




Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

Explaining adversarial vulnerability with a data sparsity hypothesis

Mahsa Paknezhad^{a,*}, Cuong Phuc Ngo^b, Amadeus Aristo Winarto^b, Alistair Cheong^b, Chuen Yang Beh^b, Jiayang Wu^b, Hwee Kuan Lee^c

^a Senior Research Fellow at Bioinformatics Institute, A*STAR, Singapore

^b Intern at Bioinformatics Institute, A*STAR, Singapore

^c Senior Principal Investigator at Bioinformatics Institute, A*STAR, National University of Singapore, Singapore Eye Research Institute, Singapore

ARTICLE INFO

Article history:

Received 15 April 2021

Revised 12 October 2021

Accepted 15 January 2022

Available online xxxx

Keywords:

Adversarial vulnerability

Adversarial attacks

Decision boundary

Regularization

Data Sparsity

Smooth SoftMax score surface

ABSTRACT

Despite many proposed algorithms to provide robustness to deep learning (DL) models, DL models remain susceptible to adversarial attacks. We hypothesize that the adversarial vulnerability of DL models stems from two factors. The first factor is data sparsity which is that in the high dimensional input data space, there exist large regions outside the support of the data distribution. The second factor is the existence of many redundant parameters in the DL models. Owing to these factors, different models are able to come up with different decision boundaries with comparably high prediction accuracy. The appearance of the decision boundaries in the space outside the support of the data distribution does not affect the prediction accuracy of the model. However, it makes an important difference in the adversarial robustness of the model. We hypothesize that the ideal decision boundary is as far as possible from the support of the data distribution. In this paper, we develop a training framework to observe if DL models are able to learn such a decision boundary spanning the space around the class distributions further from the data points themselves. Semi-supervised learning was deployed during training by leveraging unlabeled data generated in the space outside the support of the data distribution. We measured adversarial robustness of the models trained using this training framework against well-known adversarial attacks and by using robustness metrics. We found that models trained using our framework, as well as other regularization methods and adversarial training support our hypothesis of data sparsity and that models trained with these methods learn to have decision boundaries more similar to the aforementioned ideal decision boundary. We show that the unlabeled data generated by noise in our framework is almost as effective on adversarial robustness as unlabeled data sourced from existing datasets or generated by synthesis algorithms. The code for our training framework is available online.

© 2022 Elsevier B.V. All rights reserved.

1. Introduction

Deep learning (DL) models have achieved unprecedented accuracy for many visual recognition [1], speech recognition [2], and natural language processing tasks [3]. As DL models find their way to real-world applications, an important vulnerability in them was brought to light by the introduction of adversarial examples, generated by applying human imperceptible perturbations to the

original input [4]. In a white-box setup where the adversary has access to weight parameters of the DL model, different attack algorithms such as Fast Gradient Sign Method (FGSM) [5], Projected Gradient Descent (PGD) [6], and DeepFool [7] already exist. Adversarial training [5] is currently one of the most effective approaches to improve robustness. Efforts to explain the adversarial vulnerability of DL models suggest that correct classification only occurs on a thin manifold and most of the high dimensional input data space consists of adversarial examples [5]. We hypothesize that two factors play important roles in adversarial vulnerability of DL models. The first factor is that a large portion of the input data space lies outside the support of the data distributions, i.e. there exists no training data points in a large portion of the input data space. As a result, the shape of the SoftMax score surfaces learned by DL models is not known in regions outside the support of the data distribution. The second factor is that DL models have many

Abbreviations: CNN, Convolutional Neural Network; DL, Deep Learning; FGSM, Fast Gradient Sign Method; OSR, Open Set Recognition; PGD, Projected Gradient Descent; KKC, Known Known Class; SSL, Semi-supervised Learning; TRADES, TRadeoff-inspired Adversarial DEfense via Surrogate-loss minimization; UUC, Unknown Unknown Class.

* Corresponding author at: Bioinformatics Institute, A*STAR, 30 Biopolis Street, 07-01, Matrix 138671, Singapore.*

E-mail address: mahsap@bii.a-star.edu.sg (M. Paknezhad).

<https://doi.org/10.1016/j.neucom.2022.01.062>

0925-2312/© 2022 Elsevier B.V. All rights reserved.

redundant parameters that allow the models to learn different possible decision boundaries, especially outside the support of the input data distribution. These different decision boundaries may have similar prediction accuracy, but they may show major differences when it comes to adversarial robustness. We hypothesize that an ideal SoftMax score surface is smooth and its decision boundary will stay as far from the support of each class distribution as possible.

Our hypothesis about the existence of large factions in data space without training data is similar to that in the ongoing independent efforts on Open Set Recognition (OSR) [8], whose objective is to train models that can identify data points from seen classes in the training data or Known Known Classes (KCCs) accurately, while rejecting data points from unseen classes or Unknown Unknown Classes (UUCs) that are passed into the model during testing. In these algorithms, it is assumed that UUCs are located in a subspace far from KCCs in the high dimensional input data space and the model is trained to have decision boundaries that bind the support of each KCC distribution tightly in order to reserve the space outside the support of the data distribution for UUCs.

Semi-supervised learning (SSL) has been widely used to train models with higher accuracy when limited labeled data is available. SSL has also recently been deployed to train adversarially robust DL models by regularizing DL models using unlabeled datasets. These methods generate unlabeled data points using Gaussian noise and synthesis algorithms or source data points from existing datasets. The idea is that the unlabeled data provides more points in the high dimensional input data space at which the gradient of the SoftMax score surface can be regularized. In this paper, we also leverage unlabeled data, carefully generated to cover a large area around the support of the data distributions to remove the first factor in our hypothesis that is data sparsity. We then use the unlabeled data and a training framework equipped with regularization to train models with decision boundaries similar to the described ideal decision boundary. We will analyse if models trained to have decision boundaries similar to the ideal decision boundary described above are more robust to adversarial attacks.

Our contributions are listed as follows:

1. We hypothesize that adversarial vulnerability of DL models for high dimensional input data stems from two factors. The first factor is data sparsity in the high dimensional input data space. The second factor is that DL models with a large number of parameters tend to learn many different SoftMax score surfaces with decision boundaries that get close to the class distributions in the input data space. We develop a training framework for DL models to learn SoftMax score surfaces similar to what we believe is the ideal SoftMax score surface for DL models. We define the ideal SoftMax score surface as the SoftMax score surface that is smooth and its decision boundary lies far from the support of each class distribution. The training framework generates unlabeled data in a novel way to alleviate data sparsity in a large area around the labeled dataset in the high dimensional input data space. Using the generated unlabeled data and a regularization term, the training framework encourages DL models to learn SoftMax score surfaces similar to the ideal SoftMax score surface. We show that the relative distance of the decision boundary for DL models trained with our training framework using Jacobian regularizer is larger compared to DL models trained with the Jacobian regularizer alone for MNIST, CIFAR10 and Imagenette datasets.
2. We show that our training framework is able to train DL models with SoftMax score surfaces similar to the ideal SoftMax score surface for a 2D points dataset.
3. We show that unlabeled data generated by Gaussian noise generated by our training framework is almost as effective as unlabeled data generated by more costly and complex synthesis algorithms or those sourced from other existing datasets.
4. We compare adversarial robustness of DL models trained using Gaussian-generated unlabeled data and the Jacobian regularizer with intermediate noise injection and the Jacobian regularizer and show that injecting noise at the intermediate layers of the network following the same strategy proposed in our training framework improves adversarial robustness of models.
5. In an OSR setting, we show that models trained with our training framework are consistently less confident in predicting data points from unknown classes which is the desirable behaviour.

There are a number of works that utilize unlabeled or labeled data generated by Gaussian noise for training adversarially robust models. These include the work by Zhang et al. [9] where instead of generating adversarial examples using the direction from the data point to the decision boundary, perturbations are generated for adversarial training using inter-sample relationships in the input data batch. Also, in the works by Chernyak et al. [10], Zantedeschi et al. [11] and Li et al. [12] models trained with new data points drawn from a neighborhood around each training example with the same label as the original data point are found to be more robust to adversarial attacks. None of these methods ensure that the generated data distribution covers a large area around the support of the labeled data distribution. However, they do support our hypothesis that data sparsity is one of the factors that makes DL models prone to adversarial attacks. In summary, in this paper we first propose two hypotheses: 1) there exists data sparsity in high dimensional input data space. As a result, the shape of the SoftMax score surface learned by DL models is unknown in the area outside the support of the input data distribution and 2) the large number of parameters in DL models allows DL models to learn many possible SoftMax score surfaces in the area outside the support of the input data distribution. Our method is related to our hypothesis by generating unlabeled data that alleviates data sparsity in the high-dimensional input data space and training DL models to learn a unique SoftMax score surface that we believe is an ideal SoftMax score surface for DL models. Training DL models to learn the ideal SoftMax score surface was performed using the generated unlabeled data and a regularization term.

We show that a model trained with the generated unlabeled data and a regularization term can learn a SoftMax score surface similar to the ideal SoftMax score surface for a 2D points dataset. We show that the relative distance of the decision boundary of models trained with our training framework increased compared to normal training and is on par with the relative distance of the decision boundary of models trained with other regularization methods. We also show that the prediction confidence of models on UUC data points is consistently lower for models trained with our training framework compared to models trained without regularization implying that the SoftMax score surface at the location of UUC distributions in the high dimensional input data space is flattened. These results imply that the SoftMax score surface of the trained models with our training framework is similar to the ideal SoftMax score surface.

2. Related Work

Over the last few years, many algorithms have been proposed to improve adversarial robustness of DL models. A number of such algorithms propose to detect adversarial attacks. The work by Goswami et al. [13] is an example of such algorithms in which attacks are detected by observing the response of hidden layers in the net-

work or the work by Grosse et al. [14] where adversarial examples are detected by taking into account that they are drawn from a different distribution as the distribution of the original data. There are many other proposed algorithms that detect adversarial examples such as the works by Peck et al. [15], Gong et al. [16], Feinman et al. [17], Das et al. [18], Liang et al. [19] and Lu et al. [20]. Another group of methods aim to cancel the effect of adversarial perturbations on the model's output. The work by Xie et al. [21] is a good example of such methods where the authors perform low-level image transformations such as random resizing and random padding on the input image before passing the image to the classification model to destroy the structure of adversarial perturbations. The work by Liao et al. [22] proposes a denoiser based on the observation that while the difference between the clean and adversarial sample is small, this difference is large in the high-level representation of the two samples in the classifier. The proposed denoiser is trained to suppress the influence of adversarial perturbations by minimizing the distance between the high-level representation of the two images. In another work, Xie et al. [23] propose a feature denoising network that when combined with adversarial training significantly improves adversarial robustness of the trained model. Methods that are built upon adversarial training, such as the work by Baytas et al. [24] are another group of algorithms. In the work by Baytas et al. [24] adversarial training is done using perturbations that are automatically generated by a generator network or in the work by Tramèr et al. [25] perturbations are transferred from other models. Qin et al. [26] perform adversarial training while treating correctly-classified and misclassified examples differently. Apart from these methods, there are other approaches to train robust models such as the work by Theagarajan et al. [27] where adversarial-free zones around the input data distribution of the model are found and adversarial examples are projected to the adversarial-free zones. In the work by Wong et al. [28], a convex outer bound is constructed by applying norm-bounded perturbations on the training data. The authors compute and optimize over the worst case loss within this convex outer bound to train a model that is provably robust to any norm-bounded adversarial attacks. Being computationally heavy, the work by Wong et al. [28] was further modified to scale to larger networks [29,30]. Dapello et al. [31] increase adversarial robustness of DL models by replacing the first few layers of the network with neural layers with fixed weights that simulate primary visual cortex. Gittings et al. [32] propose a defence mechanism against adversarial patch attacks which leverages GANs to come up with effective adversarial patches and fine tunes the model to build resilience against adversarial patch attacks. Many other algorithms have been proposed to address adversarial vulnerability of DL models such as the works by Shaham et al. [33], Li et al. [34], Yang et al. [35], Guo et al. [36], Pang et al. [37], Zhou et al. [38], Cohen et al. [39], Goel et al. [40], Rusak et al. [41], and Liu and Sun [42]. A large number of proposed methods for training adversarially robust models are categorized as regularization methods where an additional loss term is specified for model training. Such regularization terms in the literature are designed to achieve different objectives. One objective is to smooth the decision surface of models such as the work by Simon-Gabriel et al. [43], Xu et al. [44] and Roth et al. [45]. Another objective is to keep the decision boundary away from the training data points such as the work by Singla and Feizi [46], Glimer et al. [47] and Zhang et al. [48]. In the following, we will provide a brief review of related literature on regularization approaches as well as SSL algorithms.

