

# ES Attack: Model Stealing Against Deep Neural Networks Without Data Hurdles

Xiaoyong Yuan , Member, IEEE, Leah Ding, Member, IEEE, Lan Zhang , Member, IEEE, Xiaolin Li, and Dapeng Oliver Wu, Fellow, IEEE

**Abstract**—Deep neural networks (DNNs) have become the essential components for various commercialized machine learning services, such as Machine Learning as a Service (MLaaS). Recent studies show that machine learning services face severe privacy threats - well-trained DNNs owned by MLaaS providers can be stolen through public APIs, namely model stealing attacks. However, most existing works undervalued the impact of such attacks, where a successful attack has to acquire confidential training data or auxiliary data regarding the victim DNN. In this paper, we propose *ES Attack*, a novel model stealing attack without any data hurdles. By using heuristically generated synthetic data, *ES Attack* iteratively trains a substitute model and eventually achieves a functionally equivalent copy of the victim DNN. The experimental results reveal the severity of *ES Attack*: i) *ES Attack* successfully steals the victim model without data hurdles, and *ES Attack* even outperforms most existing model stealing attacks using auxiliary data in terms of model accuracy; ii) most countermeasures are ineffective in defending *ES Attack*; iii) *ES Attack* facilitates further attacks relying on the stolen model.

**Index Terms**—Model stealing, deep neural network, knowledge distillation, data synthesis.

## I. INTRODUCTION

AS ONE of the typical business models, Machine-Learning-as-a-Service (MLaaS) provides a platform to facilitate users to use machine learning models [1]. Users can access well-trained machine learning models via public APIs provided

Manuscript received September 26, 2021; revised December 3, 2021 and January 5, 2022; accepted January 14, 2022. This work was supported by National Science Foundation under Grants CCF-2007210, CCF-2106610, and CCF-2106754. (Corresponding author: Xiaoyong Yuan.)

Xiaoyong Yuan is with the College of Computing, Michigan Technological University, Houghton, MI 49931 USA (e-mail: xyuan@mtu.edu).

Leah Ding is with the Department of Computer Science, American University, Washington, DC 20016 USA (e-mail: ding@american.edu).

Lan Zhang is with the Department of Electrical and Computer Engineering, Michigan Technological University, Houghton, MI 49931 USA (e-mail: lanzhang@mtu.edu).

Xiaolin Li is with Deep Learning, Cognization Lab, Palo Alto, CA 94301 USA (e-mail: xiaolinli@ieee.org).

Dapeng Oliver Wu is with the Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611 USA (e-mail: dpwu@ufl.edu).

We show the synthetic images using *OPT-SYN* with the best substitute model. We compare them with the victim's training data in the SVHN, CIFAR-10, MNIST and KMNIST datasets (Figs. 10, 11, , 13). First row: images in the victim's training dataset. Second row: images in the synthetic dataset. Third Row: corresponding labels of images.

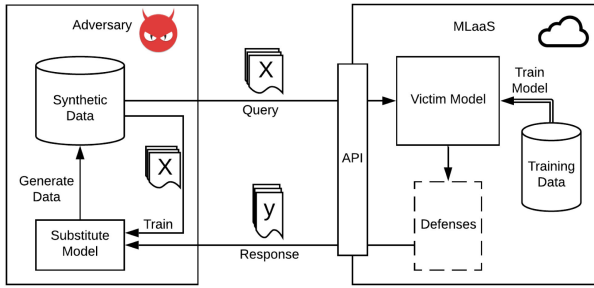
This article has supplementary downloadable material available at <https://doi.org/10.1109/TETCI.2022.3147508>, provided by the authors.

Digital Object Identifier 10.1109/TETCI.2022.3147508

by MLaaS providers, without building a model from scratch. MLaaS allows users to query machine learning models in the form of pay-per-query and get responses of the model's predictions. Recent studies show that machine learning services face severe privacy threats: model stealing attacks steal functionally equivalent copies from MLaaS providers through multiple queries [2]–[6]. Model stealing attacks exploit the tensions between queries and their corresponding feedback, *i.e.*, the output predictions. Tramer *et al.* extract machine learning model's parameters by solving equations derived from the model architecture [2]. However, it requires the exact knowledge of ML architectures and becomes difficult to scale up to steal deep neural networks (DNNs) [7]. Existing model stealing attacks against DNNs require specific knowledge of the model's training data: the exact training data [6], [8], seed samples from the training data [9], and auxiliary data that shares similar attributes as the training data or within the same task domain [3]–[5]. Most existing model stealing attacks require the knowledge of training data or auxiliary data regarding the victim deep neural networks, which undervalued the impact of model stealing attacks. In practice, these data are not always accessible. Due to recent regulations on data protection (*e.g.*, GDPR and CCPA), many types of personal data are hard to acquire, such as health data and biometric data. In many domains, companies collect data for their own business and are reluctant to share their data. Government and other organizations usually lack resources and financial supports to create open datasets. Often, the quality of public data is out-of-date and questionable without updates and maintenance. The availability of appropriate data protects the victim models from being stolen by the existing model stealing attacks.

In this paper, we introduce *ES Attack*, a new class of model stealing attacks against DNNs without data hurdles. *ES Attack* heuristically generates synthetic data to overcome the limitations of existing approaches. Fig. 1 illustrates the diagram of *ES Attack*. The adversary queries the victim model with the synthetic data  $x$  and labels the data using responses  $y$  via the MLaaS provider's APIs. The MLaaS provider may deploy defenses to prevent model leakage. A substitute model is iteratively trained using the synthetic data of  $x$  with corresponding labels  $y$  and eventually approximates the functionality of the victim model. There are two key steps in *ES Attack*: E-Step to Estimate the parameters in a substitute model and S-Step to Synthesize the dataset for attacking.

Compared to existing model stealing attacks, *ES Attack* does not require i) information about the internals of the victim's

Fig. 1. Diagram of *ES Attack*.

model (e.g., its architecture, hyper-parameters, and parameters), and ii) prior knowledge of the victim model's training data. The adversary only observes responses given by the victim model.

In summary, our contributions are fourfold:

- 1) We propose a novel model stealing framework *ES Attack* that does not require any knowledge of the victim model's training data. Compared to model stealing using auxiliary datasets, our proposed *ES Attack* improves the model accuracy by 44.57%.
- 2) *ES Attack* generates better synthetic datasets compared with the auxiliary datasets in terms of quality and diversity.
- 3) We demonstrate that the stolen model successfully facilitates black-box adversarial attacks against the victim model.
- 4) Three investigated countermeasures are not effective in preventing *ES Attack*.

## II. PROBLEM STATEMENT

In this paper, we consider a typical service in Machine-Learning-as-a-Service (MLaaS), *i.e.*, image classification. Given a task domain  $\mathcal{T}$ , an MLaaS provider (*i.e.*, victim) collects a training dataset  $\mathcal{D}_{\text{train}}$  and a test dataset  $\mathcal{D}_{\text{test}}$ , consisting of a set of images and their corresponding labels  $\{(\mathbf{x}, \mathbf{l})\}$ . Further, the MLaaS provider trains a machine learning model (*i.e.*, victim model)  $f_v$  on private training dataset  $\mathcal{D}_{\text{train}}$  and provides the image classification services to the public using the victim model. Normal users can access the trained model  $f_v$  by querying data sample  $\mathbf{x}$  and get the response from the victim model regarding the predicted probabilities of  $K$  classes  $\mathbf{y} = f_v(\mathbf{x})$ .

The goal of model stealing adversaries is to build a model (*i.e.*, stolen model)  $f_s$  that is functionally equivalent to the victim model  $f_v$ . We assume that the adversary can query the victim model by pretending themselves as normal users. Most model stealing attacks assume that the adversary has full or partial prior knowledge of the victim model's training data  $\mathcal{D}_{\text{train}}$ . In this paper, we consider a more realistic scenario that the adversary cannot access the victim's private training data. Specifically, we assume that the adversary does not know the victim's training data  $\mathcal{D}_{\text{train}}$  or any auxiliary data related to  $\mathcal{D}_{\text{train}}$ .

To evaluate the risk of model stealing attacks, we leverage the prediction accuracy on the test dataset  $\mathcal{D}_{\text{test}}$ . By mimicking the behavior of the victim model, the adversary aims to achieve good performance on the unknown test dataset  $\mathcal{D}_{\text{test}}$  using the stolen model  $f_s$ .

## III. *ES ATTACK*

In this section, we present the design of *ES Attack*, and propose two heuristic approaches for data synthesis.

