest neighbors for speaker verification," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2020, pp. 7574–7578.
- [7] M. Nakada, H. Wang, and D. Terzopoulos, "AcFR: Active face recognition using convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jul. 2017, pp. 35–40.
- [8] A. Kurakin, I. J. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," in *Artificial Intelligence Safety and Security*. London, U.K.: Chapman & Hall/CRC, 2018, pp. 99–112.
- [9] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," 2013, *arXiv:1312.6199*.
- [10] N. Papernot and P. McDaniel, "Deep k-nearest neighbors: Towards confident, interpretable and robust deep learning," 2018, *arXiv:1803.04765*.
- [11] R. Wang, T. Chen, S. M. Lindsly, C. M. Stansbury, A. Rehemtulla, I. Rajapakse, and A. O. Hero, "RAILS: A robust adversarial immune-inspired learning system," *IEEE Access*, vol. 10, pp. 22061–22078, 2022.
- [12] D. Bahri, H. Jiang, and M. Gupta, "Deep k-NN for noisy labels," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 540–550.
- [13] C. Sitawarin and D. Wagner, "On the robustness of deep K-Nearest neighbors," in *Proc. IEEE Secur. Privacy Workshops (SPW)*, May 2019, pp. 1–7.
- [14] X. Li, Y. Chen, Y. He, and H. Xue, "AdvKnn: Adversarial attacks on K-nearest neighbor classifiers with approximate gradients," 2019, *arXiv:1911.06591*.
- [15] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," 2017, *arXiv:1706.06083*.
- [16] H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan, "Theoretically principled trade-off between robustness and accuracy," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 7472–7482.
- [17] T. Plötz and S. Roth, "Neural nearest neighbors networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 31, 2018, pp. 1087–1098.
- [18] J. Zhao, "Analyzing the robustness of deep learning against adversarial examples," in *Proc. 56th Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Oct. 2018, pp. 5133–5142.
- [19] N. Frosst, N. Papernot, and G. Hinton, "Analyzing and improving representations with the soft nearest neighbor loss," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 2012–2020.
- [20] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 1597–1607.
- [21] K. Sohn, "Improved deep metric learning with multi-class N-pair loss objective," in *Proc. 30th Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 1857–1865.
- [22] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," 2018, *arXiv:1807.03748*.
- [23] Y. Bengio, I. Goodfellow, and A. Courville, *Deep Learning*, vol. 1. Cambridge, MA, USA: MIT Press, 2017.
- [24] A. Bewley and B. Upercroft, "Advantages of exploiting projection structure for segmenting dense 3D point clouds," in *Proc. Austral. Conf. Robot. Autom.*, vol. 2, 2013, pp. 1–8.
- [25] A. Rajaraman and J. D. Ullman, *Mining Massive Datasets*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [26] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [27] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms," 2017, *arXiv:1708.07747*.



REN WANG (Member, IEEE) received the bachelor's and master's degrees in electrical engineering from Tsinghua University, Beijing, China, in 2013 and 2016, respectively, and the Ph.D. degree in electrical engineering from the Rensselaer Polytechnic Institute, Troy, NY, USA, in 2020. He was a Postdoctoral Research Fellow with the Department of Electrical Engineering and Computer Science, University of Michigan. He is currently an Assistant Professor with the Department of Electrical and Computer Engineering, Illinois Institute of Technology. His research interests include trustworthy machine learning, high-dimensional data analysis, bio-inspired machine learning, and robustness/optimization on smart grid.



TIANQI CHEN received the B.S. degree in mathematics from Fudan University, Shanghai, China, in 2019, and the M.S. degree in statistics from the University of Michigan, Ann Arbor, MI, USA, in 2021. He is currently pursuing the Ph.D. degree in statistics with The University of Texas at Austin, Austin, TX, USA.



PHILIP YAO received the B.S.E. degree in computer science and electrical engineering from the University of Michigan, Ann Arbor, MI, USA, in 2020. He is interested in deep learning research.



SIJIA LIU received the Ph.D. degree in electrical and computer engineering from Syracuse University, Syracuse, NY, USA, in 2016. He was a Postdoctoral Research Fellow at the University of Michigan, Ann Arbor, in 2016 and 2017, and a Research Staff Member at the MIT-IBM Watson AI Laboratory, in 2018 and 2020. His research interests include scalable and trustworthy AI, e.g., adversarial deep learning, optimization theory and methods, computer vision, and computational biology. He received the Best Student Paper Award at the 42nd IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). His work has been published at top-tier AI conferences, such as NeurIPS, ICML, ICLR, CVPR, and AAAI. He received the All University Doctoral Prize for the Ph.D. degree.



INDIKA RAJAPAKSE is currently an Associate Professor in computational medicine and bioinformatics with the Medical School and an Associate Professor in mathematics at the University of Michigan, Ann Arbor. He is also a member of the Smale Institute. His research interests include interface of biology, engineering, mathematics, dynamical systems, networks, mathematics of data, and cellular reprogramming.



ALFRED O. HERO III (Life Fellow, IEEE) received the B.S. degree (*summa cum laude*) in electrical engineering from Boston University, in 1980, and the Ph.D. degree in electrical engineering from Princeton University, in 1984. Since 1984, he has been with the University of Michigan, Ann Arbor, where he is currently the John H. Holland Distinguished University Professor in electrical engineering and computer science and also the R. Jamison and Betty Williams Professor in engineering. His primary appointment is with the Department of Electrical Engineering and Computer Science and he also has appointments, by courtesy, with the Department of Biomedical Engineering and the Department of Statistics. His recent research interests include high dimensional spatio-temporal data, multi-modal data integration, statistical signal processing, machine learning, applications to social networks, network security and forensics, computer vision, and personalized health. He is a Section Editor of the *SIAM Journal on Mathematics of Data Science* and a Senior Editor of the *IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING*. He is on the Editorial Board of the *Harvard Data Science Review* (HDSR) and serves as a Moderator for the *Electrical Engineering and Systems Science* category of the arXiv. He was the Co-General Chair of the 2019 IEEE International Symposium on Information Theory (ISIT) and the 1995 IEEE International Conference on Acoustics, Speech and Signal Processing. He was the Founding Co-Director of the University's Michigan Institute for Data Science (MIDAS) (2015–2018). From 2008 to 2013, he held the Digeito Chaire d'Excellence at the Ecole Supérieure d'Electricité, Gif-sur-Yvette, France. He is a fellow of the Society for Industrial and Applied Mathematics (SIAM). Several of his research articles have received best paper awards. He was awarded the University of Michigan Distinguished Faculty Achievement Award (2011), the Stephen S. Attwood Excellence in Engineering Award (2017), and the H. Scott Fogler Award for Professional Leadership and Service (2018). He received the IEEE Signal Processing Society Meritorious Service Award (1998), the IEEE Third Millennium Medal (2000), the IEEE Signal Processing Society Technical Achievement Award (2014), the Society Award from the IEEE Signal Processing Society (2015), and the Fourier Award from the IEEE (2020). He was the President of the IEEE Signal Processing Society (2006–2008) and was on the IEEE Board of Directors (2009–2011), where he served as the Director of the Division IX (Signals and Applications). From 2011 to 2020, he was a member and the Chair (2017–2020) of the Committee on Applied and Theoretical Statistics (CATS) of the U.S. National Academies of Science.

• • •