2.1. Regularization Methods

Existing algorithms train DL models with decision boundaries that stay far from training data points, have smooth gradients or

small curvatures[49]. Yan et al. [50] incorporated a DeepFool attack module in training to maximize the distance between the decision boundary of the trained model and the training data points. They added the regularization term $\lambda \sum_k R(-\|\Delta_{x_k}\|_p / \|x_k\|_p)$ with a weight parameter λ to the natural error in the loss function and jointly optimized the model over the loss function during training. The regularization term is a scale of Δ_{x_k} which is the minimum adversarial perturbation found by DeepFool for an input image x_k . R is a monotonically increasing function. Zhang et al. [48] also proposed a regularization term to increase the distance to the decision boundary by adding the boundary error $\max_{x'_{k_i} \in \mathcal{B}(x_k, \epsilon)} \phi(f_{\theta}(x_k) f_{\theta}(x'_{k_i}) / \lambda)$ as a regularization term. $\mathcal{B}(x, \epsilon)$ represents the neighborhood of input x_k , and ϕ is a surrogate of 0-1 loss for binary classification. The regularization term minimizes the difference between the prediction for the natural image $f_{\theta}(x_k)$ and the adversarial example $f_{\theta}(x'_{k_i})$. Jakobovitz and Giryas [51] added the regularization term $\lambda \sqrt{\sum_{i=1}^d \sum_{j=1}^n \sum_{k=1}^N \left(\frac{\partial}{\partial x_i} z_j^{(l)}(x_k) \right)^2}$ which is the Euclidean norm of the network's Jacobian matrix evaluated on input x_k where $f_{\theta}(x_k) = \text{SoftMax}\{z^{(L)}(x_k)\}$. The proposed regularization is applied during extended training of the model after normal training where d is the dimension of the input x_k , n is the dimension of the output of $f_{\theta}(x_k)$, and N is the number of data points in the dataset. The reason for choosing a Jacobian matrix is said to be that Jacobian is related to the curvature of the model's decision boundary and the distance to the closest adversarial example. Moosavi et al. [49] provide theoretical evidence that a small curvature in the loss function increases the robustness of the model. They propose a regularization term that reduces the curvature of the loss function by minimizing eigen values for the Hessian matrix with respect to the inputs: $\lambda \frac{1}{h^2} \mathbf{E} \|\nabla l(x_k, hz) - \nabla l(x_k)\|^2$. Here, h is the discretization step over $z \sim N(0, Id)$. Ross and Doshi-velez [52] use adversarial training and distillation [53] as a regularization mechanism and use the regularization term $\lambda \sum_{i=1}^d \sum_{k=1}^N \left(\frac{\partial}{\partial x_i} \sum_{j=1}^n -y_{kj} \log f_{\theta}(x_k)_j \right)^2$ which encourages that the KL divergence between the predictions and the labels does not change significantly if the input changes slightly [54]. Carmon et al. [55] use the same regularization term proposed by Zhang et al. [48] except that they deploy unlabeled samples in addition to labeled samples during training

2.2. Semi-Supervised Learning

SSL algorithms have also been used for consistency training which refers to all the algorithms that regularize model predictions to be resilient against noise injections into input examples or hidden states. Xie et al. [56] substitute noise injection methods with advanced data augmentation methods to generate unlabeled data and use it together with labeled data for consistency training. Verma et al. [57] proposes training the network with new minibatches generated by randomly mixing two minibatches in an intermediate layer of the network and giving the new minibatch a soft label. Miyato et al. [58] extend adversarial training [5] to use both labeled and unlabeled data points for regularizing the conditional label distribution around each input data point. Zhai et al. [59] utilize SSL for adversarially robust generalization by using portions of the MNIST and CIFAR10 dataset as unlabeled data. The authors train DL models to be accurate on labeled data and robust on both labeled and unlabeled data. Uesato et al. [60] utilize the 80 Million Tiny Images dataset as unlabeled data for the labeled CIFAR10 dataset, which led to a 4% improvement in the robust accuracy of the trained model. Robust accuracy is defined as the accuracy of the model in predicting the correct classes for adversarially perturbed images. The authors also uti-

lized portions of CIFAR10 and SVHN as unlabeled data points for these datasets to improve the adversarial robustness of trained models. They showed that models trained with smaller labeled datasets, with most labels for samples in the CIFAR10 and SVHN dataset removed are more adversarially robust. Najafi et al. [61] removed labels from portions of their training dataset, assigned soft labels to the unlabeled data according to an adversarial loss, and used these soft-labeled images with the labeled images for training. Carmon et al.[55] also augmented CIFAR10 with 500 K unlabeled images from the 80 Million Tiny Images dataset. On SVHN, they used the dataset's own extra training set as unlabeled images. In these studies except for noise injection methods, unlabeled data was captured in a similar expensive manner as labeled data.

3. Method

This section presents our developed framework for training DL models to learn decision boundaries similar to our hypothesized ideal decision boundary. We will first present the overall objective. Later, we introduce the notations used in the paper. Finally, the data generation and model training steps are described.

3.1. The Ideal Decision Boundary

Fig. 1 shows a binary classification problem for a 2D points dataset. Also shown is the schematic representation of the SoftMax score surface for two trained DL models on the 2D points dataset using contour plots. Fig. 1 (Left) shows the SoftMax values for a DL model with a SoftMax score surface that is hypothesized to be non-ideal. Due to the short distance of its decision boundary to certain data points, an attack model will be able to find small perturbations such as ∇x_l and ∇x_k for data points (e.g. x_k and x_l , respectively) that will make the perturbed data points cross the decision boundary and consequently be assigned to the other class label. In contrast, the SoftMax score surface shown in Fig. 1 (Right) is hypothesized to be more robust against adversarial attacks. The SoftMax score surface spans outwards in the space around the data distribution and gradually changes from zero to one from one class to the other. This is desirable as it keeps the decision boundary as far as possible from both class distributions, making it more difficult for adversarial attacks to fool the model via small perturbations.

In Fig. 1 (Left), the SoftMax score surface learned by the model is very steep around the decision boundary and the decision boundary gets close to the data points. The hypothesized ideal SoftMax score surface (Fig. 1) has two distinct characteristics:

1. The first characteristic is the smooth gradient of the SoftMax score surface. The advantage of SoftMax score surfaces with smooth gradients has been studied and is well-known [51,52].
2. The other characteristic is spreading of the SoftMax scores in the space away from the class distributions. This characteristic is very beneficial as will be explained in the following paragraphs.

A SoftMax score surface that spreads outward in the space around the data distribution results in a decision boundary that stays far from all class distributions. Consequently, it will be more difficult for the attack model to find a small perturbation that changes the predicted label for the clean data point. Another advantage of having such a SoftMax score surface is in the behaviour of the model when it receives a data point that does not belong to any of the classes the model was trained on or unknown unknown classes (UUCs). As an example, consider the model with the non-ideal SoftMax score surface in Fig. 1 (Left). If during testing the model receives a data point from a new class that it has not been trained on, the distribution from which that data point is sampled from may be far from the data distribution the model is trained on. If the data distribution for the unseen class is at the right side of the decision boundary in Fig. 1 (Left), the model will label the new data point with the label of the class distribution on the right with a very high confidence (e.g. > 90%), whereas the proposed ideal SoftMax score surface in Fig. 1 (Right) assigns label of the class distribution on the right with a lower confidence (e.g. 60%). To train a DL model to learn such an ideal SoftMax score surface with an ideal decision boundary we sample unlabeled data points in a large neighborhood of the labeled data distribution so that the generated unlabeled dataset covers a large fraction of the high dimensional input data space outside the support of the data distribution and between class distributions. We use a regularization term to encourage small gradients on the SoftMax score surface of the DL model at the location of data points in the labeled and unlabeled datasets in the high dimensional input data space. Our experiments show that the learned SoftMax score surface by DL models using our training framework looks similar to what we define as the ideal SoftMax score surface for DL models.

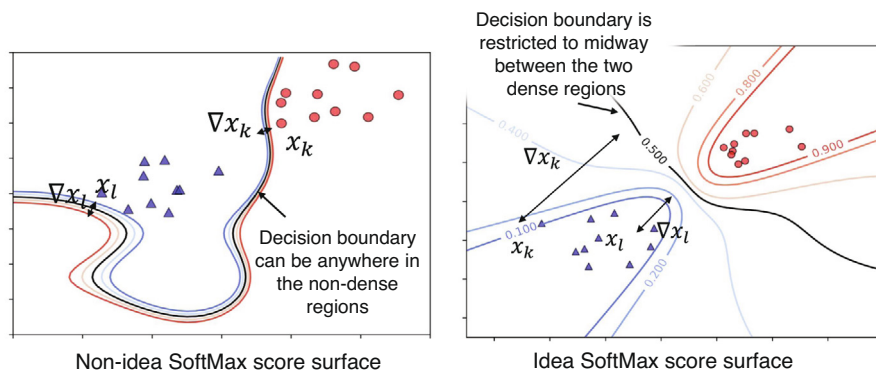


Fig. 1. (Left) shows a SoftMax score surface that a DL model can learn for classifying a 2D point dataset. A model with such a SoftMax score surface is very susceptible to adversarial attacks since the decision boundary gets close to the data distribution and small perturbations such as ∇x_k and ∇x_l in data points (e.g. x_k and x_l) can make the data points cross the decision boundary. (Right) shows what we hypothesize to be the ideal SoftMax score surface which can improve the adversarial robustness of a model. In contrast to the left figure, the decision boundary for this surface is smooth and stays far away from both class distributions. The reader is referred to the online version of this paper for the color representation of this figure.

3.2. Notations

Consider a c -class classification problem for learning a predictor f_θ with parameters θ to map inputs $x_i \in X \subset \mathbb{R}^n$ to labels $y_i \in Y = \{1, \dots, c\}$ where n is the dimension of the input data space. In this case, f_θ will be of the form: $f_\theta(x) = \arg \max_{y \in Y} p_\theta(y|x)$. The labeled dataset is specified by $D = \{(x_i, y_i) | x_i \in X, y_i \in Y, i = 1, \dots, N_l\}$, y_i is the ground-truth label for x_i , and N_l is the number of labeled data points. The generated unlabeled dataset is specified by $U = \{u_{ij} \in \mathbb{R}^n | i = 1, \dots, N_l, j = 1, \dots, N_u\}$ where u_{ij} is the j^{th} unlabeled data point generated using the i^{th} labeled data point, and N_u is the number of generated unlabeled data points using one labeled data point. No ground-truth label is known for the unlabeled data points. The unlabeled data points are sampled from the same data space as the labeled data points, hence $u_{ij} \in \mathbb{R}^n$. We concatenate the sets X and U and define $\Psi = \{\psi_k | \psi_k \in X \cup U, k = 1, \dots, N_l + N_l \times N_u\}$. Hence, a new dataset is generated:

$$D_{\text{new}} = \begin{cases} (\psi_k, y_k) & \text{if } \psi_k \in X \\ \psi_k & \text{if } \psi_k \in U. \end{cases} \quad (1)$$

In the next section, we explain how unlabeled data is generated.

