

# Improving Transferability of Adversarial Examples via Bayesian Attacks

Qizhang Li<sup>1</sup>, Yiwen Guo<sup>1</sup>, Xiaochen Yang<sup>1</sup>, Wangmeng Zuo<sup>1</sup>, *Senior Member, IEEE*,  
and Hao Chen<sup>2</sup>, *Fellow, IEEE*

**Abstract**—The transferability of adversarial examples allows for the attack on unknown deep neural networks (DNNs), posing a serious threat to many applications and attracting great attention. In this paper, we improve the transferability of adversarial examples by incorporating the Bayesian formulation into both the model parameters and model input, enabling their joint diversification. We demonstrate that combination of Bayesian formulations for both the model input and model parameters yields significant improvements in transferability. By introducing advanced approximations of the posterior distribution over the model input, adversarial transferability achieves further enhancement, surpassing all state-of-the-arts when attacking without model fine-tuning. Additionally, we propose a principled approach to fine-tune model parameters within this Bayesian framework. Extensive experiments demonstrate that our method achieves a new state-of-the-art in transfer-based attacks, significantly improving the average success rate on ImageNet and CIFAR-10. We will make our code publicly available.

**Index Terms**—Deep neural networks, adversarial examples, transferability, generalization ability.

## I. INTRODUCTION

DEEP neural networks (DNNs) have demonstrated remarkable performances in various applications, such as computer vision [1], natural language processing [2], and speech recognition [3]. Nevertheless, these models have been found to be vulnerable to adversarial examples [4], [5], [6], [7], [8], [9], which are maliciously perturbed input data that could mislead the models into making incorrect predictions. This raises serious concerns about the safety of DNNs, particularly when applying them in security-critical domains. More severely, adversarial examples are transferable [10]. An attacker can manipulate adversarial examples on a substitute model to attack unknown victim models with different architectures and parameters. An in-depth exploration of adversarial

transferability is crucial as it provides valuable insights into the nature of DNNs, enables more comprehensive evaluation of their robustness, and promotes the design of stronger defense methods and fully resistant networks.

Over the past several years, considerable effort has been devoted to improving the transferability of adversarial examples. In particular, ensemble-based methods [11], [12], [13] attract our attention as they can be readily combined with almost all other methods, such as optimizing intermediate-level representations [14], [15] or modifying the backpropagation computation [16], [17]. The effectiveness of ensemble-based attack is generally related to the number of models available in the bucket, thus we improve it to a statistical ensemble which consists of infinitely many substitute models in some sense by introducing a Bayesian formulation. To model the randomness of models comprehensively, we introduce distributions to the model parameters and model input. Parameters of these distributions can be obtained from an off-the-shelf model, and if fine-tuning is possible, these parameters can further be optimized. Adversarial examples are then crafted by maximizing the mean prediction loss averaged across the parameters and inputs sampled from their respective distributions.

We evaluate our method in attacking a variety of models on ImageNet [18] and CIFAR-10 [19]. The proposed method achieves comparable performance without fine-tuning a substitute model, and it outperforms state-of-the-arts considerably when fine-tuning is applied. We also show that our method can be readily integrated with existing methods for further improving the attack performance.

Compared with conventional input augmentation methods [20], [21], [22], [23], which apply a fixed set of transformations such as translation, scaling, or simple noise to diversify the input during the attack, our method introduces input diversity through a principled Bayesian formulation. This allows us to model input diversity via an advanced approximation of a posterior distribution learned directly from the dynamics of the iterative attack itself. Furthermore, this work extends parameter-only Bayesian attacks, such as our previous work [24] (referred to as “BasicBayesian” in this paper). Specifically, [24] exclusively modeled the diversity of the substitute model’s parameters, while the new method introduces a probability measure to model input in conjunction with the parameters of the substitute models, enabling the joint diversification of both the model and input throughout the iterative process of generating adversarial examples. Moreover,

Received 11 March 2025; revised 8 August 2025; accepted 9 September 2025. Date of publication 19 September 2025; date of current version 9 March 2026. This work was supported in part by the National Key Research and Development Program of China under Grant 2022YFA1004100. This article was recommended by Associate Editor Y. Liu. (*Corresponding author: Wangmeng Zuo.*)

Qizhang Li and Wangmeng Zuo are with the School of Computer Science and Technology, Harbin Institute of Technology, Harbin 150001, China (e-mail: liqizhang95@gmail.com; cswmzuo@gmail.com).

Yiwen Guo resides in Beijing, China (e-mail: guoyiwen89@gmail.com).

Xiaochen Yang is with the School of Mathematics and Statistics, University of Glasgow, G12 8QQ Glasgow, U.K. (e-mail: xiaochen.yang@glasgow.ac.uk).

Hao Chen is with the Department of Computer Science, University of California at Davis, Davis, CA 95616 USA (e-mail: chen@ucdavis.edu).

Digital Object Identifier 10.1109/TCSVT.2025.3609284

with the incorporation of input randomness in our Bayesian formulation, we newly derive a principled optimization objective for fine-tuning that encourages flat minima in both the parameter space and the input space. This newly proposed optimization objective enables our method to significantly outperform state-of-the-art approaches when fine-tuning the substitute model is feasible.

## II. BACKGROUND AND RELATED WORK

*Background on Adversarial Examples:* Deep neural networks are known to be vulnerable to adversarial attacks, where malicious inputs are crafted to deceive trained models. This security challenge has spurred a vast and diverse field of research. Attack paradigms, for instance, range from generating localized, perceptible patches [5], [25] to crafting the subtle, full-image perturbations which are the focus of this work. The targets of these attacks are similarly varied, encompassing not only standard static models [4], [26] but also more complex, adaptive networks with dynamic architectures [6]. Adversarial attacks have also been expanded to attack models performing various tasks, *e.g.*, facial recognition and retrieval [27], [28], [29], image compression [8], and person detection [30]. This ongoing threat has fueled an arms race between offense and defense, with novel attacks prompting the development of more generalizable defenses [31], [32], [33], [34], [35]. To push the offensive frontier, state-of-the-art methods have explored advanced techniques, including the use of generative models to manipulate latent spaces for better transferability [36]. Among the most fundamental and widely-studied of these methods are gradient-based attacks, which iteratively optimize perturbations to maximize a model's prediction loss.

Given a benign input  $\mathbf{x}$  with label  $y$  and a victim model  $f_w$  with full knowledge of architecture and parameters  $\mathbf{w}$ , gradient-based approaches generate adversarial examples  $\mathbf{x}^{\text{adv}}$  within the  $\ell_p$ -bounded neighborhood of  $\mathbf{x}$  to maximize the prediction loss  $L$ :

$$\arg \max_{\|\Delta \mathbf{x}\|_p \leq \epsilon} L(\mathbf{x} + \Delta \mathbf{x}, y, \mathbf{w}), \quad (1)$$

where  $\epsilon$  is the perturbation budget.

FGSM [26] is a simple one-step attack method to obtain adversarial examples in the  $p = \infty$  setting:

$$\mathbf{x}^{\text{adv}} = \mathbf{x} + \epsilon \cdot \text{sign}(\nabla_{\mathbf{x}} L(\mathbf{x}, y, \mathbf{w})). \quad (2)$$

The iterative variant of FGSM, I-FGSM [37], is capable of generating more powerful attacks:

$$\mathbf{x}_1^{\text{adv}} = \mathbf{x}, \mathbf{x}_{t+1}^{\text{adv}} = \text{Clip}_{\epsilon}(\mathbf{x}_t^{\text{adv}} + \eta \cdot \text{sign}(\nabla_{\mathbf{x}^{\text{adv}}} L(\mathbf{x}_t^{\text{adv}}, y, \mathbf{w}))), \quad (3)$$

where  $\eta$  is the step size and the  $\text{Clip}_{\epsilon}$  function ensures that the generated adversarial examples remain within the pre-specified range.

### A. Transfer-Based Attacks

FGSM and I-FGSM require calculating gradients of the victim model. Nevertheless, in practical scenarios, the attacker may not have enough knowledge about the victim model

for calculating gradients. To address this issue, many attacks rely on the transferability of adversarial examples, meaning that adversarial examples crafted for one classification model (using, for example, FGSM or I-FGSM) can often successfully attack other victim models. It is normally assumed to be able to query the victim model to annotate training samples, collect a set of samples from the same distribution as that modeled by the victim models, or collect a pre-trained substitute model that is trained to accomplish the same task as the victim models.

To enhance the transferability of adversarial examples, several groups of methods have been proposed. Advancements include improved optimizer and gradient computation techniques for updating  $\mathbf{x}^{\text{adv}}$  [14], [15], [17], [23], [35], [38], [39], [40], innovative strategies of training and fine-tuning substitute models [41], [42], and two groups of methods closely related to this paper: random input augmentation [20], [21], [22], [23], [43], [44], [45], [46] and substitute model augmentation (*i.e.*, ensemble) [11], [13], [47], [48]. For instance, in input augmentation methods, Wu et al. [49], Dong et al. [21], Xie et al. [20], Lin et al. [23], Wang et al. [44], and Wang et al. [46] introduced different sorts of transformations into the iterative update of adversarial examples. For instance, DI<sup>2</sup>-FGSM [20] utilizes random resizing and padding, and TI-FGSM [21] employs a translation-invariant methodology. Along similar lines, SI-FGSM [23] enforces scale invariance, and BSR [46] explores robustness to structural changes through block shuffling and rotation. The effectiveness of this paradigm is further highlighted by SIA [44], which achieves superior results by combining multiple distinct transformation types. GRA [45] also involves input augmentation strategy where multiple inputs are sampled from the neighborhood of the original image. The final update gradient is then computed by establishing a relevance between the gradients of the original and the various augmented inputs. Comparing with these methods, our method introduces input diversity via a principled Bayesian formulation and for the first time takes input diversity into account during substitute model fine-tuning. For substitute model augmentation (*i.e.*, ensemble), Liu et al. [11] proposed to generate adversarial examples on an ensemble of multiple substitute models that differ in their architectures. Additionally, Xiong et al. [13] proposed stochastic variance reduced ensemble to reduce the variance of gradients of different substitute models following the spirit of stochastic variance reduced gradient [50]. Gubri et al. [48] suggested fine-tuning with a fixed and large learning rate to collect multiple models along the training trajectory for the ensemble attack. In this paper, we consider the diversity in both the substitute models and model inputs by introducing a Bayesian approximation for achieving this.

### B. Bayesian DNNs

If a deep neural network (DNN) is considered as a probabilistic model, the process of training its parameters, denoted as  $\mathbf{w}$ , can be seen as maximum likelihood estimation or maximum a posteriori estimation (with regularization). In Bayesian deep learning, the approach involves simultaneously

estimating the posterior distribution of the parameters given the data. The prediction for any new input instance is obtained by taking the expectation over this posterior distribution. Due to the large number of parameters typically involved in DNNs, optimizing Bayesian models becomes more challenging compared to shallow models. As a result, numerous studies have been conducted to address this issue, leading to the development of various scalable approximations. Effective methods utilize variational inference [43], [51], [52], [53], [54], [55], [56], [57], [58], [59], dropout inference [60], [61], [62], Laplace approximation [63], [64], [65], or stochastic gradient descent (SGD)-based approximation [66], [67], [68], [69]. Taking SWAG [67] as an example, which is an SGD-based approximation, it approximates the posterior using a Gaussian distribution with the stochastic weight averaging (SWA) solution as its first raw moment and the composition of a low rank matrix and a diagonal matrix as its second central moment. Our method is developed in a Bayesian spirit and we shall discuss SWAG thoroughly later in this paper. In addition to approximating the posterior over model parameters, our work also involves approximating the posterior over adversarial perturbations. This is implemented in order to introduce randomization into the model input during each iteration of iterative attacks.

In recent years, there has been research on studying the robustness of Bayesian DNNs. Besides exploring the probabilistic robustness and safety measures of such models [70], [71], attacks have been adapted [72], [73] to evaluate the robustness of these models in practice. While Bayesian models are often considered to be more robust [74], [75], adversarial training has also been proposed for further safeguarding them, as can be seen in the work by Liu et al. [72]. However, these studies have not specifically focused on adversarial transferability as we do in this paper.

### III. BAYESIAN ATTACK FOR IMPROVED TRANSFERABILITY

An intuition for improving the transferability of adversarial examples suggests improving the diversity during back-propagation. The BasicBayesian [24] increases model diversity through Bayesian modeling, while this paper considers both model diversity and input diversity jointly.