### A. Design of *ES Attack*

Model stealing attacks aim to build a model  $f_s$  that is functionally equivalent to the victim model  $f_v$ . Theoretically, if the adversary can train the substitute model on all the samples in the input space of  $f_v$ , the substitute model can achieve the same performance as the victim model. However, it is infeasible to query all the samples in the input space. Therefore, a critical step for model stealing attacks is to explore the input sub-space that represents the task domain  $\mathcal{T}$ . Adversaries will mimic the behavior of victim models in the input sub-space. [3]–[5] leverage public datasets as an auxiliary dataset to train the substitute model to approximate the output of the victim model. The auxiliary data share common attributes with  $\mathcal{D}_{\text{train}}$ , which can be used to train the substitute model. However, these approaches are not practical due to many reasons: i) Data with shared attributes is not always available. Confidential data such as medical images, financial records are not publicly available. The scarcity of data is still a critical problem in many domains. ii) The relationship between the available auxiliary data on the public domain and the task domain  $\mathcal{T}$  is unclear, which brings a challenge to select a proper auxiliary dataset. The rationale for selecting a specific auxiliary dataset is missing in most of the existing approaches. In the experiments, we show that using a randomly generated dataset, a special case of an auxiliary dataset, fails to steal the victim model. iii) The quality of data used for training the substitute model cannot be improved during model stealing. The data samples are fixed in the auxiliary dataset.

Therefore, we propose an *ES Attack* to heuristically explore the potential input space related to task domain  $\mathcal{T}$  by learning from the victim model. We outline *ES Attack* in Algorithm 1. First, *ES Attack* initializes a randomly synthetic dataset  $\mathcal{D}_{\text{syn}}^{(0)}$ , which may share few attributes with  $\mathcal{D}_{\text{train}}$ , most likely fewer than  $\mathcal{D}_{\text{aux}}$ . Second, *ES Attack* trains a substitute model  $f_s$  based on the samples from the synthetic dataset and their predictions from the victim model. Then, *ES Attack* can generate a better synthetic dataset using the improved substitute model; in the meanwhile, the better synthetic dataset can help improve the substitute model. In this way, *ES Attack* iteratively synthesizes the datasets and trains the substitute model to improve the quality of the synthetic dataset and the performance of the substitute model. Eventually, the synthetic datasets will approximate the private training dataset, and the substitute model will approximate the victim model or steal the victim model.

Fig. 2 illustrates the progress of data synthesis during *ES Attack*. In Fig. 2, we compare the synthetic datasets  $\mathcal{D}_{\text{syn}}^{(t)}$  (in red) with the victim's training dataset  $\mathcal{D}_{\text{train}}$  (in blue) and the auxiliary dataset  $\mathcal{D}_{\text{aux}}$  (in green).  $\mathcal{D}_{\text{aux}}$  may share similar input space with  $\mathcal{D}_{\text{train}}$ , but in most cases, adversaries may not know the distance between the distribution of  $\mathcal{D}_{\text{aux}}$  and the distribution of  $\mathcal{D}_{\text{train}}$ . Hence,  $\mathcal{D}_{\text{train}}$  may not be fully covered by  $\mathcal{D}_{\text{aux}}$ . However,

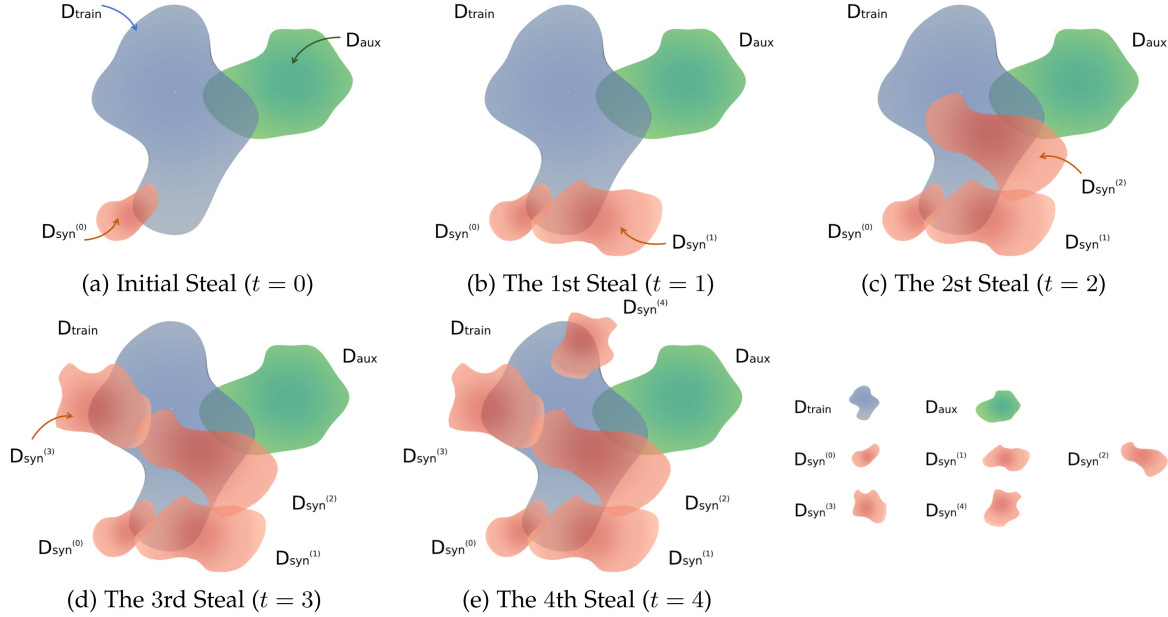


Fig. 2. Progress of data synthesis during *ES Attack*. We compare the synthetic datasets  $\mathcal{D}_{\text{syn}}^{(t)}$  (in red), generated by our proposed attack, with the victim's training dataset  $\mathcal{D}_{\text{train}}$  (in blue) and the auxiliary dataset  $\mathcal{D}_{\text{aux}}$  (in green). The auxiliary dataset  $\mathcal{D}_{\text{aux}}$  may share similar input space with the victim's training dataset  $\mathcal{D}_{\text{train}}$ , but a large space of  $\mathcal{D}_{\text{train}}$  cannot be covered by the auxiliary dataset. In our attack, we first initial a randomly generated synthetic dataset  $\mathcal{D}_{\text{syn}}^{(0)}$ .  $\mathcal{D}_{\text{syn}}^{(i)}$  share some attributes with  $\mathcal{D}_{\text{train}}$ , which might be less than  $\mathcal{D}_{\text{aux}}$ . During our attacks,  $\mathcal{D}_{\text{syn}}^{(t)}$  ( $t = 1, 2, 3$ ) get information from our substitute models and explore more space in  $\mathcal{D}_{\text{train}}$  in each iteration. The goal of the attacks is to cover the input space of  $\mathcal{D}_{\text{train}}$  as much as possible. Note that the adversary trains the substitute model on a synthetic dataset  $\mathcal{D}_{\text{syn}}^{(i)}$ , ( $i = 0, 1, \dots, N$ ) in each stealing epoch. As a sum, the substitute model is trained on all the synthetic datasets. (a) Initial Steal ( $t = 0$ ). (b) The 1st Steal ( $t = 1$ ). (c) The 2st Steal ( $t = 2$ ). (d) The 3rd Steal ( $t = 3$ ). (e) The 4th Steal ( $t = 4$ ).

after initializing the synthetic dataset  $\mathcal{D}_{\text{syn}}^{(0)}$ , *ES Attack* will iteratively improve the synthetic datasets  $\mathcal{D}_{\text{syn}}^{(t)}$  ( $t = 1, 2, 3, 4, \dots$ ), and explore more space in the training dataset  $\mathcal{D}_{\text{train}}$ .

Two key steps in *ES Attack*, E-Step and S-Step, are described as follows.

*E-Step*: Estimate parameters in the substitute model on the synthetic dataset using knowledge distillation. The knowledge distillation approach transfers the knowledge from  $f_v$  to  $f_s$  with minimal performance degradation:

$$f_s^{(t)} \leftarrow \arg \min_{f_s} \text{L}_{\text{KD}} \left( f_s, f_v; \mathcal{D}_{\text{syn}}^{(t-1)} \right), \quad (1)$$

where  $f_s^{(t)}$  denotes the substitute model at iteration  $t$  and  $\mathcal{D}_{\text{syn}}^{(t-1)}$  denotes the synthetic dataset at the previous iteration  $t - 1$ . The objective function  $\text{L}_{\text{KD}}$  is defined as knowledge distillation loss to make  $f_s^{(t)}$  approximate the victim model  $f_v$ :

$$\text{L}_{\text{KD}}(f_s, f_v; \mathcal{D}_{\text{syn}}) = \frac{1}{|\mathcal{D}_{\text{syn}}|} \sum_{\mathbf{x} \in \mathcal{D}_{\text{syn}}} \text{L}_{\text{CE}}(f_s(\mathbf{x}), f_v(\mathbf{x})), \quad (2)$$

where  $\text{L}_{\text{CE}}$  denotes the cross-entropy loss. We train the substitute model by minimizing the objective function (1) for  $M$  epochs using Adam [10].