3.3. Generating Unlabeled Data Points

Utilizing unlabeled data points to train adversarially robust neural networks is not a new idea as explained in the Section 2.2. In many domains such as cancer research, obtaining even an uncured dataset of images can be expensive and time consuming. Preparation and scanning of a single sample of cancer tissue can cost hundreds to thousands of dollars. Also, there is no guarantee that the acquired dataset will cover a large portion of the data space around the support of the labeled data distribution. In our training framework, we generate unlabeled data points very cheaply from the labeled dataset using simple additive Gaussian noise. Unlabeled data points were generated from the labeled data points by adding noise to the labeled data points sampled from an isotropic Gaussian distribution. The novelty in our training framework is in sampling Gaussian noise from a suitable Gaussian distribution that allows the generated unlabeled dataset to cover a large neighborhood around the labeled data distribution. The suitable Gaussian distribution was found by measuring the spread of the labeled data distribution in the input data space which was measured as the mean of the Euclidean distance between pairs of data points (μ_{pair}) in the labeled dataset. In other words, μ_{pair} is a parameter that we measure only once from the input labeled dataset. From the input labeled dataset X with N_l labeled data points, pairs of data points are randomly selected ($x_i, x_j \sim X$) and the Euclidean distance between each pair is measured. The mean value of these measurements is defined as μ_{pair} . We then use multiples of μ_{pair} as the standard deviation for the Gaussian distribution from which noise is sampled to generate unlabeled data from each labeled data point by adding the sampled noise to the labeled data point. Generating unlabeled data in this way increases the likelihood that the unlabeled data covers a large area around the original labeled data distribution. The standard deviation of the Gaussian distribution was specified as k -times the mean value. The Gaussian distribution was uniform in all input data dimensions. More precisely, the covariance matrix could be represented by the simplified matrix $\sigma^2 I$ with σ equal to $k \times \mu_{\text{pair}}$. For each labeled data point x_i , a series of N_u unlabeled data points $\{u_{ij} | j = 1, \dots, N_u\}$ were sampled from the area around the labeled data point in \mathbb{R}^n . We set $k = 3$ in our experiments in the paper. The effect of the value of k on the robustness of the DL models is analysed in [Appendix E of the Supplemen-](#)

[tary Materials](#). Unlabeled data points were generated once only for the labeled dataset.

Thus, the unlabeled dataset is generated to cover a large data space outside the support of the data distribution. We believe our generated unlabeled data will support not only the adversarial attack directions but many other directions in the input data space. Using this unlabeled dataset together with the labeled dataset, the DL model is trained to have a smooth SoftMax score surface using regularization. Using the generated unlabeled data and a training framework equipped with regularization, we experimentally show that models learn a decision boundary that is away from the support of each class distribution. We will show in Section 4.4 that using Gaussian noise to generate unlabeled data instead of collecting more data or using synthesis algorithms is almost as useful to achieve a smooth SoftMax score surface with a decision boundary away from the class distributions similar to the ideal decision boundary explained before. In the next section, we explain what regularization technique was deployed to train the model to have a smooth SoftMax score surface.

3.4. Training the CNN

A strong regularization term for training a model with a smooth SoftMax score surface is the Euclidean norm of the network's Jacobian matrix evaluated on each input data point as proposed by Jakobovitz et al. [51] and introduced in Section 2. We believe, regularization algorithms that impose a boundary condition such as TRADES [48] may not be as effective in training models with the aforementioned ideal decision boundary as these algorithms push the decision boundary away only in the ϵ -neighborhood of labeled data points. Minimizing the Euclidean norm of the Jacobian matrix for each input data point will make the SoftMax score surface smooth in each input dimension at that data point. Applying this regularization term not only on the labeled data points but also on the unlabeled data points will smoothen the SoftMax score surface even outside the support of the data distribution which aligns well with our objective. We believe this regularization term and the generated unlabeled data points which cover a large fraction of the high dimensional input data space around the original data distribution will train models similar to the aforementioned ideal decision boundary.

The disadvantage of this regularization term is that training models using this regularization term is very time-consuming. According to the mean value theorem, there exists two input data points $x_1, x_2 \in X$ and $x' = tx_1 + (1-t)x_2, 0 \leq t \leq 1$ for which:

$$\frac{\|f_\theta(x_1) - f_\theta(x_2)\|_2}{\|x_1 - x_2\|_2} \leq \sum_{k=1}^c \|\nabla_{x'} f_{\theta,k}(x')\|_2. \quad (2)$$

where $\|\cdot\|_2$ is the Euclidean norm function. The Jacobian regularization minimizes the term at the right side of Eq. 2. To decrease the training runtime in order to test our hypothesis rigorously, we implemented a more naive version of this regularization term by instead minimizing the term at the left side of the Eq. 2 which is the slope of the SoftMax score surface between two input data points. Fig. 2 shows a schematic representation of the training framework for a DL model for the input dataset shown on the left. The final SoftMax score surface of the trained model using this training framework is also shown at the bottom of the figure. During the training process, a data point, ψ_k , was selected randomly from the new dataset D_{new} (Eq. 1). Similar to the way unlabeled data points were generated, a specific number of neighbor data points, N_b (similar to N_u for unlabeled data points), were generated for the selected data point by adding N_b noises sampled from an isotropic Gaussian distribution with 10-times smaller standard deviation than the standard deviation for generating unlabeled data. Neighbor

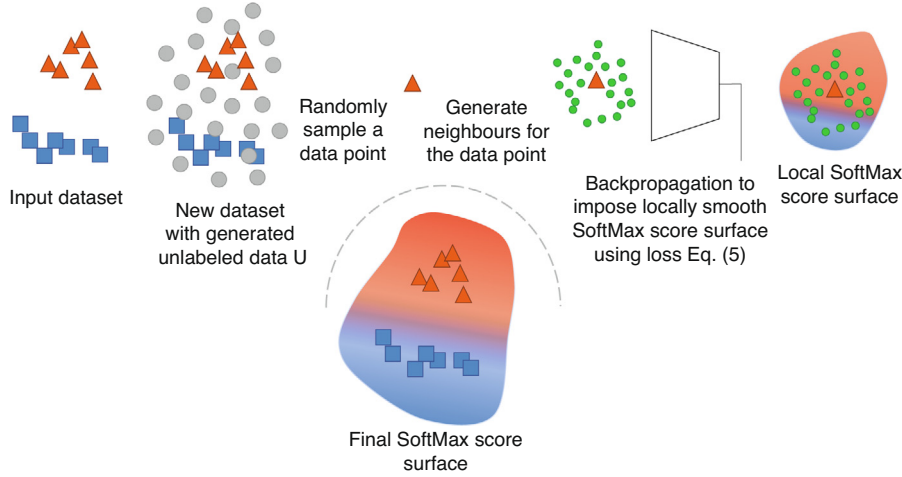


Fig. 2. Figure shows the pipeline for training a DL model for the input dataset shown on the left. For the training dataset an unlabeled dataset is generated. During training a data point is sampled from either the labeled or unlabeled dataset. The local SoftMax score surface around the sampled data point and the generated neighbors for that data point is made smooth using the loss function in Eq. 5. The final ideal SoftMax score surface of the model for the input dataset is shown at the bottom of the figure.

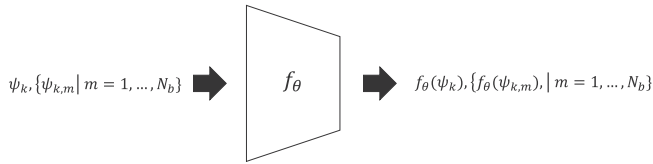


Fig. 3. The input training batch to the CNN model. In each batch, a data point, ψ_k , together with N_b generated neighbor data points, $\{\psi_{k,m} | m = 1, \dots, N_b\}$, are fed to the model. The model is trained by minimising the loss function 5.

data points were generated on the fly for regularization purposes. Hence, they were not part of the new dataset. These neighbor data points are denoted by $\{\psi_{k,m} | m = 1, \dots, N_b\}$. The selected data point was then fed into the model together with the generated neighbors, as shown in Fig. 3 as a mini-batch.

For the selected data point ψ_k and a single neighbor data point $\psi_{k,m}$, which was generated by adding small noise to ψ_k , the model with a smooth SoftMax score surface would ideally output similar predictions, denoted by $f_\theta(\psi_k)$ and $f_\theta(\psi_{k,m})$, respectively. This requirement was enforced by smoothing the slope of the SoftMax score surface between the data point ψ_k and each of its neighbor points $\psi_{k,m}$. This is our naive regularization term, $Loss_{reg}$, and is defined as:

$$Loss_{reg}(\psi_k, \theta) = \sum_{m=1}^{N_b} \frac{\|f_\theta(\psi_{k,m}) - f_\theta(\psi_k)\|_2}{\|\psi_{k,m} - \psi_k\|_2}, \quad (3)$$

$$\forall \psi_k \in X \cup U, \{\psi_{k,m} | m = 1, \dots, N_b\}$$

To ensure that the model learns to correctly classify the labeled data points $\{\psi_k \in X\}$, a supervised loss function, $Loss_{ce}$, in our case cross entropy shown as H , was considered to compare the predicted label, $f_\theta(\psi_k)$ for the labeled data point with the ground-truth label y_k .

$$Loss_{ce}(\psi_k, y_k, \theta) = H(y_k, f_\theta(\psi_k)), \forall \psi_k \in X \quad (4)$$

The overall loss function for the CNN model is a weighted combination of the supervised and the regularization loss functions, defined as below. The cross entropy loss is measured only for the data points that came from the labeled dataset X and is zero otherwise.

$$Loss = \sum_{\psi_k \in X} Loss_{ce}(\psi_k, y_k, \theta) + \lambda \times \sum_{\psi_k \in X \cup U} Loss_{reg}(\psi_k, \theta), \quad (5)$$

where λ is a hyperparameter. We use this loss function for training DL models in the coming experiments to test our hypotheses. More details on the size of the isotropic Gaussian distributions that are used to sample noise and generate unlabeled data points and neighbor data points are provided in [Appendix B of the Supplementary Materials](#). In the next section, we describe our experiment design. We assess our hypothesis using the proposed training framework and show the conclusions that were drawn from these experiments.

4. Experiments

To test our hypothesis about the proposed ideal decision boundary, different datasets comprising of both 2D points and image datasets including MNIST [62], CIFAR10 [63], and Imagenette which is a subset of the well-known ImageNet ILSVRC 2012 dataset [64] were used. The following describes the experiments carried out on each dataset.

4.1. Effect of the Training Framework on the SoftMax Score Surface for the 2D Points Dataset

This experiment was carried out to examine the feasibility of training a DL model to learn a SoftMax score surface which stays far from each class distribution similar to the example shown in Fig. 1 (right). A simple neural network consisting of 4 fully-connected layers with TanH activation functions was trained to perform binary classification on the 2D points dataset shown in Fig. 4 (a). The unlabeled data points generated for this dataset are shown in gray in Fig. 4 (b). Fig. 4 (c) shows the SoftMax score surface the model learns without utilizing regularization and unlabeled data. We also show the learned SoftMax score surface by the Jacobian Regularizer which does not use unlabeled data in Fig. 4 (d). Fig. 4 (e) shows the SoftMax score surface learned by the model using the naive regularization term and unlabeled data. These results show that the decision boundary learned by the model with normal training gets close to the support of the class distributions at certain locations and the SoftMax score surface is considerably steep around the decision boundary. The SoftMax score surface learned by the model trained by the Jacobian regularization term is slightly smoother than that of the model with normal training. However, there are certain locations outside the support of the

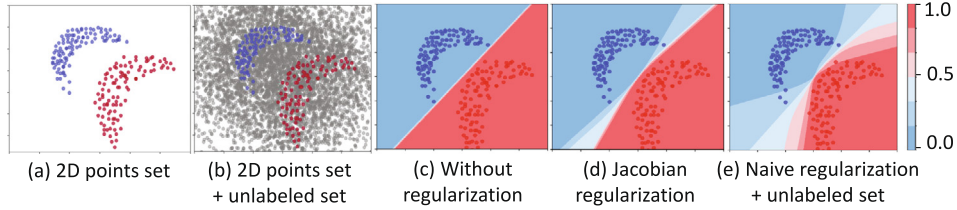


Fig. 4. (a) our 2D points dataset $(\{X, Y\})$. (b) generated unlabeled dataset for the labeled dataset is shown in gray (U). (c) the SoftMax score surface learned by a model without regularization and unlabeled data, (d) the SoftMax score surface learned by the Jacobian Regularization term [51] without unlabeled data, (e) the SoftMax score surface learned by our naive regularization term using the unlabeled dataset. The reader is referred to the online version of this paper for the color representation of this figure.

data distribution where the SoftMax score surface is steep. The model trained using the naive regularizer and unlabeled data, however, keeps the decision boundary away, far from both class distributions and the SoftMax score surface stays smooth outside the support of the data distribution. These results show that unlabeled data helps the model to learn a SoftMax score surface similar to the ideal SoftMax score surface described despite the fact that a weaker regularization term was used compared to the stronger Jacobian regularizer. More examples of different 2D datasets are provided in [Appendix A of the Supplementary Materials](#).