#### A. Generate Adversarial Examples via Bayesian Modeling

Bayesian learning aims to discover a distribution of likely models rather than a single deterministic model. Based on the obtained distribution of models, we can perform a transfer-based attack on an ensemble of infinitely many models. Let  $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$  denote a training set. Bayesian inference incorporates a prior belief about the parameters  $\mathbf{w}$  through the prior distribution  $p(\mathbf{w})$  and updates this belief after observing the data  $\mathcal{D}$  using Bayes' theorem, resulting in the posterior distribution:  $p(\mathbf{w}|\mathcal{D}) \propto p(\mathcal{D}|\mathbf{w})p(\mathbf{w})$ . For a new input  $\mathbf{x}$ , the predictive distribution of its class label  $y$  is given by the Bayesian model averaging, *i.e.*,

$$p(y|\mathbf{x}, \mathcal{D}) = \int_{\mathbf{w}} p(y|\mathbf{x}, \mathbf{w})p(\mathbf{w}|\mathcal{D})d\mathbf{w}, \quad (4)$$

where  $p(y|\mathbf{x}, \mathbf{w})$  is the conditional probability, obtained from the DNN output followed by a softmax function.

To perform attack on such a Bayesian model, a straightforward idea is to minimize the probability of the true class for a given input, as in [24]:

$$\begin{aligned} & \arg \min_{\|\Delta\mathbf{x}\|_p \leq \epsilon} p(y|\mathbf{x} + \Delta\mathbf{x}, \mathcal{D}) \\ & = \arg \min_{\|\Delta\mathbf{x}\|_p \leq \epsilon} \int_{\mathbf{w}} p(y|\mathbf{x} + \Delta\mathbf{x}, \mathbf{w})p(\mathbf{w}|\mathcal{D})d\mathbf{w}. \end{aligned} \quad (5)$$

In this paper, considering multiple-step attacks, an iterative optimizer (*e.g.*, I-FGSM) is applied, and at the  $t$ -th iteration, it seeks  $\Delta\mathbf{x}_t$  to minimize

$$\int_{\mathbf{w}} p(y|\mathbf{x} + \widehat{\Delta\mathbf{x}}_t + \Delta\mathbf{x}_t, \mathbf{w})p(\mathbf{w}|\mathcal{D})d\mathbf{w} \quad (6)$$

while ensuring  $\|\widehat{\Delta\mathbf{x}}_t + \Delta\mathbf{x}_t\|_p \leq \epsilon$ .  $\widehat{\Delta\mathbf{x}}_1 = \mathbf{0}$ , and for  $t > 1$ ,  $\widehat{\Delta\mathbf{x}}_t = \sum_{t'=1}^{t-1} \Delta\mathbf{x}_{t'}$  is the sum of all perturbations accumulated over the previous  $t-1$  iterations. Optimizing Eq. (5) or minimizing Eq. (6) can be regarded as generating adversarial examples that could succeed on a distribution of models, and it has been proved effective in our previous work [24]. For the rest of this paper, we shall use the notation  $\mathbf{x}_t^{\text{adv}}$  to represent  $\mathbf{x} + \widehat{\Delta\mathbf{x}}_t$ .

In Eqs. (5) and (6), Bayesian model averaging is used to predict the label of a deterministic model input. Such perturbation optimization processes solely consider the diversity of models, while the diversity of model inputs is overlooked. Nevertheless, considering that different DNN models may be equipped with different pre-processing operations, given the same benign input (*e.g.*, a clean image),  $\mathbf{x}$  can be different for different models after their specific pre-processing steps. Moreover, even given the same  $\mathbf{x}$  (which is the pre-processed model input), different models obtain different  $\Delta\mathbf{x}_t$ . Therefore, following the motivation of improving the model diversity in our previous work [24], introducing input diversity into Eqs. (5) and (6) may also be beneficial to the transferability of generated adversarial examples.

To incorporate such diversity, we simply add a randomness term  $\mathbf{e}$  to the input  $\mathbf{x}_t^{\text{adv}}$ . The randomness term should not significantly affect the model's prediction. Therefore, the Eq. (6) can be rewritten as:

$$\int_{\mathbf{w}, \mathbf{e}} p(y|\mathbf{x}_t^{\text{adv}} + \mathbf{e} + \Delta\mathbf{x}_t, \mathbf{w})p(\mathbf{w}|\mathcal{D})p(\mathbf{e}|\mathbf{w}, \mathbf{x}_t^{\text{adv}})d\mathbf{w}d\mathbf{e}. \quad (7)$$

Here the randomness term  $\mathbf{e}$  is added linearly to the input. We can also introduce more complex modifications, such as by using a Gaussian filter  $g(\mathbf{x}_t^{\text{adv}}, \mathbf{e})$  with a random standard deviation [76]; yet, for simplicity, we discuss the linear case in this paper. Due to the very large number of parameters, it is intractable to perform exact inference using Eq. (7). Instead, we adopt the Monte Carlo sampling to approximate the integral, where a set of  $M$  models, each parameterized by  $\mathbf{w}_j$ , are sampled from the posterior  $p(\mathbf{w}|\mathcal{D})$ , and  $S$  inputs are sampled from  $p(\mathbf{e}|\mathbf{w}, \mathbf{x}_t^{\text{adv}})$ . The optimization problem can then be cast to

$$\arg \min_{\Delta\mathbf{x}_t} \frac{1}{MS} \sum_{j=1}^M \sum_{k=1}^S p(y|\mathbf{x}_t^{\text{adv}} + \mathbf{e}_k + \Delta\mathbf{x}_t, \mathbf{w}_j)$$

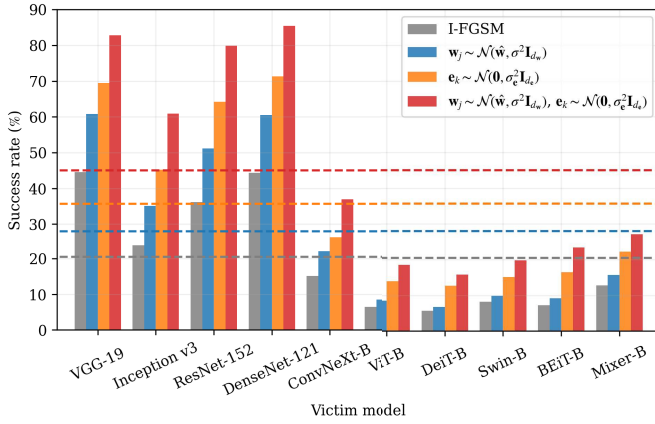


Fig. 1. Comparison of success rates in attacking 10 different victim models when adversarial examples are generated on a substitute model (ResNet-50) in a non-Bayesian manner (*i.e.*, I-FGSM) and using Bayesian modeling for model parameters  $\mathbf{w}_j$ , inputs  $\mathbf{e}_k$  (*i.e.*, vr-IGSM [49]), and both  $\mathbf{w}_j$  and  $\mathbf{e}_k$  (without fine-tuning). Dotted lines indicate the average success rate across all 10 victim models. We performed  $t_\infty$  attacks with  $\epsilon = 8/255$ . Best viewed in color.

$$= \arg \max_{\Delta \mathbf{x}_t} \frac{1}{MS} \sum_{j=1}^M \sum_{k=1}^S L(\mathbf{x}_t^{\text{adv}} + \mathbf{e}_k + \Delta \mathbf{x}_t, y, \mathbf{w}_j),$$

$$\mathbf{w}_j \sim p(\mathbf{w}|\mathcal{D}), \mathbf{e}_k \sim p(\mathbf{e}|\mathbf{w}, \mathbf{x}_t^{\text{adv}}), \quad (8)$$

where  $L(\cdot, \cdot, \mathbf{w}_j)$  is a function evaluating the prediction loss of a model parameterized by  $\mathbf{w}_j$ .

### B. Construct a Posterior Without Fine-Tuning

Given any pre-trained DNN model whose parameters are  $\hat{\mathbf{w}}$ , we can simply obtain a posterior distribution for  $\mathbf{w}_j$  without fine-tuning by assuming it follows an isotropic Gaussian [24] centered at  $\hat{\mathbf{w}}$ , *i.e.*,  $\mathbf{w}_j \sim \mathcal{N}(\hat{\mathbf{w}}, \sigma^2 \mathbf{I}_{d_w})$ , where  $\sigma$  is a positive constant that controls the diversity of distribution. Similarly, we can consider  $\mathbf{e}_k \sim \mathcal{N}(0, \sigma_e^2 \mathbf{I}_{d_e})$ , and it is equivalent to vr-IGSM [49] which suggests adding isotropic Gaussian noise into the model at each attack iteration. However, the improvement of joint diversification of model input and model parameters on adversarial transferability remains uncertain. Figure 1 compares the effectiveness of generating transferable attacks using I-FGSM (non-Bayesian), the original implementation of BasicBayesian [24] (Bayesian over  $\mathbf{w}_j$  alone), the implementation of vr-IGSM [49] (Bayesian over  $\mathbf{e}_k$  alone), and our approach in Eq. (8) (Bayesian over both  $\mathbf{w}_j$  and  $\mathbf{e}_k$ ); Gaussian posteriors were assumed for the latter three methods. The experiment was conducted on ImageNet with ResNet-50 used as the substitute model; experiment details are deferred to Section IV-A. Apparently, a joint diversification could further enhance adversarial transferability with a 17.05% increase in the average success rate compared with considering the model diversity alone [24], and a 9.37% increase compared with considering the input diversity alone [49]. Observing the success of joint diversification, we propose a fine-tuning approach in Section III-C to obtain a more suitable posterior of model parameters. Additionally, we incorporate the principles of improved distribution modeling from SWAG to further

improve the posteriors of both parameters and inputs in Section III-D.

### C. Obtain a More Suitable Posterior via Fine-Tuning

In this subsection, we explain the optimization procedure of the posterior when fine-tuning the Bayesian model from a pre-trained model is possible. Following prior work, we consider a threat model in which fine-tuning can be performed on datasets collected for the same task as the victim models.

As in Section III-B, we assume an isotropic Gaussian posterior  $\mathcal{N}(\hat{\mathbf{w}}, \sigma^2 \mathbf{I}_{d_w})$ ; however, the mean vector  $\hat{\mathbf{w}}$  is now considered as a trainable parameter. Optimization of the Bayesian model, or more specifically  $\hat{\mathbf{w}}$ , can be formulated as:

$$\max_{\hat{\mathbf{w}}} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\hat{\mathbf{w}}, \sigma^2 \mathbf{I}_{d_w}), \mathbf{e} \sim \mathcal{N}(0, \sigma_e^2 \mathbf{I}_{d_e})} P(y_i | \mathbf{x}_i, \mathbf{e}, \mathbf{w}). \quad (9)$$

By adopting Monte Carlo sampling, it can further be reformulated as:

$$\min_{\hat{\mathbf{w}}} \frac{1}{NMS} \sum_{i=1}^N \sum_{j=1}^M \sum_{k=1}^S L(\mathbf{x}_i + \mathbf{e}_{i,k}, y_i, \hat{\mathbf{w}} + \Delta \mathbf{w}_j),$$

$$\Delta \mathbf{w}_j \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{d_w}), \mathbf{e}_{i,k} \sim \mathcal{N}(0, \sigma_e^2 \mathbf{I}_{d_e}). \quad (10)$$

The computational complexity of Eq. (10) is high, thus we focus on the worst-case performance in the distributions, whose loss bounds the objective in Eq. (10) from below. The optimization problem then becomes:

$$\min_{\hat{\mathbf{w}}} \max_{\Delta \mathbf{w}, \mathbf{e}_i} \frac{1}{N} \sum_{i=1}^N L(\mathbf{x}_i + \mathbf{e}_i, y_i, \hat{\mathbf{w}} + \Delta \mathbf{w}),$$

$$\text{s.t. } \varphi_{(0, \sigma^2 \mathbf{I}_{d_w})}(\Delta \mathbf{w}) \geq \epsilon, \varphi_{(0, \sigma_e^2 \mathbf{I}_{d_e})}(\mathbf{e}_i) \geq \epsilon_e, \quad (11)$$

where  $\varphi_{(\boldsymbol{\mu}, \boldsymbol{\Sigma})}$  denotes the probability density function of a Gaussian distribution with a mean vector of  $\boldsymbol{\mu}$  and a covariance matrix of  $\boldsymbol{\Sigma}$ ;  $\epsilon$  and  $\epsilon_e$  control the confidence region of the Gaussian distributions.