*S-Step*: Synthesize the dataset  $\mathcal{D}_{\text{syn}}^{(t)} = \{\mathbf{x}\}$  consisted of the synthetic input samples.

### B. Two Approaches for Data Synthesis (S-Step)

Data synthesis (S-step in *ES Attack*) aims to explore the possible data that reveal the data distribution of the victim's

---

#### Algorithm 1: *ES Attack*.

---

##### INPUT:

- The black-box victim model  $f_v$
- Number of classes  $K$
- Number of stealing epochs  $N$
- Number of training epochs for each stealing epoch  $M$

##### OUTPUT:

- The substitute model  $f_s^{(N)}$
  - 1: Initialize a synthetic dataset  $\mathcal{D}_{\text{syn}}^{(0)}$  by randomly sampling  $\mathbf{x}$  from a Gaussian distribution.
  - 2: Construct an initial substitute model  $f_s^{(0)}$  by initializing the parameters in the model.
  - 3: **for**  $t \leftarrow 1$  to  $N$  **do**
  - 4:   **E-Step**: Estimate the parameters in the substitute model  $f_s^{(t)}$  using knowledge distillation for  $M$  epochs on the synthetic dataset  $\mathcal{D}_{\text{syn}}^{(t-1)}$ .
  - 5:   **S-Step**: Synthesize a new dataset  $\mathcal{D}_{\text{syn}}^{(t)}$  based on the knowledge of the substitute model  $f_s^{(t)}$ .
  - 6: **end for**
  - 7: **return**  $f_s^{(N)}$ .
- 

training dataset and benefit the substitute model. The most challenging problem in data synthesis is the lack of gradient of the victim model. The adversary can only get the prediction rather than the gradient from the victim model. Accordingly, the data generated by the adversary cannot be tuned by the



victim model directly. The existing approaches used in data-free knowledge distillation fail to be applied in the model stealing attacks, since they require to back-propagate the gradients from the victim model. More discussion about the difference between model stealing and data-free knowledge distillation can be found in Section VI-C. To overcome this challenge, we propose two data synthesis approaches that make use of the gradients of the substitute model as a proxy for updating the synthetic input data.

Specifically, we introduce two approaches to generate the synthetic data: namely, *DNN-SYN* and *OPT-SYN*. Both approaches start with generating a set of random labels and initial data samples. Further the approaches update the data samples based on the assigned labels and the gradients from the substitute model. Once the substitute model is close to the victim model, the synthetic data becomes close to the distribution of the victim's training dataset.

1) *DNN-SYN*: We design a DNN-based generator to synthesize images that can be classified by our substitute model with high confidence. The design of image generator  $G$  follows the major architecture of Auxiliary Classifier GAN (ACGAN) [11], a variant of Generative Adversarial Network (GAN), which can generate images with label conditioning. We refer to the data generation approach using DNN as *DNN-SYN*.

We describe the procedure of *DNN-SYN* as follows:

- 1) *Step 1*: Randomly assign a set of labels  $\mathbf{l} = \{\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_n\}$ , where  $\mathbf{l}_i$  denotes a  $K$ -dimensional one-hot vector.
- 2) *Step 2*: Train a DNN generator  $G$  with parameter  $w_G$  to generate data from a random latent vector  $\mathbf{z}_i$ .  $G$  is optimized that the generated data can be classified by  $f_s$  as assigned labels  $L$  with high confidence:

$$\min_{w_G} L_{\text{img}}(G, \mathbf{l}) \stackrel{\text{def}}{=} \sum_i^n L_{\text{CE}}(f_s(G(\mathbf{z}_i, \mathbf{l}_i)), \mathbf{l}_i). \quad (3)$$

- 3) *Step 3*: Generate a synthetic dataset using the generator trained in Step 2:  $\mathcal{D}_{\text{syn}} = \{G(\mathbf{z}_i, \mathbf{l}_i)\}$ .

In addition, mode collapse is one of the critical problems for training GANs. To avoid mode collapse in *DNN-SYN*, we use a mode seeking approach to increase the diversity of data samples [12]. Mode seeking has been shown simple yet effective in mitigating mode collapse. We generate two images  $G(\mathbf{z}_i^1)$  and  $G(\mathbf{z}_i^2)$  using latent vectors  $\mathbf{z}_i^1$  and  $\mathbf{z}_i^2$  and maximize the ratio of the distance of images to the distance of latent vectors. In other words, we minimize the mode seeking loss  $L_{\text{ms}}$ :

$$L_{\text{ms}}(G, \mathbf{l}) \stackrel{\text{def}}{=} \sum_i^n \frac{d(\mathbf{z}_i^1, \mathbf{z}_i^2)}{d(G(\mathbf{z}_i^1, \mathbf{l}_i), G(\mathbf{z}_i^2, \mathbf{l}_i))}, \quad (4)$$

where  $d(\cdot, \cdot)$  denotes a distance metric. In this paper, we use  $\ell_2$  norm distance. We sum up the original objective function  $L_{\text{img}}$  and the regularization term  $L_{\text{ms}}$  and minimize the new objective function:

$$L_{\text{DNN}} = L_{\text{img}} + \lambda L_{\text{ms}}, \quad (5)$$

where  $\lambda$  denotes the hyper-parameter to adjust the value of regularization. In the experiment, we set  $\lambda$  as 1.




Fig. 3. DNN generator  $G$  in *DNN-SYN*.

Fig. 3 illustrates the architecture of DNN-based generator in *DNN-SYN*. We follow the major design of ACGAN [11]. We feed a random latent vector  $\mathbf{z}_i$  and a generated one-hot label  $\mathbf{l}_i$  into generator  $G$ . We concatenate these two vectors and up-sample them using several transposed convolution layers. Transposed convolution layers are parameterized by learnable weights. Each transposed convolution layer is followed by a batch normalization layer and a ReLU activation function except the last transposed convolution layer. 4 transposed convolution layers are used in the model. In the final layer, we use a Tanh function to output a synthesis image within  $(-1, 1)$ .

2) *OPT-SYN*: Instead of training a generator to synthesize the dataset, we propose an optimization-based data synthesis approach, *OPT-SYN*, which operates on the input space directly and does not suffer the problem of mode collapse. In addition, *OPT-SYN* explores a more diverse label space compared to the one-hot labels used in *DNN-SYN*. *OPT-SYN* first explores the possible prediction vectors  $\{\mathbf{y}\}$  in the task domain  $\mathcal{T}$  and then minimizes the cross-entropy loss between  $\{\mathbf{y}\}$  and the substitute model's prediction on the synthetic data:

$$\min_{\mathbf{y}} L_{\text{CE}}(f_s^{(t)}(\mathbf{x}), \mathbf{y}), \quad (6)$$

where  $f_s^{(t)}$  denotes the substitute model at the  $t$ th stealing epoch. In the experiments, we find that *OPT-SYN* performs better than *DNN-SYN* does in most scenarios.

The proposed *OPT-SYN* approach is detailed in Algorithm 2. First, to explore the possible prediction vectors, we sample each random vector  $\mathbf{y} = \{y_1, y_2, \dots, y_K\}$  from a  $K$ -dimensional Dirichlet distribution with parameter  $\alpha$ . Dirichlet distribution is commonly used as conjugate prior distribution of categorical distribution. From the Dirichlet distribution, we can sample prior probabilities  $\{y_1, y_2, \dots, y_K\}$ , where  $y_i \in (0, 1)$  and  $\sum_{i=1}^K y_i = 1$ .  $\alpha$  is referred to as the concentration parameter,

**Algorithm 2:** Data Synthesis of *OPT-SYN*.**INPUT:**

The substitute model  $f_s^t$  at iteration  $t$   
 Number of synthetic data  $S$   
 Number of output classes  $K$   
 Number of optimization iterations  $m$

**OUTPUT:** A set of optimized data samples  $\mathcal{X}$ 

```

1:  $\mathcal{X} \leftarrow \emptyset$ 
2: for  $i \leftarrow 1, S$  do
3:   // Generate a  $K$ -dimensional random parameter  $\alpha$ 
   // from a Gaussian distribution
4:    $\alpha \sim \mathcal{N}(0, 1)$ 
5:   // Sample a prediction vector  $\mathbf{y}$  from a Dirichlet
   // distribution
6:    $\mathbf{y} \sim D(K, \alpha)$ 
7:   // Initialize a data sample  $\mathbf{x}$  from Gaussian
   // distribution
8:    $\mathbf{x} \sim \mathcal{N}(0, 1)$ 
9:   // Minimize (6) for  $m$  iterations
10:   $\mathbf{x}^* \leftarrow \arg \min_{\mathbf{x}} \text{LCE}(f_s^{(t)}(\mathbf{x}), \mathbf{y})$ 
11:  // Add  $\mathbf{x}^*$  to set  $\mathcal{X}$ 
12:   $\mathcal{X} \leftarrow \mathcal{X} \cup \mathbf{x}^*$ 
13: end for
14: return:  $\mathcal{X}$ .
```

which controls the distribution. The probability density function of Dirichlet distribution  $Dir(K, \alpha)$  can be calculated by:

$$f(y_1, y_2, \dots, y_K, \alpha) = \frac{1}{B(\alpha)} \prod_{i=1}^K y_i^{\alpha_i - 1}, \quad (7)$$

where  $B(\alpha)$  denotes the gamma function and  $\sum y_1, y_2, \dots, y_K = 1$ . In the experiment, we randomly sample the parameter  $\alpha$  from a Gaussian distribution:  $\alpha \sim \mathcal{N}(0, 1)$  to explore the possible Dirichlet distribution.