4.2. Numerical Results on the Distance of Decision Boundary from the Data Distribution

MNIST Dataset: The MNIST dataset consists of 60,000 monochrome images of handwritten digits in the training set. It has 10 classes which correspond to the digits 0 to 9 and an equal number of images per class. The images were resized to 32, rescaled to the range $[0, 1]$ and z-score standardized with mean and standard deviation of the training dataset. The LeNet [62] network architecture and Sigmoid activation function were used for all experiments on the MNIST dataset. Details of the training hyperparameters are given in [Appendix B of our Supplementary Materials](#). Five models were trained to perform classification on the MNIST dataset using our proposed loss function (Eq. 5). Different values of λ were tested and the model's loss curve was checked for each λ value. The highest λ value that did not decrease the models' validation accuracy by more than 5% was chosen. We show adversarial robustness of the trained models with same number of models trained using Jacobian Regularization [51], Input Gradient Regularization [52], TRADES [48] and adversarial training. [Table 1](#) shows the results

obtained using Jacobian Regularization [51], Input Gradient Regularization [52], TRADES [48], adversarial training and our naive regularizer with unlabeled data. The $\hat{\rho}_{adv}$ value defined as $\hat{\rho}_{adv}(f_0) = \sum_{x_k \in X} \frac{1}{|X|} \frac{\| \Delta x_k \|_2}{\| x_k \|_2}$ measures the average relative distance between data points and the decision boundary and was introduced by Moosavi et al. [7] as a measure of robustness of models against adversarial attacks. Our analysis shows that the $\hat{\rho}_{adv}$ values were relatively high for the model trained with the naive regularizer and unlabeled data, and when combined with adversarial training compared to other regularization methods. Jacobian Regularization [51], Input Gradient Regularization [52], and adversarial training also show larger $\hat{\rho}_{adv}$ value compared to models trained with no defence mechanism. We have also reported the test accuracy of each model in the same table. We define test accuracy as the prediction accuracy of the model on the clean test dataset. The test accuracy achieved by all the models is very close to the test accuracy for models with no defense. Finally, we have also reported the average robust accuracy of the trained DL models against FGSM and PGD attacks for epsilon value 0.7. This epsilon value is the largest epsilon value that DL models trained on MNIST dataset are evaluated for in the literature. We put '-' in the table if we could not find these measurements for those methods which we have taken from literature. Models trained with smaller values than the chosen λ value had smaller robust accuracy while models trained with larger values than the chosen λ value had smaller test accuracy.

CIFAR10 Dataset: The CIFAR10 dataset contains color images of 10 natural image classes with equal numbers of images per class. CIFAR10 images were also rescaled to the range $[0, 1]$ and z-score standardized using the mean and standard deviation of the training dataset and augmentation methods were applied during

Table 1

Column $\hat{\rho}_{adv}$ shows measured $\hat{\rho}_{adv}$ values using DeepFool attack for the DL models trained on the MNIST dataset. Test accuracy is also shown for each model. Test accuracy is defined as prediction accuracy of the model on the clean test dataset. Average robust accuracy of the models is also shown against FGSM [5] and PGD [6] attacks using epsilon value 0.7.

	Defence method	$\hat{\rho}_{adv}$ ($\times 10^{-2}$)	Test accuracy (%)	Robust accuracy (FGSM, $\epsilon = 0.7$)	Robust accuracy (PGD, $\epsilon = 0.7$)	
Literature	No defence	20.67	99.08	-	-	
	Adversarial training [5]	22.38	99.03	-	-	
	Input Gradient [52]	23.43	99.25	-	-	
	Input Gradient [52] & Adversarial training [5]	23.49	98.88	-	-	
	Jacobian [51]	34.24	98.44	-	-	
	Jacobian [51] & Adversarial training [5]	36.29	98.00	-	-	
	TRADES [48]	25.82	99.21	86.13 \pm 0.97	54.90 \pm 3.03	
	TRADES [48] & Adversarial training [5]	33.54	99.17	90.83 \pm 0.33	67.09 \pm 0.54	
	Ours	No defence	20.30	98.29	41.40 \pm 3.21	0.02 \pm 0.01
		Adversarial training [5]	61.84	98.66	89.50 \pm 3.30	43.82 \pm 24.08
Naive regularizer & Unlabeled data		46.29	98.69	62.62 \pm 2.11	4.11 \pm 1.07	
Naive regularizer, Unlabeled data & Adversarial training [5]		51.20	98.80	93.62 \pm 1.17	55.66 \pm 1.87	

Table 2

Column $\hat{\rho}_{adv}$ shows measured $\hat{\rho}_{adv}$ values using DeepFool attack for the models trained on CIFAR10 dataset. Test accuracy is also shown for each model. Test accuracy is defined as prediction accuracy of the model on the clean test dataset. Average robust accuracy of the models is also shown against FGSM [5] and PGD [6] attacks using epsilon value 15.

	Defence method	$\hat{\rho}_{adv}$ ($\times 10^{-2}$)	Test accuracy (%)	Robust accuracy (FGSM, $\epsilon = 15$)	Robust accuracy (PGD, $\epsilon = 15$)	
Literature	No defence	1.21	88.79	-	-	
	Adversarial training [5]	1.23	88.88	-	-	
	Input Gradient [52]	1.43	88.56	-	-	
	Input Gradient [52] & Adversarial training [5]	2.17	88.49	-	-	
	Jacobian [51]	3.42	89.16	-	-	
	Jacobian [51]& Adversarial training [5]	6.03	88.49	-	-	
	TRADES [48]	1.44	86.74	16.11 \mp 0.75	7.74 \mp 0.28	
	TRADES [48] & Adversarial training [5]	1.34	87.15	16.78 \mp 1.30	7.99 \mp 0.57	
	Ours	No defence	1.23	91.17	10.61 \mp 0.39	1.49 \mp 0.08
		Adversarial training [5]	1.27	89.51	80.06 \mp 1.74	2.30 \mp 0.89
Naive regularizer & Unlabeled data		3.29	83.96	47.03 \mp 0.85	43.60 \mp 0.67	
Naive regularizer, Unlabeled data & Adversarial training [5]		4.45	83.50	64.49 \mp 3.22	58.46 \mp 2.78	

training. For CIFAR10 dataset, ResNet 9 [65] network architecture was used with a CELU activation function for all layers. The lambda value was chosen similarly as for the MNIST dataset. Details of the training hyperparameters are given in [Appendix B of our Supplementary Materials](#). Five models were trained using the proposed loss function (Eq. 5), the labeled dataset and the generated unlabeled dataset. Training and comparison of the results were conducted similar to the MNIST experiments. [Table 2](#) shows the results obtained using Jacobian Regularization [51], Input Gradient Regularization [52], TRADES [48] and the naive regularizer with unlabeled data. The $\hat{\rho}_{adv}$ values achieved for the naive regularizer and unlabeled data and when combined with adversarial training are larger compared to models trained with no defense, adversarial training, Input Gradient Regularization [52] and TRADES [48]. Jacobian regularization and Jacobian Regularization with adversarial training show larger $\hat{\rho}_{adv}$ values. The test accuracy achieved was slightly lower using our training framework. We believe the test accuracy is compromised over achieving a smooth SoftMax surface outside the support of the data distribution in the high dimensional input data space. Also shown is the average robust accuracy of the trained DL models against FGSM and PGD attacks for epsilon value 15. We put '-' in the table if we could not find these measurements for those methods which we have taken from literature.

Imagenette Dataset: Imagenette is a subset of the well-known ImageNet ILSVRC 2012 dataset [64] with 10 easily classified classes. Images were rescaled to the range [0, 1] and z-score standardized using the mean and standard deviation of images in the training dataset. The XResNet 18 [66] network architecture was used to train a model to classify Imagenette dataset with MISH

activation function. Details of the training hyperparameters are given in [Appendix B of our Supplementary Materials](#). The training and attacks performed were similar to those in the experiments for MNIST and CIFAR10 datasets. [Table 3](#) shows the results obtained without regularization, with adversarial training, and the naive regularizer with unlabeled data. Also shown is the results obtained using Jacobian regularizer [51] and TRADES [48] with adversarial training. Similar to the previous experiments, our models show higher $\hat{\rho}_{adv}$ values compared to normal training and TRADES [48]. When combined with adversarial training, the model shows the highest $\hat{\rho}_{adv}$ value. The experiments on MNIST, CIFAR10 and Imagenette datasets show that a naive regularizer together with unlabeled data indeed results in models with larger $\hat{\rho}_{adv}$, which implies that the distance between data points and decision boundary increases for the models trained with unlabeled data that supports a larger area in the high dimensional input data space. We showed that even a naive regularizer can improve adversarial robustness of DL models using our generated unlabeled data. We defined and used a naive regularizer so that it would allow us to run rigorous experiments to study our hypotheses in section 4.4 to 4.7. Due to the slow speed of the Jacobian regularizer, we were not able to use this regularizer to run the experiments in section 4.4 to 4.7 in a reasonable amount of time. However, we experimentally show Jacobian regularizer's superiority over our naive regularizer in Section 4.3. We believe other regularization techniques such as Input Gradient Regularization [52] and Jacobian Regularization [51] and adversarial training are indirectly trying to achieve the ideal decision boundary explained before. The drawback of these algorithms, however, is that they do not take advantage of unlabeled data

Table 3

Column $\hat{\rho}_{adv}$ shows measured $\hat{\rho}_{adv}$ values using DeepFool attack for the models trained on Imagenette dataset. Test accuracy is also shown for each model. Test accuracy is defined as prediction accuracy of the model on the clean test dataset. Average robust accuracy of the models is also shown against FGSM [5] and PGD [6] attacks using epsilon value 19.

	Defence method	$\hat{\rho}_{adv}$ ($\times 10^{-2}$)	Test accuracy (%)	Robust accuracy (FGSM, $\epsilon = 19, \%$)	Robust accuracy (PGD, $\epsilon = 19, \%$)
Literature	Jacobian [51]	2.49	85.60	23.67 \mp 0.68	8.54 \mp 0.48
	Jacobian [51] & Adversarial training [5]	4.39	87.18	42.71 \mp 0.47	28.21 \mp 0.36
	TRADES [48]	1.20	85.97	7.02 \mp 0.23	0.08 \mp 0.01
	TRADES [48] & Adversarial training [5]	4.15	82.88	56.08 \mp 0.40	53.18 \mp 0.39
	Ours	No defence	0.941	87.18	6.48 \mp 0.37
Adversarial training [5]		4.82	83.06	69.86 \mp 0.45	67.22 \mp 0.52
Naive regularizer & Unlabeled data		2.21	84.92	12.69 \mp 0.79	0.55 \mp 0.09
Naive regularizer, Unlabeled data & Adversarial training [5]		4.96	82.44	69.59 \mp 0.39	67.68 \mp 0.37

during training. In Section 4.3, we conduct experiments to analyse how models trained with Jacobian regularization and unlabeled data perform in comparison to Jacobian regularization alone and the naive regularizer with unlabeled data. In other words, does unlabeled data help the models trained with Jacobian regularization perform even better. In the next section, we show the results of these experiments.

4.3. Comparison between Our Regularization and the Jacobian Regularization

As mentioned before, applying Jacobian regularization during the training process is very time consuming. However, this regularizer is stronger than our naive regularization term. For the purpose of confirming this point, we ran multiple experiments and compared the robust accuracy of models trained with the Jacobian regularization term and the naive regularization term. Although Jakobovitz and Gryes [51] applied Jacobian regularization as a post-processing step after training the model has finished, we applied the Jacobian regularizer to the model during the training process similar to our training framework to ensure a fair comparison. We compared the robustness of models trained with the naive regularizer and unlabeled data, the Jacobian regularizer without unlabeled data and the Jacobian regularization with unlabeled data using the FGSM [5] and PGD [6] attacks.