To optimize the min-max problem in Eq. (11), we first solve the inner maximization problem to find  $\Delta \mathbf{w}^*$  and  $\mathbf{e}_i^*$ . Based on the first-order Taylor expansion, the loss function  $L$  can be written as:

$$L(\mathbf{x}_i + \mathbf{e}_i, y_i, \hat{\mathbf{w}} + \Delta \mathbf{w})$$

$$\approx L(\mathbf{x}_i, y_i, \hat{\mathbf{w}}) + \nabla_{\hat{\mathbf{w}}} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}})^\top \Delta \mathbf{w} + \nabla_{\mathbf{x}_i} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}})^\top \mathbf{e}_i. \quad (12)$$

The inner maximization problem then simplifies to maximizing a linear function of  $\Delta \mathbf{w}$  and  $\mathbf{e}_i$ :

$$\max_{\Delta \mathbf{w}, \mathbf{e}_i} \left[ \frac{1}{N} \sum_{i=1}^N (\nabla_{\hat{\mathbf{w}}} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}}))^\top \Delta \mathbf{w} \right. \\ \left. + \frac{1}{N} \sum_{i=1}^N (\nabla_{\mathbf{x}_i} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}}))^\top \mathbf{e}_i \right],$$

$$\text{s.t. } \varphi_{(0, \sigma^2 \mathbf{I}_{d_w})}(\Delta \mathbf{w}) \geq \epsilon, \varphi_{(0, \sigma_e^2 \mathbf{I}_{d_e})}(\mathbf{e}_i) \geq \epsilon_e, \quad (13)$$

The solution to maximizing a dot product under a norm constraint is to align the vector with the gradient. Thus, the optimal solutions  $\Delta \mathbf{w}^*$  and  $\mathbf{e}_i^*$  are in the direction of

their respective gradients (*i.e.*,  $\frac{1}{N} \sum_{i=1}^N (\nabla_{\hat{\mathbf{w}}} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}}))$  and  $\nabla_{\mathbf{x}_i} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}})$ ), with magnitudes as large as possible within the specified confidence regions. These confidence regions are defined by the constraints. For example, consider the constraint  $\varphi_{\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_{d_w})}(\Delta \mathbf{w}) \geq \epsilon$ . For an isotropic Gaussian distribution, this probability density constraint is equivalent to bounding the  $\ell_2$ -norm of the  $\Delta \mathbf{w}$  within a hypersphere.  $\lambda_{\epsilon, \sigma}$  is the radius of this hypersphere, *i.e.*,  $\|\Delta \mathbf{w}\|_2 \leq \lambda_{\epsilon, \sigma}$ . It is determined by the probability  $\epsilon$  and the variance  $\sigma^2$ . Specifically, this radius can be formally derived by solving the constraint inequality for  $\|\Delta \mathbf{w}\|_2$ . Starting with the constraint and the definition of the Gaussian probability density, we have

$$\varphi_{\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_{d_w})}(\Delta \mathbf{w}) = \frac{1}{(2\pi\sigma^2)^{d_w/2}} \exp\left(-\frac{\|\Delta \mathbf{w}\|_2^2}{2\sigma^2}\right) \geq \epsilon. \quad (14)$$

By taking the logarithm and rearranging the terms to solve for the norm, we find the bound for its squared value:

$$\|\Delta \mathbf{w}\|_2^2 \leq 2\sigma^2 \ln\left(\frac{1}{\epsilon(2\pi\sigma^2)^{d_w/2}}\right). \quad (15)$$

Therefore, the radius  $\lambda_{\epsilon, \sigma}$  of the hypersphere is the square root of this bound:

$$\lambda_{\epsilon, \sigma} = \sqrt{2\sigma^2 \ln\left(\frac{1}{\epsilon(2\pi\sigma^2)^{d_w/2}}\right)}. \quad (16)$$

The calculation of the radius  $\lambda_{\epsilon_e, \sigma_e}$  of the confidence region defined by  $\varphi_{(\mathbf{0}, \sigma_e^2 \mathbf{I}_{d_e})}(\mathbf{e}_i) \geq \epsilon_e$  proceeds analogously. Recall that the magnitude of  $\Delta \mathbf{w}^*$  and  $\mathbf{e}_i^*$  should be as large as possible within the confidence regions, we have  $\|\Delta \mathbf{w}^*\| = \lambda_{\epsilon, \sigma}$  and  $\|\mathbf{e}_i^*\|_2 = \lambda_{\epsilon_e, \sigma_e}$ . Then we obtain the approximate solution of the inner maximization in Eq. (11):

$$\begin{aligned} \Delta \mathbf{w}^* &\approx \lambda_{\epsilon, \sigma} \frac{\frac{1}{N} \sum_{i=1}^N \nabla_{\hat{\mathbf{w}}} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}})}{\left\| \frac{1}{N} \sum_{i=1}^N \nabla_{\hat{\mathbf{w}}} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}}) \right\|_2}, \\ \mathbf{e}_i^* &\approx \lambda_{\epsilon_e, \sigma_e} \frac{\nabla_{\mathbf{x}_i} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}})}{\left\| \nabla_{\mathbf{x}_i} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}}) \right\|_2}. \end{aligned} \quad (17)$$

Next, for the outer minimization in Eq. (11), we have:

$$\min_{\hat{\mathbf{w}}} \left[ \frac{1}{N} \sum_{i=1}^N L(\mathbf{x}_i, y_i, \hat{\mathbf{w}}) + \nabla_{\hat{\mathbf{w}}} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}})^\top \Delta \mathbf{w}^* + \nabla_{\mathbf{x}_i} L(\mathbf{x}_i, y_i, \hat{\mathbf{w}})^\top \mathbf{e}_i^* \right]. \quad (18)$$

The gradient for solving this minimization problem is:

$$\nabla_{\hat{\mathbf{w}}} \frac{1}{N} \sum_{i=1}^N L(\mathbf{x}_i, y_i, \hat{\mathbf{w}}) + \mathbf{H}_{\hat{\mathbf{w}}, \hat{\mathbf{w}}} \Delta \mathbf{w}^* + \mathbf{H}_{\hat{\mathbf{w}}, \mathbf{x}_i} \mathbf{e}_i^*, \quad (19)$$

which involves second-order partial derivatives in the Hessian matrix, and it can be approximately calculated using the finite difference method. That is, we use:

$$\begin{aligned} \mathbf{H}_{\hat{\mathbf{w}}, \hat{\mathbf{w}}} \Delta \mathbf{w}^* &\approx \frac{1}{\gamma} \left( \nabla_{\hat{\mathbf{w}}} \frac{1}{N} \sum_{i=1}^N L(\mathbf{x}_i, y_i, \hat{\mathbf{w}} + \gamma \Delta \mathbf{w}^*) \right. \\ &\quad \left. - \nabla_{\hat{\mathbf{w}}} \frac{1}{N} \sum_{i=1}^N L(\mathbf{x}_i, y_i, \hat{\mathbf{w}}) \right), \end{aligned} \quad (20)$$




Fig. 2. Comparison of adversarial attacks with and without the proposed model fine-tuning, both using isotropic Gaussian posteriors over model parameters and inputs. Dotted lines indicate the average success rates over all 10 victim models. We performed  $\ell_\infty$  attacks with  $\epsilon = 8/255$ . Best viewed in color.

where  $\gamma$  is a small positive constant.  $\mathbf{H}_{\hat{\mathbf{w}}, \mathbf{x}_i} \mathbf{e}_i^*$  is approximated similarly as:

$$\begin{aligned} \mathbf{H}_{\hat{\mathbf{w}}, \mathbf{x}_i} \mathbf{e}_i^* &\approx \frac{1}{\gamma} \left( \nabla_{\hat{\mathbf{w}}} \frac{1}{N} \sum_{i=1}^N L(\mathbf{x}_i + \gamma \mathbf{e}_i^*, y_i, \hat{\mathbf{w}}) \right. \\ &\quad \left. - \nabla_{\hat{\mathbf{w}}} \frac{1}{N} \sum_{i=1}^N L(\mathbf{x}_i, y_i, \hat{\mathbf{w}}) \right). \end{aligned} \quad (21)$$

By introducing  $\mathbf{H}_{\hat{\mathbf{w}}, \hat{\mathbf{w}}} \Delta \mathbf{w}^*$  and  $\mathbf{H}_{\hat{\mathbf{w}}, \mathbf{x}_i} \mathbf{e}_i^*$  in Eq. (19), we encourage flat minima in the parameter space and the input space, respectively. The former is known to be beneficial to the generalization ability of DNNs [77], while the later has been paid little attention during training/fine-tuning.

To evaluate the effectiveness of fine-tuning, an experiment was carried out on ImageNet using ResNet-50 as the substitute model, and results are shown in Figure 2. It clearly suggests that fine-tuning leads to more significant adversarial transferability. More detailed comparison results are provided in Table I.

#### D. Improved Distribution Modeling

In Sections III-B and III-C, we have demonstrated the superiority of adopting the Bayesian formulation for generating transferable adversarial examples. However, it should be noted that our approach relies on a relatively strong assumption that the posterior follows an isotropic distribution. Taking one step further, we remove the assumption about the covariance matrix and try to learn it from data in this subsection.

Numerous methods have been proposed to learn covariance matrices, but in this paper we opt for SWAG [67] due to its simplicity and scalability. SWAG offers an enhanced Gaussian approximation to the distribution of  $\mathbf{w}$  and  $\mathbf{e}_i$  (for brevity, the subscript  $i$  will be dropped in this subsection). Specifically, we adopt the SWA solution [88] as the mean of  $\mathbf{w}$ , and decompose the covariance matrix into a diagonal term, a low-rank term, and a scaled identity term, *i.e.*,  $\mathbf{w} \sim \mathcal{N}(\mathbf{w}_{\text{SWA}}, \Sigma_{\mathbf{w}})$ , where

$$\Sigma_{\mathbf{w}} = \alpha(\Sigma_{\text{diag}} + \Sigma_{\text{low-rank}}) + \beta \mathbf{I}, \quad (22)$$

TABLE I

COMPARING TRANSFERABILITY OF FGSM AND I-FGSM ADVERSARIAL EXAMPLES GENERATED ON A DETERMINISTIC SUBSTITUTE MODEL AND THE BAYESIAN SUBSTITUTE MODEL (WITH ISOTROPIC OR IMPROVED GAUSSIAN POSTERIORIS AND WITH OR WITHOUT FINE-TUNE) UNDER THE  $\ell_\infty$  CONSTRAINT WITH  $\epsilon = 8/255$  ON IMAGENET. THE ARCHITECTURE OF THE SUBSTITUTE MODELS IS RESNET-50, AND ‘‘AVERAGE’’ WAS CALCULATED OVER ALL TEN VICTIM MODELS. THE BEST RESULTS ARE MARKED IN BOLD

		fine-tune	VGG -19 [78]	Inception v3 [79]	ResNet -152 [80]	DenseNet -121 [81]	ConvNeXt -B [82]	ViT -B [83]	DeiT -B [84]	Swin -B [85]	BEiT -B [86]	Mixer -B [87]	Average
FGSM	-	$\times$	41.26%	28.08%	33.64%	41.20%	14.50%	9.04%	8.04%	9.52%	9.30%	17.72%	21.23%
	Isotropic	$\times$	52.02%	38.10%	46.40%	53.84%	17.06%	12.50%	11.90%	13.00%	13.62%	23.18%	28.16%
		$\checkmark$	75.50%	57.14%	66.08%	76.78%	24.94%	26.02%	24.00%	23.58%	30.12%	38.04%	44.22%
Improved	$\times$	-	-	-	-	-	-	-	-	-	-	-	-
	$\checkmark$	<b>84.50%</b>	<b>68.04%</b>	<b>77.10%</b>	<b>85.94%</b>	<b>31.06%</b>	<b>33.32%</b>	<b>30.76%</b>	<b>29.42%</b>	<b>38.62%</b>	<b>45.06%</b>	<b>52.38%</b>	
I-FGSM	-	$\times$	44.52%	24.06%	36.14%	44.40%	15.16%	6.50%	5.68%	8.24%	7.24%	12.76%	20.47%
	Isotropic	$\times$	82.86%	60.96%	79.94%	85.42%	36.92%	18.52%	15.86%	19.82%	23.46%	27.14%	45.09%
		$\checkmark$	97.32%	89.94%	96.58%	98.84%	49.44%	43.68%	35.00%	34.56%	59.62%	49.70%	65.47%
Improved	$\times$	92.40%	77.22%	91.42%	94.48%	45.44%	29.20%	26.32%	31.36%	36.74%	39.32%	56.39%	
	$\checkmark$	<b>99.12%</b>	<b>95.64%</b>	<b>99.12%</b>	<b>99.66%</b>	<b>64.28%</b>	<b>58.20%</b>	<b>51.66%</b>	<b>49.26%</b>	<b>73.98%</b>	<b>61.04%</b>	<b>75.20%</b>	

### Algorithm 1 Fine-Tuning of Substitute Model

- 1: **Input:** Pre-trained substitute model parameters  $\hat{\mathbf{w}}$ , training data  $\mathcal{D}$ , number of iterations  $T$ , SWAG start iteration  $T_{\text{SWA}}$ , SWAG posterior update interval  $c$
- 2: **Output:** Posterior distribution over model parameters
- 3: **for**  $t = 1$   $T$  **do**
- 4: Draw a mini-batch from training set  $\{\mathbf{x}_i, y_i\}_i^N \sim \mathcal{D}$
- 5: Calculate gradient using Eqs. (19)–(21), and update parameters  $\hat{\mathbf{w}}$  with the SGD optimizer
- 6: Set  $r = t - T_{\text{SWA}}$
- 7: **if** adopt SWAG and  $r > 0$  and  $r \bmod c = 0$  **then**
- 8: Update posterior over model parameters:
 
$$p(\mathbf{w}|\mathcal{D}) \leftarrow \mathcal{N}(\mathbf{w}_{\text{SWA}}, \Sigma_{\mathbf{w}})$$
- 9: **end if**
- 10: **end for**
- 11: **if** adopt SWAG **then**
- 12: **return**  $\mathcal{N}(\mathbf{w}_{\text{SWA}}, \Sigma_{\mathbf{w}})$
- 13: **else**
- 14: **return**  $\mathcal{N}(\hat{\mathbf{w}}, \sigma^2 \mathbf{I}_{d_w})$
- 15: **end if**

$\alpha \geq 0$  represent the scaling factor of SWAG for disassociating the learning rate of the covariance [67], and  $\beta \geq 0$  controls the covariance matrix of the isotropic Gaussian distribution.  $\mathbf{w}_{\text{SWA}}$ ,  $\Sigma_{\text{diag}}$ , and  $\Sigma_{\text{low-rank}}$  are obtained after fine-tuning converges, thus we keep the fine-tuning loss and fine-tuning mechanism as in Section III-C even with the improved distribution modeling. Algorithm 1 describes the fine-tuning procedure for the substitute model, optionally incorporating SWAG.

For multi-step attacks, e.g., in I-FGSM,  $\mathbf{e}$  can be similarly obtained from the distribution of perturbations. Specifically, we can obtain a Gaussian distribution of  $\mathbf{x}^{\text{adv}}$  using the information contained in the trajectory of a multi-step attack by the approach of SWAG, i.e.,  $\mathcal{N}(\mathbf{x}_{\text{SWA}}^{\text{adv}}, \Sigma_{\mathbf{x}^{\text{adv}}})$ . The  $\mathbf{x}^{\text{adv}}$  from this distribution would elicit similar predictions by being fed into the model. Therefore, for the model input at the  $t$ -th iteration, i.e.,  $\mathbf{x}_t^{\text{adv}}$ , we can readily obtain a posterior distribution of  $\mathbf{e}$  as  $\mathcal{N}(\mathbf{e}_{\text{SWA}}, \Sigma_{\mathbf{e}}) = \mathcal{N}(\mathbf{x}_{\text{SWA}}^{\text{adv}} - \mathbf{x}_t^{\text{adv}}, \Sigma_{\mathbf{x}^{\text{adv}}})$ . This posterior distribution

### Algorithm 2 Generation of Adversarial Examples

- 1: **Input:** Benign input  $\mathbf{x}$  and label  $y$ , posterior over model parameters  $p(\mathbf{w}|\mathcal{D})$ , posterior over input randomness terms  $p(\mathbf{e}|\mathbf{w}, \mathbf{x}^{\text{adv}})$ , number of iterations  $T$ , step size  $\eta$ , perturbation budget  $\epsilon$ , number of model samples  $M$ , number of input samples  $S$ , SWAG start iteration  $T_{\text{SWA}}$ , SWAG posterior update interval  $c$
- 2: **Output:** An adversarial example
- 3: **Initialize:**  $\mathbf{x}_1^{\text{adv}} \leftarrow \mathbf{x}$ ,  $p(\mathbf{e}|\mathbf{w}, \mathbf{x}^{\text{adv}}) \leftarrow \mathcal{N}(\mathbf{0}, \sigma_{\mathbf{e}}^2 \mathbf{I}_{d_e})$
- 4: **for**  $t = 1$   $T$  **do**
- 5: Sample model parameters  $\mathbf{w}_j$  and input randomness terms  $\mathbf{e}_k$ :
 
$$\begin{aligned} \mathbf{w}_j &\sim p(\mathbf{w}|\mathcal{D}) && \text{for } j = 1, \dots, M \\ \mathbf{e}_k &\sim p(\mathbf{e}|\mathbf{w}, \mathbf{x}_t^{\text{adv}}) && \text{for } k = 1, \dots, S \end{aligned}$$
- 6: Compute input gradient of Eq. (8):
 
$$\mathbf{g}_t = \frac{1}{MS} \sum_{j=1}^M \sum_{k=1}^S \nabla_{\mathbf{x}_t^{\text{adv}}} L(\mathbf{x}_t^{\text{adv}} + \mathbf{e}_k, y, \mathbf{w}_j)$$
- 7: Update the adversarial example:
 
$$\mathbf{x}_{t+1}^{\text{adv}} = \text{Clip}_{\epsilon}(\mathbf{x}_t^{\text{adv}} + \eta \cdot \text{sign}(\mathbf{g}_t))$$
- 8: Let  $r = t - T_{\text{SWA}}$
- 9: **if** adopt SWAG and  $r > T_{\text{SWA}}$  and  $r \bmod c = 0$  **then**
- 10: Update posterior over randomness term:
 
$$p(\mathbf{e}|\mathbf{w}, \mathbf{x}_t^{\text{adv}}) \leftarrow \mathcal{N}(\mathbf{x}_{\text{SWA}}^{\text{adv}} - \mathbf{x}_t^{\text{adv}}, \Sigma_{\mathbf{x}^{\text{adv}}})$$
- 11: **end if**
- 12: **end for**
- 13: **return**  $\mathbf{x}_{T+1}^{\text{adv}}$

is preferable to the isotropic distribution since it constrains the effect of the randomness term  $\mathbf{e}$  on the prediction. Unlike for  $\mathbf{w}_{\text{SWA}}$  and  $\Sigma_{\mathbf{w}}$ , it does not necessarily require model fine-tuning to obtain the mean vector and covariance matrix of the distribution, as they model the randomness in model input instead of in model parameters and the learning dynamics of the adversarial inputs are obtained during I-FGSM. We

can start modeling the distribution and sampling from the distribution at the 10-th iteration out of 50. Yet, for single-step attacks like FGSM, such a distribution modeling boils down to be equivalent to the isotropic case. Our approach to the diversity of model inputs differs from existing input augmentation techniques. Instead of simply adding noise, cropping, or translating images, we establish the posterior distribution over the input through the learning dynamics of iterative attacks and sample from this posterior distribution. Moreover, our approach can be easily combined with existing input augmentation methods to further improve performance. Experimental results are shown in Section IV-C. Algorithm 2 outlines the iterative generation of adversarial examples using the joint Bayesian framework.

We compare performance of methods with and without such improved distribution, in single-step and multi-step attacks, in Table I. Approximating the covariance matrix of parameters through SWAG improves the attack success rates on all victim models, leading to an average increase of 8.16% (from 44.22% to 52.38%) when using single-step attacks based on FGSM. This indicates that the more general distributional assumption of model parameters aligns better with the distribution of victim parameters in practice. A similar benefit is observed when applying the improved distribution modeling to the model input, where the average success rate increases from 45.09% to 56.39% in using I-FGSM. The best performance is achieved when both improved posteriors and fine-tuning are adopted. However, note that with fine-tuning, we suggest not to adopt such improved distribution modeling of model input and model parameters simultaneously, as the two distributions will not be independent under such circumstances and there will be more hyper-parameters to be tuned. Following this suggestion, in Table I and subsequent experiments, when our fine-tuning, which considers both model parameters and model input, is performed, we only adopt SWAG for the model parameters and use the isotropic formulation for the model inputs during attack.

#### IV. EXPERIMENTS

We evaluate the effectiveness of our method by comparing it to recent state-of-the-arts in this section.

##### A. Experimental Settings

To be consistent with [24], we focused on untargeted  $\ell_\infty$  attacks to study the adversarial transferability. All experiments were conducted on the same set of ImageNet [18] models collected from the timm repository [89], *i.e.*, ResNet-50 [80], VGG-19 [78], Inception v3 [79], ResNet-152 [80], DenseNet-121 [81], ConvNeXt-B [82], ViT-B [83], DeiT-B [84], Swin-B [85], BEiT-B [86], and MLP-Mixer-B [87]. These models are well-known and encompass CNN, transformer, and MLP architectures, making the experiments more comprehensive. We randomly sampled 5000 test images that can be correctly classified by all these models from the ImageNet validation set for evaluation. Since some victim models are different from those in [24], the test images are also different. The ResNet-50 was chosen as the substitute model, same as in

[24]. For the experiments conducted on CIFAR-10, we adhered to the settings established in prior work [24]. Specifically, we performed attacks on VGG-19 [78], WRN-28-10 [90], ResNeXt-29 [91], DenseNet-BC [81], PyramidNet-272 [92], and GDAS [93], employing ResNet-18 [80] as the substitute model. The test images used are the entire test set of the CIFAR-10 dataset. We ran 50 iterations with a step size of  $1/255$  for all the iterative attacks.

In this paper, we set  $\gamma$  to a fixed constant, which is slightly different from the approach described in [24]. In [24],  $\gamma$  was set to be  $0.1/\|\Delta\mathbf{w}^*\|_2$ , dependent on the value of  $\|\Delta\mathbf{w}^*\|_2$ . However, for our experiments on ImageNet and CIFAR-10, we set  $\gamma$  to 0.1 with  $\lambda_{\varepsilon,\sigma} = 1$ , and  $\gamma$  to 0.5 with  $\lambda_{\varepsilon,\sigma} = 0.2$ , respectively. These hyper-parameters match the ones actually used in [24]. We set  $\lambda_{\varepsilon_e,\sigma_e}$  to be 1 and 0.01 for ImageNet and CIFAR-10, respectively. We used a learning rate of 0.05, an SGD optimizer with a momentum of 0.9 and a weight decay of 0.0005, a batch size of 1024 for ImageNet and 128 for CIFAR-10, and a number of epochs of 10. We fixed  $\sigma = 0.006$  and 0.012 for ImageNet and CIFAR-10, respectively. The  $\sigma_e$  was set to be 0.01 and 0.05 with and without fine-tuning, respectively. For simplicity and convenient hyper-parameter tuning, we directly set the values of  $\lambda_{\varepsilon,\sigma}$  and  $\lambda_{\varepsilon_e,\sigma_e}$ , instead of deriving them from the confidence levels  $\varepsilon$  and  $\varepsilon_e$  according Eq. (16). Considering Eq. (22), we used  $\alpha = 1$  for the posterior distribution over model parameters. For the posterior distribution over model input, we set  $\alpha = 100$  and 25 for ImageNet and CIFAR-10, respectively. Due to the negligible difference in success rates observed between using a diagonal matrix and a combination of diagonal and low-rank matrices as the covariance for the SWAG posterior, we opted for simplicity and consistently employed the diagonal matrix. We also set  $\beta = 0$  in the experiments for the same reason. For our current method with fine-tuning, we used SWAG to approximate the posterior over model parameters, while we used an isotropic Gaussian distribution for the posterior over model inputs. This choice was made because modeling the posterior over model inputs becomes challenging in situations where there is a high degree of randomness in the model parameters. We set  $M = 5$  and  $S = 5$  unless otherwise specified.

