

# Exploring Non-additive Randomness on ViT against Query-Based Black-Box Attacks

Jindong Gu<sup>1</sup>  
jindong.gu@outlook.com  
Fangyun Wei<sup>2</sup>  
fawe@microsoft.com  
Philip Torr<sup>1</sup>  
philip.torr@eng.ox.ac.uk  
Han Hu<sup>2</sup>  
hanhu@microsoft.com

<sup>1</sup>Torr Vision Group  
Department of Engineering Science  
University of Oxford  
Oxford, UK  
<sup>2</sup>Microsoft Research Asia  
Beijing, China

---

## Abstract

Deep Neural Networks can be easily fooled by small and imperceptible perturbations. The query-based black-box attack (QBBA) is able to create the perturbations using model output probabilities of image queries requiring no access to the underlying models. QBBA poses realistic threats to real-world applications. Recently, various types of robustness have been explored to defend against QBBA. In this work, we first taxonomize the stochastic defense strategies against QBBA. Following our taxonomy, we propose to explore non-additive randomness in models to defend against QBBA. Specifically, we focus on underexplored Vision Transformers based on their flexible architectures. Extensive experiments show that the proposed defense approach achieves effective defense, without much sacrifice in performance.

## 1 Introduction

The decisions of deep neural networks (DNNs) can be misled by imperceptible perturbations. The adversary perturbations can be created only using the model output scores (i.e., softmax probabilities) of image queries and require no access to the underlying models [1, 2]. As a result, DNN-based real-world applications face the potential threats posed by query-based black-box attacks (QBBA).

**Threat Model of QBBA:** For this type of attack, attackers can only access model output scores, while they cannot get access to the model architecture, model parameters, input gradients, or defense strategies. Many efficient and effective such attacks have raised the attention of the community [3, 4, 5]. Recent work explores various types of randomness to defend against QBBA [6, 7, 8, 9]. For example, the randomness in inputs or models can mislead the gradient estimation or search process of QBBA.

To articulate different types of defensive randomness, in this work, we first taxonomize the defense strategies against QBBA. Following our taxonomy, we propose to explore non-additive randomness in models to defend against QBBA. Specifically, we focus on Vision

Transformers (ViT) [10] models for the following reasons: 1) As an alternative to CNNs, ViTs have received great attention; 2) The non-additive randomness on ViT is underexplored; 3) The network architecture is flexible.

Different from CNNs, ViT represents an input image as a sequence of image patches and learns image representations by applying the self-attention module to the patches. We propose three ways to integrate non-additive robustness in the self-attention module of ViT and provide analyses of their effectiveness against QBBA. We verify our approach on popular QBBA and show that ours can defend against them without sacrificing clean performance.

## 2 Related Work

**Query-based Black-box Attacks.** The popular black-box attacks include transfer-based black-box attacks and query-based black-box attacks (QBBA). The former leverages the observation that adversary examples created on one model can be transferred to different model architectures, e.g., Vision Transformers [6, 37], Capsule Networks [18], Spiking Neural Networks [47]. QBBA can be categorized into score-based attacks and decision-based ones according to the availability of model outputs. The score-based attacks are able to access class-wise output probabilities [0, 0, 3, 7, 22, 21, 23, 28, 29, 38, 41, 42], while decision-based ones assume only the top-1 class index is available [9, 11, 14, 15, 25, 34]. On the other hand, according to the method to create adversarial examples, query-based attacks can be also categorized into optimization-based attacks and search-based attacks. The former estimates gradients of the loss with respect to the input to optimize a pre-defined adversary's objective loss [0, 22, 14, 15, 23, 38, 42]. In contrast, the latter repeatedly searches for the desired images from randomly perturbed ones so that the desired ones get closer to adversary's objective [0, 0, 3, 3, 9, 11, 21, 25, 28, 29, 34].

**Defense against Query-based Black-box Attacks.** Most defense strategies are designed for defending against white-box attacks. For example, adversarial training [27]. The adversarially trained models also show certain robustness against QBBA, but often with a sacrifice in clean performance [27, 46]. Randomness has also been intensively explored to defend against both white-box and black-box attacks. The works [27, 24] propose to add Gaussian noise to model parameters and intermediate activations to boost model robustness. Random transformation of input, as a pre-process defense strategy is proposed in the work [45]. The works [10, 33] show that simply adding small noise to inputs can effectively defend against black-box attacks. The relationship between these stochastic defense strategies remains to be articulated.