For each dataset, three models were trained using the naive regularizer and unlabeled data and three models were trained using the Jacobian regularizer without unlabeled data. Their robust accuracy evaluated against the FGSM [5] and PGD [6] attacks with different epsilon perturbation values (ϵ) is outlined in Fig. 5. The lambda values used for the Jacobian Regularizer were set to 0.01 for the MNIST dataset, 0.01 for the CIFAR10 dataset and 0.005 for the Imagenette dataset and were found as specified in the previous section.

As shown in Fig. 5, the proposed training framework performs better for the MNIST dataset, whereas the Jacobian regularizer

performs better for the CIFAR10 and Imagenette dataset. This improved performance has two potential explanations:

1. The data sparsity hypothesis is proven wrong for the CIFAR10 and Imagenette datasets.
2. The Jacobian regularizer is a stronger regularizer than our naive regularizer, regardless of the extent to which the data distribution supports the high dimensional input data space.

To find out which explanation is true, we compared the robustness of models trained using just the Jacobian regularizer versus models trained with both Jacobian regularization and unlabeled data. Three models were trained per dataset. If the robustness does not significantly improve after the Jacobian regularizer was deployed with unlabeled data, we could conclude that unlabeled data does not provide significant benefits for training models on CIFAR10 and Imagenette datasets. In other words, the first explanation holds true. Conversely, if the robustness does significantly improve, the second explanation holds true. The $\hat{\rho}_{adv}$ values and test accuracy of the trained models are shown in the rows "Jacobian", and "Jacobian + Unlabeled" in Table 4. The robust accuracy is also shown with labels "Jacobian" and "Jacobian + Unlabeled" in Fig. 6. As seen in Table 4 and Fig. 6, when the Jacobian regularizer is used in conjunction with unlabeled data, robustness indeed improves. Therefore, the second explanation holds true. Jacobian Regularization is indeed more powerful than our naive regularizer. The proposed naive regularizer uses neighbor points to minimize the slope of the SoftMax score surface. Depending on the relative position of the neighbor data point with respect to the labeled or unlabeled data point, our naive regularizer smooths the SoftMax score surface in not all but a limited number of dimensions in the input data space in exchange for a shorter training runtime whereas the Jacobian regularizer smooths the SoftMax score surface in all dimensions in the input data space at the location of the input data point. A comparison of the runtime for the naive regularizer with unlabeled data, the Jacobian regularizer and the

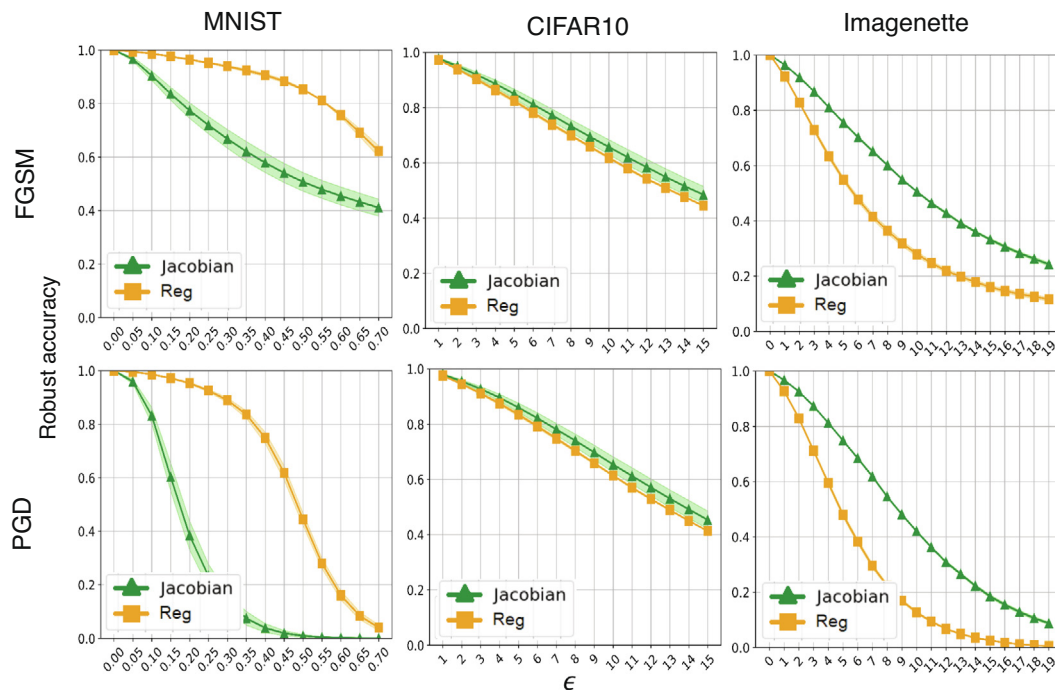


Fig. 5. Figure shows the robust accuracy of models trained using the naive regularizer (Reg) with unlabeled data and the Jacobian regularization (Jacobian) without unlabeled data [51] for MNIST, CIFAR10 and Imagenette datasets. The robust accuracy is measured using FGSM [5], and PGD [6] attacks with different epsilon perturbation values (ϵ). The reader is referred to the online version of this paper for the color representation of this figure.

Table 4

Columns $\hat{\rho}_{adv}(\times 10^{-2})$ and Test Accuracy (%) show measured $\hat{\rho}_{adv}$ values from the DeepFool attack [7] and test accuracy for the models trained on MNIST, CIFAR10, and Imagenette datasets. This table compares the robustness of models trained using the Jacobian regularization (Jacobian), Jacobian regularization and unlabeled data (Jacobian + Unlabeled), as well as Jacobian regularization with Gaussian noise injected to the intermediate model outputs (Jacobian + Inject).

Defence method	MNIST		CIFAR10		Imagenette	
	$\hat{\rho}_{adv}$ ($\times 10^{-2}$)	Test accuracy (%)	$\hat{\rho}_{adv}$ ($\times 10^{-2}$)	Test accuracy (%)	$\hat{\rho}_{adv}$ ($\times 10^{-2}$)	Test accuracy (%)
Jacobian	561	98.75	3.56	88.01	2.57	85.52
Jacobian + Unlabeled	1083	98.84	3.98	87.78	2.73	85.29
Jacobian + Inject	1582	98.70	4.42	87.24	4.21	86.93

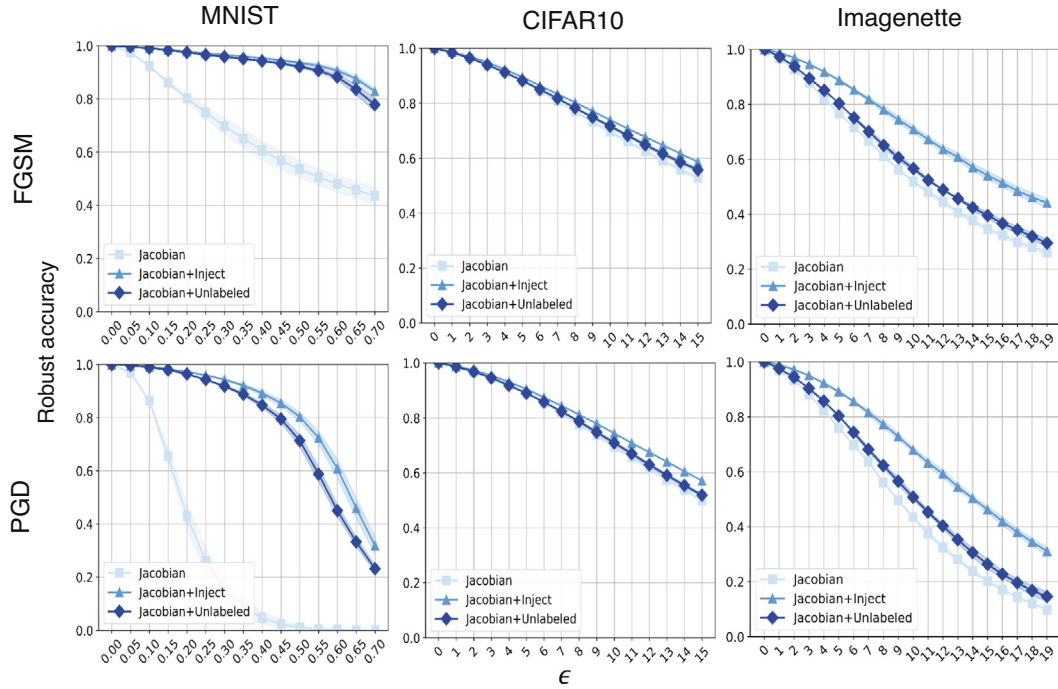


Fig. 6. Figure shows the robust accuracy of models trained using the Jacobian regularization term (Jacobian) versus models trained using both Jacobian regularization and unlabeled data (Jacobian + Unlabeled) and the Jacobian regularization with Gaussian noise injected to the intermediate model outputs (Jacobian + Inject). The robust accuracy is measured using FGSM [5] and PGD [6] attacks with different epsilon perturbation values (ϵ). It can be seen that training models using both Jacobian regularization and unlabeled data and Jacobian regularization with injected Gaussian noise to the intermediate model outputs are more resilient to adversarial perturbations. The reader is referred to the online version of this paper for the color representation of this figure.

Table 5

Table shows the average runtime (in minutes) for training a LeNet network [62] with MNIST dataset for 1 epoch, the runtime for training a ResNet 9 network [65] on CIFAR10 dataset and a XResNet 18 network [66] on Imagenette dataset using our naive regularization term with unlabeled data, the Jacobian regularization term and the Jacobian regularization term with unlabeled data. All measurements were done on a Google Colab T4 Tesla GPU for a single run.

Training runtime	Proposed + Unlabeled (min)	Jacobian (min)	Jacobian + Unlabeled (min)
MNIST	5.82	4.05	8.72
CIFAR10	11.83	12.47	25.75
Imagenette	8.48	16.80	34.55

Jacobian regularizer with unlabeled data is shown in Table 5 where the average training runtime for 1 epochs is reported in minutes for each dataset. One can see that the runtime for the naive regularizer with unlabeled data is less than the training time for the Jacobian regularizer with unlabeled data by 1.5–4 times.

The reader may find certain differences in the robust accuracy plots in our paper and the plots in the literature. There are two reasons for the differences seen in the plots in our paper and the plots in the literature. The first reason is that unlike the papers in the

literature where robust accuracy is reported for all samples in the test dataset, we show adversarial robustness of a model for samples that the model originally classified correctly. In other words, we only perturb images that the model classifies correctly and see if the model can still assign the perturbed images to the correct class label. We believe showing how much a model's performance deteriorates on correctly classified images is a more accurate measurement of robust accuracy of models.

Another reason for the difference seen in the plots can be due to differences in the data preprocessing pipeline of our experiments and the experiments conducted in the literature. In our preprocessing pipelines, we followed standard preprocessing steps in the literature for the three datasets: MNIST, CIFAR10, and Imagenette. We z-score standardized image intensities using the mean and standard deviation of all the images in the training dataset. We also applied augmentation methods for CIFAR10 and Imagenette datasets. The preprocessing pipelines are included in our code that is available online. However, we were not able to find out what preprocessing steps were performed for the algorithms in the literature such as Input Gradient [52] and Jacobian [51] regularization algorithms. We used the standard Foolbox library [67] to measure robust accuracy for our models. In the next section, the effect of

different types of unlabeled data on the adversarial robustness of DL models is analysed.

4.4. Gaussian Noise is as Effective as Costly and More Complex Unlabeled Datasets

An important question is whether models trained with unlabeled datasets widely utilized in the literature for SSL show a higher adversarial performance compared to models trained with unlabeled data generated with Gaussian noise. To answer this question we trained two models for CIFAR10 dataset using our naive regularizer with unlabeled data extracted from the 80 Million Tiny Images dataset following the same automatic filtering technique used by Xie et al. [56]. We also prepared another unlabeled dataset by applying the synthesis technique proposed by Uesato et al. [60]. Similar to many papers on semi-supervised learning and regularization [58,68,59], we ran experiments for different numbers of labeled data points per class ($N_l = 128, 256, 512$) and the same number of unlabeled data points per labeled data point ($N_u = 1$) using these two unlabeled datasets, and compared the robust accuracy of the trained models with those of models trained with unlabeled data generated from Gaussian noise. Fig. 7 (a) and (b) compare robust accuracy of models trained with unlabeled data extracted from the 80 Million Tiny Images dataset and our models using FGSM and PGD attacks. Similarly, Fig. 7 (c) and (d) compare robust accuracy of models trained with unlabeled datasets generated using the synthesis technique proposed in the work by Uesato et al. [60] with our models for FGSM and PGD attacks. Our models show slightly less robust accuracy for the same number of unlabeled data points, suggesting that generating unlabeled datasets using Gaussian noise as explained in this paper is almost as effective as more complex and sometimes expensive unlabeled datasets.