Given the prediction vector  $\mathbf{y}$ , we synthesize data  $\mathbf{x}$  by iteratively minimizing the objective function 6. The goal is to generate a data sample  $\mathbf{x}^*$  that  $f_s^{(t)}$  predicts  $\mathbf{x}^*$  close to  $\mathbf{y}$ . An adaptive gradient-based optimization algorithm, Adam [10], is applied to optimize the objective function iteratively.

#### IV. EVALUATION OF *ES ATTACK*

In this section, we evaluate our proposed *ES Attack* on three different neural networks and four image classification datasets. We compare our results with two baseline attacks. Moreover, we investigate the data synthesized during the attacks in terms of data quality and diversity.

##### A. Experiment Setup

1) *Victim Models and Datasets:* In our experiments, we evaluate model stealing attacks on four image classification datasets: MNIST, KMNIST, SVHN, CIFAR-10. The MNIST dataset [13] contains 70,000 28-by-28 Gy images of 10 digits. We use 60,000 images for training and 10,000 images for testing following the original train/test split in the MNIST dataset. Kuzushiji-MNIST

(KMNIST) [14] is a similar dataset to MNIST, containing 70,000 28-by-28 grey images of 10 Hiragana characters. We use 60,000 images for training and 10,000 images for testing. The SVHN [15] dataset consists of 60,000 32-by-32 RGB images from house numbers (10 classes from 0 to 9) in the Google Street View dataset. The CIFAR10 [16] dataset contains 60,000 32-by-32 RGB images with 10 classes.

We train three types of DNN models on four datasets and use them as the victim models in our experiments. We train LeNet5 [13] on the MNIST [13] and KMNIST [14] datasets. We train ResNet18 [17] and ResNet34 [17] on the SVHN [15] and CIFAR10 [16] datasets. LeNet5 is trained for 30 epochs using an SGD optimizer with a learning rate of 0.1 on the MNIST and KMNIST datasets. We train ResNet18 and ResNet34 for 200 epochs with an initial learning rate of 0.1 on the SVHN and CIFAR10 datasets. We reduce the learning rate by 10 after 80 and 120 epochs. We select the models with the highest test accuracies as the victim models.

2) *Settings of ES Attack:* For *DNN-SYN*, we input a 100-dimensional random latent vector and a one-hot label vector into DNN-based generator  $G$ . The substitute model  $f_s$  and DNN-based generator  $G$  is trained by an Adam optimizer with a learning rate of 0.001.  $f_s$  and  $G$  are trained alternatively for 2,000 epochs each on the MNIST, KMNIST, and SVHN dataset ( $N = 2000, M = 1$ ), and 15,000 epochs each on the CIFAR10 dataset ( $N = 15000, M = 1$ ).

For *OPT-SYN*, we synthesize data for 30 iterations ( $M = 30$ ) in each stealing epoch using an Adam optimizer with a learning rate of 0.01. We train the adversary model for 10 epochs on the synthetic dataset ( $M = 10$ ). We repeat the stealing for 200 epochs on the MNIST, KMNIST, and SVHN dataset ( $N = 200$ ), and 1,500 epochs on the CIFAR10 dataset ( $N = 1500$ ). To speed up the stealing process, we augment the synthetic dataset by random horizontal flip, horizontal shift, and adding Gaussian noise.

3) *Baseline Model Stealing Attacks:* We compare *ES Attack* with two baseline attacks - model stealing using randomly generated data and auxiliary data. First, if the adversary has no knowledge of the victim's training data, randomly generated data could be the only dataset the adversary can leverage. We form a random dataset with the random data sampled from a Gaussian Distribution  $\mathcal{N}(0, 1)$  and their prediction from the victim model. We train our substitute model using the random dataset iteratively. Second, we consider public data as an auxiliary dataset. We use data samples from other public datasets and query the victim model with them. We construct an auxiliary dataset and train the substitute model on it. To make a fair comparison, we make sure that all the model stealing attacks, including two baseline attacks and two *ES Attacks* (*DNN-SYN* and *OPT-SYN*), train their substitute models for the same epochs.

##### B. Performance Evaluation

We evaluate the performance of *ES Attacks* using two data synthesis approaches and compare them with two baseline attacks. We report the accuracy of model stealing attacks in Table I. We compare the results with two baseline attacks that use randomly

TABLE I  
PERFORMANCE COMPARISON OF MODEL STEALING ATTACKS

Dataset	Model	Victim accuracy (%)	Attacks	Substitute accuracy (%)
SVHN	ResNet18	95.40	Random	50.71
			Auxiliary	74.84
			DNN-SYN	93.95
			OPT-SYN	<b>93.97</b>
SVHN	ResNet34	95.94	Random	60.95
			Auxiliary	82.00
			DNN-SYN	<b>93.34</b>
			OPT-SYN	93.19
CIFAR10	ResNet18	91.12	Random	11.72
			Auxiliary	48.73
			DNN-SYN	33.44
			OPT-SYN	<b>84.60</b>
CIFAR10	ResNet34	91.93	Random	14.45
			Auxiliary	38.55
			DNN-SYN	12.69
			OPT-SYN	<b>80.79</b>
MNIST	LeNet5	99.10	Random	72.18
			Auxiliary	<b>98.96</b>
			DNN-SYN	91.02
			OPT-SYN	92.03
KMNIST	LeNet5	95.67	Random	56.39
			Auxiliary	59.43
			DNN-SYN	90.37
			OPT-SYN	<b>90.37</b>

generated data (Random) and auxiliary data (Auxiliary) to steal the victim model.

From the evaluation, we observe that *OPT-SYN* can successfully steal the victim models over all the datasets and model architectures. Our proposed attacks achieve better performance compared with two baseline attacks. **On average, *OPT-SYN* improves the best accuracy by 44.57% compared to the best results of two baseline attacks.** *OPT-SYN* performs as well as *DNN-SYN* on the SVHN, MNIST, and KMNIST datasets. However, *DNN-SYN* cannot achieve a good performance on the CIFAR10 dataset, which is a more complex dataset and the generator  $G$  in *DNN-SYN* may still cause the mode collapse problem. Both our proposed attacks perform worse than the attacks using auxiliary data (KMNIST) on the MNIST dataset, which suggests that the auxiliary data can be used in the model stealing if the auxiliary data well-represent the target task and the data are available to the adversary. Note that we assume that the adversary has no knowledge of any victim's data, which means the adversary cannot evaluate the substitute model on a validation dataset and select the best substitute model during the attack. If the performance of the stealing attacks fluctuates, then the adversary cannot guarantee the best performance of the substitute model. The convergence of the substitute model is essential for stealing attacks without a validation dataset. Our experiments show that the performance of the substitute model converges after a few stealing epochs (Fig. 4). If the adversary has the knowledge of a validation dataset or the victim's

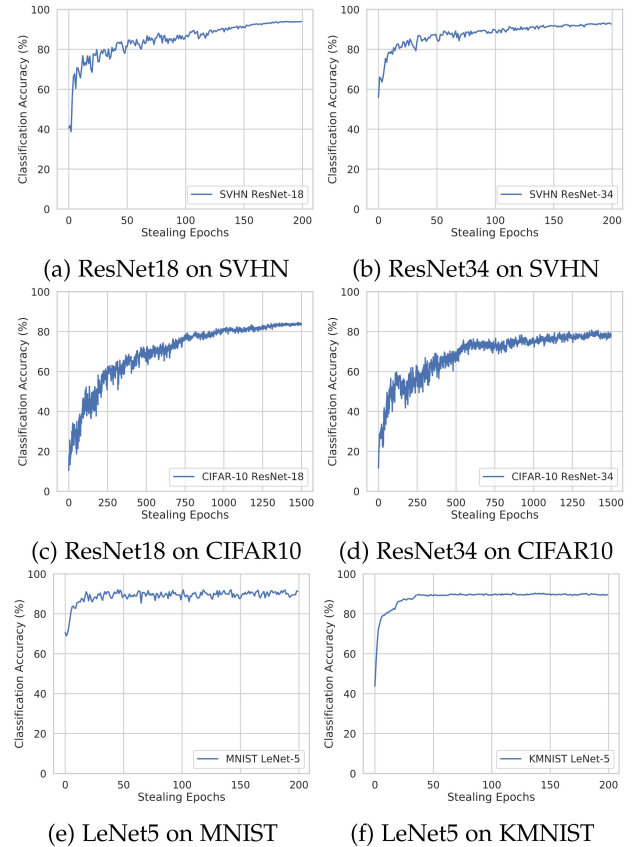


Fig. 4. Substitute model accuracy during attacks. (a) ResNet18 on SVHN. (b) ResNet34 on SVHN. (c) ResNet18 on CIFAR10. (d) ResNet34 on CIFAR10. (e) LeNet5 on MNIST. (f) LeNet5 on KMNIST.

