

Towards Achieving Adversarial Robustness by Enforcing Feature Consistency Across Bit Planes

Sravanti Addepalli*, Vivek B.S.*, Arya Baburaj, Gaurang Sriramanan, R.Venkatesh Babu
Video Analytics Lab, Department of Computational and Data Sciences
Indian Institute of Science, Bangalore, India

Abstract

As humans, we inherently perceive images based on their predominant features, and ignore noise embedded within lower bit planes. On the contrary, Deep Neural Networks are known to confidently misclassify images corrupted with meticulously crafted perturbations that are nearly imperceptible to the human eye. In this work, we attempt to address this problem by training networks to form coarse impressions based on the information in higher bit planes, and use the lower bit planes only to refine their prediction. We demonstrate that, by imposing consistency on the representations learned across differently quantized images, the adversarial robustness of networks improves significantly when compared to a normally trained model. Present state-of-the-art defenses against adversarial attacks require the networks to be explicitly trained using adversarial samples that are computationally expensive to generate. While such methods that use adversarial training continue to achieve the best results, this work paves the way towards achieving robustness without having to explicitly train on adversarial samples. The proposed approach is therefore faster, and also closer to the natural learning process in humans.

1. Introduction

Deep Neural Networks have been used to achieve remarkable performance in many Computer Vision tasks, such as Classification [14], Segmentation [19] and Object recognition [24]. While these networks achieve near-human accuracy on many benchmark datasets, they are far from being as robust as the human visual system. Deep Networks are known to be vulnerable to carefully crafted imperceptible noise known as Adversarial Perturbations [30], which could have disastrous implications in critical applications such as autonomous navigation and surveillance systems. The compelling need of securing these systems, coupled with the goal of improving the worst-case robustness of

Deep Networks has propelled research in the area of Adversarial Robustness over the last few years. While adversarial training methods [21, 40] have led to significant progress in improving adversarial robustness, these methods are computationally expensive and also non-intuitive when compared to the learning process in humans.

Humans perceive images based on features of large magnitude and use finer details only to enhance their impressions [29, 28]. This *background knowledge* of giving higher importance to information present in higher bit planes naturally equips the human visual system to develop resistance towards adversarial perturbations, which are of relatively lower magnitude. On the contrary, these adversarial perturbations can arbitrarily flip the predictions of Deep Networks to completely unrelated classes, suggesting that such *background knowledge* of giving hierarchical importance to different bit planes is missing in these networks. In this work, we propose to equip Deep Networks with such knowledge, and demonstrate that this improves their robustness to adversarial examples.

We propose a novel *Bit Plane Feature Consistency (BPFC)* regularizer, which can significantly improve adversarial robustness of models, without exposure to adversarial samples during training. The proposed method is considerably faster than methods that require multi-step adversarial samples for training [21], and is therefore scalable to large datasets such as ImageNet. Through this work, we hope to pave the path towards training robust Deep Networks without using adversarial samples, similar to the learning process that exists in human beings.

The organization of this paper is as follows: The subsequent section presents a discussion on the existing literature related to our work. Section-3 lays out the preliminaries related to notation and threat model. This is followed by details and analysis of our proposed approach in Section-4. We present the experiments performed and an analysis on the results in Section-5, followed by our concluding remarks in Section-6.

The code and pretrained models are available at:
<https://github.com/val-iisc/BPFC>.

*Equal contribution

2. Related Works

2.1. Adversarial Training methods

The most popular methods of improving adversarial robustness of Deep Networks involve Adversarial Training (AT), where clean data samples are augmented with adversarial samples during training. Early formulations such as FGSM-AT [9] proposed training on adversarial samples generated using single-step optimization, that assumes a first-order linear approximation to the loss function. This was later shown to be ineffective against multi-step attacks by Kurakin *et al.* [16], wherein the effect of gradient masking was identified. Gradient masking, first identified by Papernot *et al.* [23], is the phenomenon where the trained network yields masked gradients, thereby resulting in the generation of weak adversaries, leading to a false sense of robustness. In a wide variety of settings, numerous countermeasures have been developed that can produce strong adversaries to circumvent gradient masking [31].

Madry *et al.* [21] proposed Projected Gradient Descent (PGD) based training, that employs an iterative procedure for finding strong adversaries that maximise training loss under the given norm constraint. Crucially, PGD trained models are robust against various gradient-based iterative attacks, as well as several variants of non-gradient based attacks. However, the process of PGD Adversarial Training (PGD-AT) is computationally expensive.

In order to address this, Vivek *et al.* [35] revisited single-step adversarial training, and introduced a regularizer that helps mitigate the effect of gradient masking. This regularizer penalizes the ℓ_2 distance between logits of images perturbed with FGSM and R-FGSM [31] attacks. In the proposed method, we achieve adversarial robustness without using adversarial samples during training, and hence achieve a further reduction in computation time.

