Efficient Defense Against Model Stealing Attacks on Convolutional Neural Networks

Kacem Khaled^{*}, Mouna Dhaouadi[†], Felipe Gohring de Magalhães^{*} and Gabriela Nicolescu^{*}

*Department of Computer Engineering and Software Engineering, Polytechnique Montreal, Canada

[†]Department of Computer Science and Operations Research, University of Montreal, Canada

Email: kacem.khaled@polymtl.ca

Abstract-Model stealing attacks have become a serious concern for deep learning models, where an attacker can steal a trained model by querying its black-box API. This can lead to intellectual property theft and other security and privacy risks. The current state-of-the-art defenses against model stealing attacks suggest adding perturbations to the prediction probabilities. However, they suffer from heavy computations and make impracticable assumptions about the adversary. They often require the training of auxiliary models. This can be timeconsuming and resource-intensive which hinders the deployment of these defenses in real-world applications. In this paper, we propose a simple yet effective and efficient defense alternative. We introduce a heuristic approach to perturb the output probabilities. The proposed defense can be easily integrated into models without additional training. We show that our defense is effective in defending against three state-of-the-art stealing attacks. We evaluate our approach on large and quantized (i.e., compressed) Convolutional Neural Networks (CNNs) trained on several vision datasets. Our technique outperforms the state-ofthe-art defenses with a $\times 37$ faster inference latency without requiring any additional model and with a low impact on the model's performance. We validate that our defense is also effective for quantized CNNs targeting edge devices.

Index Terms—Deep learning, Privacy, Security, Model stealing attacks, Quantization

I. INTRODUCTION

Deep Learning (DL) achieved human-level performance in several computer vision tasks [1]. Developing such highperformance models is a costly process for companies [2]. They often provide their models as a service through an Application Programming Interface (API). This enables users to have predictions for their queries. Besides, state-of-the-art Deep Neural Networks (DNNs) are more and more computationally expensive and require intensive memory, which in some cases hinder their deployment in embedded devices. Therefore, during the past years, researchers proposed to *quantize* DNNs without significantly affecting their performance [3]. This enables a DL model to run efficiently on edge devices to enable a service for a user (e.g. smartphones) or other applications in the system (e.g. autonomous vehicle).

Sharing models as a black box to users or between companies creates a dilemma for the model owner since there might be a financial risk of Intellectual Property (IP) theft. Some research work focuses on obfuscating the model architecture to hinder attackers from retrieving the original model architecture. Nevertheless, the model obfuscation is useless when the attacker can retrieve a surrogate model using only the prediction API [4]. An attacker can act as a normal user and send input queries to the DL model. Leveraging the output predictions, a malicious user can train a high-accuracy and high-fidelity copy of the original model. In addition, model stealing attacks can be used to facilitate other attacks such as adversarial attacks where the attacker creates malicious examples to evade the model classification [5].

Previous work proposed different defenses against APIbased model stealing attacks. Some work proposes to withhold information about probabilities, but this reduces the model's transparency and utility [4]. The prediction confidence is important information, especially for critical systems, e.g. in an autonomous vehicle the system might take an action only when a confidence threshold is reached. Other work proposes to add perturbations to the prediction probabilities in a way that reduces the attacker performance [6]-[8]. However, these techniques often inject high perturbation which hurts the model transparency or even reduces its performance. Additionally, state-of-the-art defense techniques are computationally expensive. Despite the focus of previous work on maintaining the defender model's performance by adding a minimal perturbation, none of them measure their defenses' impact on the model inference latency [6]-[8]. In critical systems, inference latency is very important which makes state-of-theart defenses infeasible. When developing DL models targeting edge devices, it is crucial to have efficient algorithms that neither overload the system nor reduce its utility [3].

To overcome these issues, we propose a novel efficient perturbation-based defense alternative called *Deception* (*DCP*). Our *heuristic* approach aims to deceive the attacker and decrease the accuracy of his stolen model by adding an adaptive noise that deviates the probability scores without increasing the inference time. In order to keep the model's transparency, we aim to decrease the attack performance with a low amount of perturbation. To the best of our knowledge, our approach is the first to propose a defense that takes into account the *attack performance*, the *defender's accuracy*, the *perturbation level*, the *inference time*, the *energy consumption*

This research was funded by Synopsys Inc. and the Natural Sciences and Engineering Research Council of Canada (NSERC).