4.5. Effect of Regularization and Unlabeled Data on the Trained DL Models

This section analyses the effect of our naive regularizer and the amount of unlabeled data on adversarial robustness of DL models. Fig. 8(a) (Left) shows robust accuracy for models trained with different number of unlabeled data per labeled data points ($N_u = 0, 5, 10$) using our training framework ($\lambda = 100k$), $N_l = 64$, for the MNIST dataset while other hyperparameters were fixed. Five models were trained for each set of parameters. Three different attack methods (FGSM, PGD, and DeepFool (I^2)) were conducted on the models, and the mean and standard error of the robust accuracy were measured. The plots show that models trained with larger

N_u show higher robust accuracy. Models trained with and without our naive regularizer ($\lambda = 100k, 0$) were also compared for different numbers of labeled data points per class ($N_l = 8, 64, 256$) with other hyperparameters fixed. Five different models were trained similarly per set of parameters. All trained models had similar validation accuracy. More details on the experiment design and hyperparameters are provided in Appendix C of the Supplementary Materials. Fig. 8 (b) (Left) shows the mean and standard error of the robust accuracy values measured for the trained models suggesting that models trained with our regularizer were more robust compared to models trained without our regularizer. We ran similar experiments on the CIFAR10 dataset. Fig. 8 (a) (Right) shows the robust accuracy of models trained with $N_u = 1, 5, 10$, $\lambda = 10k, N_l = 128$ with other hyperparameters fixed for the CIFAR10 dataset. All trained models had similar validation accuracy. As shown, the models trained with larger N_u show higher robust accuracy. Fig. 8 (b) (Right) shows the mean and standard error of the robust accuracy values measured for models trained with and without our regularizer ($\lambda = 10k, 0$) and different number of labeled data points per class ($N_l = 128, 256, 512$) with other hyperparameters fixed. For all three parameter settings, models trained with our naive regularizer were more robust than models trained without our naive regularizer. The experiments show that regularization and unlabeled data both have positive influence on adversarial robustness of trained models. We only show the adversarial robustness of the models using DeepFool (I^2) attack because the models show a very small decrease in adversarial robustness when attacked with DeepFool (I^∞).

4.6. The Effect of the Network Layer where Gaussian Noise is Applied

In the previous experiments, the unlabeled data was generated once and passed through the model for regularization. In this section, we conduct experiments to ascertain if applying Gaussian noise to intermediate layers in the network would train models with almost the same adversarial robustness or training models with the generated unlabeled data results in more adversarially robust models. Noise injection in the intermediate layers was conducted in the same way the unlabeled dataset was generated. The mean pair-wise Euclidean distance (μ_{pair}) between the outputs of the previous layer in the network was measured for the labeled dataset and multiples of μ_{pair} was used as the standard deviation of the Gaussian distribution from which noise was sampled. We trained three models with intermediate noise injection explained in detail in the Appendix D of the Supplementary Materials for each dataset. Table 6 shows the measured $\hat{\rho}_{adv}$ values, and test accuracy for these models and Fig. 9 shows the robust accuracy

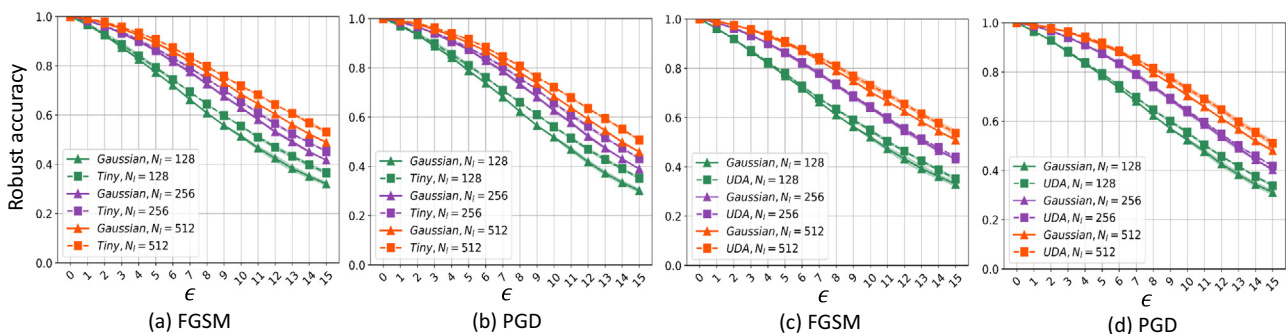


Fig. 7. Figure shows robustness of two models trained with regularization using a Gaussian generated dataset as the unlabeled data and using the 80 m@200k dataset as the unlabeled data against FGSM (a) and PGD (b) attacks for models trained with 128, 256, and 512 labeled data points per class (N_l) from the CIFAR10 dataset. (c) and (d) show robustness of two models trained with regularization using a Gaussian generated dataset as the unlabeled and using the UDA generated dataset [56] as the unlabeled data against FGSM and PGD attacks respectively. The reader is referred to the online version of this paper for the color representation of this figure.

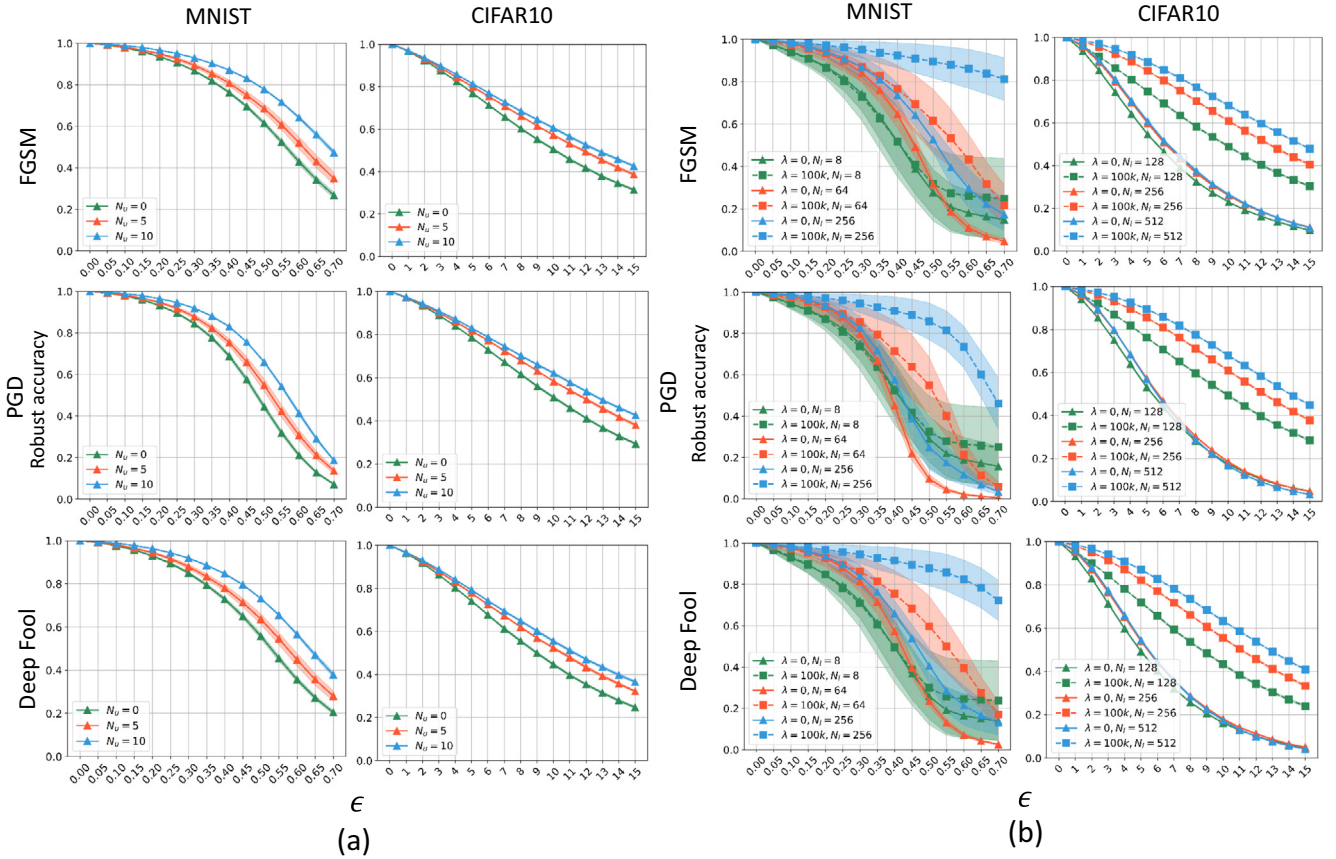


Fig. 8. (a) Shows the robust accuracy for models trained on the MNIST dataset using regularization with $\lambda = 100k$ and $N_l = 64$ (left), and robust accuracy for models trained on the CIFAR10 dataset using regularization with $\lambda = 10k$ and $N_l = 128$ for different numbers of unlabeled data points per labeled data points, N_u using FGSM, PGD, and DeepFool (l^2) attacks (right). (b) Shows the robust accuracy for models trained on the MNIST dataset with and without regularization ($\lambda = 100k, 0$) with different numbers of labeled data points, $N_l = 8, 64, 256$ (left). (Right) Shows the robust accuracy for models trained on the CIFAR10 dataset with and without regularization ($\lambda = 10k, 0$) with different numbers of labeled data points, N_l , using FGSM, PGD, and DeepFool (l^2) attacks. The reader is referred to the online version of this paper for the color representation of this figure.

Table 6

Columns $\hat{\rho}_{adv}(\times 10^{-2})$ and Test Accuracy (%) show measured $\hat{\rho}_{adv}$ values using DeepFool attack [7] and test accuracy for the models trained on the MNIST, CIFAR10, and Imagenette datasets. This table compares the robustness of models trained using the naive regularizer with unlabeled data (Reg + Unlabeled) with models trained using the naive regularizer and intermediate noise injection (Reg + Inject).

Defence method	MNIST		CIFAR10		Imagenette	
	$\hat{\rho}_{adv}$ ($\times 10^{-2}$)	Test accuracy (%)	$\hat{\rho}_{adv}$ ($\times 10^{-2}$)	Test accuracy (%)	$\hat{\rho}_{adv}$ ($\times 10^{-2}$)	Test accuracy (%)
Reg + Unlabeled	46.29	98.69	3.29	83.96	2.21	84.92
Reg + Inject	23.51	98.41	1.39	91.02	2.29	86.43

for these models compared with models trained with our naive regularizer and unlabeled data. The results in Table 6 show that the $\hat{\rho}_{adv}$ values are lower when noise is injected to the intermediate layers of the models during training compared to when models are trained with unlabeled data for the MNIST and CIFAR10 datasets. The $\hat{\rho}_{adv}$ value is higher when noise is injected to the intermediate layers of the models during training for the Imagenette dataset. Fig. 9 also shows that the robust accuracy of models trained with noise injection decreases faster than models trained with unlabeled data when epsilon perturbation value, ϵ , increases except for models trained on the Imagenette dataset and evaluated against FGSM attack.

We also compared models trained with Jacobian regularization and unlabeled data and Jacobian regularization and

intermediate noise injection to investigate if the adversarial robustness would similarly decrease faster for models trained with intermediate noise injection as observed for our naive regularizer. Three models were trained for each dataset in the same way. The $\hat{\rho}_{adv}$ values and test accuracy is shown in Table 4 with labels ‘‘Jacobian + Unlabeled’’ and ‘‘Jacobian + Inject’’, respectively. The corresponding robust accuracy is also shown in Fig. 6. The results show models trained with intermediate noise injection are more robust against adversarial attacks compared to models trained with unlabeled data when the Jacobian regularization is utilized. This suggests that noise injection at intermediate layers in the same way the unlabeled data was generated trains more adversarially robust models when using Jacobian regularization term.