**Adversarial Robustness of ViT.** The robustness of ViT have achieved great attention due to its great success in many vision tasks [6, 6, 8, 19, 30, 31, 32, 36, 37]. The works [6, 8, 36, 44] demonstrate that ViT achieves higher adversarial robustness than CNNs under white adversarial attacks. The work [19, 32] shows a different observation in a test setting, e.g., under distribution shift and patch attack. The work [6, 37] makes the exploration of QBBA attacks and transfer-based attacks on ViTs. They show ViT also suffers from the basic QBBA, such as boundary attacks and Bandits-based attacks. The works [36, 39, 44] show adversarial training can also be applied to boot the adversarial robustness of ViT. The simple patch permutation of ViT has also been explored to study its robustness [17, 19, 30, 36]. In this work, we systematically explore non-additive randomness of ViT against QBBA.




Figure 1: Taxonomy of Stochastic Defense Strategy. The randomness can be categorized into additive randomness and non-additive one. Each of the two types randomness can be integrated either into the inputs or into the model.

## 3 Stochastic Defense of ViT against Query-Based Black-box Attacks

In this work, we first present our taxonomy the stochastic defense strategies. Following our taxonomy, we propose non-additive randomness on ViT to defend against QBBA. Then, we show our method performs better than the rest of the methods in our taxonomy.

### 3.1 Taxonomy of Stochastic Defense Strategy

Randomness has been explored intensively to defend against attacks. The earlier exploration in [24] adds trainable gaussian noise into intermediate activations to build a robust model, which was shown to be robust to white-box attacks. A later work [45] shows the random transformation of input can be used to defend against both white-box and black-box attacks. The recent work [10, E3] shows simply adding small noise to inputs is surprisingly effective to defend against QBBA. Without a doubt, all these methods can be applied to defend against QBBA. In this work, we attempt to articulate the relationship among them and taxonomize the existing stochastic defense strategies, as shown in Fig. 1.

In our taxonomy, we categorize randomness into two types, namely, additive randomness and non-additive one. The additive randomness is often implemented by adding small noise to the input or the model parameters. In contrast, the non-additive randomness integrated non-additive perturbation into the input or the model. For example, the affine transformation of inputs is one type of non-additive perturbation. As shown in taxonomy, the non-additive randomness on the model has not been explored yet to defend against QBBA. In the next subsection, we describe our proposal to fill the vacancy.

### 3.2 Non-additive Randomness in ViT for Defense

We first review the background knowledge of ViT and then present our three ways to integrate non-additive randomness into ViT model.

**Background knowledge on ViT** Different from CNNs, ViT represents an input image as a sequence of image patches. Then, a list of self-attention modules is applied to the sequence of image patches sequentially. We now introduce the details of the primary ViT architecture in [10]. The input image  $X \in \mathbb{R}^{(H \times W \times C)}$  is first split into image patches  $\{x_i \in \mathbb{R}^{P \times P \times C} | i \in (1, 2, 3, \dots, HW/P^2)\}$  where  $P$  is the patch size.

The embedding of each patch is extracted from the raw image patch with linear projection. Before the application of self-attention module, the position information of image patches is also integrated into the patch embedding. The embedding of the image patches




Figure 2: Non-additive Randomness in ViT. The color marks denote our operations to introduce randomness into ViT. The blue arrows in (a) are the random permutation operation. The red crosses in (b) are the random drop operation. The green arrows in (c) are the random dimension reduction operation. All three operations introduce non-additive randomness.

in the input layer is denoted as  $\{Z_i^0 \in \mathbb{R}^D | i \in (1, 2, 3, \dots, HW/P^2)\}$ . A learnable class-token embedding  $Z_0^0$  is added into the list of patch embeddings. The class embedding in the last layer is taken as the image embedding for classification.

We now introduce the transformer encoder where the list of blocks is applied to transform the input embeddings. Each block consists of two main modules, namely, a multi-head self-attention module  $MSA(\cdot)$  to model the inter-patch relationship and an MLP module  $MLP(\cdot)$  to project patch embeddings into a new space, respectively.

When the self-attention module in  $(l+1)$ -th layer is applied to input patches  $\{Z_i^l \in \mathbb{R}^D | i \in (0, 1, 2, \dots, HW/P^2)\}$  in the  $l$ -th layer, the attention of  $i$ -th patch in a single head is computed as follows  $K_i^{l+1} = W_k^{l+1} \cdot Z_i^l$ ,  $Q_i^{l+1} = W_q^{l+1} \cdot Z_i^l$ ,  $A_i^{l+1} = \text{Softmax}(Q_i^{l+1}(K^{l+1})^T/\sqrt{D})$ , The output embedding of the  $i$ -th patch  $Z_i^l$  is  $V_i^{l+1} = W_v^{l+1} \cdot Z_i^l$ ,  $Z_i^{l+1} = \sum_{j=0}^{HW/P^2} A_{ij}^{l+1} \cdot V_j^{l+1}$ .