To appear in Proceedings of ICMLA, Florida, USA ©2023 IEEE

TABLE I: Our approach considers multiple aspects while defending against model stealing attacks

Approach	Fast inference	No auxiliary model	Low perturbation
RS [8]	1	✓	×
AM [16]	✓	×	1
MAD [6]	×	×	1
$GRAD^2$ [7]	×	×	1
Ours	1	\checkmark	1

and the number of required models to perform the defense (Table I). We test our approach on several vision datasets: CIFAR-10, CIFAR-100 [9], SVHN [10], GTSRB [11] and CUB200 [12]. Our evaluation proves that our technique is effective in defending against three state-of-the-art model stealing attacks KnockoffNets [13], MAZE [14] and DFME [15]. We find better or comparable results with state-of-the-art defenses in the matter of the perturbation budget and the attack performance. Furthermore, our technique outperforms the stateof-the-art defenses $GRAD^2$ [7] and MAD [6] with a faster inference latency by consuming less energy and without relying on other surrogate models. We validate our approach on large (i.e., original architecture) and quantized Convolutional Neural Networks (CNNs). We show that the latter is as vulnerable to model stealing attacks as their corresponding original models. Our technique succeeds in defending against these attacks in the quantization setting which ensures that our defense extends to real-world applications in edge devices. Our code is available at: https://github.com/KacemKhaled/defending-extraction

The remainder of the paper is organized as follows: Section II reviews the state-of-the-art that relates to model stealing attacks and defenses, and CNNs quantization; Section III details our methodology; we include our experiments and obtained results in Section IV; and Section V concludes the paper.

II. RELATED WORK

A. Model stealing attacks

Model stealing attacks also referred to as model extraction attacks, can be categorized into two groups: attacks that exploit hardware access and attacks that leverage API query access. We refer to the latter category with API-based attacks. In APIbased attacks, the adversary observes the prediction outputs of his queries, then uses them to steal the functionality of a model [4]. Prior works [5], [17] propose extraction attacks using a jacobian-based data augmentation approach to produce synthetic data. However, their success is limited to small datasets. Reference [13] proposes KnockoffNets attack where the adversary aims to obtain a functionally equivalent clone of the victim model using random and publicly available data. The attack is easy and simple to execute, yet very effective in outperforming other state-of-the-art stealing attacks. Recently, other works in extraction attacks namely MAZE [14] and DFME [15] focus on modern techniques to generate the adversary's dataset instead of relying on public datasets. They leverage generative models to generate data with an objective that enables a successful extraction. However, they require millions of queries instead of thousands compared to *KnockoffNets* [13]. Other stealing attacks in the literature assume that a target model can only output hard labels instead of probabilities [18], [19].

B. Model stealing defenses

To enable model owners to claim their models if stolen, previous work proposes watermarking the model [20], [21]. Watermarking can be done by changing the output probabilities for a small subset of queries. However, this technique does not prevent model stealing. Other defense approaches protect against attacks by perturbing the posterior prediction. These perturbation-based defenses can be categorized into two categories: *proactive* and *reactive* techniques.

Reactive techniques are triggered with the detection of an ongoing extraction attack. They continuously monitor the interactions with the user and the defended model to be able to assess the knowledge stolen by the attacker [22] or to detect an ongoing attack through detecting a deviation from the normal distribution [17] or a large number of Out-of-Distribution (OOD) queries [16]. To detect OOD queries, [16] proposes a technique called *Adaptive Misinformation (AM)* where they flag as malicious queries the ones with low prediction scores. Then they use another model for "misinformation" that generates noise vectors which will be combined with the prediction posteriors. However, this technique is sensitive to false positives which deteriorates the model's performance. Reactive defenses help reduce the efficiency of attacks, but they cause inference delays and increase computational costs.

Proactive defenses are always perturbing the posterior predictions, regardless an attack is detected or not. Some techniques propose to truncate the probability scores or to add random noise [13]. But, these techniques are not efficient against state-of-the-art attacks. The Reverse Sigmoid (RS) defense [8] leverages the last layer in a neural network, i.e., the one that outputs the prediction probabilities of each class, and changes its activation function to a reverse sigmoid. This technique uses a high perturbation level which hurts the model's transparency. To find the suitable perturbation per query, MAD [6] tries to solve an optimization problem where it simulates the attacker by another network and tries to add the perturbation that maximizes the error of the simulated surrogate model. It adds targeted noise to the posterior probabilities to poison the adversary's gradient signal. $GRAD^2$ [7] is a similar approach to MAD that relies on redirecting the gradient signal. Besides, it proposes to train a surrogate model on the attacker's queries to simulate the attacker's knowledge. GRAD² always assumes that the attacker will send a large batch of queries that it can use for training the simulated attacker. This hinders its practicality since it assumes that the system, where the model is implemented to perform predictions, has sufficient hardware to perform the training.

In general, existing perturbation-based defenses are computationally expensive or they use high perturbations that hurt the model's transparency. However, in real-world systems that make decisions based on predictions (e.g., a computer vision system in an autonomous vehicle), the model has to provide a level of confidence in its predictions. The techniques that provide high perturbations make the model useless for these applications. In addition, optimization-based techniques require a lot of gradient calculations, which results in enormous inference delays and an increase in computational costs. Our proposed defense overcomes these issues through a heuristic approach that aims to add perturbations without requiring more hardware resources, as shown in Table I. We offer better tradeoffs in satisfying the perturbation budget vs. the defender's error constraint.