test dataset  $\mathcal{D}_{\text{test}}$ , the adversary will achieve the best accuracy. Otherwise, the adversary will use the substitute model in the last stealing epoch ( $t = N$ ). We observe the subtle difference between the best accuracy and the last accuracy achieved by the substitute model (0.79% difference on average for *OPT-SYN* and 1.53% for *DNN-SYN*). The stable convergence suggests that our proposed attacks do not rely on a validation dataset.

We find that model stealing attacks do not require a large query budget in real-world settings. The query to the victim model only occurs in the E-Step, where the adversary uses the synthetic data to get the model prediction. In our experiments, to steal the victim model trained on the MNIST, KMNIST, and SVHN dataset using *OPT-SYN*, the adversary only needs to pay \$30 K for all the required queries (around 120 M queries) according to the pricing of Amazon AWS [18]. For the CIFAR10 dataset, it costs \$187.5 K to steal the victim model (around 750 M queries). The expenses are much less than hiring ML experts and collecting data from scratch. This indicates that the adversary can replicate the function of the existing MLaaS models with much less cost than the victim's actual cost of establishing a machine learning service.

TABLE II  
ES ATTACK USING DIFFERENT DNN ARCHITECTURES

Dataset	Victim model	Substitute model	Victim accuracy (%)	Substitute accuracy (%)
MNIST	ResNet18	LeNet5	99.31	97.28
	LeNet5	ResNet18	92.03	98.13
SVHN	ResNet18	ResNet34	95.40	94.64
	ResNet34	ResNet18	95.94	94.03
CIFAR10	ResNet18	ResNet34	91.12	82.54
	ResNet34	ResNet18	91.93	62.73

### C. Sensitivity Analysis of DNN Architectures

In this section, we consider a more realistic scenario where the adversary has no knowledge of the victim model’s architecture. The adversary may choose a different architecture of the substitute model from that of the victim model. Thus, we investigate when the adversary uses a different neural network architecture from the victim. For the MNIST dataset, we use ResNet18 for the victim model and LeNet5 for the substitute model, and vice versa. For the SVHN and CIFAR10 datasets, we consider ResNet18 and ResNet34 as the victim model and the substitute model’s architecture. Because *OPT-SYN* outperforms *DNN-SYN* in most model stealing attacks, we use *OPT-SYN* to evaluate the sensitivity of DNN architectures.

From Table II, we do not observe a significant performance loss due to different DNN architectures for the MNIST and SVHN dataset, compared with using the same architecture. For the CIFAR10 dataset, we observe subtle performance loss using ResNet34 to steal ResNet18 models. The only degradation of performance occurs when the adversary uses a small model ResNet18 to steal a large model ResNet34. We believe the degradation is due to the gap between the size of the victim model and the substitute model. We find similar performance degradation in many other tasks using knowledge distillation. The performance of the student model (substitute model in our paper) will be degraded if there is a gap between student and teacher (victim) [19]. The adversary can easily avoid performance degradation by selecting a large model. From our experiments, if the adversary chooses a DNN model with the same size or the large size compared with the victim model, the adversary will be able to steal a substitute model with high accuracy.

### D. Convergence of ES Attack

Fig. 4 illustrates the convergence of *ES Attack*. We observe that the accuracy of the substitute model always converges at the end of the stealing. We observe the subtle difference between the best accuracy and the last accuracy achieved by the substitute model (0.79% difference on average for *OPT-SYN* and 1.53% for *DNN-SYN*). The stable convergence at the end of the attacks suggests that our proposed attacks do not rely on a test dataset. Hence, the adversary can successfully steal the victim model even without knowing test data, which suggests the practicality

of *ES Attack* and substantially raises the severity of model stealing attacks.

### E. Quality Analysis of Synthetic Data

In this section, we investigate the quality of synthetic data. Fig. 5 shows examples of the synthetic data used in model stealing. We compare them with the victim’s training data. Humans cannot recognize these images, yet our substitute model can be well-trained using these synthetic data.

Therefore, we further investigate synthetic data in terms of quality and diversity. Inspired by the measurements for GANs, we use Inception Score (IS) [20] and Fréchet Inception Distance (FID) [21] to evaluate the synthetic data. In the experiments, we observe that the synthetic data achieves better quality and higher diversity compared to the auxiliary data.

**Inception Score (IS)** was originally proposed to measure the quality of generated images using a pre-trained Inception-V3 network. In our experiments, we replaced the Inception-V3 network with the victim models. To be consistent with the concept, we keep Inception Score as the name of our metric. Given the prediction provided by the victim models, Inception Score compares the conditional prediction distribution with the marginal prediction distribution:

$$\text{IS} = \exp(\mathbb{E}_{\mathbf{x}} D_{\text{KL}}(p(\mathbf{y}|\mathbf{x})||p(\mathbf{y}))), \quad (8)$$

where  $D_{\text{KL}}$  denotes the KL divergence. A high Inception Score indicates: 1) generated images are meaningful and predictable to the victim model, so that  $p(\mathbf{y}|\mathbf{x})$  has low entropy; 2) generated images are diverse, so that  $p(\mathbf{y})$  has high entropy.

**Fréchet Inception Distance (FID)** was proposed to improve IS by comparing the intermediate features of the victim model. FID models the real images and the synthesis data as two multivariate Gaussian distributions and calculates the distance between the two distributions:

$$\text{FID} = \|\mu_t - \mu_s\|_2^2 + \text{Tr} \left( \Sigma_t + \Sigma_s - 2(\Sigma_t \Sigma_s)^{\frac{1}{2}} \right), \quad (9)$$

where  $(\mu_t, \Sigma_t)$  and  $(\mu_s, \Sigma_s)$  denote the mean and covariance matrix of intermediate features predicted using training data and synthesis data.  $\text{Tr}(\cdot)$  denotes the trace of a matrix. A low FID indicates better image quality and diversity. FID is shown to be more consistent with human perception than IS [21] and more robust to mode collapse [22]. In our experiments, we used the features from the layer before the last linear layer and compared our synthesis data with the training data using FID. These two metrics are widely used to measure the quality of generated images.

We compare the four types of data and report the average values of IS and FID in Table III: 1) victim’s training dataset, 2) auxiliary dataset used in the baseline attack, 3) random data generated in the first initialization epoch ( $t = 0$ ), 4) the synthetic data generated in the last stealing epoch ( $t = N$ ). The value of FID is evaluated by comparing the data with the victim’s training data. From our analysis, we find that synthetic data usually achieves better quality and high diversity than the auxiliary data (higher IS value and lower FID value). On average over six settings, synthetic data  $\mathcal{D}_{syn}^{(1)}$  achieves 60.58% higher IS values



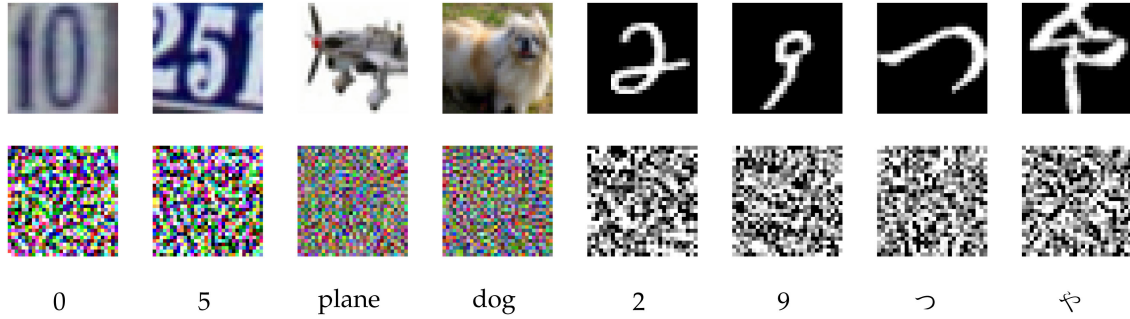


Fig. 5. Victim’s training data vs. synthetic data. We synthesize images using *OPT-SYN* with the best substitute model. We compare them with the victim’s training data. First row: images in the victim’s training dataset. Second row: images in the synthetic dataset. Third Row: corresponding labels of images. Other examples are presented in Appendix (Fig. 10, 11, 12, 13).