For compared competitors, we followed their official implementations for the most of them. It is worth noting that recent research has found that, for input augmentation methods, aggregating more input gradients can lead to the creation of more transferable adversarial examples [97]. To ensure a fair comparison focused on the quality of different augmentation strategies, we standardized the computational budget for this class of methods by strictly controlling the number of gradient computations to be 25 per iteration. Specifically, for our method, we used a combination of  $M = 5$  model samples and  $S = 5$  input noise samples per iteration, resulting in an average of  $5 \times 5 = 25$  gradients. For methods involving random input transformations (*i.e.*, vr-IGSM [49], DI<sup>2</sup>-FGSM [20], SI-FGSM [23], Admix [22], PGN [95], GRA [45], BSR [46], and SIA [44]), we execute the transformation-and-gradient-computation process 25 independent times per iteration and aggregate the resulting gradients by either averaging them or integrating them with specific approaches. For TI-FGSM

TABLE II

SUCCESS RATES OF TRANSFER-BASED ATTACKS ON IMAGENET USING RESNET-50 AS SUBSTITUTE ARCHITECTURE AND I-FGSM AS THE BACK-END ATTACK, UNDER THE  $\ell_\infty$  CONSTRAINT WITH  $\epsilon = 8/255$  IN THE UNTARGETED SETTING. "AVERAGE" WAS CALCULATED OVER ALL TEN VICTIM MODELS. DETAILED SETTINGS AND STRATEGIES FOR ENSURING FAIRNESS CAN BE FOUND IN SECTION IV-A. THE BEST RESULTS ARE MARKED IN BOLD

Method	VGG -19 [78]	Inception v3 [79]	ResNet -152 [80]	DenseNet -121 [81]	ConvNeXt -B [82]	ViT -B [83]	DeiT -B [84]	Swin -B [85]	BEiT -B [86]	Mixer -B [87]	Average
I-FGSM	44.52%	24.06%	36.14%	44.40%	15.16%	6.50%	5.68%	8.24%	7.24%	12.76%	20.47%
MI-FGSM (2018) [38]	54.98%	34.08%	46.76%	55.48%	20.22%	10.14%	8.82%	11.24%	10.94%	17.82%	27.05%
vr-IGSM (2018) [49]	72.80%	50.02%	68.24%	75.96%	29.48%	15.18%	13.74%	16.84%	18.76%	24.44%	38.55%
TI-FGSM (2019) [21]	88.10%	63.02%	87.46%	90.96%	48.22%	19.26%	20.54%	25.94%	24.12%	27.66%	49.53%
DI <sup>2</sup> -FGSM (2019) [20]	85.78%	63.74%	86.84%	90.50%	41.80%	20.56%	20.62%	23.54%	26.26%	28.02%	48.77%
SI-FGSM (2019) [23]	56.32%	36.40%	50.34%	60.70%	18.58%	9.16%	8.18%	9.86%	10.00%	16.64%	27.62%
NI-FGSM (2019) [23]	54.80%	34.52%	46.92%	55.30%	20.46%	10.34%	8.68%	10.86%	11.04%	17.70%	27.06%
ILA (2019) [14]	75.30%	53.54%	71.40%	76.88%	35.34%	15.02%	14.04%	19.10%	17.34%	23.00%	40.10%
SGM (2020) [16]	73.02%	47.40%	62.22%	70.72%	34.74%	17.22%	15.22%	19.60%	16.92%	26.44%	38.35%
LinBP (2020) [17]	77.84%	51.00%	63.70%	75.66%	24.58%	10.82%	8.42%	12.74%	13.38%	20.88%	35.90%
Admix (2021) [22]	75.54%	56.14%	72.32%	80.78%	26.84%	13.12%	11.20%	14.78%	17.24%	22.26%	39.02%
NAA (2022) [94]	79.00%	63.78%	72.94%	82.78%	27.44%	13.70%	12.48%	16.96%	17.14%	26.28%	41.25%
ILA++ (2022) [40]	78.22%	59.16%	75.46%	80.44%	41.30%	17.26%	16.96%	21.60%	21.06%	25.80%	43.73%
PGN (2023) [95]	72.52%	49.36%	67.52%	75.16%	27.78%	14.94%	13.92%	15.94%	18.46%	24.50%	38.01%
GRA (2023) [45]	75.70%	55.10%	72.34%	78.94%	31.80%	17.56%	17.16%	19.34%	21.58%	27.60%	41.71%
ILPD (2023) [96]	83.82%	71.26%	83.30%	87.14%	45.04%	24.30%	26.24%	30.48%	31.12%	34.60%	51.73%
SIA (2023) [44]	98.12%	86.50%	97.80%	98.98%	73.60%	33.36%	32.78%	40.08%	42.50%	42.24%	64.60%
BSR (2024) [46]	98.12%	85.92%	96.42%	98.60%	46.86%	20.06%	28.86%	37.34%	45.34%	40.80%	60.43%
Ours (w/o fine-tune)	92.40%	77.22%	91.42%	94.48%	45.44%	29.20%	26.32%	31.36%	36.74%	39.32%	56.39%
DRA (2022) [42]	87.10%	76.40%	86.60%	92.02%	39.12%	45.26%	43.94%	35.76%	55.98%	59.12%	62.13%
LGV (2022) [48]	96.50%	81.76%	95.00%	98.20%	44.84%	28.42%	23.28%	27.96%	36.60%	35.58%	56.81%
Ours BasicBayesian (2023) [24]	98.30%	89.92%	97.64%	99.38%	47.40%	24.78%	24.78%	33.04%	47.40%	39.30%	61.12%
Ours	<b>99.12%</b>	<b>95.64%</b>	<b>99.12%</b>	<b>99.66%</b>	<b>64.28%</b>	<b>58.20%</b>	<b>51.66%</b>	<b>49.26%</b>	<b>73.98%</b>	<b>61.04%</b>	<b>75.20%</b>

TABLE III

SUCCESS RATES OF TRANSFER-BASED ATTACKS ON CIFAR-10 USING RESNET-18 AS SUBSTITUTE ARCHITECTURE AND I-FGSM AS THE BACK-END ATTACK, UNDER THE  $\ell_\infty$  CONSTRAINT WITH  $\epsilon = 4/255$  IN THE UNTARGETED SETTING. "AVERAGE" WAS CALCULATED OVER ALL SIX VICTIM MODELS. DETAILED SETTINGS AND STRATEGIES FOR ENSURING FAIRNESS CAN BE FOUND IN SECTION IV-A. THE BEST RESULTS ARE MARKED IN BOLD

Method	VGG -19 [78]	WRN -28-10 [90]	ResNeXt -29 [91]	DenseNet -BC [81]	PyramidNet -272 [92]	GDAS [93]	Average
I-FGSM	38.47%	57.43%	58.61%	53.23%	12.82%	40.52%	43.51%
MI-FGSM (2018) [38]	42.42%	62.25%	63.84%	57.45%	14.15%	44.21%	47.39%
vr-IGSM (2018) [49]	59.85%	80.85%	79.84%	76.48%	20.26%	62.09%	63.23%
TI-FGSM (2019) [21]	50.65%	74.23%	74.77%	69.23%	18.74%	53.31%	56.82%
DI <sup>2</sup> -FGSM (2019) [20]	56.99%	71.50%	72.90%	66.51%	21.57%	52.22%	56.95%
SI-FGSM (2019) [23]	47.10%	63.83%	68.48%	62.26%	16.99%	43.49%	50.36%
NI-FGSM (2019) [23]	42.38%	62.23%	63.50%	57.05%	14.10%	44.23%	47.25%
ILA (2019) [14]	56.16%	76.37%	76.78%	72.78%	22.49%	56.26%	60.14%
SGM (2020) [16]	40.85%	59.61%	62.50%	55.48%	13.84%	45.17%	46.24%
LinBP (2020) [17]	58.51%	78.83%	81.52%	76.98%	27.44%	61.34%	64.10%
Admix (2021) [22]	56.93%	75.60%	78.60%	73.70%	21.65%	52.98%	59.91%
NAA (2022) [94]	39.40%	57.94%	60.26%	56.29%	13.74%	45.47%	45.52%
ILA++ (2022) [40]	60.31%	77.82%	78.43%	74.71%	26.25%	59.44%	62.83%
PGN (2023) [95]	58.29%	79.84%	78.52%	74.98%	19.98%	60.61%	62.04%
GRA (2023) [45]	61.30%	80.56%	79.20%	75.86%	21.15%	63.06%	63.52%
ILPD (2023) [96]	62.19%	81.87%	80.50%	78.81%	26.07%	63.98%	65.57%
SIA (2023) [44]	73.12%	88.27%	91.55%	86.66%	36.23%	77.83%	75.61%
BSR (2024) [46]	39.01%	59.64%	60.96%	55.22%	16.82%	42.14%	45.63%
Ours (w/o fine-tune)	67.44%	83.31%	84.21%	80.25%	27.83%	65.15%	68.03%
LGV (2022) [48]	83.22%	93.73%	94.22%	92.10%	44.11%	80.65%	81.34%
Ours BasicBayesian (2023) [24]	86.60%	94.84%	<b>95.73%</b>	93.78%	50.11%	84.97%	84.34%
Ours	<b>88.04%</b>	<b>95.00%</b>	95.62%	<b>93.93%</b>	<b>55.80%</b>	<b>86.07%</b>	<b>85.74%</b>

[21], we modified its standard implementation. Instead of convolving the input gradient with a kernel, we generated 25 randomly translated input images and averaged their respective gradients, a process we found to be more effective. In contrast, methods that do not inherently involve input augmentation (*i.e.*, MI-FGSM [38], NI-FGSM [23], ILA [14], SGM [16], LinBP [17], NAA [94], ILA++ [40], and ILPD [96]) were run in their standard, single-gradient configuration. Combining them with input augmentation would turn them

into a different, hybrid method, making it difficult to isolate the true benefit of their core innovations. When comparing our method with the methods innovated by substitute model fine-tuning (*i.e.*, LGV [48], DRA [42], and our BasicBayesian [24]), in order to fairly demonstrate the improvements achieved by our fine-tuning framework, for the competitors, we additionally introduced Gaussian noise into the model input multiple times and then averaged the obtained gradients. For LGV [48] and our BasicBayesian [24], during each iteration,

we sampled 5 models from the collected model set and sampled 5 Gaussian noise into the model input for each model. For DRA [42], we tested it on ImageNet using the ResNet-50 model provided by the authors, and we averaged 25 input gradients obtained by adding Gaussian noise into the model input per iteration. We chose the best standard deviation of Gaussian noise from {0.01, 0.03, 0.05, 0.07, 0.09} for these methods, *i.e.*, 0.01 for LGV and our BasicBayesian [24], and 0.05 for DRA. By doing so, the performance of the competitors will be better compared with their default setting. All experiments were performed on an NVIDIA V100 GPU.

### B. Comparison With State-of-the-Arts

We compared our method with recent state-of-the-arts in Tables II and III. A variety of methods were included in the comparison, including methods that adopt advanced optimizers (MI-FGSM [38] and NI-FGSM [23]), increase input diversity (vr-IGSM [49], TI-FGSM [21], DI<sup>2</sup>-FGSM [20], SI-FGSM [23], and Admix [22]), use advanced gradient computations (ILA [14], SGM [16], LinBP [17], NAA [94], ILA++ [40], PGN [95], and ILPD [96]), and employ substitute model fine-tuning (LGV [48] and DRA [42]). Detailed settings and strategies for ensuring fairness can be found in Section IV-A. The performance of all compared methods was evaluated in the task of attacking 10 victim models on ImageNet and 6 victim models on CIFAR-10.

It can be observed from the tables that the proposed method outperforms all compared methods. Among the methods without fine-tuning (*i.e.*, the upper half in Table II and Table III), our method (w/o fine-tune) achieves the best average success rate of 56.39% on ImageNet and 68.03% on CIFAR-10. When comparing with the substitute model fine-tuning methods (*i.e.*, the lower half in Table II and Table III), even after combining these competitors with Gaussian noise input augmentation and controlling the number of input gradients averaged per iteration, our method still achieves the best average success rate, *i.e.*, 75.20% on ImageNet and 85.74% on CIFAR-10. This shows that our fine-tuning framework can indeed boost the transferability of adversarial examples.

### C. Combination With Other Methods

We would also like to mention that it is possible to combine our method with other attack methods to further enhance the transferability. In Table IV, we report the attack success rate of our method in combination with MI-FGSM, DI<sup>2</sup>-FGSM, SGM, LinBP, ILA++, and ILPD. Note that the number of input gradient computations is controlled to be the same in the baseline methods and after combining with our method (as detailed in Section IV-A). It can be seen that the transferability to all victim models improves substantially, compared with both the original methods and our BasicBayesian [24]. The best performance is obtained when combining our current method with ILPD, improving over the original ILPD by 28.38% and achieving a record-high success rate of 80.11%.

### D. Attacking Defensive Models

It is also of interest to evaluate the transferability of adversarial examples to robust models, and we compared the performance of competitive methods in this setting in Table V. The victim models used in this study were collected from RobustBench [98]. All these models were trained using some sorts of advanced adversarial training [31], [99], [100], and they exhibit high robust accuracy against AutoAttack [101] on the official ImageNet validation set. These models included a robust ConvNeXt-B [102], a robust Swin-B [102], and a robust ViT-B-CvSt [103]. Following the setting in [24], two models from Bai et al.'s open-source repository [104], namely a robust ResNet-50-GELU and a robust DeiT-S, were also adopted as the robust victim. The tested robust ConvNeXt-B, robust Swin-B, and robust ViT-B-CvSt show higher prediction accuracy (*i.e.*, 55.82%, 56.16%, and 54.66%, respectively, against AutoAttack) than that of the robust ResNet-50-GELU and robust DeiT-S (35.51% and 35.50%, respectively).

We still used the ResNet-50 substitute model which was trained just as normal and not robust to adversarial examples at all. From Table V, we can observe that our newly proposed method improves the transferability of adversarial examples to these defensive models, compared with the basic Bayesian formulation in [24].

Since the objective of adversarial training is different from that of normal training, in the sense that the distribution of input is different, we suggest increasing  $\lambda_{\epsilon_e, \sigma_e}$  to achieve better alignment between the distributions and to further enhance transferability. When we increase  $\lambda_{\epsilon_e, \sigma_e}$  from 1 to 5 or even 10, we indeed obtain even better results in attacking robust models on ImageNet, as will be discussed in the ablation study.

In addition to adversarial training, we evaluated our method against a diverse set of defenses: a certified defense, namely randomized smoothing (RS) [105]), which was applied to a ResNet-50 victim model, three input transformation defenses, namely Bit-Red [33], JPEG [33], and R&P [34], all of which pre-process the adversarial input before classification, and a purification-based defense, namely NRP [106], which uses a dedicated model to purify the input. In this challenging scenario, we compare our method against the high-performing competitors from our main experiments, specifically those with over a 40% success rate shown in Table II). As shown in Table VI, our method consistently achieves the highest success rates, confirming its superior performance against a variety of defense strategies.

### E. Ablation Study

We conduct a series of ablation experiments to study the impact of different hyper-parameters.

1) *The Effect of  $\lambda_{\epsilon, \sigma}$  and  $\lambda_{\epsilon_e, \sigma_e}$* : When adopting fine-tuning, we have two main hyper-parameters that have an effect, namely  $\lambda_{\epsilon, \sigma}$  and  $\lambda_{\epsilon_e, \sigma_e}$ . As introduced in Section III-C, these hyperparameters govern the breadth of the model parameter and input distributions for fine-tuning. Consequently, larger values of  $\lambda_{\epsilon, \sigma}$  and  $\lambda_{\epsilon_e, \sigma_e}$  represent training across a wider region of their respective distributions. We conducted an

TABLE IV

SUCCESS RATES OF COMBINING OUR METHOD WITH OTHER STATE-OF-THE-ARTS ON IMAGENET USING RESNET-50 AS SUBSTITUTE ARCHITECTURE AND I-FGSM AS THE BACK-END ATTACK, UNDER THE  $\ell_\infty$  CONSTRAINT WITH  $\epsilon = 8/255$  IN THE UNTARGETED SETTING. "AVERAGE" WAS CALCULATED OVER ALL TEN VICTIM MODELS. DETAILED SETTINGS AND STRATEGIES FOR ENSURING FAIRNESS CAN BE FOUND IN SECTION IV-A. THE BEST RESULTS ARE MARKED IN BOLD

Method	VGG -19 [78]	Inception v3 [79]	ResNet -152 [80]	DenseNet -121 [81]	ConvNeXt -B [82]	ViT -B [83]	DeiT -B [84]	Swin -B [85]	BEiT -B [86]	Mixer -B [87]	Average
Our BasicBayesian [24]	98.30%	89.92%	97.64%	99.38%	47.40%	34.06%	24.78%	33.04%	47.40%	39.30%	61.12%
Ours (w/o fine-tune)	92.40%	77.22%	91.42%	94.48%	45.44%	29.20%	26.32%	31.36%	36.74%	39.32%	56.39%
Ours	99.12%	95.64%	99.12%	99.66%	64.28%	58.20%	51.66%	49.26%	73.98%	61.04%	75.20%
MI-FGSM (2018) [38]	54.98%	34.08%	46.76%	55.48%	20.22%	10.14%	8.82%	11.24%	10.94%	17.82%	27.05%
+Our BasicBayesian [24]	98.90%	92.38%	97.92%	<b>99.48%</b>	47.10%	39.12%	27.62%	35.26%	52.92%	45.36%	63.61%
+Ours (w/o fine-tune)	92.58%	78.64%	91.42%	94.72%	46.38%	29.98%	26.82%	31.98%	36.98%	39.68%	56.92%
+Ours	<b>99.02%</b>	<b>95.52%</b>	<b>98.82%</b>	<b>99.48%</b>	<b>62.46%</b>	<b>59.62%</b>	<b>53.80%</b>	<b>50.80%</b>	<b>74.40%</b>	<b>65.44%</b>	<b>75.94%</b>
DI <sup>2</sup> -FGSM (2019) [20]	85.78%	63.74%	86.84%	90.50%	41.80%	20.56%	20.62%	23.54%	26.26%	28.02%	48.77%
+Our BasicBayesian [24]	<b>99.40%</b>	95.98%	99.36%	<b>99.82%</b>	56.58%	53.38%	39.38%	45.96%	69.36%	54.80%	71.40%
+Ours (w/o fine-tune)	96.62%	87.70%	96.90%	98.16%	55.62%	40.24%	38.34%	41.82%	50.40%	49.02%	65.48%
+Ours	99.32%	<b>96.58%</b>	<b>99.42%</b>	99.70%	<b>63.40%</b>	<b>66.66%</b>	<b>59.08%</b>	<b>53.54%</b>	<b>81.62%</b>	<b>68.02%</b>	<b>78.73%</b>
SGM (2020) [16]	73.02%	47.40%	62.22%	70.72%	34.74%	17.22%	15.22%	19.60%	16.92%	26.44%	38.35%
+Our BasicBayesian [24]	98.72%	91.30%	97.76%	99.26%	49.82%	36.10%	26.46%	34.28%	49.92%	43.06%	62.67%
+Ours (w/o fine-tune)	94.40%	81.66%	92.34%	95.74%	50.50%	33.74%	32.30%	35.54%	41.34%	47.50%	60.51%
+Ours	<b>99.16%</b>	<b>95.76%</b>	<b>98.92%</b>	<b>99.64%</b>	<b>63.12%</b>	<b>58.28%</b>	<b>51.96%</b>	<b>49.96%</b>	<b>74.10%</b>	<b>64.04%</b>	<b>75.49%</b>
LinBP (2020) [17]	77.84%	51.00%	63.70%	75.66%	24.58%	10.82%	8.42%	12.74%	13.38%	20.88%	35.90%
+Our BasicBayesian [24]	98.66%	92.38%	97.98%	99.48%	48.62%	36.26%	24.42%	33.68%	51.02%	40.90%	62.34%
+Ours (w/o fine-tune)	94.30%	80.94%	92.68%	95.76%	48.56%	31.58%	28.50%	34.14%	40.52%	40.58%	58.76%
+Ours	<b>99.16%</b>	<b>95.72%</b>	<b>99.14%</b>	<b>99.68%</b>	<b>62.94%</b>	<b>60.84%</b>	<b>51.74%</b>	<b>49.88%</b>	<b>76.10%</b>	<b>61.78%</b>	<b>75.70%</b>
ILAA (2022) [40]	78.22%	59.16%	75.46%	80.44%	41.30%	17.26%	16.96%	21.60%	21.06%	25.80%	43.73%
+Our BasicBayesian [24]	<b>99.02%</b>	93.80%	<b>98.40%</b>	<b>99.52%</b>	57.84%	50.58%	45.54%	50.80%	62.74%	56.30%	71.45%
+Ours (w/o fine-tune)	92.04%	80.60%	91.62%	93.72%	48.46%	32.20%	31.38%	33.80%	41.06%	39.78%	58.47%
+Ours	98.58%	<b>94.30%</b>	98.26%	99.18%	<b>63.06%</b>	<b>63.28%</b>	<b>58.56%</b>	<b>57.12%</b>	<b>74.96%</b>	<b>67.40%</b>	<b>77.47%</b>
ILPD (2023) [96]	83.82%	71.26%	83.30%	87.14%	45.04%	24.30%	26.24%	30.48%	31.12%	34.60%	51.73%
+Our BasicBayesian [24]	<b>99.18%</b>	96.04%	98.86%	99.54%	60.90%	54.02%	45.66%	52.50%	68.46%	59.16%	73.43%
+Ours (w/o fine-tune)	95.58%	87.64%	95.88%	97.02%	61.00%	41.34%	40.72%	44.48%	50.34%	51.04%	66.50%
+Ours	98.98%	<b>96.06%</b>	<b>98.88%</b>	<b>99.56%</b>	<b>66.12%</b>	<b>67.46%</b>	<b>62.94%</b>	<b>60.18%</b>	<b>80.50%</b>	<b>70.38%</b>	<b>80.11%</b>
SIA (2023) [44]	98.12%	86.50%	97.80%	98.98%	60.04%	27.62%	32.78%	40.08%	42.50%	42.24%	62.67%
+Our BasicBayesian [24]	99.72%	96.58%	99.64%	<b>99.90%</b>	62.46%	34.32%	37.78%	46.70%	60.40%	50.54%	68.80%
+Ours (w/o fine-tune)	98.56%	91.68%	98.36%	99.22%	61.92%	32.56%	38.90%	45.66%	51.98%	49.80%	66.86%
+Ours	<b>99.76%</b>	<b>96.88%</b>	<b>99.78%</b>	99.82%	<b>64.54%</b>	<b>59.56%</b>	<b>56.18%</b>	<b>54.36%</b>	<b>81.46%</b>	<b>68.96%</b>	<b>78.13%</b>
BSR (2024) [46]	98.12%	85.92%	96.42%	98.60%	46.86%	26.06%	28.86%	37.34%	45.34%	40.80%	60.43%
+Our BasicBayesian [24]	99.34%	94.50%	98.42%	99.62%	40.30%	30.96%	32.20%	39.76%	57.18%	49.12%	64.14%
+Ours (w/o fine-tune)	99.72%	96.98%	<b>99.32%</b>	<b>99.84%</b>	49.66%	38.58%	39.08%	46.42%	66.98%	54.64%	69.12%
+Ours	<b>99.78%</b>	<b>97.40%</b>	99.22%	99.66%	<b>65.40%</b>	<b>58.98%</b>	<b>55.40%</b>	<b>51.84%</b>	<b>79.84%</b>	<b>64.66%</b>	<b>77.22%</b>

TABLE V

SUCCESS RATES OF ATTACKING ADVERSARIALLY TRAINED MODELS ON IMAGENET USING RESNET-50 AS SUBSTITUTE ARCHITECTURE AND I-FGSM AS THE BACK-END ATTACK, UNDER THE  $\ell_\infty$  CONSTRAINT WITH  $\epsilon = 8/255$  IN THE UNTARGETED SETTING. "AVERAGE" WAS CALCULATED OVER ALL FIVE VICTIM MODELS. THE BEST RESULTS ARE MARKED IN BOLD

Method	ConvNeXt -B [102]	Swin -B [102]	ViT -B [103]	ResNet -50 [104]	DeiT -S [104]	Average
Ours BasicBayesian [24]	6.80%	7.08%	7.22%	16.86%	17.64%	11.12%
Ours (w/o fine-tune)	6.98%	7.22%	7.32%	17.40%	18.10%	11.40%
Ours	<b>8.14%</b>	<b>8.04%</b>	<b>8.74%</b>	<b>21.00%</b>	<b>20.74%</b>	<b>13.33%</b>

empirical study to demonstrate how the performance of our method varies with the values of these two hyper-parameters on ImageNet. We varied  $\lambda_{\epsilon,\sigma}$  from the set  $\{0, 0.1, 0.2, 0.5, 1, 2\}$  and varied  $\lambda_{\epsilon_e,\sigma_e}$  from the set  $\{0, 0.5, 1, 5, 10\}$  for attacking models which are the same as in Table II and Table V. The average success rates of attacking these normally trained models and robust models are given in Table VII. To achieve the best performance for each  $\lambda_{\epsilon,\sigma}$  and  $\lambda_{\epsilon_e,\sigma_e}$ , we tuned the other hyper-parameters using 500 randomly selected images

from the validation set. These images show no overlap with the 5000 test images.

From the table, it can be observed that increasing the values of  $\lambda_{\epsilon,\sigma}$  and  $\lambda_{\epsilon_e,\sigma_e}$  from 0 to 0.5 enhances the transferability of adversarial examples in attacking normally trained models. However, excessively large values of these hyper-parameters can lead to inferior performance. The best performance of attacking normally trained models is achieved by setting  $\lambda_{\epsilon,\sigma} = 0.5$  and  $\lambda_{\epsilon_e,\sigma_e} = 1$ , which results in an average success rate of 75.60%. However, despite achieving optimal performance, we do not adopt this setting in our main ImageNet experiments (presented in Tables I, II, IV, V, and VI). Instead, we use  $\lambda_{\epsilon,\sigma} = 1$  and  $\lambda_{\epsilon_e,\sigma_e} = 1$  to ensure a direct and fair comparison with our prior work, BasicBayesian [24]. Considering the performance on attacking robust models, it can be observed that the average success rate peaks a larger value of  $\lambda_{\epsilon_e,\sigma_e}$ . This is partially because that reducing the prediction loss of the perturbed inputs during fine-tuning in Eq. (11) resembles performing adversarial training, and larger  $\lambda_{\epsilon_e,\sigma_e}$  implies making the substitute model robust in a larger neighborhood of benign inputs. A recent related method, namely DRA [42],

TABLE VI

ATTACK SUCCESS RATES OF ATTACKING AGAINST DEFENSE METHODS UNDER THE  $\ell_\infty$  CONSTRAINT WITH  $\epsilon = 8/255$  ON IMAGENET IN THE UNTARGETED SETTING. THE BEST RESULTS ARE MARKED IN BOLD

Method	RS	Bit-Red	JPEG	R&P	NRP
I-FGSM	12.24%	15.68%	18.03%	20.31%	13.38%
TI-IGSM [21]	14.72%	39.24%	37.96%	49.19%	21.68%
DI <sup>2</sup> -FGSM [20]	15.80%	40.05%	41.25%	48.13%	20.14%
ILA [14]	13.86%	32.55%	32.60%	40.10%	16.02%
NAA [94]	17.88%	35.79%	37.85%	39.34%	13.38%
ILA++ [40]	15.30%	36.67%	36.11%	40.60%	18.12%
GRA [45]	19.40%	38.54%	40.98%	38.40%	15.23%
ILPD [96]	20.26%	47.35%	47.74%	51.23%	23.19%
SIA [44]	17.06%	55.32%	53.28%	62.03%	22.40%
BSR [46]	15.44%	51.90%	52.43%	56.11%	20.92%
Ours (w/o fine-tune)	25.68%	52.85%	54.80%	55.90%	22.75%
DRA [42]	60.94%	60.24%	67.42%	65.70%	10.58%
LGV [48]	22.28%	52.92%	55.15%	58.27%	16.21%
Ours-BasicBayesian [24]	26.76%	57.57%	60.39%	63.40%	22.92%
Ours	<b>85.04%</b>	<b>72.79%</b>	<b>74.07%</b>	<b>74.24%</b>	<b>26.01%</b>

TABLE VII

AVERAGE SUCCESS RATES OF ATTACKING NORMALLY TRAINED VICTIM MODELS AND ADVERSARIALY TRAINED VICTIM MODELS ON IMAGENET, WITH DIFFERENT CHOICES OF  $\lambda_{\epsilon, \sigma}$  AND  $\lambda_{\epsilon_e, \sigma_e}$ . THE SUBSTITUTE MODEL IS RESNET-50. THE NORMALLY TRAINED VICTIM MODELS ARE THE SAME AS THOSE IN TABLE II, AND THE ADVERSARIALY TRAINED MODELS ARE THE SAME AS THOSE IN TABLE V. THE SYMBOL “-” SIGNIFIES THAT THE TRAINING DOES NOT CONVERGE. THE BEST RESULTS ARE MARKED IN BOLD

$\lambda_{\epsilon, \sigma}$	$\lambda_{\epsilon_e, \sigma_e}$				
	0	0.5	1	5	10
0	47.66%/10.86%	71.22%/12.77%	74.33%/14.11%	66.00%/20.47%	47.43%/ <b>22.72%</b>
0.1	50.74%/10.93%	73.19%/12.83%	74.08%/14.04%	66.14%/20.38%	48.23%/22.70%
0.2	51.97%/11.04%	73.62%/12.80%	74.19%/13.84%	66.65%/19.94%	49.57%/22.22%
0.5	54.24%/11.14%	73.78%/12.64%	<b>75.60%</b> /13.56%	67.51%/19.09%	49.09%/21.70%
1	55.58%/11.26%	72.81%/12.54%	75.20%/13.33%	65.84%/18.48%	-
2	51.67%/11.48%	66.02%/12.66%	70.78%/13.36%	-	-

also suggests that performing regularized fine-tuning before attacking is beneficial to attacking defensive models, and it achieves an average success rate of 62.13% and 17.42% in attacking normally trained models and adversarialy trained models, respectively. By setting  $\lambda_{\epsilon_e, \sigma_e} = 5$ , our method has an average success rate ranging from 65.84% to 67.51% in attacking normally trained models and ranging from 18.48% to 20.47% in attacking adversarialy trained models, which outperforms DRA considerably.

2) *The Effect of M and S*: In Figure 3, we conducted experiments to evaluate the transferability of adversarial examples crafted using different choices of the number of model parameters and input noise sampled at each attack iteration, *i.e.*,  $M$  and  $S$ , in the cases of with and without fine-tuning. The average success rates were obtained by attacking the same victim models as in Table II. The results apparently show that sampling more substitute models and more input noise can indeed enhance the transferability of adversarial examples, just as expected. Note that our fine-tuning framework which considers the joint diversification of both the model input and model parameters can outperform our previous work [24] even when setting  $M = 1$  and  $S = 1$ . Our method with this configuration achieves a higher attack success rate (72.63%) compared to the previous approach that used 5 model parameters and

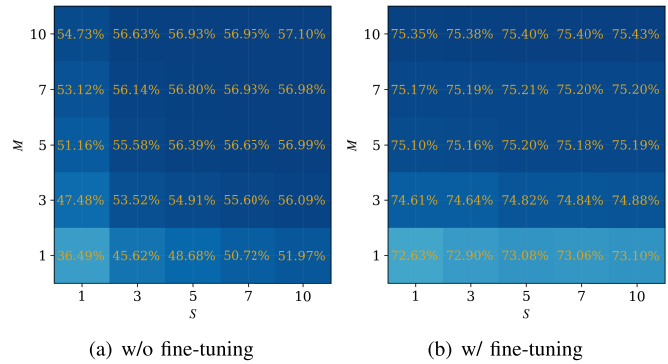


Fig. 3. Average success rates of attacking 10 victim models on ImageNet with varying  $M$  and  $S$ . Darker cubic indicates better performance. Best viewed in color.

5 input noise samples at each iteration (61.12%), indicating that our method can generate more transferable adversarial examples with lower computational complexity compared to our previous work [24].

## V. CONCLUSION

In this paper, we aim at improving the transferability of adversarial examples. We have developed a Bayesian formulation for performing attacks, which can be equivalently regarded as generating adversarial examples on a set of infinitely many substitute models with input augmentations. We also advocated possible fine-tuning and advanced posterior approximations for improving the Bayesian model. Extensive experiments have been conducted on ImageNet and CIFAR-10 to demonstrate the effectiveness of the proposed method in generating transferable adversarial examples. It has been shown that our method outperforms recent state-of-the-arts by large margins in attacking more than 10 DNNs, including convolutional networks and vision transformers and MLPs, as well as in attacking defensive models. We have also showcased the compatibility of our method with existing transfer-based attack methods, leading to even more powerful adversarial transferability.

## REFERENCES


- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with deep convolutional neural networks,” *Commun. ACM*, vol. 60, no. 6, pp. 84–90, May 2017.
- [2] R. Collobert and J. Weston, “A unified architecture for natural language processing: Deep neural networks with multitask learning,” in *Proc. 25th Int. Conf. Mach. Learn.*, Helsinki, Finland, 2008, pp. 160–167.
- [3] O. Abdel-Hamid, A.-R. Mohamed, H. Jiang, L. Deng, G. Penn, and D. Yu, “Convolutional neural networks for speech recognition,” *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 22, no. 10, pp. 1533–1545, Oct. 2014.
- [4] C. Szegedy et al., “Intriguing properties of neural networks,” in *Proc. ICLR*, 2014, pp. 1–10.
- [5] Y. Ran, W. Wang, M. Li, L.-C. Li, Y.-G. Wang, and J. Li, “Cross-shaped adversarial patch attack,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 34, no. 4, pp. 2289–2303, Apr. 2024.
- [6] A. Tao, Y. Duan, Y. Wang, J. Lu, and J. Zhou, “Dynamics-aware adversarial attack of adaptive neural networks,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 34, no. 7, pp. 5505–5518, Jul. 2024.
- [7] C. Yang, Y. Liu, D. Li, and T. Jiang, “Exploring vulnerabilities of no-reference image quality assessment models: A query-based black-box method,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 34, no. 12, pp. 12715–12729, Dec. 2024.

- [8] T. Chen and Z. Ma, "Toward robust neural image compression: Adversarial attack and model finetuning," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 33, no. 12, pp. 7842–7856, Dec. 2023.
- [9] A. Zou, Z. Wang, N. Carlini, M. Nasr, J. Z. Kolter, and M. Fredrikson, "Universal and transferable adversarial attacks on aligned language models," 2023, [arXiv:2307.15043](https://arxiv.org/abs/2307.15043).
- [10] N. Papernot, P. McDaniel, and I. Goodfellow, "Transferability in machine learning: From phenomena to black-box attacks using adversarial samples," 2016, [arXiv:1605.07277](https://arxiv.org/abs/1605.07277).
- [11] Y. Liu, X. Chen, C. Liu, and D. Song, "Delving into transferable adversarial examples and black-box attacks," in *Proc. ICLR*, 2016, pp. 1–14.
- [12] Y. Li, S. Bai, Y. Zhou, C. Xie, Z. Zhang, and A. Yuille, "Learning transferable adversarial examples via ghost networks," in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, 2020, pp. 11458–11465.
- [13] Y. Xiong, J. Lin, M. Zhang, J. E. Hopcroft, and K. He, "Stochastic variance reduced ensemble adversarial attack for boosting the adversarial transferability," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 14963–14972.
- [14] Q. Huang, I. Katsman, Z. Gu, H. He, S. Belongie, and S.-N. Lim, "Enhancing adversarial example transferability with an intermediate level attack," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 4732–4741.
- [15] Q. Li, Y. Guo, and H. Chen, "Yet another intermediate-level attack," in *Proc. ECCV*, 2020, pp. 241–257.
- [16] D. Wu, Y. Wang, S.-T. Xia, J. Bailey, and X. Ma, "Rethinking the security of skip connections in resnet-like neural networks," in *Proc. ICLR*, 2020, pp. 1–15.
- [17] Y. Guo, Q. Li, and H. Chen, "Backpropagating linearly improves transferability of adversarial examples," in *Proc. NeurIPS*, 2020, pp. 85–95.
- [18] O. Russakovsky et al., "ImageNet large scale visual recognition challenge," in *Proc. IJCV*, vol. 115, 2015, pp. 211–252.
- [19] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," Dept. Comput. Sci., Univ. Toronto, Toronto, ON, Canada, 2009.
- [20] C. Xie et al., "Improving transferability of adversarial examples with input diversity," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2725–2734.
- [21] Y. Dong, T. Pang, H. Su, and J. Zhu, "Evading defenses to transferable adversarial examples by translation-invariant attacks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 4307–4316.
- [22] X. Wang, X. He, J. Wang, and K. He, "Admix: Enhancing the transferability of adversarial attacks," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 16158–16167.
- [23] J. Lin, C. Song, K. He, L. Wang, and J. E. Hopcroft, "Nesterov accelerated gradient and scale invariance for adversarial attacks," 2019, [arXiv:1908.06281](https://arxiv.org/abs/1908.06281).
- [24] Q. Li, Y. Guo, W. Zuo, and H. Chen, "Making substitute models more Bayesian can enhance transferability of adversarial examples," in *Proc. Int. Conf. Learn. Represent.*, 2023, pp. 1–16.
- [25] T. B. Brown, D. Mané, A. Roy, M. Abadi, and J. Gilmer, "Adversarial patch," 2017, [arXiv:1712.09665](https://arxiv.org/abs/1712.09665).
- [26] I. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *Proc. ICLR*, 2014, pp. 1–11.
- [27] L. Tang, Y. Lv, D. Ye, Y. He, Z. Liu, and C. Xie, "Toward a universal, transferable, and robust adversarial perturbation framework against deep hashing-based facial image retrieval," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 35, no. 8, pp. 7423–7438, Aug. 2025.
- [28] S. Komkov and A. Petiushko, "AdvHat: Real-world adversarial attack on ArcFace face ID system," in *Proc. 25th Int. Conf. Pattern Recognit. (ICPR)*, Jan. 2021, pp. 819–826.
- [29] Y. Dong et al., "Efficient decision-based black-box adversarial attacks on face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 7714–7722.
- [30] K. Xu et al., "Adversarial T-shirt! Evading person detectors in a physical world," in *Proc. Eur. Conf. Comput. Vis.*, 2019, pp. 665–681.
- [31] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *Proc. ICLR*, 2017, pp. 1–23.
- [32] H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan, "Theoretically principled trade-off between robustness and accuracy," in *Proc. Int. Conf. Mach. Learn.*, Jordan, 2019, pp. 7472–7482.
- [33] C. Guo, M. Rana, M. Cisse, and L. van der Maaten, "Countering adversarial images using input transformations," 2017, [arXiv:1711.00117](https://arxiv.org/abs/1711.00117).
- [34] C. Xie, J. Wang, Z. Zhang, Z. Ren, and A. Yuille, "Mitigating adversarial effects through randomization," 2017, [arXiv:1711.01991](https://arxiv.org/abs/1711.01991).
- [35] X. Fu, L. Ma, and L. Zhang, "Remove to regenerate: Boosting adversarial generalization with attack invariance," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 35, no. 3, pp. 1999–2012, Mar. 2025.
- [36] Y. Gan, C. Wu, D. Ouyang, S. Tang, M. Ye, and T. Xiang, "LESEP: Boosting adversarial transferability via latent encoding and semantic embedding perturbations," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 35, no. 4, pp. 3368–3382, Apr. 2025.
- [37] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial machine learning at scale," in *Proc. ICLR*, 2016, pp. 1–17.
- [38] Y. Dong et al., "Boosting adversarial attacks with momentum," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 9185–9193.
- [39] Y. Huang and A. Wai-Kin Kong, "Transferable adversarial attack based on integrated gradients," 2022, [arXiv:2205.13152](https://arxiv.org/abs/2205.13152).
- [40] Y. Guo, Q. Li, W. Zuo, and H. Chen, "An intermediate-level attack framework on the basis of linear regression," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 3, pp. 2726–2735, Mar. 2023.
- [41] J. M. Springer, M. Mitchell, and G. T. Kenyon, "A little robustness goes a long way: Leveraging robust features for targeted transfer attacks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, pp. 9759–9773.
- [42] Y. Zhu et al., "Toward understanding and boosting adversarial transferability from a distribution perspective," *IEEE Trans. Image Process.*, vol. 31, pp. 6487–6501, 2022.
- [43] A. Wu, S. Nowozin, E. Meeds, R. E. Turner, J. M. Hernandez-Lobato, and A. L. Gaunt, "Deterministic variational inference for robust Bayesian neural networks," 2018, [arXiv:1810.03958](https://arxiv.org/abs/1810.03958).
- [44] X. Wang, Z. Zhang, and J. Zhang, "Structure invariant transformation for better adversarial transferability," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2023, pp. 4584–4596.
- [45] H. Zhu, Y. Ren, X. Sui, L. Yang, and W. Jiang, "Boosting adversarial transferability via gradient relevance attack," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2023, pp. 4741–4750.
- [46] K. Wang, X. He, W. Wang, and X. Wang, "Boosting adversarial transferability by block shuffle and rotation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 24336–24346.
- [47] M. Gubri, M. Cordy, M. Papadakis, Y. L. Traon, and K. Sen, "Efficient and transferable adversarial examples from Bayesian neural networks," in *Proc. Uncertainty Artif. Intell.*, 2020, pp. 738–748.
- [48] M. Gubri, M. Cordy, M. Papadakis, Y. Le Traon, and K. Sen, "LGV: Boosting adversarial example transferability from large geometric vicinity," 2022, [arXiv:2207.13129](https://arxiv.org/abs/2207.13129).
- [49] L. Wu, Z. Zhu, and C. Tai, "Understanding and enhancing the transferability of adversarial examples," 2018, [arXiv:1802.09707](https://arxiv.org/abs/1802.09707).
- [50] R. Johnson and T. Zhang, "Accelerating stochastic gradient descent using predictive variance reduction," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 26, 2013, pp. 315–323.
- [51] A. Graves, "Practical variational inference for neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 24, 2011, pp. 2348–2356.
- [52] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight uncertainty in neural network," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1613–1622.
- [53] D. P. Kingma, T. Salimans, and M. Welling, "Variational dropout and the local reparameterization trick," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 28, 2015, pp. 2575–2583.
- [54] M. E. Khan, D. Nielsen, V. Tangkaratt, L. Wu, Y. Gal, and A. Srivastava, "Fast and scalable Bayesian deep learning by weight-perturbation in Adam," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 2611–2620.
- [55] G. Zhang, S. Sun, D. Duvenaud, and R. Grosse, "Noisy natural gradient as variational inference," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 5852–5861.
- [56] J. Xu, B. Liu, and Y. Xiao, "A variational inference method for few-shot learning," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 33, no. 1, pp. 269–282, Jan. 2023.
- [57] Z. Gao et al., "Multi-domain adversarial variational Bayesian inference for domain generalization," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 35, no. 4, pp. 3081–3093, Apr. 2025.
- [58] K. Osawa et al., "Practical deep learning with Bayesian principles," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 4287–4299.
- [59] M. Dusenberry et al., "Efficient and scalable Bayesian neural nets with rank-1 factors," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 2782–2792.

- [60] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. 33rd Int. Conf. Mach. Learn.*, vol. 48, Jun. 2016, pp. 1050–1059.
- [61] A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5574–5584.
- [62] Y. Gal, J. Hron, and A. Kendall, "Concrete dropout," in *Proc. Adv. neural Inf. Process. Syst.*, vol. 30, 2017, pp. 3581–3590.
- [63] J. Kirkpatrick et al., "Overcoming catastrophic forgetting in neural networks," *Proc. Nat. Acad. Sci. USA*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [64] H. Ritter, A. Botev, and D. Barber, "A scalable Laplace approximation for neural networks," in *Proc. 6th Int. Conf. Learn. Represent.*, 2018, pp. 1–15.
- [65] D. X. Li, "On default correlation: A copula function approach," *J. Fixed Income*, vol. 9, no. 4, pp. 43–54, 2000.
- [66] S. Mandt, M. D. Hoffman, and D. M. Blei, "Stochastic gradient descent as approximate Bayesian inference," 2017, *arXiv:1704.04289*.
- [67] W. J. Maddox, P. Izmailov, T. Garipov, D. Vetrov, and A. G. Wilson, "A simple baseline for Bayesian uncertainty in deep learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 13132–13143.
- [68] W. J. Maddox, S. Tang, P. G. Moreno, A. G. Wilson, and A. Damianou, "Fast adaptation with linearized neural networks," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2021, pp. 2737–2745.
- [69] A. G. Wilson and P. Izmailov, "Bayesian deep learning and a probabilistic perspective of generalization," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 4697–4708.
- [70] L. Cardelli, M. Kwiatkowska, L. Laurenti, N. Paoletti, A. Patane, and M. Wicker, "Statistical guarantees for the robustness of Bayesian neural networks," 2019, *arXiv:1903.01980*.
- [71] M. Wicker, L. Laurenti, A. Patané, and M. Kwiatkowska, "Probabilistic safety for Bayesian neural networks," in *Proc. Conf. Uncertainty Artif. Intell.*, 2020, pp. 1198–1207.
- [72] X. Liu, Y. Li, C. Wu, and C.-J. Hsieh, "Adv-BNN: Improved adversarial defense through robust Bayesian neural network," 2018, *arXiv:1810.01279*.
- [73] M. Yuan, M. Wicker, and L. Laurenti, "Gradient-free adversarial attacks for Bayesian neural networks," 2020, *arXiv:2012.12640*.
- [74] G. Carbone, M. Wicker, L. Laurenti, A. Patané, L. Bortolussi, and G. Sanguinetti, "Robustness of Bayesian neural networks to gradient-based attacks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 15602–15613.
- [75] Y. Li, J. Bradshaw, and Y. Sharma, "Are generative classifiers more robust to adversarial attacks," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 3804–3814.
- [76] Z. Wang, Y. Guo, and W. Zuo, "Deepfake forensics via an adversarial game," *IEEE Trans. Image Process.*, vol. 31, pp. 3541–3552, 2022.
- [77] P. Foret, A. Kleiner, H. Mobahi, and B. Neyshabur, "Sharpness-aware minimization for efficiently improving generalization," 2020, *arXiv:2010.01412*.
- [78] Simonyan, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, 2015, pp. 1–14.
- [79] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2818–2826.
- [80] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [81] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2261–2269.
- [82] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A ConvNet for the 2020s," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 11976–11986.
- [83] A. Dosovitskiy et al., "An image is worth 16×16 words: Transformers for image recognition at scale," 2020, *arXiv:2010.11929*.
- [84] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, "Training data-efficient image transformers & distillation through attention," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 10347–10357.
- [85] Z. Liu et al., "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 10012–10022.
- [86] H. Bao, D. Li, and F. Wei, "BEiT: BERT pre-training of image transformers," in *Proc. Int. Conf. Learn. Represent.*, 2021, pp. 1–18.
- [87] I. Tolstikhin et al., "MLP-mixer: An all-MLP architecture for vision," in *Proc. 35th Conf. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 24261–24272.
- [88] P. Izmailov, D. Podoprikin, T. Garipov, D. Vetrov, and A. Gordon Wilson, "Averaging weights leads to wider optima and better generalization," 2018, *arXiv:1803.05407*.
- [89] R. Wightman. (2019). *Pytorch Image Models*. [Online]. Available: <https://github.com/rwightman/pytorch-image-models>
- [90] S. Zagoruyko and N. Komodakis, "Wide residual networks," in *Proc. Brit. Mach. Vis. Conf.*, 2016, pp. 1–12.
- [91] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 5987–5995.
- [92] D. Han, J. Kim, and J. Kim, "Deep pyramidal residual networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 6307–6315.
- [93] X. Dong and Y. Yang, "Searching for a robust neural architecture in four GPU hours," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 1761–1770.
- [94] J. Zhang et al., "Improving adversarial transferability via neuron attribution-based attacks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 14973–14982.
- [95] Z. Ge, F. Shang, H. Liu, Y. Liu, and X. Wang, "Boosting adversarial transferability by achieving flat local maxima," in *Proc. Adv. Neural Inf. Process. Syst.*, 2023, pp. 70141–70161.
- [96] Q. Li, Y. Guo, W. Zuo, and H. Chen, "Improving adversarial transferability via intermediate-level perturbation decay," in *Proc. 37th Conf. Neural Inf. Process. Syst.*, 2023, pp. 32900–32912.
- [97] Z. Zhao, H. Zhang, R. Li, R. Sicre, L. Amsaleg, and M. Backes, "Towards good practices in evaluating transfer adversarial attacks," 2022, *arXiv:2211.09565*.
- [98] F. Croce et al., "RobustBench: A standardized adversarial robustness benchmark," 2020, *arXiv:2010.09670*.
- [99] E. Wong, L. Rice, and J. Z. Kolter, "Fast is better than free: Revisiting adversarial training," 2020, *arXiv:2001.03994*.
- [100] C. Xie, M. Tan, B. Gong, J. Wang, A. L. Yuille, and Q. V. Le, "Adversarial examples improve image recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 819–828.
- [101] F. Croce and M. Hein, "Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks," in *Proc. ICML*, 2020, pp. 2206–2216.
- [102] C. Liu et al., "A comprehensive study on robustness of image classification models: Benchmarking and rethinking," 2023, *arXiv:2302.14301*.
- [103] N. D. Singh, F. Croce, and M. Hein, "Revisiting adversarial training for ImageNet: Architectures, training and generalization across threat models," 2023, *arXiv:2303.01870*.
- [104] Y. Bai, J. Mei, A. Yuille, and C. Xie, "Are transformers more robust than CNNs," in *Proc. 35th Conf. Neural Inf. Process. Syst.*, 2021, pp. 26831–26843.
- [105] J. M. Cohen, E. Rosenfeld, and J. Z. Kolter, "Certified adversarial robustness via randomized smoothing," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 1310–1320.
- [106] M. Naseer, S. Khan, M. Hayat, F. S. Khan, and F. Porikli, "A self-supervised approach for adversarial robustness," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 259–268.



**Qizhang Li** received the B.E. and M.E. degrees from Harbin Engineering University, Harbin, China, in 2018 and 2021, respectively. He is currently pursuing the Ph.D. degree in computer science and technology with the School of Computer Science and Technology, Harbin Institute of Technology, Harbin. His current research interests include computer vision, natural language processing, and machine learning.



**Yiwen Guo** received the B.E. degree from Wuhan University in 2011 and the Ph.D. degree from Tsinghua University in 2016. He was a Research Scientist with the ByteDance AI Laboratory, Beijing. Prior to this, he was a Staff Research Scientist with Intel Labs China. His current research interests include computer vision, pattern recognition, and machine learning.



**Wangmeng Zuo** (Senior Member, IEEE) received the Ph.D. degree in computer application technology from Harbin Institute of Technology, Harbin, China, in 2007. He is currently a Professor with the School of Computer Science and Technology, Harbin Institute of Technology. He has published over 100 papers in top tier academic journals and conferences. His current research interests include image enhancement and restoration, image and face editing, object detection, visual tracking, and image classification. He has served as a Tutorial Organizer for ECCV 2016 and an Associate Editor for IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, *The Visual Computer*, and *Journal of Electronic Imaging*.



**Xiaochen Yang** received the Ph.D. degree in statistical science from University College London, London, in 2020. She is currently a Lecturer with the School of Mathematics and Statistics, University of Glasgow, Glasgow, U.K. Her research interests include statistical classification, metric learning, and medical image analysis. She is an Associate Editor of *Neurocomputing* and IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY.



**Hao Chen** (Fellow, IEEE) received the B.S. and M.S. degrees from Southeast University and the Ph.D. degree from the Computer Science Division, University of California at Berkeley, Berkeley, CA, USA. He is currently a Professor with the Department of Computer Science, University of California at Davis, Davis, CA. His research interests are computer security, machine learning, and mobile computing.