2.2. Attempts of Adversary-Free Training

In this section, we discuss existing training methods that do not utilize adversarial samples during training. Works such as Mixup [39] and Manifold-Mixup [34] propose training methods to learn better feature representations. In Mixup, the network is trained to map a random convex combination of the input data to the corresponding convex combination of their one-hot encodings. This work is extended further in Manifold Mixup, where the network is trained to map a convex combination of intermediate hidden-layers generated by two different data points to the corresponding convex combination of their one-hot encodings. Hence, these methods encourage the network to behave in a linearized manner between input data points, or between hidden-layers deeper in the network. While these methods resulted in improved performance against single-step FGSM attacks, they were susceptible to stronger multi-

step attacks.

Another attempt of adversary-free training to achieve robustness utilized input transformations for defense. In the work by Guo *et al.* [10], the effect of various input transformations such as bit-depth reduction, JPEG compression, total variation minimisation and image quilting was studied. The robustness from these techniques primarily originated from the non-differentiable pre-processing steps, in order to possibly thwart gradient-based iterative attacks. This method, along with a few others [2, 20, 6, 37, 27], were broken in the work by Athalye *et al.* [1], where it was identified that obfuscated gradients do not provide reliable security against adversaries.

Yet another avenue pursued was towards the detection of adversarial samples. Feature Squeezing, proposed by Xu *et al.* [38], used transformations such as reduction of color bit depth, spatial smoothing with a median filter and a combination of both, in order to generate a feature-squeezed image from a given input image. By thresholding the ℓ_1 distance between logits of an input image and its feature-squeezed counterpart, the image was classified to be either adversarial or legitimate in nature. However, in the work by He *et al.* [12], it was shown that an adaptive attacker cognizant of this defense strategy could fool the model by constructing attacks that retain adversarial properties even after feature-squeezing is applied, thereby evading detection.

While we use the concept of quantization to defend against adversarial attacks in this work, we do not introduce any pre-processing blocks that lead to obfuscated or shattered gradients.

3. Preliminaries

3.1. Notation

In this paper, we consider $f(\cdot)$ as the function mapping of a classifier C , from an image x , to its corresponding softmax output $f(x)$. The predicted class label, which is an argmax over the softmax output, is denoted by $c(x)$. The ground truth label corresponding to x is denoted by y . The image is said to be correctly classified when $c(x) = y$. The pre-softmax output of the classifier C is denoted by $g(x)$. We define $\mathcal{A}(x)$ to be the set of all *Adversarial Samples* corresponding to x . A specific adversarial sample corresponding to a clean sample x is denoted by x' .

3.2. Threat Model

In this paper, we consider the task of improving the worst-case robustness of Deep Networks. The goal of an adversary is to cause an error in the prediction of the classifier. We define an *Adversarial Sample* x' , as one that causes the output of the network to be different from the ground truth label y . We do not restrict the adversary to flip labels from a specific source class, or to a specific target class. We



Figure 1: (a) Original 8-bit image (b) Weighted sum of (higher) bit planes 7, 6 and 5 (c) Weighted sum of (higher) bit planes 7 and 6 (d) Bit plane 7 - Most significant bit plane (e) Weighted sum of (lower) bit planes 4, 3, 2, 1 and 0.

restrict x' to be in the ℓ_∞ -ball of radius ε around x . The set of *Adversarial Samples* can be formally defined as follows:

$$\mathcal{A}(x) = \{x' : c(x') \neq y, \|x - x'\|_\infty \leq \varepsilon\} \quad (1)$$

We therefore impose a constraint that any individual pixel in the image x cannot be perturbed by more than ε .

Since the goal of this work is to improve worst-case robustness, we do not impose any restrictions on the access to the adversary. We consider that the adversary has complete knowledge of the model architecture, weights and the defense mechanism employed.

4. Proposed Method

In this section, we first present the motivation behind our proposed method, followed by a detailed discussion of the proposed algorithm. We further describe local properties of networks trained using the proposed regularizer, which lead to improved robustness.

4.1. Hierarchical Importance of Bit Planes

Bit planes of an image are the spatial maps (of the same dimension as the image) corresponding to a given bit position. For an n -bit representation of an image, bit plane $n - 1$ corresponds to the most significant bit (MSB), and bit plane 0 corresponds to the least significant bit (LSB). An n -bit image can be considered as the sum of n bit planes weighted by their relative importance. The importance of features embedded within lower bit planes is significantly lower than that of features embedded within higher bit planes, both in terms of pixel value, and information content [26]. Steganography methods [8] utilize lower bit planes to embed crucial copyright information that needs to be visually imperceptible. However, information content in natural images decreases from the most significant bit (MSB) to the least significant bit (LSB). A weighted sum of the five least significant bit planes of the image in Fig.1(a) is shown in Fig.1(e), from which it is evident that lower bit planes contribute only to fine details. Fig.1(b), (c) and (d) show images ranging from fine to coarse structure, with different levels of quantization. The difference between Fig.1(a) and Fig.1(b) is Fig.1(e). While the addition of Fig.1(e) certainly improves the information content, it is not as crucial as the higher bit planes for interpreting the image.

The human visual system is known to give higher importance to global information when compared to fine details [28]. Sugase *et al.* [29] demonstrate that global information is used for coarse classification in early parts of the neural response, while information related to fine details is perceived around 51ms later. This demonstrates a hierarchical classification mechanism, where the response to an image containing both coarse and fine information is aligned with that containing only coarse information.