TABLE III  
ANALYSIS OF SYNTHETIC DATA USING IS AND FID

Dataset	Model	Victim $\mathcal{D}_{\text{train}}$		Auxiliary $\mathcal{D}_{\text{aux}}$		Random $\mathcal{D}_{\text{syn}}^{(0)}$		Synthetic $\mathcal{D}_{\text{syn}}^{(N)}$	
		IS	FID	IS	FID	IS	FID	IS	FID
MNIST	LeNet5	9.86	4.22	274.80	1.22	500.16	4.63	257.55	
KMNIST	LeNet5	9.96	4.75	1073.92	1.78	1498.75	4.70	962.47	
CIFAR10	ResNet18	6.89	2.58	5.34	3.45	5.82	4.32	3.31	
CIFAR10	ResNet34	7.64	4.16	18.16	6.75	14.88	6.65	18.29	
SVHN	ResNet18	6.89	2.24	7.62	3.45	5.82	4.32	3.31	
SVHN	ResNet34	7.64	4.16	18.16	6.75	14.88	6.65	18.29	

and 27.64% lower FID values than the auxiliary data  $\mathcal{D}_{\text{aux}}$ , which suggests better quality and higher diversity of synthetic images. The victim’s training data always achieve the highest IS value: the training data is the best representation of the input space among data we investigate. The Random Data are always the worst data due to the low IS values and high FID values.

#### F. Further Attacks: A Case Study on Black-Box Adversarial Attacks

DNNs are vulnerable to adversarial examples, a slight modification on the original data sample that can easily fool DNNs [23]–[26]. Black-box adversarial attacks assume that the adversary can only access the output of the model victim instead of its internal information, which is an emerging topic in adversarial attacks. After *ES Attack*, the adversary has full knowledge of the stolen substitute model. Hence, in this section, we try to answer the following question: Can the adversary leverage the knowledge to conduct adversarial attacks against the victim model?

Transferability is commonly used to conduct black-box adversarial attacks, by training a surrogate model to transfer the adversarial attacks [27], [28]. In the experiments, we demonstrate how model stealing facilitates black-box adversarial attacks through transferability. The evaluation of the black-box adversarial attack is outlined as follows: 1) Steal the victim model using *ES Attack* and get a substitute model; 2) Perform a white-box  $\ell_\infty$ -PGD attack against the substitute model and generate adversarial examples; 3) Evaluate the generated adversarial examples on the victim model.

We evaluate the black-box adversarial attack against the victim model using the substitute model. We implement a white-box  $\ell_\infty$ -PGD attacks [25] and leverage the transferability of adversarial examples to conduct the attack. PGD attack is an iterative gradient-based attack in a white-box setting and has been proven effective in many machine learning models and tasks.

In the experiment, we use the test dataset  $\mathcal{D}_{\text{test}}$  as our evaluation dataset. We follow the adversarial settings in [25] and consider the untargeted adversarial attacks, where adversarial examples can be classified as any classes other than the ground-truth class. For the MNIST and KMNIST dataset, we run 40 iterations of  $\ell_\infty$ -PGD attack with a step size of 0.01. We set the maximal perturbation size as 0.3. For the SVHN and CIFAR10 dataset, we run 20 iterations with a step size of  $2/255$ . The maximal perturbation size is set as  $8/255$ .

We report the success rate of adversarial attacks against our substitute model and the victim model (transferring attack) in Table IV. We compare the success rate of three adversarial attacks: 1) white-box attacks against the victim model, 2) white-box attacks against the substitute model, and 3) black-box attacks against the victim model via transferring. For the third attack, we evaluate the adversarial examples generated against substitute model (white-box attacks) on the victim model. We show the performance of the white-box  $\ell_\infty$ -PGD attack against the victim model as well.

From the experimental results, the black-box adversarial attacks using the substitute model can achieve the same success rate as the white-box, which suggests that the substitute models can transfer the adversarial examples to the victim model successfully. Almost all black-box adversarial attacks achieve



TABLE IV  
SUCCESS RATES OF ADVERSARIAL ATTACKS

Dataset	Model	Attack success rate (%)		
		white-box victim model	white-box substitute model	black-box victim model
SVHN	ResNet18	99.95	99.94	98.71
SVHN	ResNet34	99.93	99.90	98.21
CIFAR10	ResNet18	100.00	100.00	93.60
CIFAR10	ResNet34	100.00	100.00	100.00
MNIST	LeNet5	86.07	99.57	92.14
KMNIST	LeNet5	66.44	99.72	98.99




Fig. 6. Evaluation of rounding prediction on MNIST.

high accuracy rates (over 90%). We observe that the success rates of the black-box attacks against the victim models are less than that of the white-box attacks against the substitute models, but the change is subtle. Hence, most adversarial examples can be transferred from the substitute model to the victim model. Surprisingly, the black-box attacks against the victim model perform even better than the white-box attacks against the victim model on the MNIST and KMNIST dataset.

## V. COUNTERMEASURES OF MODEL STEALING

In this section, we discuss the defense strategies of MLaaS providers and evaluate their effectiveness. Given the good performance of *OPT-SYN* on all the datasets, we evaluate three countermeasures against *OPT-SYN*. We find that the countermeasures are ineffective in defending or detecting proposed *OPT-SYN*.

### A. Rounding Prediction

The MLaaS providers fix the decimals of the output prediction and zero-out the rest to provide only the necessary information. For example, if we round the prediction with 2 decimals, then  $\text{round}(0.2474, r = 2) = 0.25$ . We deploy rounding predictions with 2 decimals as a defensive strategy. Our experiments show that none of the model stealing attacks are affected by rounding to two decimals (Table V). On average, the best accuracy of the substitute model even increases by 0.55% and the last accuracy only decreases by 0.30%.

We further investigate the impact of rounding decimals on the *ES Attack*. Figs. 6 and 7 show the results of experiments with class probabilities rounded to 0-5 decimals. We compare the after-rounding classification accuracy of the substitute model




Fig. 7. Evaluation of rounding prediction on KMNIST.

and the victim model. Class probabilities rounded to 2 to 5 decimals have no effect on the adversary's success. When rounding further to 1 decimal, the attack is weakened, but still successful. When we round the precision to 0 decimal - the victim model only outputs 0 or 1 - the attack is further weakened, but still can predict in most cases (over 80% accuracy). We observe that rounded information brings the uncertainty of the prediction, while this uncertainty sometimes will slightly improve the training.

### B. Top-K Prediction

Instead of providing the predictions of all the classes, MLaaS providers can provide partial information - predictions of  $K$  classes with the highest probabilities. Since this defense can be easily detected by the adversary, we assume the adversary is aware of the top- $K$  defenses equipped by the MLaaS provider and will try to circumvent such defense. Therefore, the adversary can slightly change the attack by making up the missing probabilities of the rest classes. The adversary remains the probabilities of the Top- $K$  classes and fills up the rest classes with the same probabilities. For example, given a prediction output of  $[0.5, 0.02, 0.3, 0.02, 0.15, 0.01]$ , by using Top-2 defense, MLaaS provider can hide the predictions of eight classes and only respond with the prediction of  $[0.5, 0.0, 0.3, 0.0, 0.0, 0.0]$ . By knowing the top- $k$  defense, The adversary can then convert the predictions to  $[0.5, 0.05, 0.3, 0.05, 0.05, 0.05]$  and resume *ES Attack*.

From the experiments, we observe that Top-1 prediction will not affect much on most datasets (Table V). For the MNIST and KMNIST datasets, we find that the accuracy of the substitute model even gets improved. Top-1 prediction is only effective in preventing model stealing on the CIFAR10 dataset. However, we believe Top-1 prediction is a very strong defense, which will also affect normal users by providing very limited information. On average, the best accuracy of the substitute model with Top-1 prediction is only decreased by 2.86%. In addition, we investigate the impact of probabilities with different numbers ( $K$ ) of classes on our model stealing attacks (Figs. 8 and 9). The performance of model stealing attacks is not decreased with fewer classes providing probabilities (small  $K$ ). We find our attacks are minimally impacted by reducing the informativeness of black-box predictions in the response.

TABLE V  
EVALUATION OF DEFENSES AGAINST MODEL STEALING

Dataset	Model	Accuracy without defense (%)	Accuracy with rounding (%)	Accuracy with Top-K (%)
SVHN	ResNet34	93.19	92.58	88.76
CIFAR10	ResNet34	80.79	80.31	69.64
MNIST	LeNet5	92.03	96.69	95.43
KMNIST	LeNet5	90.37	90.03	91.11




Fig. 8. Evaluation of top-K prediction on MNIST.