C. CNNs Quantization

Quantization is a technique used to reduce the memory footprint and computation cost of CNNs by representing their weights and activations using lower precision data types. In other words, quantizing the weights involves reducing the number of bits used to represent the parameters, typically from the standard 32-bit floating-point representation to 8bit or even lower precision [23], [24]. In this paper, we consider two of the most used quantization techniques for CNNs: Post Training Quantization (PTQ) and Quantization Aware Training (QAT) [3]. PTQ finds the quantization parameters without retraining the model. It computes the resulting distributions of the different activation functions after feeding some batches of data to the trained model, to determine the specific quantization of each activation at inference time. Then the model is quantized based on the calibration results. This technique is fast and simple, however, the quantized model may have lower accuracy than the original one [3]. During QAT, float parameter values are rounded to mimic int8 values, so the weight adjustments are made while being aware that the model will be quantized. The calibration is performed while re-training the model. This technique often results in quantized models with higher accuracy than with PTQ [3].

III. METHODOLOGY

In this section, we present the threat model, the attack strategy, and our proposed defense approach.

A. Threat Model

We assume that the model owner provides the DL model as a black box. The attacker has an API access to the victim model without internal knowledge about the model parameters nor the training data. The attacker is able to act as a normal user by sending queries to the model and receiving prediction probabilities. Fig. 1 shows an API-based model stealing attack. In these attacks, the attacker attempts to steal the functionality of the victim model by labeling a dataset with the prediction outputs obtained from the attacked model. The goal of the attacker is to obtain a clone with high fidelity to the attacked model and high accuracy on the victim's prediction task.



Fig. 1: In an API-based model stealing attack, the attacker sends queries to the victim model through its prediction API, then he uses the predictions received to create a labeled dataset. The attacker then trains a new model, called the *Adversary* Model, on that dataset. This model represents the stolen model from the victim.

B. Attack strategy

Our aim is to comprehensively evaluate our perturbationbased defense approach against diverse strategies for model stealing. We consider attacks that differ in the query distributions and strategies (random vs. adaptive querying). Our defense aims to perturb the prediction probabilities, thus, we use the state-of-the-art model stealing attacks that leverage the output probabilities of target victim models: *KnockoffNets* [13], *MAZE* [14] and *DFME* [15].

Let $F_V(x;\theta)$ be the victim model provided by the prediction API, with x representing the input query and θ the model parameters. The stealing attack approach has three steps: (i) selecting a query set Q, (ii) constructing a transfer set, and (iii) training a stolen model (the adversary's model) F_A .

MAZE [14] and DFME [15] attacks rely on generative models to create the query set Q. They leverage the victim's outputs and estimate their gradients to adaptively construct suitable image queries. In these attacks steps (i) and (ii) happen simultaneously. KnockoffNets [13] attack selects publicly available datasets to construct the query set Q. The attacker selects an image distribution and randomly samples images from it. We work with different datasets to simulate two scenarios: a knowledgeable adversary and a limited-knowledge adversary. The former is an attacker with knowledge about the distribution of the training data. In this case, we select Q that comes from the same distribution as the victim's dataset \mathcal{D} , so there might be overlapping data between both datasets. We call this a distribution-aware attack. For knowledge-limited adversary, we select a different dataset from the victim $\mathcal{Q} \neq \mathcal{D}$. The adversary uses the obtained output predictions to label the data Q. Finally, using the obtained transfer set, the adversary trains the stolen model F_A using a reasonably complex architecture (e.g., ResNets [25]). The adversary's model F_A can be trained to substitute F_V by minimizing the cross-entropy loss, defined with $y_i = F_V(x; \theta)_i$ as:

$$L(y, F_A(x; \theta')) = -\sum_i y_i \log F_A(x; \theta')_i \tag{1}$$

C. Defense strategy

The defender's goal is to mitigate the extraction attack threat. Since the attacker relies on the prediction probabilities for a successful attack, in order to defend against it, the defender can carefully modify the *posteriors* (i.e., the prediction probabilities) to maximize the attacker's loss Eq. (1). However, the prediction probabilities y have to keep a certain level of confidence score to maintain the model's utility and transparency. Hence, we constrain the perturbation magnitude on the posteriors y by budget ϵ defined as the maximum ℓ_1 distance between the clean posteriors y and the perturbed ones y', defined as:

$$\ell_1 = ||y' - y||_1 \le \epsilon \tag{2}$$

D. Approach

We propose a new heuristic approach to defend against model stealing attacks that could be applied in both large CNNs and Quantized CNNs for edge device implementations. Our defense is characterized by a low computational overhead that does not slow down the model's inference while maintaining a low perturbation of the prediction probabilities. In contrast to the state-of-the-art defenses GRAD² [7] and MAD [6], our technique is not computationally expensive since it does not rely on auxiliary models to simulate attackers. Our technique is motivated by the small inference times of both RS [8] and AM [16] (Table I). We revisit these defenses and we propose a better alternative called Deception (DCP) to overcome their limitations such as high perturbation levels for RS or requiring an auxiliary model and deteriorating the model's performance with false positives for AM.

Through comparing previous defenses, we describe a general defense mechanism that perturbs the prediction probabilities $F_V(x;\theta)$ in a classification problem by:

$$y' = N(a \cdot F_V(x;\theta) + b \cdot r) \tag{3}$$

where $a, b \leq 1$ are linear combination parameters to control the proportions of the noise function r and the prediction probabilities $F_V(x;\theta)$; N(.) is a sum-to-one normalizer so that y' represents the probability values for the different classes in the prediction problem.

Our strategy in injecting the noise is to decrease the uncertain probabilities. Our rationale is that probabilities with low values are of less importance. We assume that if a *welltrained* model is uncertain about the prediction of an input, it is highly probable that the model made a wrong prediction (false positive). Hence, modifying these probabilities will not affect the defender's accuracy $Acc(F_V)$. However, if the model predicts an output with a high probability, it is more likely that it is a correct prediction (true positive). Thus, the injected noise is minimal without affecting the defender's performance.

In order to heuristically find which posteriors to perturb, i.e., low probability values, and switch between clean posteriors $F_V(x;\theta)$ and perturbed ones r(y), we leverage a detector function α inspired from AM [16] and defined as:

$$\alpha = S(\nu \cdot (y_{max} - \tau)) \tag{4}$$

where S is the sigmoid function $S(z) = 1/(1 + e^{-z})$; τ is a threshold that pushes α towards 1 when $y_{max} > \tau$ and towards 0 otherwise (Fig. 2). Similar to [16], we set ν to 1000



Fig. 2: The detector function α that switches between the clean and perturbed posteriors using a threshold τ , which pushes α towards 1 when $y_{max} > \tau$ and towards 0 otherwise, where y_{max} represents the highest value in a probability vector y.

in order to have a non-smooth switch between clean and perturbed predictions.

In contrast to AM that multiplies F_V by $(1 - \alpha)$ in Eq. (7) to create a combination between the clean posteriors and the noise, which results in deteriorating the defender's performance in case of false positives, we set in Eq. (3) the parameter $a \leftarrow 1$ and $b \leftarrow -\beta \cdot \alpha$ to obtain:

$$y' = N(F_V(x;\theta) - \beta \cdot S(\nu \cdot (y_{max} - \tau)) \cdot r(y))$$
 (5)

To compute the perturbation, AM leverages another previously trained model, called *misinformation model* $F_M(x;\theta)$. The latter is trained on the same data as the victim but in a way that maximizes the loss $L(y, F_V(x;\theta))$. However, this deteriorates the performance of the original model because of false positives.

Instead of using an extra model to inject the noise, our defense leverages a noise function r(y) inspired from RS [8]. The latter maintains the defender's accuracy in most cases, without being computationally expensive. However, its limitation is that it uses a high perturbation, increasing the $\ell 1$ distance between the clean and perturbed predictions.

For simplicity, to avoid estimating the adversary's loss through auxiliary models, similar to RS [8], we assume that the attacker has perfectly duplicated the victim model. Thus, we use the victim's predictions $y = F_V(x;\theta)$ as starting point to compute the perturbations r(y). RS uses a reverse sigmoid function $S^{-1}(y) = \log(y/(1-y))$ on the predictions y in order to retrieve the initial logits (i.e., the obtained vectors from the last layer). Then it selects a perturbation $z = \gamma S^{-1}(y)$, where γ is a dataset-specific constant that gives more control over the defense's convergence [8]. After that, a sigmoid function is applied on $z: S(z) = 1/(1 + e^{-z})$. Finally, the heuristically obtained perturbation r is:

$$r(y) = S(z) - 1/2 = S(\gamma S^{-1}(y)) - 1/2$$
(6)

Our defense's final prediction output is obtained through substituting r(y) in Eq. (5), to finally obtain:

$$y' = N(F_V(x;\theta) - \beta \cdot S(\nu \cdot (y_{max} - \tau)) \cdot (S(\gamma S^{-1}(y)) - 1/2))$$
(7)

where S and S^{-1} are respectively the sigmoid and reverse sigmoid functions. We set $\tau \leftarrow \beta$ to control the perturbation magnitude. With this setting we have $y' \rightarrow F_V(x;\theta)$ when $\alpha \rightarrow 0$ and $y' \rightarrow RS$ when $\alpha \rightarrow 1$.

IV. EXPERIMENTS

In this section, we detail our setup: we explain the datasets, the training process, the quantization techniques, the evaluation metrics as well as the baselines defenses. Then, we report and discuss our results.

A. Setup

1) Datasets: We tackle CNNs trained on 5 benchmark vision datasets with diverse resolutions and different numbers of classes. We leverage CIFAR-10 [9], a dataset of images of animals and vehicles with 10 classes; CIFAR-100 [9], a diverse real-world images dataset with 100 classes; SVHN [10] a street view house numbers image dataset with 10 classes; and GTSRB [11], a German traffic signs dataset with 43 classes. All images are sized to 32×32 pixels. Furthermore, we work with CUB-200 [12] dataset which contains images of 200 classes of bird species with 224×224 image size.

Since the KnockoffNets [13] attack requires access to an adversary dataset, we simulate two attack scenarios: a) *knowledgeable adversaries* and b) *knowledge-limited adversaries*. Therefore, to perform the first one, since CIFAR-10 and CIFAR-100 come from the same distribution TinyImages [9], we select queries Q from CIFAR-10 to attack models trained on CIFAR-100 and vice-versa. Second, for knowledge-limited adversaries we choose unrelated datasets, we select queries from CIFAR-10 as an adversary dataset for victims trained on GTSRB and SVHN. To attack the CUB200 victim, we select images from Pascal-VOC [26] which contains a diverse set of objects, such as animals, vehicles, and household items.

2) Training: In order to have a fair comparison with the state-of-the-art techniques and the other baseline defenses, we follow the same training process using the same hyperparameters by the latest technique: $GRAD^2$ [7]. For the CUB-200 dataset, we leverage a ResNet-34 [25] pretrained CNN on ImageNet [27] and we fine-tune it for 50 epochs to obtain our victim model. For the other datasets, we train our victim models from scratch for 50 epochs using the ResNet-18 [25] architecture. In KnockoffNets attack [13], the adversary architecture is a pretrained WideResNet [28] for an attack against a CUB-200 dataset victim and ResNet-34 against the small size datasets. In this attack, we limit the query budget to 50k queries. Similar to [7] in training, we use an SGD optimizer with an initial learning rate of 0.01 that is annealed with a cosine schedule. We use a Nesterov Momentum of 0.9 and a weight decay of $5 \cdot 10^{-4}$. In DFME [15] and MAZE [14] attacks we train adversaries with ResNet-34-8x for small size datasets and WideResNet for CUB-200 dataset. We train each model using 20M queries for 200 epochs. Similar to [15], we use SGD with an initial learning rate of 0.1 and a weightdecay of $5 \cdot 10^{-4}$. Additionally, a learning rate scheduler was applied that multiplies the learning rate by a factor of 0.3 at $0.1 \times$, $0.3 \times$, and $0.5 \times$ the total training epochs.

3) Quantization: We train CNNs for computer vision tasks in our work. Hence, we use two quantization techniques from the literature that are most suited to quantizing our models [3]: PTQ and QAT. We leverage our previously trained



Fig. 3: State-of-the-art attacks vs. our defense. Curves are obtained by varying the perturbation magnitude parameter β (Eq. 7) between 0 and 1.5. The attack performance is represented on the *y*-axis by the *Adversary's Error*. The *x*-axis considers the *Defender's Error*. Our defense is effective in defending against state-of-the-art stealing attacks.

victim models to quantize them with PTQ. We use the same previously described architectures to retrain quantized models with QAT.

4) Evaluation metrics:

a) Classification Error and ℓ_1 distance: In order to assess the performance of a defense against stealing attacks, we compute the adversary's classification error (Adv. Clf.Error) on the victim's test set. This metric measures the stolen functionality from the victim, i.e., the performance of the extracted model on the same task. Similar to [7], [13], we use two utility metrics: the defender's classification error (Def. Clf.Error) and the ℓ_1 distance between the clean predictions and the perturbed predictions by each defense technique. The ℓ_1 distance metric measures the magnitude of the perturbation that impacts the prediction probabilities (Eq. 2). A high ℓ_1 distance (e.g. $\ell_1 > 1.2$) entails a high perturbation that might harm the model's transparency and reduces its utility when users require confident probabilities. This helps in identifying the trade-off between the defender's performance and the perturbation budget for each defense. Both accuracies are calculated using a held-out labeled test set that was not seen before by the victim (during the training) nor the stolen model.

b) Hardware-related metrics: We compute the average inference latency per query and the average GPU energy consumption for all queries using Weights and Biases [29] platform. This helps perceive the impact of the computation cost of each defense technique in real-world applications. To have a fair evaluation, we run all experiments on a Linux CentOS 7 computer with 1 GPU NVIDIA Tesla P100-PCIE-12GB and 24 CPU cores Intel E5-2650 v4 Broadwell @ 2.2GHz.

5) Defenses: We use several baselines to compare our proposed defenses and validate our approach. No Defense means that we attack an undefended model. To the best of our knowledge, the best defenses in the literature are: $GRAD^2$ [7], MAD [6], AM [16], and RS [8]. First, we compare our defense using large (i.e., normally trained) models with all the other defenses. After that, for the *quantized* models, we compare our defenses only to the techniques that are more suitable for edge device implementations. In other words, we select the defense that maintains the inference time without requiring additional models: RS.



Fig. 4: Impact of the KnockoffNets stealing attack [13] on our defense vs. other baselines. Curves are obtained by varying the perturbation magnitude parameter β (Eq. 7) between 0 and 1.5. The attack performance is represented on the *y*-axis by the *Adversary's Classification Error*. The *x*-axis considers the *Defender's Classification Error*. Our defense outperforms the state-of-the-art defenses on knowledgeable adversaries (a) and knowledge-limited attackers (b).

TABLE II: Impact of the stealing attack on the defenses with constraints on both the variation of the Defender's Classification Error " Δ Def.Clf.Error" and the ℓ_1 distance between clean and perturbed predictions. We report the maximum attainable value under both constraints. Δ Def.Clf.Error ranges from 1% to 5% and the ℓ_1 distance is maintained at a budget of 0.9. **Bold** denotes the best value and <u>underlined</u> denotes the second-best.

	CIFAR-10→CIFAR-100		CIFAR-100→CIFAR-10		CIFAR-10→SVHN		CIFAR-10→GTSRB			Pascal-VOC→CUB200					
	Δ Def.Clf.Error		Δ Def.Clf.Error		Δ Def.Clf.Error		Δ Def.Clf.Error		Δ Def.Clf.Error						
Method	1%	2%	5%	1%	2%	5%	1%	2%	5%	1%	2%	5%	1%	2%	5%
No defense	56.89	56.89	56.89	18.03	18.03	18.03	30.92	30.92	30.92	7.18	7.18	7.18	61.14	61.14	61.14
RS	64.20	64.20	64.20	19.22	19.22	19.22	51.07	51.07	51.07	14.23	14.23	14.23	83.80	83.80	83.80
AM	59.81	61.27	<u>65.87</u>	19.88	21.34	25.65	58.29	58.29	58.29	12.04	15.46	15.46	88.92	88.92	88.92
MAD (pre-sota)	60.27	61.51	64.94	21.54	23.08	27.24	40.42	43.55	51.98	<u>14.84</u>	<u>17.65</u>	23.52	77.11	78.45	80.61
GRAD ² (sota)	60.77	62.06	65.20	21.56	23.37	27.25	43.73	47.27	52.51	11.15	12.86	15.02	82.04	83.49	85.36
DCP (ours)	70.31	70.31	70.31	21.87	23.78	23.78	59.73	59.73	59.73	21.69	21.69	21.69	85.62	85.62	85.62

B. Results

We present the results of our experiments in Fig. 3, Fig. 4, Fig. 5, Fig. 6, Fig. 7, Table II, and Table III. In the figures, we visualize two utility metrics: the defender's error and the ℓ_1 distance between the original predictions y and the perturbed ones y'. The attack performance is measured by the adversary's classification error. The higher this error, the stronger the defense.

1) DCP defense vs. state-of-the-art attacks: Fig. 3 presents an evaluation of our defense DCP (Eq. 7) against the three attacks KnockoffNets [13], DFME [15], and MAZE [14]. We observe the performance of these attacks on undefended and defended models. A good adversary produces a model with a low classification error. As shown in Fig. 3, using our defense approach for all datasets and attack models can significantly reduce the effectiveness of the attacker by increasing its error. Our defense produces reasonable operating points across all models, with a minimal impact on the defender's error and a high impact on the adversary's error. For instance, a defended victim trained on the SVHN dataset can multiply the adversary's error by up to $\times 3.3$. From Fig. 3 we notice that the strongest model stealing attack is KnockoffNets as it produces the lowest adversary's error for most datasets.

2) DCP defense vs. baselines: We evaluate our defense and the baselines against the strongest stealing attack in the attack models evaluation: KnockoffNets. Fig. 4 shows that for knowledgeable and knowledge-limited attackers, our defense technique outperforms the state-of-the-art defenses in most cases. It increases the adversary's classification error significantly before it starts impacting the defender's classification error. For instance, for CIFAR-10 (Fig. 4a, right column), our defense DCP outperforms other techniques in increasing the adversary's classification error by 22% with a small impact on the defender's performance (1%). Most baseline defenses fail to decrease the adversary's performance without deteriorating the defender's performance. For a victim trained on GTSRB, when we constrain only the defender's error to be lower than 1%, our defense DCP increases the adversary's classification error by more than 109%. The second-best is RS which increases the adversary's classification error up to 79%.

In order to balance both metrics measuring the impact on the defender's performance and the ℓ_1 distance between clean and perturbed predictions, we constrain both the ℓ_1 distance with a maximum budget of 0.9 and the defender's error variation to be lower than 1%, 2% or 5%. Table II shows the results of this experiment. In most cases, our defense DCP yields comparable to or better results than the state-of-the-art defenses while offering reasonable trade-offs that prioritize the defender's performance.

Through these extensive experiments, we prove that our results are significantly better than the state-of-the-art for knowledgeable and knowledge-limited adversaries. We pro-



Fig. 5: Our defense vs. Reverse Sigmoid. The Adversary's classification error with relation to the ℓ_1 distance between the perturbed predictions and the clean predictions. Our defense achieves higher performance using lower perturbation budget.



Fig. 6: Performance of the defenses against extraction attacks relying on different numbers of queries. Each column considers ℓ_1 budget constraint. The attacks are performed on a victim trained on the SVHN dataset. In most cases of a low number of queries, our defense outperforms other techniques in increasing the adversary's classification error.

vide better robustness against stealing attacks while maintaining a reasonable perturbation.

3) Ablation study:

a) Comparing ℓ_1 distance of DCP and RS: In the previous experiments, we observe that RS [8] and our defense DCP sometimes achieve the same attack performance, i.e., they both increase the adversary's error to comparable values with a fast inference latency. This could be explained by the fact that our defense uses a noise function that draws some inspiration from RS. We evaluate both defenses from a perturbation magnitude standpoint. Fig. 5 shows that our defense outperforms RS in the matter of ℓ_1 budget. In fact, our defense uses less perturbation to increase the adversary's classification error. Hence, DCP overcomes the limitation of RS that requires a high perturbation to perform well.

b) The impact of the number of queries used in the stealing attack: We explore the resilience of our defense in protecting models from multiple stealing attacks with different budgets of queries. For each attack, we select a budget for the number of queries ranging from 10k to 50k images. Each victim model is protected with a perturbation-based defense constrained by an $\ell 1$ budget for the distance between the clean and perturbed predictions. Fig. 6 shows the results of this experiment on the SVHN dataset. Our defense increases the adversary's classification error significantly higher than other defenses, especially for a lower number of queries. This entails that our defense slows down the adversary which makes him use a larger number of queries to perform the attack.

TABLE III: The average inference latency per query and the average energy consumption required to generate all perturbed queries, **bold** means best.

	CIFAR10	CIFAR100	SVHN	GTSRB	CUB200				
Defense	Inference latency (ms)								
No defense	0.10	0.09	0.09	0.07	0.09				
RS [8]	0.10	0.10	0.09	0.13	0.10				
AM [16]	0.14	0.14	0.13	0.16	0.17				
MAD [6]	58.02	499.09	59.12	225.50	121.4				
$GRAD^2$ [7]	4.09	21.83	4.39	14.47	44.17				
DCP (ours)	0.11	0.09	0.09	0.13	0.09				
Defense	Energy consumption (Wh)								
No defense	0.13	0.15	0.20	0.12	0.43				
RS [8]	0.18	0.15	0.20	0.12	0.90				
AM [16]	0.21	0.18	0.23	0.17	0.93				
MAD [6]	2.79	22.18	6.96	12.66	80.06				
$GRAD^2$ [7]	0.42	1.20	0.94	1.19	12.15				
DCP (ours)	0.17	0.17	0.20	0.12	0.59				

4) Inference latency and energy consumption: To validate the efficiency of our proposed defense, we report in Table III the hardware-related evaluation metrics for our defense and the baselines. For example, for the CIFAR-10 dataset, we obtain $\times 37$ faster inference than the state-of-the-art technique GRAD² and more than $\times 527$ faster inference than MAD. In addition, we show that our defenses have comparable minimal energy consumption for each dataset compared to other defenses. Finally, we only require half the memory used by most of the baselines, since our technique does not require an extra model to compute the perturbations. To conclude, our defense is practical for edge device implementations: in general, we report comparable-to-best results in inference latency and energy consumption.

5) Model stealing defenses on quantized models: We validate the extension of our techniques on quantized CNNs. Fig. 7 shows that in most cases our defense outperforms or matches the results of RS. For example, in PTQ results (Fig. 7a), for a selected perturbation budget (e.g. $\ell_1 = 0.5$), our technique DCP provides better protection against model stealing attacks through increasing the adversary's classification error to between 5% and 20% while maintaining the defender's performance. Our technique sometimes achieves the same performance while having lower perturbations. This shows that our defense can be used in DL models targeting embedded systems applications.

V. CONCLUSION

The motivation for our work is the non-efficiency of previous defense techniques against model stealing attacks. This hinders their deployment in edge devices. We overcame the heavy computation problem with a novel heuristic defense. We empirically proved that it works efficiently on large CNNs and their corresponding quantized versions targeting edge devices. We found that our approach achieves comparable-to-better performance with the state-of-the-art while maintaining low perturbations and fast inference time. Our defense proved to



Fig. 7: Performance of the defense against stealing attack on Quantized CNNs with two quantization techniques PTQ (a) and QAT (b). Our defense outperforms the baseline defense that does not require an auxiliary model, nor increases inference delays.

be effective in amplifying the adversary's classification error under various attack models.