We take motivation from this aspect of the human visual system, and enforce Deep Networks to maintain consistency across decisions based on features in high bit planes alone (quantized image) and all bit planes (normal image). Such a constraint will ensure that Deep Networks give more importance to high bit planes when compared to lower bit planes, similar to the human visual system. Adversarial examples constrained to the ℓ_∞ -ball utilize low bit planes to transmit information which is inconsistent with that of higher bit planes. The fact that Deep Networks are susceptible to such adversarial noise demonstrates the weakness of these networks, which emanates from the lack of consistency between predictions corresponding to coarse information and fine details. Therefore, enforcing feature consistency across bit planes results in a significant improvement in adversarial robustness when compared to conventionally trained networks.

While we use the base-2 (binary) representation of an image to illustrate the concept of ignoring low magnitude additive noise, the same can be formulated in terms of any other representation (in any other base) as well. Secondly, low magnitude noise does not always reside in low bit planes. It can overflow to MSBs as well, based on the pixel values in the image. We introduce pre-quantization noise in our proposed approach to mitigate these effects. This is illustrated in the following section, where we explain our proposed method in greater detail.

4.2. Proposed Training Algorithm

We present the proposed training method in Algorithm-1. Motivated by the need to learn consistent representations for coarse and fine features of an image, we introduce a regularizer that imposes feature consistency between each image and its quantized counterpart.

Algorithm 1: Bit Plane Feature Consistency

Input: Network f with parameters θ , fixed weight λ , training data $\mathcal{D} = \{(x_i, y_i)\}$ of n -bit images, quantization parameter k , learning rate η , minibatch size M

for minibatch $B \subset \mathcal{D}$ **do**

 Set $L = 0$

for $i = 1$ **to** M **do**

$x_{pre} = x_i + \mathcal{U}(-2^{k-2}, 2^{k-2})$ // Add noise

$x_q = x_{pre} - (x_{pre} \bmod 2^k)$ // Quantization

$x_q = x_q + 2^{k-1}$ // Range Shift

$x_q = \min(\max(x_q, 0), 2^n - 1)$ // Clip

$L = L + ce(f(x_i), y_i) + \lambda \|g(x_i) - g(x_q)\|_2^2$

end for

$\theta = \theta - \frac{1}{M} \cdot \eta \cdot \nabla_{\theta} L$ // SGD update

end for




Figure 2: Quantization of a given 8-bit pixel to 3-bits ($n = 8$, $k = 5$): The probability P of a pixel i being assigned to the quantized values $q(i) = 48, 80$ and 112 is shown here.

4.2.1 Quantization

The steps followed for generating a coarse image are described in this section. The input image x_i is assumed to be represented using n -bit quantization. The intensity of pixels is hence assumed to be in the range $[0, 2^n)$. We generate an $n - k + 1$ bit image using the quantization process described here. The allowed range of k is between 1 and $n - 1$.

- **Pre-quantization noise:** Initially, uniform noise sampled independently from $\mathcal{U}(-2^{k-2}, 2^{k-2})$ is added to each pixel in the image x_i , to generate x_{pre} .
- **Quantization step:** Next, each pixel is quantized to $n - k$ bits, by setting the last k bits to 0.
- **Range Shift:** The intensity of all pixels is shifted up by 2^{k-1} . This shifts the range of quantization error (w.r.t. x_{pre}) from $[0, 2^k)$ to $[-2^{k-1}, 2^{k-1})$.
- **Clip:** Finally, the quantized image is clipped to the original range, $[0, 2^n)$.

Fig.2 illustrates the probability of a pixel i being assigned to different quantization levels using the above method, when $n = 8$ and $k = 5$. As shown in the figure, addition of pre-quantization noise introduces randomness in the quantized value of each pixel. The probability of being assigned the nearest quantization level is 1 when the input

pixel intensity is close to the quantized value, and decays linearly to 0 after a fixed distance. This helps mitigate unwanted non-linear effects at the edges of each quantization bin, due to specific pixel intensity values of the original image.

We do an ablation experiment (in Section-5) called Simple Quantization, where the pre-quantization noise is not added. Simple Quantization can be also viewed as addition of correlated low magnitude (quantization) noise, where the noise pattern depends on local pixel intensity levels. This noise is the difference between the original image, and the image subject to simple quantization. Since the pixel intensities can be assumed to be locally correlated in space, the noise is also correlated locally. The correlated nature of noise differentiates quantization noise from random noise, and also brings it closer to the properties of adversarial perturbations. We also consider an ablation experiment of replacing the quantization step with addition of random noise sampled from a uniform distribution in Section-5.

While pre-quantization noise disturbs the local correlation properties of the quantization noise for some of the pixels, it is crucial to mitigate the bin edge effects discussed above. We demonstrate through experiments that the proposed solution is better than both the ablation experiments discussed above.