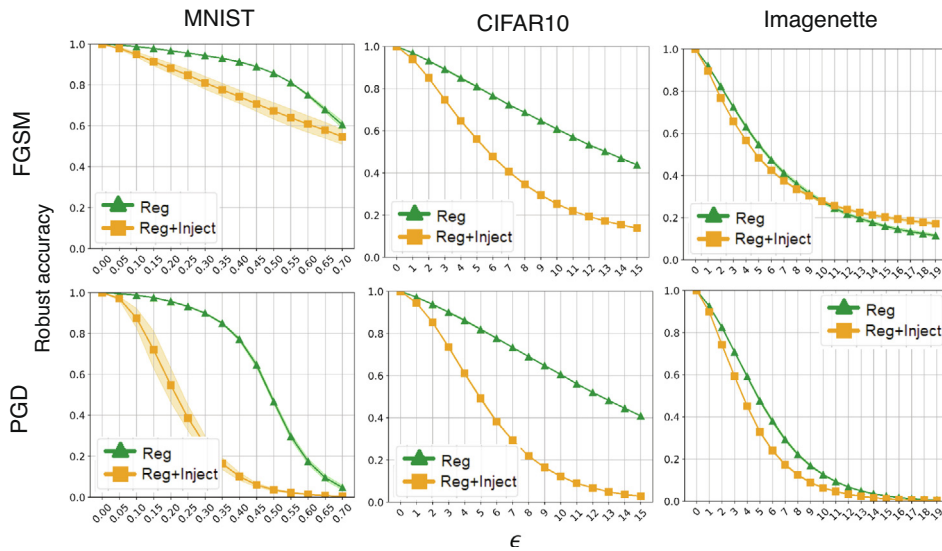
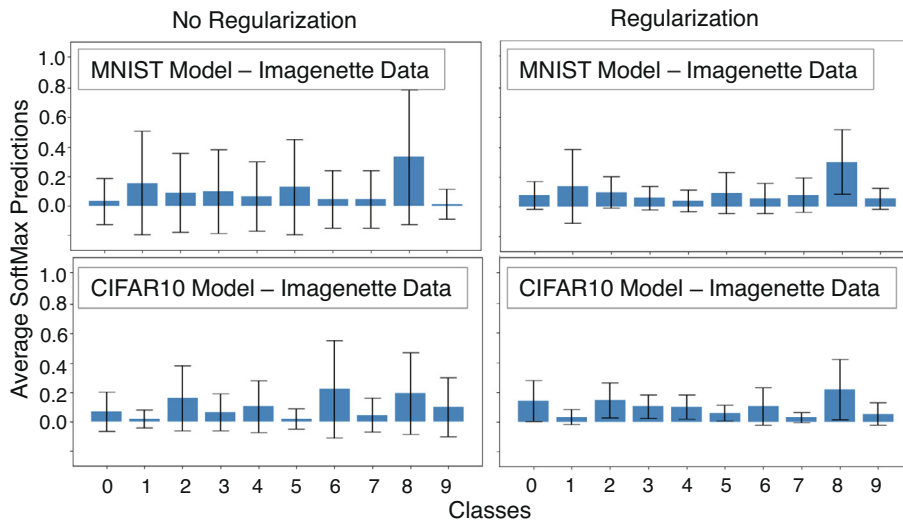


Fig. 9. Figure shows the robust accuracy of models trained using the naive regularizer with unlabeled data (Reg) and the naive regularizer with noise injection (Reg + Inject). The robust accuracy is measured using FGSM [5] and PGD [6] attacks with different epsilon perturbation values (ϵ). The reader is referred to the online version of this paper for the color representation of this figure.

Table 7

Table compares prediction confidence of DL models on UUC data points taken from Imagenette dataset. The DL models are trained on MNIST and CIFAR10 datasets using our training framework and without regularization for five different seeds.



4.7. Prediction Confidence Decreases for Unknown Unknown Classes

Open Set Recognition (OSR) describes the scenario where a DL model is trained with data points from the training dataset known as known known data points (KCCs) similar to normal training. However, during testing, new classes which are not seen during training also known as unknown unknown data points (UUCs) appear among KCC data points. The model is expected to correctly identify KCC data points and reject UUC data points [69]. Many proposed methods train DL models with an additional class label for the UUCs and aim for a decision boundary that tightly binds the support for each class distribution. The work by Shao et al. [70] showed that such OSR methods are vulnerable to adversarial attacks and proposed a defense mechanism to detect and remove adversarial noise. Other OSR algorithms do not add a new class label, but instead train models to assign UUC data points to every class with low probability [71]. Since UUC distributions are

hypothesized to be away from the class distributions seen in the training dataset, we evaluated whether our trained models with the proposed training framework are able to provide OSR by decreasing the prediction confidence for UUC data points.

We fed the models trained on the MNIST dataset with the test set for Imagenette dataset and measured the mean and standard deviation of the predictions for each class. We performed the same experiment for the models trained on the CIFAR10 dataset. The results were compared with the mean and standard deviation of predictions for models trained using normal training. Table 7 shows the average class prediction and the standard deviation of predictions for each class. As shown, the standard deviation of predictions for each class decreases and the predictions are spread more evenly among the classes for models trained using regularization and unlabeled data. In other words, the prediction confidence of models on UUC data points is consistently lower for models trained with regularization compared to model trained

without regularization implying that the SoftMax score surface at the location of UUC distributions in the high dimensional input data space is flattened.

5. Discussion and Conclusion

In this paper, we hypothesized that data sparsity and large numbers of redundant parameters in DL models allow different DL models to learn decision boundaries with different appearances in the space outside the support of the data distribution. Such models can have equally high prediction accuracy on clean datasets, however, their robustness against adversarial attacks can be very different. We also hypothesized that a model with a smooth SoftMax score surface whose decision boundary stays far from each class distribution is ideal in terms of adversarial robustness. Following these hypotheses, we automatically generated an unlabeled dataset to cover the space outside the support of the data distribution using Gaussian noise. We then used a regularization term to train DL models with smooth SoftMax score surfaces using both the original and unlabeled dataset. Using the generated unlabeled data, we were able to train DL models to have smooth SoftMax score surfaces outside the support of the input data distribution. We empirically showed that the proposed training framework indeed increases the distance from data points to the decision boundary for models trained on MNIST, CIFAR10 and Imagenette datasets. We also showed that the proposed training framework trains models with SoftMax score surfaces similar to the hypothesized ideal SoftMax score surface for a 2D Points dataset. Although we believe the Jacobian regularizer is a strong regularization term which can train DL models with smooth SoftMax score surfaces, we deployed a more naive regularization term with shorter runtime for our rigorous experiments. However, in multiple occasions we deployed the Jacobian regularizer and showed that the latter generally performs better than the former. We also showed that using an unlabeled dataset generated by Gaussian noise while considering the spread of the labeled data distribution in the high dimensional input data space is almost as effective as unlabeled datasets sourced from expensive and manually collected datasets or generated using complex synthesis algorithms. Finally, we demonstrated that models trained with our framework also show consistently less confidence in classifying samples from unknown classes as expected in the open set recognition domain. This paper empirically tests our hypothesis about data sparsity, DL models and their decision boundaries. Our future objective is to theoretically prove that Jacobian regularization and our unlabeled data indeed train models with smooth SoftMax score surface with decision boundaries far from each class distribution.

CRedit authorship contribution statement

Mahsa Paknezhad: Conceptualization, Methodology, Validation, Formal analysis, Writing - original draft, Writing - review & editing. **Cuong Phuc Ngo:** Methodology, Validation, Formal analysis. **Amadeus Aristo Winarto:** Methodology, Validation, Formal analysis. **Alistair Cheong:** Validation, Formal analysis, Visualization. **Beh Chuen Yang:** Validation, Formal analysis, Visualization. **Wu Jiayang:** Validation, Formal analysis, Visualization. **Lee Hwee Kuan:** Conceptualization, Methodology, Resources, Supervision.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.neucom.2022.01.062>.

References

- [1] W. Rawat, Z. Wang, Deep convolutional neural networks for image classification: A comprehensive review, *Neural computation* 29 (9) (2017) 2352–2449, https://doi.org/10.1162/neco_a_00990.
- [2] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen, et al., Deep speech 2: End-to-end speech recognition in english and mandarin, in: *International conference on machine learning PMLR*, 2016, pp. 173–182.
- [3] T. Young, D. Hazarika, S. Poria, E. Cambria, Recent trends in deep learning based natural language processing, *IEEE Computational Intelligence Magazine* 13 (3) (2018) 55–75, <https://doi.org/10.1109/MCI.2018.2840738>.
- [4] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, R. Fergus, Intriguing properties of neural networks, *arXiv preprint arXiv:1312.6199*.
- [5] I.J. Goodfellow, J. Shlens, C. Szegedy, Explaining and harnessing adversarial examples, *arXiv preprint arXiv:1412.6572*.
- [6] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, A. Vladu, Towards deep learning models resistant to adversarial attacks, *arXiv preprint arXiv:1706.06083*.
- [7] S.M. Moosavi-Dezfooli, A. Fawzi, 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.
- [8] W.J. Scheirer, A. de Rezende Rocha, A. Sapkota, T.E. Boult, Toward open set recognition, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35 (7) (2012) 1757–1772, <https://doi.org/10.1109/TPAMI.2012.256>.
- [9] H. Zhang, J. Wang, Defense against adversarial attacks using feature scattering-based adversarial training, *Advances in Neural Information Processing Systems* 32 (2019) 1831–1841.
- [10] B.R. Chernyak, B. Raj, T. Hazan, J. Keshet, Constant random perturbations provide adversarial robustness with minimal effect on accuracy, *arXiv preprint arXiv:2103.08265*.
- [11] V. Zantedeschi, M.-I. Nicolae, A. Rawat, Efficient defenses against adversarial attacks, in: *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, 2017, pp. 39–49.
- [12] Y. Li, B. Wu, Y. Feng, Y. Fan, Y. Jiang, Z. Li, S. Xia, Toward adversarial robustness via semi-supervised robust training, *arXiv preprint arXiv:2003.06974*.
- [13] G. Goswami, N. Ratha, A. Agarwal, R. Singh, M. Vatsa, Unravelling robustness of deep learning based face recognition against adversarial attacks, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32, 2018.
- [14] K. Grosse, P. Manoharan, N. Papernot, M. Backes, P. McDaniel, On the (statistical) detection of adversarial examples, *arXiv preprint arXiv:1702.06280*.
- [15] J. Peck, B. Goossens, Y. Saeys, Detecting adversarial manipulation using inductive venn-abers predictors, *Neurocomputing* 416 (2020) 202–217, <https://doi.org/10.1016/j.neucom.2019.11.113>.
- [16] Z. Gong, W. Wang, W.-S. Ku, Adversarial and clean data are not twins, *arXiv preprint arXiv:1704.04960*.
- [17] R. Feinman, R.R. Curtin, S. Shintre, A.B. Gardner, Detecting adversarial samples from artifacts, *arXiv preprint arXiv:1703.00410*.
- [18] N. Das, M. Shanhogue, S.-T. Chen, F. Hohman, L. Chen, M.E. Kounavis, D.H. Chau, Keeping the bad guys out: Protecting and vaccinating deep learning with jpeg compression, *arXiv preprint arXiv:1705.02900*.
- [19] B. Liang, H. Li, M. Su, X. Li, W. Shi, X. Wang, Detecting adversarial image examples in deep neural networks with adaptive noise reduction, *IEEE Transactions on Dependable and Secure Computing* doi:10.1109/TDSC.2018.2874243.
- [20] J. Lu, T. Issaranon, D. Forsyth, Safetynet: Detecting and rejecting adversarial examples robustly, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 446–454.
- [21] C. Xie, J. Wang, Z. Zhang, Z. Ren, A. Yuille, Mitigating adversarial effects through randomization, *arXiv preprint arXiv:1711.01991*.
- [22] F. Liao, M. Liang, Y. Dong, T. Pang, X. Hu, J. Zhu, Defense against adversarial attacks using high-level representation guided denoiser, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1778–1787.
- [23] C. Xie, Y. Wu, L.V.D. Maaten, A.L. Yuille, K. He, Feature denoising for improving adversarial robustness, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 501–509.
- [24] I.M. Baytas, D. Deb, Robustness-via-synthesis: Robust training with generative adversarial perturbations, *arXiv preprint arXiv:2108.09713*.
- [25] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, P. McDaniel, Ensemble adversarial training: Attacks and defenses, *arXiv preprint arXiv:1705.07204*.
- [26] C. Qin, J. Martens, S. Gowal, D. Krishnan, K. Dvijotham, A. Fawzi, S. De, R. Stanforth, P. Kohli, Adversarial robustness through local linearization, *arXiv preprint arXiv:1907.02610*.
- [27] R. Theagarajan, M. Chen, B. Bhanu, J. Zhang, Shieldnets: Defending against adversarial attacks using probabilistic adversarial robustness, in: *Proceedings*