With this paper, we shed light on the need for considering the hardware limitations during the development of security and privacy algorithms for DL applications. We strongly encourage the community to prioritize the practicality of their proposed defenses, not as an undesired side-effect, but as a crucial requirement for real-world applications. For future work, we intend to work on developing efficient defenses specifically tailored for edge devices and real-time systems.

REFERENCES

- A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.
- [2] E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for deep learning in NLP," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 3645–3650.
- [3] A. Gholami, S. Kim, Z. Dong, Z. Yao, M. W. Mahoney, and K. Keutzer, "A survey of quantization methods for efficient neural network inference," in *Low-Power Computer Vision*. Chapman and Hall/CRC, 2022, pp. 291–326.
- [4] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Stealing machine learning models via prediction apis," in 25th {USENIX} Security Symposium ({USENIX} Security 16), 2016, pp. 601–618.
- [5] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, "Practical black-box attacks against machine learning," in *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, 2017, pp. 506–519.

- [6] T. Orekondy, B. Schiele, and M. Fritz, "Prediction poisoning: Towards defenses against dnn model stealing attacks," in *International Conference on Learning Representations*, 2019.
- [7] M. Mazeika, B. Li, and D. Forsyth, "How to steer your adversary: Targeted and efficient model stealing defenses with gradient redirection," in *International Conference on Machine Learning*. PMLR, 2022, pp. 15 241–15 254.
- [8] T. Lee, B. Edwards, I. Molloy, and D. Su, "Defending against neural network model stealing attacks using deceptive perturbations," in 2019 IEEE Security and Privacy Workshops (SPW). IEEE, 2019, pp. 43–49.
- [9] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," 2009.
- [10] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, "Reading digits in natural images with unsupervised feature learning," 2011.
- [11] J. Stallkamp, M. Schlipsing, J. Salmen, and C. Igel, "The German Traffic Sign Recognition Benchmark: A multi-class classification competition," in *IEEE International Joint Conference on Neural Networks*, 2011, pp. 1453–1460.
- [12] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The Caltech-UCSD Birds-200-2011 Dataset," California Institute of Technology, Tech. Rep. CNS-TR-2011-001, 2011.
- [13] T. Orekondy, B. Schiele, and M. Fritz, "Knockoff nets: Stealing functionality of black-box models," in *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, 2019, pp. 4954–4963.
- [14] S. Kariyappa, A. Prakash, and M. K. Qureshi, "Maze: Data-free model stealing attack using zeroth-order gradient estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 13814–13823.
- [15] J.-B. Truong, P. Maini, R. J. Walls, and N. Papernot, "Data-free model extraction," in *Proceedings of the IEEE/CVF conference on computer* vision and pattern recognition, 2021, pp. 4771–4780.
- [16] S. Kariyappa and M. K. Qureshi, "Defending against model stealing attacks with adaptive misinformation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 770– 778.
- [17] M. Juuti, S. Szyller, S. Marchal, and N. Asokan, "Prada: protecting against dnn model stealing attacks," in 2019 IEEE European Symposium on Security and Privacy (EuroS&P). IEEE, 2019, pp. 512–527.
- [18] M. Zhou, J. Wu, Y. Liu, S. Liu, and C. Zhu, "Dast: Data-free substitute training for adversarial attacks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [19] S. Sanyal, S. Addepalli, and R. V. Babu, "Towards data-free model stealing in a hard label setting," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 15 284–15 293.
- [20] H. Jia, C. A. Choquette-Choo, V. Chandrasekaran, and N. Papernot, "Entangled watermarks as a defense against model extraction," in 30th USENIX Security Symposium (USENIX Security 21), 2021.
- [21] S. Szyller, B. G. Atli, S. Marchal, and N. Asokan, "Dawn: Dynamic adversarial watermarking of neural networks," in *Proceedings of the 29th* ACM International Conference on Multimedia, 2021, pp. 4417–4425.
- [22] M. Kesarwani, B. Mukhoty, V. Arya, and S. Mehta, "Model extraction warning in mlaas paradigm," in *Proceedings of the 34th Annual Computer Security Applications Conference*, 2018, pp. 371–380.
- [23] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding," arXiv preprint arXiv:1510.00149, 2015.
- [24] V. Vanhoucke, A. Senior, and M. Z. Mao, "Improving the speed of neural networks on cpus," 2011.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision* and pattern recognition, 2016, pp. 770–778.
- [26] M. Éveringham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International Journal of Computer Vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [27] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in 2009 IEEE conference on computer vision and pattern recognition. IEEE, 2009, pp. 248–255.
- [28] S. Zagoruyko and N. Komodakis, "Wide residual networks," arXiv preprint arXiv:1605.07146, 2016.
- [29] L. Biewald, "Experiment tracking with weights and biases," 2020, software available from wandb.com. [Online]. Available: https: //www.wandb.com/