Then, an MLP module with two MLP layers is applied to project the final embedding of each patch into a new feature space. The embeddings of  $i$ -th patch in different heads are concatenated as its final embedding in the  $(l+1)$ -th layer. The embedding of the class-token patch in the last block is taken as the image representation to classify the image. A linear classifier maps the features to output space.

### 3.2.1 Patch Sequence Random Permutation in ViT

The self-attention module in ViT models sequence to sequence. In the standard ViT, the input sequence of image patches is  $x = \{x_i | i \in (1, 2, 3, \dots, HW/P^2)\}$ . To encode positional information of input patches, ViT adds explicit positional embedding to the corresponding patch embedding. The positional embedding can be  $p = \{p_i | i \in (1, 2, 3, \dots, HW/P^2)\}$ . The initial patch embedding of the  $i$ -th patch is composed of the embedding of a raw patch  $x_i$  and the corresponding positional embedding of  $p_i$ , specifically, the sum of the embeddings  $E(x_i) + E(p_i)$ . The sum operation encodes the positional information into patch embedding.

We introduce randomness into the position information encoding process. We first randomly sample parts of patches and permute their correspondence to their positional embedding. In other words, once selected, the initial patch embedding of the  $i$ -th patch is  $E(x_i) + E(p_j)$  where  $p_j$  is the positional embedding of the  $j$ -th patch. Both the patch selection and the correspondence permutation are different from run to run.

The self-attention operation itself is invariant to sequence order. Hence, our patch permutation is equivalent to the permutation of the input patches of ViT model, as shown in Fig. 2a. The probability of selecting patches to permute controls the trade-off between the defensive effect and the clean performance.

### 3.2.2 Patch Random Drop in ViT

It has been said that ViT has less inductive bias than traditional CNNs [14]. This is because only grid-structured input is acceptable for standard CNNs in image classification tasks. Contrarily, given its architectural traits, the self-attention module of ViT can accept sequences with different lengths. Given the architectural flexibility, we introduce non-additive randomness by randomly sampling the elements of the input sequence. Namely, the input sequence of image patches is  $x' = \{\phi_r(x_i, p) | i \in (1, 2, 3, \dots, HW/P^2)\}$ , where  $p$  is a pre-defined probability and the function  $f_r(x_i, p)$  is defined as

$$\phi_r(x_i, p) = \begin{cases} x_i & \text{if } u > p \\ \{\} & \text{if } u \leq p, \end{cases}$$

where  $u$  is a value sampled from the uniform distribution  $\mathcal{U}(0, 1)$ . In the non-additive randomness above, different image patches are sampled in every forward pass, as shown in Fig. 2b. Namely, only part of the inputs are taken to make a prediction, meanwhile, the selected parts are random in each prediction. Given its remarkable modeling capacity, ViT still performs well when such non-additive randomness is introduced.

### 3.2.3 Patch Attention Perturbation by Sampling Keys and Queries

Another non-additive randomness we explore is in the process to create attention. The attention of the  $i$ -th patch is computed with the  $i$ -th query  $Q_i^{l+1}$  and the keys  $K^{l+1}$  of all patches. We introduce randomness into attention by injecting it into the query and the keys.

The key and the query of a patch are represented by activation vectors of the same dimension. Given the dimension  $D$  of the query and the keys, their dimension indices are  $I = (1, 2, \dots, D)$ . As illustrated in Fig. 2c, we first sample the dimension indices to obtain  $I'$  where  $|I| \leq |I'|$ . We drop each dimension with a probability of  $\alpha \in [0, 1)$ . We propose to formulate the new keys and the queries with the sampled dimensions  $I'$ .  $K'_i = \psi(K_i, I')$ ,  $Q'_i = \psi(Q_i, I')$ , where  $\psi(\cdot)$  reduces dimension by dropping the specified dimensions. Namely, only the sampled dimensions are considered. By doing this, we introduce non-additive randomness into attention. The new attention is computed with the new keys and the queries.

The non-additive model randomness is different from the input with additive robustness where the input sequence is  $x = \{x_i + \mathcal{N}(0, \sigma^2 I) | i \in (1, 2, 3, \dots, HW/P^2)\}$ , where  $\sigma$  is the variance of gaussian noise.

## 3.3 Analysis of (Non)-additive Randomness for Adversarial Defense

In this subsection, we first show why stochastic defense strategies are effective against QBBA. Furthermore, we provide an analysis to show why non-additive randomness on the model performs better than other stochastic defense strategies.

### 3.3.1 Defense against Optimization-based Black-box Attacks

Optimization-based black-box attacks [14, 14, 15] create adversarial examples using estimated gradients. The gradient estimation is often implemented with finite difference [14]. Given the model  $f(x, \theta)$  with the input  $x \in \mathbb{R}^{(H \times W \times C)}$  and the model parameters  $\theta$ , the input gradients  $G \in \mathbb{R}^{(H \times W \times C)}$  can be estimated as follows:

$$G = \frac{1}{N} \sum_{i=0}^N \frac{f(x + \beta u_i, \theta) - f(x - \beta u_i, \theta)}{2\beta} \quad (1)$$

where  $u_i \sim \mathcal{N}(0, \sigma^2 I)$  and  $G = \frac{1}{N} \sum_{i=0}^N G_i$ . The stochastic defense strategies introduce randomness to inputs or models. The gradient estimation based on the perturbed inference can differ from the original one. Note that the randomness  $\eta$  and  $\eta'$  in the forward inference are different from run to run. The introduced randomness leads to inaccurate gradients:

$$\hat{G}_i = \frac{f(x + \beta u_i, \theta, \eta) - f(x - \beta u_i, \theta, \eta')}{2\beta} \neq G_i. \quad (2)$$

Concretely, when the additive randomness is added to input, the forward inference is  $f(x + \beta u + \eta, \theta)$ . In case  $\eta = -2\beta u$ ,  $\eta' = 2\beta u$ , the gradients can even be  $\text{sign}(\hat{G}_i) = -\text{sign}(G_i)$ , which leads to total wrong estimations. In practice, the estimated gradients are often opposite to the original ones in some dimensions of  $G_i$ . When the additive randomness is added to model parameters, the output of the perturbed layer is  $(w + \eta)x + b = wx + \eta x + b$ , which is equivalent to adding perturbations to the input of the next layer.

When the non-additive randomness is added to input (e.g. affine transformation), the forward inference is  $f(g(x), \theta)$ . The finite difference does work well in this case since the two inputs in two forward inferences differ from each other significantly. Our proposed method is to add non-additive randomness to the model, i.e., sampling model functions in forward inferences. The forward inference  $f'(x, \theta)$  makes the gradient estimation inaccurate by changing the model function. The architectural flexibility of ViT models makes it feasible to introduce non-additive randomness to the model.

### 3.3.2 Defense Against Search-based Black-box Attacks

Search-based black-box attacks first add random noises  $u = \{u_k | k \in (1, 2, \dots, N)\}$  to the current input, then select the noise that leads to the maximal loss  $L_k = \ell(f(x + u_k, \theta))$ , and repeat the first two steps until an effective adversarial perturbation is found. When the additive robustness  $\eta = \{\eta_i | i \in (1, 2, \dots, M)\}$  is added to the input, the loss is  $L'_k = \ell(f(x + u_k + \eta_i, \theta))$  where  $i$  is randomly selected from  $\{1, 2, \dots, M\}$  in each forward pass. The new loss makes search-based methods infeasible since  $L'_k$  is not fixed from run to run due to  $\eta_i$ . Note that  $\theta_i$  is introduced by the defense side, which is unavailable to the adversary.

Similarly, when the non-additive randomness is introduced to input, the loss becomes  $L'_k = \ell(f(g(x + u), \theta))$ . The selected adversarial perturbation  $u_i$  does not necessarily lead to a higher loss in the next search iteration. Furthermore, the non-additive randomness on the model changes the model function. The  $x + u$  is able to fool the model function  $f'$ , but not necessarily another one  $f''$ . This error in the adversarial perturbation selection makes the searching process almost similar to a random walk.

### 3.3.3 Adaptive Attacks on Stochastic Defense

In the setting above, the adversary has no access to the model details as well as the defense strategies for a black-box attack setting. We now consider the black-box attack attacks that also work on the model with stochastic defense, i.e. defense-aware attacks. As suggested in [9], a defense-aware attack is to combine the Expectation of Transform strategy (EOT) with the existing black-box attacks. We argue that defense-aware attacks do not work well for the following two reasons: The estimated gradients are still wrong under EOT attack

strategy since the models are highly non-linear [10]; The application of EOT requires a large number of queries on the model, which is what the adversary tries to avoid in the first place when QBBA are designed.

### 3.3.4 Efficiency of Defense Strategy

Besides the effectiveness of the defense strategies, their efficiency is also important since some defense strategies can bring large extra computational costs. Concretely, adversarial training is very computationally expensive in the training stage [27]. The ensemble-based defense strategies increase the cost many times more [24]. The recently proposed randomness-based defense strategies only bring tiny extra computational costs. Different from the previous ones, our approach can even accelerate the forward inference of ViT, such as Patch Random Drop and Patch Attention Perturbations. Our non-additive randomness on the model also enjoys the advantage of high efficiency.

## 4 Experiments

We first briefly describe the baselines and our approach. Then, we show the experimental setting as well as evaluation metrics. Next, we show our experimental results and analysis.

**Methods.** The following defense methods have been applied in our experiments: Small Noise Defense (SND) [10], Parametric Noise Injection (PNI) [22], Random Resizing and Padding (R&P) [45], Patch Random Permutation (PRPerm), Patch Random Drop (PRDrop), and Patch Attention Perturbation (PAttnPert). We select representative QBBA, such as Square Attack [9], Hop Skip Jump Attack (HSJA) [10], and GeoDA [64]. The details of all attack and defense methods can be found in Supplement B and C.

**Experimental Setting.** We use the most primary ViT models [17, 40]. The comparison to various CNN models (e.g. ResNet) and the generalization on other ViT models (e.g. SwinTransformer) are also studied in our experiments. All models we used are from Timm library [43], which are pre-trained on the ImageNet dataset [43]. Following most previous work on black-box attacks, we take 1k images from the validation dataset as the test set.

We select the popular QBBA to compare different defense approaches. The hyper-parameters of each attack are shown in Supplement B. Instead of carefully selecting a hyper-parameter, we report the robust accuracy and clean accuracy under different hyper-parameters of defense strategies. By doing this, we compare the trade-off achieved by different approaches.

**Evaluation Metrics.** As discussed before, there is a trade-off between robust accuracy and clean accuracy. For a fair comparison, we report the clean accuracy (**Accu** in %) on the whole validation dataset as well as Attack Failure Rate (**AFR** in %) on the selected 1k image set. When the attack is unable to fool the target model under the allowed perturbation range within the given number of attack iterations, the attack is counted as a failure.

### 4.1 Stochastic Defense

We equip the ViT-small models with different defense methods respectively and apply the state-of-the-art QBBA to the models. We report the results in Fig. 3. Each subfigure shows the results of one type of attack. For example, the subfigure 3a presents the results on Square Attack. The x-axis and y-axis are the clean accuracy (Accu in %) and the attack failure rate






Figure 3: Stochastic Defense on ViT against Query-based Black-box Attacks. In each subfigure, the x-axis and y-axis are the clean accuracy (Accu in %) and the attack failure rate (AFR in %), respectively. Each of the lines corresponds to a type of defense. The three red lines are our non-additive randomness defense on the model. Each point in the line corresponds to a trade-off point between Accu and AFR. Our methods can achieve a better trade-off than others.

(AFR in %), respectively. In each subfigure, there are six lines, each of which corresponds to a type of defense method. Concretely, the blue, green, and yellow lines correspond to SND, R&P, and PNI defense methods, respectively. The three red lines are our method, where each mark corresponds to a method under our defense based on the non-additive randomness of the model. Each point in the line corresponds to a trade-off point between Accu and AFR.

In the subfigure 3a, we can observe that our method can achieve better defense (i.e. the higher AFR) when keeping the same clean accuracy. Overall, our methods marked with red lines can achieve a better trade-off than the others. The claim is also true when different QBBA are applied. The experimental results on more QBBA are in Supplement A, such as NES Attack [23], SimBA [20], and Boundary [9].

**Stochastic Defense on Different Models.** We also investigate the impact of model sizes and model architectures on the defense methods. Besides the ViT-small, we test the defense methods on ViT-tiny, ViT-base, and ViT-huge with the state-of-the-art Square Attack. The results are in Fig. 4. Our defense methods outperform others on ViT models with different sizes.

There are several popular ViT variants. In this experiment, we select some representative architectures, namely, DeiT, CaiT, and SwinTransformer. We apply various defenses on the selected models and test the defended models with Square Attack. Note that the patch random drop approach cannot be directly applied to Swin Transformer since the shifting and merging operating therein make the patch drop challenging. The results are reported in Fig. 5. Our methods perform better than others on different models with different architectures.

**Combination of Defense Strategies.** We also investigate the combinations of different stochastic defense methods. As shown in Supplement D, the combination is still effective when we combine three types of the proposed methods. In other words, the defense is similarly effective when we randomly drop and permute the patches and perturb the attention at the same time.

## 4.2 Trade-off between Robust Accuracy and Clean Accuracy

The main limitation of the randomization defense method is the decreased model performance on clean examples. If the injected randomness is too small, the defensive effect is






Figure 4: Stochastic Defense on ViTs with different model sizes. Besides ViT-small, we also conduct experiments on ViT-tiny, ViT-base, and ViT-large. Our non-additive stochastic defense on the model (represented by the three red lines) can achieve a better trade-off than other defense methods.




Figure 5: We apply the stochastic defense to models with different model architectures. On the popular versions of Vision Transformers, our methods (marked with red lines) are superior to others.

low. In contrast, the large noise injection leads to better defensive effectiveness, but low model performance on clean examples. Our method achieves a good trade-off, where the non-additive randomness-based defense strategy hardly reduces the clean performance of ViT. For example, the ViT-small has a clean accuracy of 81.4%. When equipped with our defense strategy, it shows a robust accuracy of 92.6% and also the same clean accuracy of 81.4%.

We now discuss the possibility of further improving the trade-off. The intuitive solution is the finetuning method. A pre-trained model is finetuned to improve clean performance where the model is modified with the defense strategy. In other words, we use the same model during finetuning as in the inference. The finetuning strategy can be applied to boost the model for both additive and non-additive randomization. The finetuned models with improved clean accuracy can achieve a better trade-off than the un-fintuned counterparts. For example, the ViT-small with our non-additive randomness can achieve the robust accuracy of 98.4% and the clean accuracy of 81.4% after finetuning.

In the defense of our non-additive randomness, the model performance will be reduced due to the reduced number of input patches or less accurate attention. Besides the finetuning method, another way to alleviate the model degradation is with scheduled non-additive randomness. Note that it is also possible to apply the proposed non-additive randomness to a specific layer instead of all layers. We leave more explorations of more sophisticated fine-tuning methods in future work.

Models	Defense	Accu(%)	AFR(%)	Models	Defense	Accu(%)	AFR(%)
ResNet18	No	69.55	0.0	ResNet50	No	75.86	0.7
	SND [14]	69.21	43.1		SND [14]	75.79	36.2
	PNI [22]	69.45	38.0		PNI [22]	75.84	32.3
	R&P [15]	69.35	68.3		R&P [15]	75.22	78.5
ViT-tiny	No	75.48	0.9	ViT-small	No	81.40	1.2
	SND [14]	75.18	41.14		SND [14]	81.38	72.0
	PNI [22]	74.81	72.2		PNI [22]	81.40	33.5
	R&P [15]	74.19	63.2		R&P [15]	80.97	89.5
	PRDrop	75.09	70.0		PRDrop	81.39	90.6
	PAPert	74.55	72.1		PAPert	80.98	95.4
	PRPerm	75.48	64.2		PRPerm	81.40	92.6

Table 1: Comparison of ViT with ResNet for Defending Against Squared Attack. ViT with the randomness-based defense achieves a better trade-off than the counterpart ResNet (e.g., ResNet50 vs. ViT-small). Besides, our non-additive randomness-based defense on ViT achieves a better trade-off than the other randomness-based defense methods.

### 4.3 Comparison with ResNet for Defense

We now study the question of whether ViT achieves a better trade-off than ResNet with or without our defense. The results are reported in Tab. 1. ViT achieves a better trade-off than ResNet to defend against the state-of-the-art QBBA.

Given the flexibility of ViT architectures, we can drop and permute the input and apply non-additive randomness to its attention. However, it is less easy to do so in CNNs since CNN requires input with strict grid data structures. One of our non-additive randomness-based strategies is to permute the input patches of each self-attention module. A similar defense strategy can also be applied to ResNet where we randomly select input patches and permute them. Our experiments show patch permutation does not work well on ResNet. We conjecture that the defense pattern of a patch is local, which is still effective when placed in different places.

## 5 Conclusion

Random-based defense strategies have been intensively explored to defend against various black-box attack methods. In this work, we taxonomize the defensive randomness from the perspective of defense against query-based black-box attacks. Following our taxonomy, we propose non-additive randomness on ViT. Specifically, we propose three ways to integrate non-additive robustness in the self-attention module of ViT. Our experiments verify that our defense method can achieve better trade-offs than the rest of the stochastic methods in our taxonomy. Based on the non-additive randomness, this work boosts the robustness of ViT to query-based black-box attacks. Almost all recent foundation models are transformer architecture-based. Extending our proposed techniques from classifiers to those foundation models remains to explore for improving their robustness [13, 20]. We leave further exploration in future work.

**Acknowledgements** This work is partially supported by the UKRI grant: Turing AI Fellowship EP/W002981/1 and EPSRC/MURI grant: EP/N019474/1. We would also like to thank the Royal Academy of Engineering and FiveAI.

## References

- [1] Abdullah Al-Dujaili and Una-May O'Reilly. There are no bit parts for sign bits in black-box attacks. *arXiv preprint arXiv:1902.06894*, 2019.
- [2] Moustafa Alzantot, Yash Sharma, Supriyo Chakraborty, Huan Zhang, Cho-Jui Hsieh, and Mani B Srivastava. Genattack: Practical black-box attacks with gradient-free optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1111–1119, 2019.
- [3] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *ECCV*, 2020.
- [4] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International conference on machine learning*, pages 274–283. PMLR, 2018.
- [5] Yutong Bai, Jieru Mei, Alan Yuille, and Cihang Xie. Are transformers more robust than CNNs? In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=hbHkvGBZB9>.
- [6] Philipp Benz, Soomin Ham, Chaoning Zhang, Adil Karjauv, and In So Kweon. Adversarial robustness comparison of vision transformer and mlp-mixer to cnns. *arXiv preprint arXiv:2110.02797*, 2021.
- [7] Arjun Nitin Bhagoji, Warren He, Bo Li, and Dawn Song. Practical black-box attacks on deep neural networks using efficient query mechanisms. In *ECCV*, 2018.
- [8] Srinadh Bhojanapalli, Ayan Chakrabarti, Daniel Glasner, Daliang Li, Thomas Unterthiner, and Andreas Veit. Understanding robustness of transformers for image classification. *arXiv:2103.14586*, 2021.
- [9] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. *arXiv preprint arXiv:1712.04248*, 2017.
- [10] Junyoung Byun, Hyojun Go, and Changick Kim. On the effectiveness of small input noise for defending against query-based black-box attacks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3051–3060, 2022.
- [11] Jianbo Chen, Michael I Jordan, and Martin J Wainwright. Hopskipjumpattack: A query-efficient decision-based attack. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 1277–1294. IEEE, 2020.
- [12] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM workshop on artificial intelligence and security*, 2017.

- [13] Shuo Chen, Jindong Gu, Zhen Han, Yunpu Ma, Philip Torr, and Volker Tresp. Benchmarking robustness of adaptation methods on pre-trained vision-language models. *arXiv preprint arXiv:2306.02080*, 2023.
- [14] Minhao Cheng, Thong Le, Pin-Yu Chen, Jinfeng Yi, Huan Zhang, and Cho-Jui Hsieh. Query-efficient hard-label black-box attack: An optimization-based approach. *arXiv preprint arXiv:1807.04457*, 2018.
- [15] Minhao Cheng, Simranjit Singh, Patrick Chen, Pin-Yu Chen, Sijia Liu, and Cho-Jui Hsieh. Sign-opt: A query-efficient hard-label adversarial attack. *arXiv preprint arXiv:1909.10773*, 2019.
- [16] Khoa D Doan, Yingjie Lao, Peng Yang, and Ping Li. Defending backdoor attacks on vision transformer via patch processing. *arXiv preprint arXiv:2206.12381*, 2022.
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv:2010.11929*, 2020.
- [18] Jindong Gu, Baoyuan Wu, and Volker Tresp. Effective and efficient vote attack on capsule networks. In *International Conference on Learning Representations (ICLR)*, 2021.
- [19] Jindong Gu, Volker Tresp, and Yao Qin. Are vision transformers robust to patch perturbations? In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XII*, pages 404–421. Springer, 2022.
- [20] Jindong Gu, Ahmad Beirami, Xuezhi Wang, Alex Beutel, Philip Torr, and Yao Qin. Towards robust prompts on vision-language models. *arXiv preprint arXiv:2304.08479*, 2023.
- [21] Chuan Guo, Jacob Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Weinberger. Simple black-box adversarial attacks. In *International Conference on Machine Learning*, pages 2484–2493. PMLR, 2019.
- [22] Zhezhi He, Adnan Siraj Rakin, and Deliang Fan. Parametric noise injection: Trainable randomness to improve deep neural network robustness against adversarial attack. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 588–597, 2019.
- [23] Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. Black-box adversarial attacks with limited queries and information. In *International Conference on Machine Learning*, pages 2137–2146. PMLR, 2018.
- [24] Xuanqing Liu, Minhao Cheng, Huan Zhang, and Cho-Jui Hsieh. Towards robust neural networks via random self-ensemble. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 369–385, 2018.
- [25] Yujia Liu, Seyed-Mohsen Moosavi-Dezfooli, and Pascal Frossard. A geometry-inspired decision-based attack. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4890–4898, 2019.

- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *arXiv:2103.14030*, 2021.
- [27] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *International Conference on Learning Representations (ICLR)*, 2017.
- [28] Laurent Meunier, Jamal Atif, and Olivier Teytaud. Yet another but more efficient black-box adversarial attack: tiling and evolution strategies. *arXiv preprint arXiv:1910.02244*, 2019.
- [29] Seungyong Moon, Gaon An, and Hyun Oh Song. Parsimonious black-box adversarial attacks via efficient combinatorial optimization. In *International Conference on Machine Learning*, pages 4636–4645. PMLR, 2019.
- [30] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *Advances in Neural Information Processing Systems*, 34, 2021.
- [31] Sayak Paul and Pin-Yu Chen. Vision transformers are robust learners. *arXiv:2105.07581*, 2021.
- [32] Francesco Pinto, Philip HS Torr, and Puneet K. Dokania. An impartial take to the cnn vs transformer robustness contest. In *European Conference on Computer Vision*, pages 466–480. Springer, 2022.
- [33] Zeyu Qin, Yanbo Fan, Hongyuan Zha, and Baoyuan Wu. Random noise defense against query-based black-box attacks. *Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- [34] Ali Rahmati, Seyed-Mohsen Moosavi-Dezfooli, Pascal Frossard, and Huaiyu Dai. Geoda: a geometric framework for black-box adversarial attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8446–8455, 2020.
- [35] Jathushan Rajasegaran, Vinoj Jayasundara, Sandaru Jayasekara, Hirunima Jayasekara, Suranga Seneviratne, and Ranga Rodrigo. Deepcaps: Going deeper with capsule networks. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10725–10733, 2019.
- [36] Rulin Shao, Zhouxing Shi, Jinfeng Yi, Pin-Yu Chen, and Cho-Jui Hsieh. On the adversarial robustness of visual transformers. *arXiv:2103.15670*, 2021.
- [37] Yucheng Shi and Yahong Han. Decision-based black-box attack against vision transformers via patch-wise adversarial removal. *arXiv preprint arXiv:2112.03492*, 2021.
- [38] Satya Narayan Shukla, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Black-box adversarial attacks with bayesian optimization. *arXiv preprint arXiv:1909.13857*, 2019.

- [39] Shiyu Tang, Ruihao Gong, Yan Wang, Aishan Liu, Jiakai Wang, Xinyun Chen, Fengwei Yu, Xianglong Liu, Dawn Song, Alan Yuille, et al. Robustart: Benchmarking robustness on architecture design and training techniques. *arXiv preprint arXiv:2109.05211*, 2021.
- [40] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning (ICML)*, pages 10347–10357. PMLR, 2021.
- [41] Chun-Chen Tu, Paishun Ting, Pin-Yu Chen, Sijia Liu, Huan Zhang, Jinfeng Yi, Chou-Jui Hsieh, and Shin-Ming Cheng. Autozoom: Autoencoder-based zeroth order optimization method for attacking black-box neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 742–749, 2019.
- [42] Jonathan Uesato, Brendan O’donoghue, Pushmeet Kohli, and Aaron Oord. Adversarial risk and the dangers of evaluating against weak attacks. In *International Conference on Machine Learning*, pages 5025–5034. PMLR, 2018.
- [43] Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019.
- [44] Boxi Wu, Jindong Gu, Zhifeng Li, Deng Cai, Xiaofei He, and Wei Liu. Towards efficient adversarial training on vision transformers. In *European Conference on Computer Vision*, pages 307–325. Springer, 2022.
- [45] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. Mitigating adversarial effects through randomization. *International Conference on Learning Representations (ICLR)*, 2017.
- [46] Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang, Alan L Yuille, and Quoc V Le. Adversarial examples improve image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 819–828, 2020.
- [47] Wenqian Yu, Jindong Gu, Zhijiang Li, and Philip Torr. Reliable evaluation of adversarial transferability. *arXiv preprint arXiv:2306.08565*, 2023.