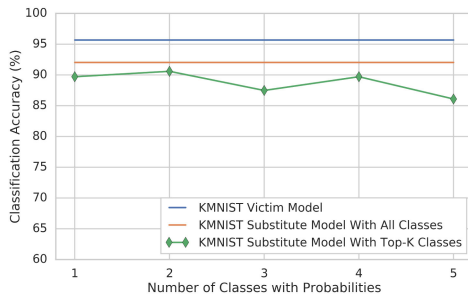


Fig. 9. Evaluation of top-K prediction on KMNIST.

### C. Anomaly Detection

Anomaly detection identifies the abnormal queries sent from users and detects the abnormal behavior that deviates from normal ones. For example, PRADA assumes that the distance between normal queries follows Gaussian distribution and detects the abnormal queries [9]. By evaluating how deviated the distance from Gaussian distribution, PRADA detects model stealing attacks. For the details of PRADA, we refer readers to [9]. We evaluate the effectiveness of anomaly detection using PRADA against *OPT-SYN*. We analyzed 300,000 image samples from the first five stealing epochs on the MNIST dataset. None of the synthetic images can be detected by PRADA. We believe that because the images are generated starting from the Gaussian distribution, the distances between queried images are too small to be detected by PRADA. Moreover, we find it is not practical for MLaaS providers to deploy a PRADA detector due to its high response time. In our experiments, it takes about 33 hours to process 300,000 images (2.46 images per second on average). With more images to be detected, the average response time will be further increased. Therefore, PRADA is ineffective and infeasible to detect the proposed *ES Attack*.

## VI. RELATED WORK

### A. Model Stealing Attacks

Several studies have been proposed for model stealing attacks. Tramer *et al.* investigated stealing model parameters using equation solving [2]. However, this approach is hard to extend to DNNs, which contains a larger number of than conventional machine learning models do. Papernot *et al.* proposed a similar framework to steal DNNs by training a substitute model [7]. Their goal is to approximate the victim model's decision boundaries to facilitate the adversarial attacks rather than to maximize the substitute model's accuracy. Thus their substitute model achieves a much lower classification accuracy compared to our work. In addition, to generate adversarial examples, their approach requires a small set of inputs that represents the input domain. In our work, we eliminate this requirement, where the adversary does not need to have prior knowledge, making the attacks more feasible in the real world. From the experimental results, the stolen model from *ES Attack* achieves a higher accuracy compared to that from [7]. Existing model stealing attacks against DNNs require an auxiliary dataset. Orekondy *et al.* proposed stealing attacks that assume access to a large dataset and use active learning to select the best samples to query [3]. Correia-Silva *et al.* leveraged public datasets from the same task domain but with different distributions to steal DNNs. Different from these works, we assume the adversary does not have any auxiliary data related to the task domain. The experiments show that *ES Attack* can achieve comparable performance, compared with the attacks using auxiliary datasets.

Zhou *et al.* used the same assumption with our work - unknown victim's training data and leveraged a generative model to synthesize the training data [29]. However, the goal of DaST is to perform a successful adversarial attack, which is different from ours. In this paper, we aim to improve the prediction performance of the substitute model. Accordingly, the substitute model trained by DaST achieves much lower accuracy than *ES Attack*. MAZE investigated a similar problem with our work, namely data-free model stealing attack [30]. To address the problem of unknown victim's training data, MAZE tried to solve the same challenge as our work, that is in generating synthetic data, the gradient for updating the synthetic data cannot be backpropagated using the victim model. MAZE addressed this issue by approximating the gradients from the victim model using zeroth-order gradient estimation, which is widely used in black-box adversarial attacks, whereas in our work, we generate the gradients by using the substitute model as a proxy of the victim model. Both approaches achieve comparable attacking




Fig. 10. The SVHN dataset.

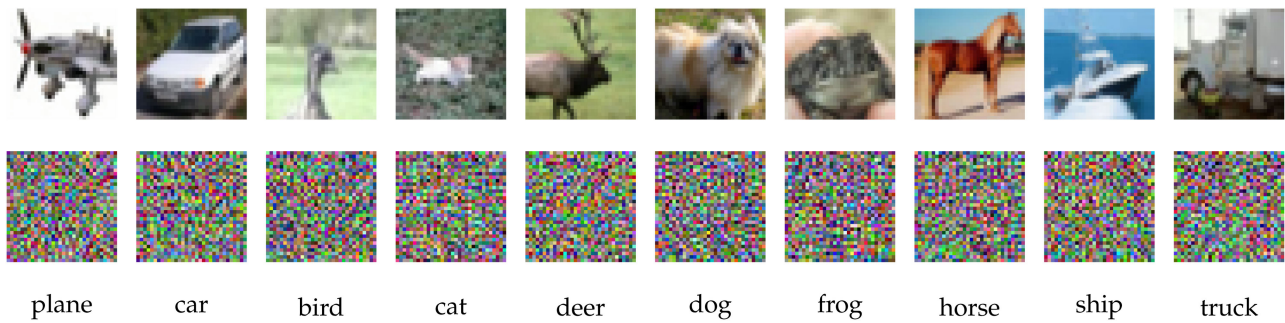


Fig. 11. The CIFAR-10 dataset.




Fig. 12. The MNIST dataset.



Fig. 13. The KMNIST dataset.

performance. In addition, the two approaches are orthogonal and could be further integrated together for better performance. We will explore the new approach benefiting from *ES Attack* and *MAZE* in the future.

### B. Model Stealing Defenses

Several detection approaches have been proposed for model stealing attacks. Juuti *et al.* detected the deviated distribution of queries from normal behavior [9]. Similarly, Kesarwani



*et al.* proposed a detection tool that uses information gain to measure the model learning rate by users with the increasing number of queries [31]. The learning rate is measured to the coverage of the input feature space in the presence of collusion. We evaluate [9] in our experiments and find that the detection approach is ineffective for our model stealing attacks.

### C. Knowledge Distillation

In our model stealing attacks, we use distillation to transfer knowledge from the victim model to the substitute model. Knowledge distillation is widely used in model compression by transferring the knowledge from one model (teacher model) to another (student model) [32], [33]. Most knowledge distillation approaches require the knowledge of training data. Recently, knowledge distillation without training data has recently been investigated [34]–[38] when the training data is infeasible due to large data size or privacy concerns. However, these data-free knowledge distillation approaches cannot be used for model stealing since the adversary cannot acquire the required information. For example, the model gradient is required to update the parameters of the generator in [34]. Similarly, model parameters are required to calculate class similarity in [36] or to calculate the feature map statistics in [37] and batch normalization statistics in [38]. Therefore, beyond data-free knowledge distillation, we introduce two data synthesis approaches in *ES Attack* to construct a novel data-free model stealing attack.

## VII. CONCLUSION

We demonstrated that our attacks successfully stole various DNNs from the MLaaS providers without any data hurdles. Even without the knowledge on the victim’s dataset, *ES Attack* outperforms the two baseline attacks by 44.57% on average of four datasets in terms of best accuracy. Our experimental results illustrated the better quality and higher diversity of the generated synthetic data compared with the auxiliary data, which benefits *ES Attack*. In addition, most existing defenses are ineffective to prevent *ES Attack*, where new countermeasures should be provided. Moreover, the stolen model can be used to conduct black-box adversarial attacks against the victim model, and sometimes the black-box attack achieves higher success rates compared with the white-box attack. In this paper, we only target image classification tasks and small datasets. We will extend our work to other machine learning tasks and more complex datasets such as image segmentation and ImageNet in the future.


Further, this paper investigated a critical challenge in the existing machine learning services. By successfully launching the model stealing attack, the adversary can provide the same machine learning service with a much lower price compared to MLaaS provider due to the low cost of model stealing attacks. Moreover, model stealing attacks may facilitate further serious attacks that have been found in recent machine learning security

research. We hope the severity of model stealing attacks can attract the attention of the community and encourage researchers in academia and industry to investigate effective countermeasures.