- of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 6988–6996.
- [28] E. Wong, Z. Kolter, Provable defenses against adversarial examples via the convex outer adversarial polytope, *International Conference on Machine Learning*, PMLR (2018) 5286–5295.
- [29] E. Wong, F.R. Schmidt, J.H. Metzen, J.Z. Kolter, Scaling provable adversarial defenses, arXiv preprint arXiv:1805.12514.
- [30] S. Wang, Y. Chen, A. Abdou, S. Jana, Mixtrain: Scalable training of verifiably robust neural networks, arXiv preprint arXiv:1811.02625.
- [31] J. Dapello, T. Marquesand, M. Schrimpf, F. Geiger, D.D.C.J.J. DiCarlo, Simulating a primary visual cortex at the front of cnns improves robustness to image perturbations, *BioRxiv* doi:10.1101/2020.06.16.154542.
- [32] T. Gittings, S. Schneider, J. Collomosse, Vax-a-net: Training-time defence against adversarial patch attacks, in: *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [33] U. Shaham, Y. Yamada, S. Negahban, Understanding adversarial training: Increasing local stability of supervised models through robust optimization, *Neurocomputing* 307 (2018) 195–204, <https://doi.org/10.1016/j.neucom.2018.04.027>.
- [34] B. Li, C. Chen, W. Wang, L. Carin, Certified adversarial robustness with additive noise, arXiv preprint arXiv:1809.03113.
- [35] Y. Yang, G. Zhang, D. Katabi, Z. Xu, Me-net: Towards effective adversarial robustness with matrix estimation, arXiv preprint arXiv:1905.11971.
- [36] Y. Guo, C. Zhang, C. Zhang, Y. Chen, Sparse dnns with improved adversarial robustness, arXiv preprint arXiv:1810.09619.
- [37] T. Pang, K. Xu, C. Du, N. Chen, J. Zhu, Improving adversarial robustness via promoting ensemble diversity, in: *International Conference on Machine Learning* PMLR, 2019, pp. 4970–4979.
- [38] X. Zhou, I.W. Tsang, J. Yin, Latent adversarial defence with boundary-guided generation, arXiv preprint arXiv:1907.07001.
- [39] J. Cohen, E. Rosenfeld, Z. Kolter, Certified adversarial robustness via randomized smoothing, *International Conference on Machine Learning*, PMLR (2019) 1310–1320.
- [40] A. Goel, A. Agarwal, M. Vatsa, R. Singh, N.K. Ratha, Dndnet: Reconfiguring cnn for adversarial robustness, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 22–23.
- [41] E. Rusak, L. Schott, R.S. Zimmermann, J. Bitterwolf, O. Bringmann, M. Bethge, W. Brendel, A simple way to make neural networks robust against diverse image corruptions, in: *European Conference on Computer Vision*, Springer, 2020, pp. 53–69.
- [42] X. Liu, S. Sun, Alpha-divergence minimization with mixed variational posterior for bayesian neural networks and its robustness against adversarial examples, *Neurocomputing* 423 (2021) 427–434, <https://doi.org/10.1016/j.neucom.2020.10.087>.
- [43] C.-J. Simon-Gabriel, Y. Olivier, L. Bottou, B. Schölkopf, D. Lopez-Paz, First-order adversarial vulnerability of neural networks and input dimension, in: *International Conference on Machine Learning* PMLR, 2019, pp. 5809–5817.
- [44] J. Xu, Y. Li, Y. Jiang, S.-T. Xia, Adversarial defense via local flatness regularization, in: *2020 IEEE International Conference on Image Processing (ICIP) IEEE*, 2020, pp. 2196–2200.
- [45] K. Roth, A. Lucchi, S. Nowozin, T. Hofmann, Adversarially robust training through structured gradient regularization, arXiv preprint arXiv:1805.08736.
- [46] S. Singla, S. Feizi, Second-order provable defenses against adversarial attacks, *International Conference on Machine Learning*, PMLR (2020) 8981–8991.
- [47] J. Gilmer, L. Metz, F. Faghri, S.S. Schoenholz, M. Raghu, M. Wattenberg, I. Goodfellow, Adversarial spheres, arXiv preprint arXiv:1801.02774.
- [48] H. Zhang, Y. Yu, J. Jiao, E. Xing, L.E. Ghaoui, M. Jordan, Theoretically principled trade-off between robustness and accuracy, in: *International Conference on Machine Learning* PMLR, 2019, pp. 7472–7482.
- [49] S.M. Moosavi-Dezfooli, A. Fawzi, J. Uesato, P. Frossard, Robustness via curvature regularization, and vice versa, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9078–9086.
- [50] Z. Yan, Y. Guo, C. Zhang, Deep defense: Training dnns with improved adversarial robustness, arXiv preprint arXiv:1803.00404.
- [51] D. Jakubovitz, R. Giryas, Improving dnn robustness to adversarial attacks using jacobian regularization, in: *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 514–529.
- [52] A. Ross, F. Doshi-Velez, Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32, 2018.
- [53] J. Ba, R. Caruana, Do deep nets really need to be deep?, in: Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, K.Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems*, Vol. 27, Curran Associates Inc, 2014.
- [54] H. Drucker, Y.L. Cun, Improving generalization performance using double backpropagation, *IEEE Transactions on Neural Networks* 3 (6) (1992) 991–997, <https://doi.org/10.1109/72.165600>.
- [55] Y. Carmon, A. Raghunathan, L. Schmidt, P. Liang, J.C. Duchi, Unlabeled data improves adversarial robustness, arXiv preprint arXiv:1905.13736.
- [56] Q. Xie, Z. Dai, E. Hovy, M.T. Luong, Q.V. Le, Unsupervised data augmentation for consistency training, arXiv preprint arXiv:1904.12848.
- [57] V. Verma, A. Lamb, C. Beckham, A. Najafi, I. Mitliagkas, D. Lopez-Paz, Y. Bengio, Manifold mixup: Better representations by interpolating hidden states, in: *International Conference on Machine Learning* PMLR, 2019, pp. 6438–6447.
- [58] T. Miyato, S.I. Maeda, M. Koyama, S. Ishii, Virtual adversarial training: a regularization method for supervised and semi-supervised learning, *IEEE transactions on pattern analysis and machine intelligence* 41 (8) (2018) 1979–1993, <https://doi.org/10.1109/TPAMI.2018.2858821>.
- [59] R. Zhai, T. Cai, D. He, C. Dan, K. He, J. Hopcroft, L. Wang, Adversarially robust generalization just requires more unlabeled data, arXiv preprint arXiv:1906.00555.
- [60] J. Uesato, J.B. Alayrac, P.S. Huang, R. Stanforth, A. Fawzi, P. Kohli, Are labels required for improving adversarial robustness?, arXiv preprint arXiv:1905.13725.
- [61] A. Najafi, S.I. Maeda, M. Koyama, T. Miyato, Robustness to adversarial perturbations in learning from incomplete data, arXiv preprint arXiv:1905.13021.
- [62] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, *Proceedings of the IEEE* 86 (11) (1998) 2278–2324, <https://doi.org/10.1109/5.726791>.
- [63] A. Krizhevsky, G. Hinton, et al., Learning multiple layers of features from tiny images doi:10.1.1.222.9220.
- [64] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al., Imagenet large scale visual recognition challenge, *International journal of computer vision* 115 (3) (2015) 211–252, <https://doi.org/10.1007/s11263-015-0816-y>.
- [65] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [66] T. He, Z. Zhang, H. Zhang, Z. Zhang, J. Xie, M. Li, Bag of tricks for image classification with convolutional neural networks, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 558–567.
- [67] J. Rauber, R. Zimmermann, M. Bethge, W. Brendel, Foolbox native: Fast adversarial attacks to benchmark the robustness of machine learning models in pytorch, tensorflow, and jax, *Journal of Open Source Software* 5 (53) (2020) 2607, <https://doi.org/10.21105/joss.02607>.
- [68] J.T. Springenberg, Unsupervised and semi-supervised learning with categorical generative adversarial networks, arXiv preprint arXiv:1511.06390.
- [69] C. Geng, S.J. Huang, S. Chen, Recent advances in open set recognition: A survey, *IEEE transactions on pattern analysis and machine intelligence* doi:10.1109/TPAMI.2020.2981604.
- [70] R. Shao, P. Perera, P. Yuen, V.M. Patel, Open-set adversarial defense, arXiv preprint arXiv:2009.00814.
- [71] W.J. Scheirer, L.P. Jain, T.E. Boult, Probability models for open set recognition, *IEEE transactions on pattern analysis and machine intelligence* 36 (11) (2014) 2317–2324, <https://doi.org/10.1109/TPAMI.2014.2321392>.



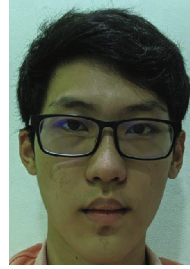
Mahsa Paknezhad is an active researcher in deep learning, reinforcement learning and medical image processing at the Bioinformatics Institute where she has gained invaluable research experience since 2018. She has several publications in machine learning, deep learning and medical image processing. Mahsa did her PhD from National University of Singapore (NUS) in computer science in 2018. She did a doctoral thesis on developing robust and automatic machine learning algorithms for analysing cardiac MRI images. Mahsa has M.Eng. and B.Eng. degrees from Shiraz University of Technology in IT engineering in 2011 and 2013, respectively.




Ngo Phuc Cuong is an undergraduate in Computer Science at the National University of Singapore from 2020. His research interests are focused on Adversarial Attacks, Unsupervised Domain Adaptation and 3D Object Detection. Cuong was a research assistant at the Bioinformatics Institute, A*STAR in 2020 where he worked on multiple deep learning projects.



Amadeus Aristo Winarto is an undergraduate in Computer Science at the National University of Singapore from 2020. He has worked with GANs and anomaly detection using GANs. His research interests are focused on Machine Learning applications, robust classification, and computer vision. Amadeus was a research assistant at the Bioinformatics Institute, A*STAR in 2020 where he worked on multiple deep learning projects.



Wu Jiayang is a Junior College student studying in Hwa Chong Institution, Singapore. His interests include web security, the internet of things, and reinforcement learning. Jiayang was an intern at the Bioinformatics Institute, A*STAR in 2020.



Alistair Cheong Liang Chuen is a Junior College student studying in Hwa Chong Institution, Singapore. He is interested in meta-learning and attempts to create artificial general intelligence capable of learning multi-modal tasks. Alistair was an intern at the Bioinformatics Institute, A*STAR in 2020.



Hwee Kuan Lee is a Senior Principal Investigator of the Imaging Informatics division in Bioinformatics Institute. Hwee Kuan obtained his Ph.D. in 2001 in Theoretical Physics from Carnegie Mellon University with a thesis on liquid-liquid phase transitions and quasicrystals. He then held a joint postdoctoral position with Oak Ridge National Laboratory and University of Georgia. In 2003, Hwee Kuan moved to Tokyo Metropolitan University. In 2005, he joined the Data Storage Institute. In 2006, he joined Bioinformatics Institute as a Principle Investigator in the Imaging Informatics Division. His current research work involves developing computer vision algorithms for clinical studies.



Beh Chuen Yang is a Junior College student studying in Hwa Chong Institution, Singapore. He is interested in Mathematics, Cybersecurity and Foreign Languages, and likes to work on Cryptography problems. Chuen Yang was an intern at the Bioinformatics Institute, A*STAR in 2020.