4.2.2 Bit Plane Feature Consistency Regularizer

The loss function used for training is shown below:

$$L = \frac{1}{M} \sum_{i=1}^M ce(f(x_i), y_i) + \lambda \|g(x_i) - g(q(x_i))\|_2^2 \quad (2)$$

For a given image x_i , the first term of Eq. (2) is the cross-entropy (ce) loss obtained from the softmax output of the network $f(x_i)$, and the corresponding ground truth label y_i . The second term is the squared ℓ_2 distance between the pre-softmax activation of the image x_i , and that of the corresponding quantized image $q(x_i)$ (generated using the process described in Section-4.2.1). We call this squared ℓ_2 loss term as the *Bit Plane Feature Consistency (BPFC)* regularizer, as it ensures that the network learns consistent feature representations across the original image as well as the coarse quantized image. The loss for each minibatch of size M is an average over all samples in the minibatch.

The cross-entropy term on original images ensures that a combination of coarse and fine features is used to learn the overall function mapping $g(\cdot)$. This helps preserve the accuracy on clean images, while the *BPFC* regularizer helps improve the adversarial robustness of the model.

4.3. Local Properties of BPFC Trained Networks

In this section, we examine local properties of the function $g(\cdot)$ learned using the proposed *BPFC* regularizer.

Let x_i denote an n -bit image sampled from the data distribution \mathbb{P}_D with pixel intensities in the range $[0, 2^n)$, and let $q(x_i)$ denote a quantized image corresponding to x_i . We assume that $q(x_i)$ is not identically equal to x_i . For a fixed value of λ , let $\Theta_{g(\lambda)}$ denote the set of parameters corresponding to a family of functions that lead to the cross-entropy term in Eq. (2) being below a certain threshold. Minimization of *BPFC* loss among the family of functions parameterized by $\Theta_{g(\lambda)}$ is shown in Eq. (3):

$$\min_{\theta_g \in \Theta_{g(\lambda)}} \mathbb{E}_{x_i \sim \mathbb{P}_D} \mathbb{E}_{q(x_i)} \|g(x_i) - g(q(x_i))\|_2^2 \quad (3)$$

$$\min_{\theta_g \in \Theta_{g(\lambda)}} \mathbb{E}_{x_i \sim \mathbb{P}_D} \mathbb{E}_{q(x_i)} \frac{\|g(x_i) - g(q(x_i))\|_2^2}{\|x_i - q(x_i)\|_2^2} \quad (4)$$

The expression in Eq. (3) can be lower bounded by the expression in Eq. (4), which is equivalent to minimizing the local Lipschitz constant of the network at each sample x_i . The denominator of the objective function in Eq. (4) is the ℓ_2 -norm between each image and its quantized counterpart, and is thus independent of θ_g . Therefore, minimization of *BPFC* loss in Eq. (3) can be viewed as minimization of the local Lipschitz constant at each sample x_i , weighted by an ℓ_2 -norm of its deviation from the quantized image. An expectation of this ℓ_2 -norm term over all $q(x_i)$ (with differently sampled pre-quantization noise) converges to a constant value for all samples, thereby establishing an equivalence between the minimization of *BPFC* loss and the minimization of local Lipschitz constant of the network.

Hence, imposing *BPFC* regularizer encourages the network to be locally Lipschitz continuous with a reduced Lipschitz constant. While the *BPFC* regularizer imposes local smoothness, the cross-entropy term in Eq. (2) requires $g(\cdot)$ to be a complex mapping for better accuracy on clean images. The final selection of θ_g would depend on λ , which is typically selected based on the amount by which clean accuracy can be traded-off for adversarial accuracy [40, 32]. During the initial epochs of training, the function learned is relatively smooth. Hence, we start with a low value of λ and step it up during training.

Therefore, the *BPFC* formulation leads to functions with improved local properties, which is closely related to adversarial robustness as explained by Szegedy *et al.* [30].

5. Experiments and Analysis

In this section, we discuss the experiments done to verify the robustness of our proposed approach. We first describe the datasets used and details on the training methodology in Section-5.1, followed by an overview of the experiments conducted in Section-5.2. We further present details on each experiment and our analysis on the results in Sections-5.3 to 5.6. We follow the guidelines laid out by Athalye *et al.* [1] and Carlini *et al.* [3] to ascertain the validity of our claim on the achieved robustness.

5.1. Preliminaries

We use the benchmark datasets, CIFAR-10 [13], Fashion-MNIST (F-MNIST) [36] and MNIST [17] for validating our proposed approach. CIFAR-10 is a ten-class dataset with RGB images of dimension 32×32 . The number of images in the training set and test set are 50,000 and 10,000 respectively. These images are equally distributed across all classes. We set aside 10,000 images from the training set as the validation set. Fashion-MNIST and MNIST are ten-class datasets with gray-scale images of dimension 28×28 . The datasets are composed of 60,000 training samples and 10,000 test samples each. We further split each of the training datasets into 50,000 training samples and 10,000 validation samples.

We use ResNet-18 [11] architecture for CIFAR-10, and a modified LeNet (M-LeNet) [18] architecture with two additional convolutional layers (details in Table-1 of the Supplementary) for MNIST and Fashion-MNIST. We train CIFAR-10 models for 100 epochs, MNIST and Fashion-MNIST models for 50 epochs each. The minibatch size is set to 128 for CIFAR-10 and 64 for Fashion-MNIST and MNIST. We use SGD optimizer with momentum of 0.9 and weight decay of $5e-4$. We use an initial learning rate of 0.1 for CIFAR-10, and 0.01 for MNIST and Fashion-MNIST. We reduce the initial learning rate by a factor of 5, three times during the training process. We use early stopping based on I-FGSM [15] accuracy on the validation split in the last 20 epochs for CIFAR-10, and last 30 epochs for MNIST and Fashion-MNIST.

The hyperparameters to be selected for training are: k , which is the number of bits to be eliminated during the quantization step in Section-4.2.1, and λ , which is the weighting factor for *BPFC* loss in Eq. (2). We set k to 5 for CIFAR-10, 6 for Fashion-MNIST and 7 for MNIST. The value of k can be selected based on the ε value of the attack to be defended. λ is selected to achieve the desired trade-off between accuracy on clean and adversarial samples (details in Section-1.2 of the Supplementary). As explained in Section-4.3, we start with a low value of λ and step it up over epochs. This helps achieve better accuracy on clean samples. For CIFAR-10, we start with λ of 1 and step it up by a factor of 9 every 25 epochs. Since the clean accuracy on Fashion-MNIST and MNIST datasets increases within very few epochs, we use high λ values from the beginning (without a step-up factor). We thus use a λ value of 30 for MNIST and 25 for Fashion-MNIST.

5.2. Overview of Experiments

We compare the proposed approach with Normal Training (NT), FGSM-AT [9], PGD-AT [21] and Regularized Single-Step Adversarial Training (RSS-AT) [35] across all three datasets. We report results on single-step (FGSM) and multi-step (I-FGSM, PGD) attacks, epsilon-bounded and

Table 1: **CIFAR-10**: Recognition accuracy (%) of models in a white-box attack setting.

Training method	Clean	FGSM	IFGSM 7 steps	PGD (n-steps)		
				7	20	1000
FGSM-AT	92.9	96.9	0.8	0.4	0.0	0.0
RSS-AT	82.3	55.0	50.9	50.0	46.2	45.8
PGD-AT	82.7	54.6	51.2	50.4	47.4	47.0
NT	92.3	16.0	0.0	0.0	0.0	0.0
Mixup	90.3	27.4	1.6	0.6	0.1	0.0
BPFC (Ours)	82.4	50.1	44.1	41.7	35.7	34.4
Ablations of the proposed approach (BPFC)						
A1: Simple quant	82.6	49.2	41.4	38.8	31.6	30.1
A2: Uniform noise	82.6	48.7	42.3	40.0	33.3	31.9
A3: ℓ_1 norm ¹	92.1	68.3	60.8	57.1	46.8	35.9

unbounded (DeepFool [22], Carlini-Wagner (C&W) [4]) attacks, untargeted and targeted attacks, and gradient-free attacks (random attacks, SPSA [33]). We consider attacks in white-box and black-box settings. We also consider adaptive attacks that are specific to the defense mechanism used.

As explained in Section-3.2, we restrict the adversary to be in the ℓ_∞ -ball of radius ε around each data point. We refer to the work by Madry *et al.* [21] for attack parameters and number of iterations for PGD attack. For an image with pixel intensities in the range $[0, 1]$, we consider an ε value of $8/255$ for CIFAR-10, 0.3 for MNIST and 0.1 for Fashion-MNIST. We consider ε_{step} to be $2/255$, 0.01 and 0.01 for CIFAR-10, MNIST and Fashion-MNIST respectively. These restrictions do not apply to the unbounded attacks, DeepFool and C&W.

We present our experiments, results and analysis for each attack in the following subsections.

5.3. Performance against White-box Attacks

As explained in Section-3.2, we consider that the adversary has access to the network architecture and weights. In this scenario, white-box attacks are expected to be stronger than black-box attacks (unless the model merely appears to be robust due to gradient masking). In this section, we consider the following types of white-box attacks: untargeted and targeted ε -bounded attacks, and unbounded attacks.

5.3.1 Bounded Attacks: Untargeted

The results of various single-step and multi-step white-box attacks for CIFAR-10 dataset are presented in Table-1. FGSM-AT achieves the best robustness to single-step attacks. However, it is not robust to multi-step attacks as explained by Kurakin *et al.* [16]. PGD-AT and RSS-AT show the best accuracy of around 45% for 1000-step PGD attack.

¹A3: The 500-step worst case PGD accuracy goes down from 37.5% to 24.8% with 100 random restarts (over 1000 test samples)

Table 2: **Computational complexity** measured in terms of absolute training time per epoch (seconds) and ratio w.r.t. the proposed method (BPFC). This experiment is run on a single Nvidia Titan-X GPU card.

Training method	CIFAR-10		F-MNIST		MNIST	
	seconds	ratio	seconds	ratio	seconds	ratio
RSS-AT	127.2	1.8	23.8	2.0	24.1	1.7
PGD-AT	257.8	3.7	199.6	16.9	199.2	14.2
NT	39.6	0.6	9.3	0.8	8.9	0.6
BPFC (Ours)	69.4	1.0	11.8	1.0	14.0	1.0

Mixup [39] does not use Adversarial Training and achieves an improvement over Normal Training (NT) in robustness towards FGSM attack. However, it is not robust to PGD attacks. The proposed method achieves a significant improvement over Normal Training and Mixup in robustness to both single-step and multi-step attacks, despite not being exposed to adversarial samples during training. As shown in Table-2, the proposed method is faster than methods that are robust to multi-step attacks (PGD-AT and RSS-AT).

As explained in Section-4.2.1, we consider ablation experiments of Simple Quantization (A1) and addition of Uniform Noise (A2). The proposed method (BPFC) achieves an improvement over these two baselines, indicating the significance of the proposed formulation. Adding uniform random noise in the range $(-8/255, 8/255)$ produces an effect similar to that of quantization by reducing the importance given to LSBs for the classification task. Hence, we see comparable results even for this ablation experiment.

We also consider an ablation experiment of using ℓ_1 -norm instead of ℓ_2 -norm in Eq. (2). While the results using ℓ_1 -norm (Table-1) show an improvement over the proposed method, the 500-step worst-case PGD accuracy goes down from 37.5% to 24.8% with 100 random restarts (over 1000 test samples represented equally across all classes), indicating that it achieves robustness due to gradient masking. For the proposed approach, the PGD accuracy with 50 steps (34.68%) is similar to that with 1000 steps (34.44%). Hence, we check 50-step PGD accuracy with 1000 random restarts (for 1000 test samples) and find that the drop in accuracy over multiple random restarts is negligible. The accuracy drops from 35.6% to 34.9% over 1000 random restarts, verifying that the robustness in the proposed approach is not due to gradient masking.

Table-3 shows the consolidated white-box results for all datasets. The proposed method has significantly better robustness to multi-step attacks when compared to methods that do not use Adversarial training (NT and Mixup). We also achieve results comparable to PGD-AT and RSS-AT, while being significantly faster. Detailed results with Fashion-MNIST and MNIST datasets are reported in Tables-2 and 3 of the Supplementary.

Table 3: **White-box setting:** Recognition accuracy (%) of different models on clean samples and adversarial samples generated using PGD-1000 step attack.

Training method	CIFAR-10		F-MNIST		MNIST	
	Clean	PGD	Clean	PGD	Clean	PGD
FGSM-AT	92.9	0.0	93.1	15.1	99.4	3.7
RSS-AT	82.3	45.8	87.7	71.8	99.0	90.4
PGD-AT	82.7	47.0	87.5	79.1	99.3	94.1
NT	92.3	0.0	92.0	0.3	99.2	0.0
Mixup	90.3	0.0	91.0	0.0	99.4	0.0
BPFC (Ours)	82.4	34.4	87.2	67.7	99.1	85.7

Table 4: Recognition accuracy (%) of the proposed method (BPFC) on different 1000-step PGD attacks.

Attack	CIFAR-10	F-MNIST	MNIST
Untargeted	34.4	67.7	85.7
Targeted (Least Likely target)	65.2	85.5	95.6
Targeted (Random target)	63.1	83.5	94.8

Table 5: **DeepFool and C&W attacks (CIFAR-10):** Average ℓ_2 norm of the generated adversarial perturbations is reported. Higher ℓ_2 norm implies better robustness. Fooling rate (FR) represents percentage of test set samples that are misclassified.

Training method	DeepFool		C&W	
	FR (%)	Mean ℓ_2	FR (%)	Mean ℓ_2
FGSM-AT	95.12	0.306	100	0.078
PGD-AT	90.78	1.098	100	0.697
RSS-AT	89.75	1.362	100	0.745
NT	94.66	0.176	100	0.108
Mixup	93.37	0.168	100	0.104
BPFC (Ours)	89.51	2.755	100	0.804

5.3.2 Bounded attacks: Targeted

We evaluate the robustness of BPFC trained models against targeted attacks of two types, as discussed in this section. In the first attack (Least Likely target), we set the target class to be the least likely predicted class of a given image. In the second variant (Random target), we assign random targets to each image. We use 1000-step PGD attacks for both these evaluations, and compare the robustness against an untargeted PGD attack in Table-4. As expected, model trained using the proposed approach is more robust to targeted attacks, when compared to an untargeted attack.

5.3.3 Unbounded Attacks

We evaluate robustness of BPFC trained models to the unbounded attacks, DeepFool and Carlini-Wagner (C&W). The goal here is to find the lowest ℓ_2 -norm bound on perturbations that can result in 100% fooling rate for all sam-

Table 6: **Black-box setting:** Recognition accuracy (%) of different models on FGSM black-box adversaries. Columns represent the source model used for generating the attack.

Training method	CIFAR-10		Fashion-MNIST		MNIST	
	VGG19	ResNet18	Net-A	M-LeNet	Net-A	M-LeNet
FGSM-AT	78.67	77.58	94.36	90.76	87.99	85.68
RSS-AT	79.80	79.99	84.99	84.16	95.28	95.19
PGD-AT	80.24	80.53	84.99	85.68	95.75	95.36
NT	36.11	15.97	34.71	16.67	29.94	16.60
Mixup	42.67	43.41	54.65	66.31	58.47	69.46
BPFC (Ours)	78.92	78.98	81.38	83.46	94.17	94.56

ples. We select the following hyperparameters for C&W attack: search steps = 9, max iterations = 200, learning rate = 0.01. For DeepFool attack, we set the number of steps as 100. With these settings, we achieve 100% fooling rate for C&W attack with all the training methods. DeepFool does not achieve 100% fooling rate for any of the methods. However, these results are consistent with the results reported in literature [25]. The performance of these models is measured in terms of average ℓ_2 -norm of the generated perturbations. A higher value on the bound implies that the model has better robustness. The results on CIFAR-10 are presented in Table-5. It can be observed that the BPFC trained model is more robust to C&W attack when compared to all other methods, including PGD-AT. The DeepFool results of the proposed method can be compared directly only with PGD-AT and RSS-AT, as they achieve similar fooling rates. The proposed method achieves significantly improved robustness when compared to both these methods. Results on Fashion-MNIST and MNIST datasets are presented in Section-2.1.2 of the Supplementary.

5.4. Performance against Black-box Attacks

We report accuracy against FGSM black-box attacks in Table-6. We consider two source models for generation of black-box attacks on each dataset; the first is a model with a different architecture, and the second is a model with the same architecture as that of the target model. In both cases, and across all datasets, black-box accuracies are significantly better with the proposed approach when compared to other non-adversarial training methods (NT and Mixup). Further, our results are comparable with those of adversarial training methods. Results on multi-step black-box attacks are presented in Section-2.2 of the Supplementary.

5.5. Performance against Gradient-free Attacks

We check the robustness of our proposed approach against the following gradient-free attacks on the CIFAR-10 dataset, to ensure that there is no gradient masking: attack with random noise [3] and SPSA attack [33].

For the attack with random noise, we consider a random sample of 1000 images from the test set of CIFAR-10, such

that all ten classes are equally represented. We randomly select 10^5 samples from an ℓ_∞ -ball of radius ε around each data point (each pixel is an *i.i.d.* sample from a Uniform distribution) and compute the accuracy of these samples. We find that the accuracy on these random samples is 79.76%, which is slightly less than the accuracy on clean samples (82.4%). We run another experiment to verify that every image that is robust to PGD attack is also robust to random noise. We run PGD attack for 50 steps and 100 random restarts and identify the images that are robust to the attack. We attack these images with 10^5 random noise-perturbations each, and find that we achieve the expected accuracy of 100%. Hence, we conclude that the attack with random noise is not stronger than a gradient-based attack.

The SPSA attack [33] is a gradient-free attack that computes a numerical approximation of the gradient along multiple random directions and approximates the final gradient to be an average over these gradients. The attack becomes stronger as more directions are used. We use the following hyperparameters to generate the attack: $\delta = 0.01$, learning rate = 0.01, batch size = 128 and iterations = 5. We get an accuracy of 70.5% against the SPSA attack using the proposed approach. For the same attack, the accuracy of a PGD trained model is 70.8%.

Therefore, we verify that gradient-based attacks are stronger than gradient-free attacks, thereby confirming the absence of gradient masking.

5.6. Performance against Adaptive Attacks

In this section, we consider methods that utilize knowledge of the defense mechanism in creating stronger attacks. We explore maximizing loss functions that are different from the standard cross-entropy loss for generating adversarial samples. We consider the CIFAR-10 dataset for this experiment. Maximizing the same loss that is used for training, with the same hyperparameters, gives a slightly lower accuracy (34.52%) when compared to PGD (34.68%) for a 50-step attack. However, this difference is not statistically significant, and it may be due to the random nature of the PGD attack. The worst accuracy across different hyperparameters in the loss function is 34.41%.

We also explore adding another term to the loss that is maximized during a PGD attack. In addition to maximizing the training loss, we minimize the magnitude of k ($= 5$) LSBs in the generated samples. This would encourage the adversaries to have low magnitude LSBs, which could possibly be the samples where the defense was less effective. However, even with this change, we get the same accuracy as that of a standard PGD attack.

Therefore, the adaptive attacks are only as strong as a PGD attack. We include more details on adaptive attacks in Section-2.3 of the Supplementary.

5.7. Basic Sanity Checks to Verify Robustness

In this section, we present results on the basic sanity checks listed by Athalye *et al.* [1] to ensure that the model's robustness is not due to gradient masking.

- Results in Table-1 illustrate that iterative attacks (PGD and I-FGSM) are stronger than an FGSM attack.
- White-box attacks are stronger than black-box attacks based on results in Tables-3 and 6.
- We note that unbounded attacks reach 100% success rate, and increasing the distortion bound increases the success rate of the attack (Fig.2 of the Supplementary).
- As discussed in Section-5.5, gradient-based attacks are stronger than gradient-free attacks.
- We note that cross-entropy loss on FGSM samples increases monotonically with an increase in perturbation size. (Fig.3 of the Supplementary).

5.8. Scalability of the Proposed Method to ImageNet

We present results on ImageNet [5], which is a 1000-class dataset with 1.2 million images in the training set and 50,000 images in the validation set. The accuracy on a targeted PGD 20-step attack is 32.91% with the proposed approach and 43.43% with a PGD-AT model [7]. The trend in robustness when compared to PGD-AT is similar to that of CIFAR-10 (Table-1), thereby demonstrating the scalability of the proposed approach to large scale datasets. We present detailed results in Section-2.5 of the Supplementary.

6. Conclusions

We have proposed a novel Bit Plane Feature Consistency (*BPFC*) regularizer, which improves the adversarial robustness of models using a normal training regime. Results obtained using the proposed regularizer are significantly better than existing non-adversarial training methods, and are also comparable to adversarial training methods. Since the proposed method does not utilize adversarial samples, it is faster than adversarial training methods. We demonstrate through extensive experiments that the robustness achieved is indeed not due to gradient masking. Motivated by human vision, the proposed regularizer leads to improved local properties, which results in better adversarial robustness. We hope this work would lead to further improvements on the front of non-adversarial training methods to achieve adversarial robustness in Deep Networks.

7. Acknowledgements

This work was supported by RBCCPS, IISc and Uchatar Avishkar Yojana (UAY) project (IISC-10), MHRD, Govt. of India. We would like to extend our gratitude to all the reviewers for their valuable suggestions.

References

- [1] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *arXiv preprint arXiv:1802.00420*, 2018. 2, 5, 8
- [2] Jacob Buckman, Aurko Roy, Colin Raffel, and Ian Goodfellow. Thermometer encoding: One hot way to resist adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [3] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, and Aleksander Madry. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019. 5, 7
- [4] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017. 6
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 8
- [6] Guneet S. Dhillon, Kamyar Azizzadenesheli, Jeremy D. Bernstein, Jean Kossaifi, Aran Khanna, Zachary C. Lipton, and Animashree Anandkumar. Stochastic activation pruning for robust adversarial defense. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [7] Logan Engstrom, Andrew Ilyas, Shibani Santurkar, and Dimitris Tsipras. Robustness (python library), 2019. 8
- [8] Jessica Fridrich, Miroslav Goljan, and Rui Du. Reliable detection of lsb steganography in color and grayscale images. In *Proceedings of the 2001 Workshop on Multimedia and Security: New Challenges*. 3
- [9] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 2, 5
- [10] Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens van der Maaten. Countering adversarial images using input transformations. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [12] Warren He, James Wei, Xinyun Chen, Nicholas Carlini, and Dawn Song. Adversarial example defense: Ensembles of weak defenses are not strong. In *11th USENIX Workshop on Offensive Technologies (WOOT)*, 2017. 2
- [13] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). 5
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2012. 1
- [15] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016. 5
- [16] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016. 2, 6
- [17] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>. 5
- [18] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 5
- [19] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 1
- [20] Xingjun Ma, Bo Li, Yisen Wang, Sarah M. Erfani, Sudanthi Wijewickrema, Grant Schoenebeck, Michael E. Houle, Dawn Song, and James Bailey. Characterizing adversarial subspaces using local intrinsic dimensionality. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [21] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Tsipras Dimitris, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations (ICLR)*, 2018. 1, 2, 5, 6
- [22] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [23] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the ACM Asia Conference on Computer and Communications Security*, 2017. 2
- [24] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1
- [25] Jérôme Rony, Luiz G Hafemann, Luiz S Oliveira, Ismail Ben Ayed, Robert Sabourin, and Eric Granger. Decoupling direction and norm for efficient gradient-based l_2 adversarial attacks and defenses. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 7
- [26] Zhao Shan and Wang Hai-Tao. Image retrieval based on bit-plane distribution entropy. In *International Conference on Computer Science and Software Engineering (CSSE)*. IEEE, 2008. 3
- [27] Yang Song, Taesup Kim, Sebastian Nowozin, Stefano Ermon, and Nate Kushman. Pixeldefend: Leveraging generative models to understand and defend against adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [28] Arun P Sripati and Carl R Olson. Representing the forest before the trees: a global advantage effect in monkey inferotemporal cortex. *Journal of Neuroscience*, 29(24):7788–7796, 2009. 1, 3

- [29] Yasuko Sugase, Shigeru Yamane, Shoogo Ueno, and Kenji Kawano. Global and fine information coded by single neurons in the temporal visual cortex. *Nature*, 400(6747):869, 1999. 1, 3
- [30] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2013. 1, 5
- [31] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017. 2
- [32] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. *arXiv preprint arXiv:1805.12152*, 2018. 5
- [33] Jonathan Uesato, Brendan O’Donoghue, Aaron van den Oord, and Pushmeet Kohli. Adversarial risk and the dangers of evaluating against weak attacks. *arXiv preprint arXiv:1802.05666*, 2018. 6, 7, 8
- [34] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, Aaron Courville, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. *arXiv preprint arXiv:1806.05236*, 2018. 2
- [35] B. S. Vivek, Arya Baburaj, and R. Venkatesh Babu. Regularizer to mitigate gradient masking effect during single-step adversarial training. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019. 2, 5
- [36] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017. 5
- [37] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. Mitigating adversarial effects through randomization. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [38] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155*, 2017. 2
- [39] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 2, 6
- [40] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P Xing, Laurent El Ghaoui, and Michael I Jordan. Theoretically principled trade-off between robustness and accuracy. *arXiv preprint arXiv:1901.08573*, 2019. 1, 5