## REFERENCES

- [1] T. Hunt, C. Song, R. Shokri, V. Shmatikov, and E. Witchel, “Chiron: Privacy-preserving machine learning as a service,” 2018, *arXiv:1803.05961*.
- [2] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, “Stealing machine learning models via prediction apis,” in *Proc. 25th USENIX Secur. Symp.*, 2016, pp. 601–618.
- [3] T. Orekondy, B. Schiele, and M. Fritz, “Knockoff nets: Stealing functionality of black-box models,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4954–4963.
- [4] J. R. C. da Silva, R. F. Berriel, C. Badue, A. F. de Souza, and T. Oliveira-Santos, “Copycat CNN: Stealing knowledge by persuading confession with random non-labeled data,” in *Proc. Int. Joint Conf. Neural Netw.*, 2018, pp. 1–8.
- [5] S. Pal, Y. Gupta, A. Shukla, A. Kanade, S. K. Shevade, and V. Ganapathy, “A framework for the extraction of deep neural networks by leveraging public data,” 2019, *arXiv:1905.09165*.
- [6] H. Yu, K. Yang, T. Zhang, Y. Tsai, T. Ho, and Y. Jin, “CloudLeak: Large-scale deep learning models stealing through adversarial examples,” in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2020.
- [7] N. Papernot, P. D. McDaniel, I. J. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical black-box attacks against machine learning,” in *Proc. ACM Asia Conf. Comput. Commun. Secur.*, 2017, pp. 506–519.
- [8] B. Wang and N. Z. Gong, “Stealing hyperparameters in machine learning,” in *Proc. IEEE Symp. Secur. Privacy*, 2018, pp. 36–52.
- [9] M. Juuti, S. Szyller, S. Marchal, and N. Asokan, “PRADA: Protecting against deep neural network model stealing attacks,” in *Proc. IEEE Eur. Symp. Secur. Privacy*, 2019, pp. 512–527.
- [10] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. 3rd Int. Conf. Learn. Representations*, 2015.
- [11] A. Odena, C. Olah, and J. Shlens, “Conditional image synthesis with auxiliary classifier gans,” in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 2642–2651.
- [12] Q. Mao, H. Lee, H. Tseng, S. Ma, and M. Yang, “Mode seeking generative adversarial networks for diverse image synthesis,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 1429–1437.
- [13] Y. LeCun *et al.*, “Gradient-based learning applied to document recognition,” *Proc. IEEE*, Nov. 1998, vol. 86, no. 11, pp. 2278–2324.
- [14] T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, and D. Ha, “Deep learning for classical japanese literature,” 2018, *arXiv:1812.01718*.
- [15] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, “Reading digits in natural images with unsupervised feature learning,” *NIPS Workshop Deep Learn. Unsupervised Feature Learn.*, 2011.
- [16] A. Krizhevsky *et al.*, “Learning multiple layers of features from tiny images,” *Tech. Rep.*, 2009.
- [17] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [18] Amazon, “Amazon rekognition pricing,” 2020. [Online]. Available: <https://aws.amazon.com/rekognition/pricing/>
- [19] S. Mirzadeh, M. Farajtabar, A. Li, and H. Ghasemzadeh, “Improved knowledge distillation via teacher assistant: Bridging the gap between student and teacher,” 2019, *arXiv:1902.03393*.
- [20] T. Salimans, I. J. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, “Improved techniques for training gans,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 2234–2242.
- [21] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, “Gans trained by a two time-scale update rule converge to a local nash equilibrium,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 6626–6637.
- [22] M. Lucic, K. Kurach, M. Michalski, S. Gelly, and O. Bousquet, “Are GANs created equal? A large-scale study,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 700–709.
- [23] C. Szegedy *et al.*, “Intriguing properties of neural networks,” in *Proc. 2nd Int. Conf. Learn. Representations*, 2014.
- [24] N. Carlini and D. A. Wagner, “Towards evaluating the robustness of neural networks,” in *Proc. IEEE Symp. Secur. Privacy*, 2017, pp. 39–57.

- [25] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *Proc. 6th Int. Conf. Learn. Representations*, 2018.
- [26] X. Yuan, P. He, Q. Zhu, and X. Li, "Adversarial examples: Attacks and defenses for deep learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 9, pp. 2805–2824, Jan. 2019.
- [27] N. Papernot, P. D. McDaniel, and I. J. Goodfellow, "Transferability in machine learning: From phenomena to black-box attacks using adversarial samples," 2016, *arXiv:1605.07277*.
- [28] Y. Liu, X. Chen, C. Liu, and D. Song, "Delving into transferable adversarial examples and black-box attacks," in *Proc. 5th Int. Conf. Learn. Representations*, 2017.
- [29] M. Zhou, J. Wu, Y. Liu, S. Liu, and C. Zhu, "Dast: Data-free substitute training for adversarial attacks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 234–243.
- [30] S. Kariyappa, A. Prakash, and M. K. Qureshi, "MAZE: Data-free model stealing attack using zeroth-order gradient estimation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 13814–13823.
- [31] M. Kesarwani, B. Mukhoty, V. Arya, and S. Mehta, "Model extraction warning in MLaaS paradigm," in *Proc. 34th Annu. Comput. Secur. Appl. Conf.*, 2018, pp. 371–380.
- [32] G. E. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015, *arXiv:1503.02531*.
- [33] C. Buciluă, R. Caruana, and A. Niculescu-Mizil, "Model compression," in *Proc. 12th ACM Int. Conf. Knowl. Discov. Data Mining*, 2006, pp. 535–541.
- [34] H. Chen *et al.*, "Data-free learning of student networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 3513–3521.
- [35] R. G. Lopes, S. Fenu, and T. Starner, "Data-free knowledge distillation for deep neural networks," 2017, *arXiv:1710.07535*.
- [36] G. K. Nayak, K. R. Mopuri, V. Shaj, V. B. Radhakrishnan, and A. Chakraborty, "Zero-shot knowledge distillation in deep networks," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 4743–4751.
- [37] H. Yin *et al.*, "Dreaming to distill: Data-free knowledge transfer via deepinversion," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 8712–8721.
- [38] M. Haroush, I. Hubara, E. Hoffer, and D. Soudry, "The knowledge within: Methods for data-free model compression," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 8491–8499.



**Xiaoyong Yuan** (Member, IEEE) received the B.S. degree from Fudan University, Shanghai, China, in 2012, the M.E. degree from Peking University, Beijing, China, in 2015, and the Ph.D. degree from the University of Florida, Gainesville, FL, USA, in 2020.

Dr. Yuan is currently an Assistant Professor with the College of Computing, Michigan Technological University, Houghton, MI, USA. His research interests mainly include the fields of machine learning, deep learning, security and privacy, and cloud computing.



**Leah Ding** (Member, IEEE) received the Ph.D. degree in electrical engineering from the University at Buffalo, Buffalo, NY, USA, in 2013. She is currently an Associate Professor with the Department of Computer Science, American University, Washington, DC, USA. Her research interests include cyber security, machine learning, wireless and emerging networks.



**Lan Zhang** (Member, IEEE) received the B.S. and M.S. degrees from the University of Electronic Science and Technology of China, Chengdu, China, in 2013 and 2016, respectively, and the Ph.D. degree from the University of Florida, Gainesville, FL, USA, in 2020.

Dr. Zhang is currently an Assistant Professor with the Department of Electrical and Computer Engineering, Michigan Technological University, Houghton, MI, USA. Her research interests mainly include wireless communications and distributed machine learning in various cyber-physical systems and Internet of Things applications.



**Xiaolin Li** received the Ph.D. degree in computer engineering from Rutgers University, New Brunswick, NJ, USA.

Dr. Li is a Partner of Tongdun Technology, heading the AI Institute, a Chair Professor and the Chief Scientist of IBMC, Chinese Academy of Sciences. He was a Professor with the University of Florida, Gainesville, FL, USA, and founding Center Director of NSF Center for Big Learning. His research interests include deep learning, cloud computing, security & privacy, IoT, and intelligent medicine.



**Dapeng Oliver Wu** (Fellow, IEEE) received the Ph.D. degree in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, USA, in 2003. He is currently a Professor with the Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL, USA, and the Director of NSF Center for Big Learning. His research interests include networking, communications, signal processing, computer vision, machine learning, smart grid, and information and network security. He was the recipient of the University of Florida Term Professorship Award in 2017, University of Florida Research Foundation Professorship Award in 2009, AFOSR Young Investigator Program (YIP) Award in 2009, ONR Young Investigator Program (YIP) Award in 2008, NSF CAREER award in 2007, the IEEE Circuits and Systems for Video Technology (CSVT) Transactions Best Paper Award for Year 2001, and the best paper awards in IEEE GLOBECOM 2011 and International Conference on Quality of Service in Heterogeneous Wired/Wireless Networks (QShine) 2006.

He was the Editor in Chief of IEEE TRANSACTIONS ON NETWORK SCIENCE AND ENGINEERING, Editor-at-Large for IEEE OPEN JOURNAL OF THE COMMUNICATIONS SOCIETY, founding Editor-in-Chief of *Journal of Advances in Multimedia* and an Associate Editor for IEEE TRANSACTIONS ON CLOUD COMPUTING, IEEE TRANSACTIONS ON COMMUNICATIONS, IEEE TRANSACTIONS ON SIGNAL AND INFORMATION PROCESSING OVER NETWORKS, *IEEE Signal Processing Magazine*, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS and IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY.