

Exploring Query Efficient Data Generation Towards Data-Free Model Stealing in Hard Label Setting

Gaozheng Pei¹, Shaojie Lyu³, Ke Ma^{1*}, Pinci Yang¹, Qianqian Xu^{2*}, Yingfei Sun¹

¹School of Electronic, Electrical and Communication Engineering, UCAS, Beijing, China

²Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, CAS, Beijing, China

³Tencent Corporate

peigaozheng23@mails.ucas.ac.cn, shaojielyu@tencent.com, make@ucas.ac.cn, yangpinci23@mails.ucas.ac.cn, xuqianqian@ict.ac.cn, yfsun@ucas.ac.cn

Abstract

Data-free model stealing involves replicating the functionality of a target model into a substitute model without accessing the target model’s structure, parameters, or training data. The adversary can only access the target model’s predictions for generated samples. Once the substitute model closely approximates the behavior of the target model, attackers can exploit its white-box characteristics for subsequent malicious activities, such as adversarial attacks. Existing methods within cooperative game frameworks often produce samples with high confidence for the prediction of the substitute model, which makes it difficult for the substitute model to replicate the behavior of the target model. This paper presents a new data-free model stealing approach called Query Efficient Data Generation (QEDG). We introduce two distinct loss functions to ensure the generation of sufficient samples that closely and uniformly align with the target model’s decision boundary across multiple classes. Building on the limitation of current methods, which typically yield only one piece of supervised information per query, we propose the query-free sample augmentation that enables the acquisition of additional supervised information without increasing the number of queries. Motivated by theoretical analysis, we adopt the consistency rate metric, which more accurately evaluates the similarity between the substitute and target models. We conducted extensive experiments to verify the effectiveness of our proposed method, which achieved better performance with fewer queries compared to the state-of-the-art methods on the real MLaaS scenario and five datasets.

Introduction

The widespread adoption of cloud services by Microsoft (Kang et al. 2024), Google (Borra 2024), and Amazon (Wittig and Wittig 2023) suggest that machine learning as a service (MLaaS) (Ribeiro, Grolinger, and Capretz 2015) may become the predominant mode of interaction between human and models in the future. From the user’s perspective, these models operate as black boxes, providing feedback based exclusively on user inputs. This interaction method safeguards the privacy of MLaaS training data and ensures the security of the model. This prompts the inquiry: has the

security concern regarding MLaaS for classification been effectively mitigated?

The answer is negative. One severe security threat is model stealing, also known as model extraction (Liang et al. 2022b; Tramèr et al. 2016; Truong et al. 2021; Sanyal, Addepalli, and Babu 2022). In this process, an unauthorized party aims to replicate the decision of a target model by creating a substitute model. Attackers generate samples using a generative model, then acquire the target model’s predictions for these samples and use them as supervisory information to iteratively refine the substitute model, until either the query budget is exhausted or the behavior of the substitute model closely matches that of the target model. Once the substitute model emulates the behavior of the target model, attackers can utilize it for malicious purposes, including adversarial attacks. The current methods within cooperative game frameworks for stealing models still have a significant issue (Zhang et al. 2022; Zhang, Chen, and Lyu 2022). This can be understood by examining the entire process of model stealing. The attacker synthesizes features using pseudo labels via generator (Goodfellow et al. 2020). As shown in the upper panel of Figure 1, the generator mainly produces synthetic features by minimizing the classification loss function of the substitute model. This strategy confers excessive confidence in the substitute model. The target and substitute models would have the same prediction of these synthetic features. Correspondingly, the synthetic features with target model predictions could not guide the update of the substitute model effectively. As a result, simulating the target model’s decision-making requires a large number of synthetic features. Blindly increasing the number of queries to receive more feedback is unrealistic in real confrontation scenarios due to the risk of malicious behavior being detected (Juuti et al. 2019). Therefore, attackers should aim to generate the minimum number of synthetic features necessary, using these limited features to identify discrepancies between the target and substitute models. Ideally, these synthetic features should mainly occupy the disputed area, where the target and substitute models have different predictions of the same synthetic features.

To this end, we dissect the generation process of synthetic features and introduce a novel data-free model stealing method named Query Efficient Data Generation (QEDG). Firstly, we implement a harmony loss to ensure the gener-

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

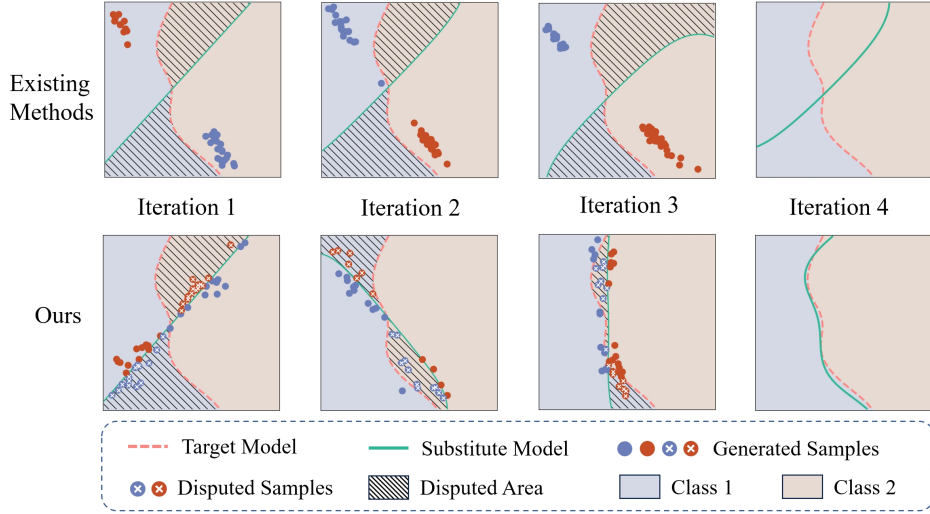


Figure 1: Compared to previous methods (top), our proposed approach (bottom) guides the generator to produce samples that are sufficiently close to the decision boundary of the substitute model. This results in more samples falling into the disputed area, thereby directing the substitute model to approach the target model in the correct direction.

ated samples closely approach the classification boundary of multiple classes. Secondly, we apply a diversity loss by enlarging the gap between samples within the same class, thereby achieving a uniform distribution of generated samples along the decision boundary of the substitute model. Furthermore, we propose a query-free augmentation, which allows for the acquisition of more diverse supervised information without increasing the number of queries to the target model. The main contributions of this paper are summarized as follows:

- By dissecting the traditional synthetic process, we reconstruct the principal of the generator to ensure that the generated samples are close to and uniformly distributed along the decision boundary of the substitute model which not only causes some samples to fall into disputed areas, but also ensures that the generated samples cover as large a decision boundary surface as possible.
- In order to make full use of the supervisory information obtained by interacting with the target model, a query-free sample augmentation is proposed to obtain multiple pieces of supervisory information through single query. Such a mechanism increases the diversity of samples.
- Motivated by theoretical analysis, we introduce the consistency rate metric, which more accurately evaluates the similarity between the substitute and target models. We perform extensive experiments on four datasets using three metrics, where the results show that our method achieves better performance with a smaller number of queries.

Proposed Framework

Threat Modelling

Here we will fully identify the potential threats of the adversary in data-free model stealing.

Goal. To implement the adversarial attack with partial knowledge of the victim, the adversary aims to acquire a substitute model that could imitate the inference of the target model. Let $\theta_T \in \mathbb{R}^{D_T}$ and $\theta_S \in \mathbb{R}^{D_S}$ be the target and substitute model parameters, where $D_T \neq D_S$. Given $x \in \mathbb{R}^D$ of a test sample $z = \{x, y^*\}$, the adversary hopes

$$T(x, \theta_T) = S(x, \theta_S), \quad (1)$$

where $T : \mathbb{R}^D \times \mathbb{R}^{D_T} \rightarrow \mathbb{R}$ is the decision function of the target model, and $S : \mathbb{R}^D \times \mathbb{R}^{D_S} \rightarrow \mathbb{R}$ corresponds to the substitute model. It is noteworthy that the predictions of the target and substitute model could be inconsistent with the true label as $y^* \neq T(x, \theta_T) = S(x, \theta_S)$.

Knowledge. Although θ_T is agnostic to the model stealing attackers, he/she can get the feedback of the target model:

$$\hat{y}_g \stackrel{\text{def}}{=} T(x_g, \theta_T) \quad (2)$$

with the synthetic sample x_g and construct $z_g = \{x_g, \hat{y}_g\}$ to train the substitute model. Note that here we are considering a stricter scenario where the target model can only return the label, rather than a probability vector. The data-free attacker does not possess more knowledge than normal users in MLaaS scenario.

Ability. The existing attackers (Zhang et al. 2022; Zhang, Chen, and Lyu 2022) generate x_g by minimizing the classification loss of the substitute model. This strategy confers excessive confidence in the substitute model. Such a behavior turns out to be a mirage. The target and substitute model could have the same prediction of x_g

$$S(x_g, \theta_S) = \hat{y}_g. \quad (3)$$




Figure 2: The confidence vectors of the points in the shaded area are displayed above. Our goal is to ensure that the generated sample points are as close to the decision boundary as possible while also minimizing intra-class similarity. This approach aims to distribute the generated samples as closely as possible along the decision boundary.

Yet the predictions of the substitute model would be still inconsistent with the target model on the true test samples

$$S(\mathbf{x}, \theta_S) \neq T(\mathbf{x}, \theta_T). \quad (4)$$

This issue requires a large number of synthetic samples $\{z_g\}$ to intimate the decision boundary of the target model. The “optimistic” attackers need to continuously interact with the target model to query $\{\hat{y}_g\}$ and construct $\{z_g\}$. However, excessive interactions with the target model can lead to the attacker’s samples x_g being identified as illegal inputs. In a real confrontation scenario, the adversary faces a paradox: he/she needs more feedback from the target model but the risk of his/her malicious behavior being detected increases. Different from the existing “optimistic” attackers, the proposed “pessimistic” adversary makes every effort to minimize the number of queries, *i.e.* generate as few synthetic samples as possible. Consequently, the adversary pays attention to the synthetic x_g with different predictions from the target and substitute models

$$S(x_g, \theta_S) \neq \hat{y}_g. \quad (5)$$

A Parsimonious Generation Process

The above threat modeling shows that minimizing the classification loss function of the substitute model is the main reason for the inefficiency of existing attack methods. Specifically, at the t step, with the given pseudo label $\mathbf{y} \in \{0, 1\}^K$ where K is the number of classes, the existing methods adopt a classification loss function (\mathcal{L}_{clf}) like the cross-entropy to generate the synthetic feature

$$\mathbf{x}_g^{(t)} \in \arg \min_{\mathbf{x}} \mathcal{L}_{\text{clf}} \left(S(\mathbf{x}, \theta_S^{(t)}), \mathbf{y} \right). \quad (6)$$

Then the attacker queries the target model $T(\cdot)$ about $\mathbf{x}_g^{(t)}$. The parameter of the substitute model will be updated by $\mathbf{z}_g^{(t)} = \{\mathbf{x}_g^{(t)}, \mathbf{y}_g^{(t)}\}$ as

$$\theta_S^{(t+1)} \in \arg \min_{\theta} \mathcal{L}_S \left(S(\mathbf{x}_g^{(t)}, \theta), \hat{\mathbf{y}}_g^{(t)} \right), \quad (7)$$

where \mathcal{L}_S is the objective function of the substitute model. When (3) holds with some occasional good synthetic samples, (7) would update θ_S very slowly. To reduce the number

of interactions, the adversary needs to efficiently discover the differences between the target model and the substitute model. This inspires us to dissect the generation process of the synthetic features. The synthetic features generated by the cross-entropy function are separated by the substitute model $\theta_S^{(t)}$. The corresponding predictions of the target and substitute models could be consistent (see the upper panel of Fig. 1). As a consequence, the “pessimistic” attacker tries to generate the disputed synthetic features with the given labels \mathbf{y} and the current model parameter $\theta_S^{(t)}$. The synthetic features $\mathbf{x}_g^{(t)}$ should lie as close as possible to the decision boundary of $S(\cdot, \theta_S^{(t)})$. Such synthetic features have a high probability in the disputed area as (5) if the substitute model can not imitate the target model.

We analyze the requirements of the desirable synthetic features through Figure 2. Generally speaking, the disputed example $\mathbf{x}_g^{(t)}$ for $S(\cdot)$ means that the prediction

$$\begin{aligned} S(\mathbf{x}_g^{(t)}) &\stackrel{\text{def}}{=} S(\mathbf{x}_g^{(t)}, \theta_S^{(t)}) \\ &= [s_1(\mathbf{x}_g^{(t)}), s_2(\mathbf{x}_g^{(t)}), \dots, s_K(\mathbf{x}_g^{(t)})] \end{aligned} \quad (8)$$

lacks the discriminative capability. Treating $S(\mathbf{x}_g^{(t)})$ as a group of random variables, we hope it holds a small variance. Let $\text{Var}(S(\mathbf{x}_g^{(t)}))$ be the variance of $S(\mathbf{x}_g^{(t)})$:

$$\begin{aligned} \text{Var} \left(S(\mathbf{x}_g^{(t)}) \right) \\ = \frac{1}{K} \sum_{k=1}^K \left(s_k(\mathbf{x}_g^{(t)}) - \bar{S}(\mathbf{x}_g^{(t)}) \right)^2, \end{aligned} \quad (9)$$

where $\bar{S}(\mathbf{x}_g^{(t)})$ be the mean of $S(\mathbf{x}_g^{(t)})$. Thus, we introduce (9) as the loss function to emphasize the harmony of the synthetic features, as comes below:

$$\mathcal{L}_{\text{harm}} = \frac{1}{N} \sum_{\mathbf{x}} \text{Var}(S(\mathbf{x})), \quad (10)$$

where N is the total number of the synthetic features. However, simply minimizing (10) can only generate the synthetic

features that lie in the junction of different classes (the center of Figure 2). Furthermore, we need to generate the synthetic features along the decision boundaries between any two classes. To extend the range of synthetic features, we try to increase the diversity of the synthetic features belonging to the same class. Specifically, the other loss function is designed as follows:

$$\mathcal{L}_{\text{div}} = -\frac{2}{N(N-1)} \sum_{i=1}^N \sum_{j=1, i \neq j}^N d(\mathbf{x}_i, \mathbf{x}_j), \quad (11)$$

where $d : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}_+$ is the distance or dissimilarity function of the synthetic features. Here we choose to use cosine similarity. The final object function of the generator is a combination of (6), (10), and (11) as

$$\mathcal{L}_G = \mathcal{L}_{\text{clf}} + \alpha \cdot \mathcal{L}_{\text{harm}} + \beta \cdot \mathcal{L}_{\text{div}}, \quad (12)$$

where α and β are hyper-parameters. Our strategy for selecting hyper-parameters is to ensure that the magnitudes of the three different loss functions are as close to the same scale as possible.

Query-free Augmentation and Substitute Model Training




Figure 3: The consistency rate of predicted labels by the target model for non-disputed samples before and after data augmentation.

Once we obtain the synthetic data \mathbf{x}_g using (12), we expect the outputs of the $T(\mathbf{x}_g)$ and $S(\mathbf{x}_g)$ to be as consistent as possible. However, using only one batch of features synthesized in the first stage of the current epoch, it will suffer from catastrophic forgetting (Binici et al. 2022; Do et al. 2022). The memory bank (**MB**) can effectively address this issue. Different from (Yuan et al. 2024) which only stores all previously synthesized features in the first stage, we also store the supervisory information from the target model. In the hard-label setting, the supervisory information refers to the label predicted by the target model for the features \mathbf{x}_g :

$$\mathbf{MB} = \mathbf{MB} \cup \{(\mathbf{x}_g, \hat{\mathbf{y}}_g)\}, \text{ where } \hat{\mathbf{y}}_g = T(\mathbf{x}_g). \quad (13)$$

We observe that HEE (Yuan et al. 2024) applies strong data augmentation to the features \mathbf{x}_g stored in the memory bank. However, strong data augmentation can change the original labels. This raises the question: is it possible to achieve feature diversity without altering the original labels? We conducted experiments on both grayscale datasets like MNIST

and colored images like CIFAR-10. For non-disputed samples, we applied three simple data augmentations, including horizontal flipping, vertical flipping, and rotation by a certain angle. In Figure 3, we found that the labels of non-disputed samples remained high consistent rate before and after these data augmentations. The intuition behind this experimental phenomenon is that since non-disputed samples are correctly classified, they are further away from the decision boundary of the target model, making them more robust to simple data augmentations. Therefore, we retain the original labels of the non-disputed samples after data augmentation, thereby introducing more supervised information without additional query costs. It is worth noting that these augmented samples are not stored in the memory bank.

With the samples stored in the memory bank and the samples after random augmentation, we can update the parameters θ_S of the substitute model. We assign different weights to the loss of disputed samples and non-disputed samples, respectively. The loss function of the substitute model $S(\cdot)$ at t step is as follows:

$$\begin{aligned} \mathcal{L}_S^{(t)}(\mathbf{x}_g, \mathbf{y}_g; \theta) \\ = \begin{cases} \gamma \cdot \mathcal{L}(\mathbf{x}_g, \mathbf{y}_g; \theta), & \text{if } S(\mathbf{x}_g, \theta_S^{(t-1)}) \neq \hat{\mathbf{y}}_g, \\ \mathcal{L}(\mathbf{x}_g, \mathbf{y}_g; \theta), & \text{otherwise,} \end{cases} \end{aligned} \quad (14)$$

where

$$\mathcal{L}(\mathbf{x}_g, \mathbf{y}_g, \theta) = \mathcal{L}_{\text{clf}}(S(\mathbf{x}_g, \theta), \hat{\mathbf{y}}_g), \quad (15)$$

and γ is the hyper-parameter to control the importance of the disputed samples. The whole process of the proposed method is summarized as Algorithm 1.

Algorithm 1: Query-efficient Data Generation

Input: random noise \mathbf{z} , generator $G(\cdot)$, target model $T(\cdot)$, memory bank (**MB**), maximal number of queries $Q \geq 1$, maximal iteration of generator $E \geq 1$.

Output: θ_S .

Initialization: $q = 0$, $\mathbf{MB} = \emptyset$;

while $q \leq Q$ **do**

Given pseudo labels of a batch $\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\}$;

for $t = 1, 2, \dots, E$ **do**

$\mathbf{x}_g = G(\mathbf{z}, \theta_G)$;

Update θ_G using (12);

end

$\hat{\mathbf{y}}_g = T(\mathbf{x}_g)$ and Update q ;

$\mathbf{MB} = \mathbf{MB} \cup \{(\mathbf{x}_g, \hat{\mathbf{y}}_g)\}$;

for $(\mathbf{x}_g, \hat{\mathbf{y}}_g) \in \mathbf{MB}$ **do**

if $S(\mathbf{x}_g) == \hat{\mathbf{y}}_g$ **then**

$\mathbf{x}'_g = \text{Augmentation}(\mathbf{x}_g)$;

Update θ_S using (14) with $(\mathbf{x}'_g, \hat{\mathbf{y}}_g)$;

else

Update θ_S using (14) with $(\mathbf{x}_g, \hat{\mathbf{y}}_g)$;

end

end

end

Theoretical Analysis

This part establishes the query complexity of the proposed active model stealing attack. Compared to the existing passive attackers, the following theorem states that the active attacker can easily steal any target model whose correct probability is greater than or equal to $1/2 + c$ with some constant $c > 0$ for all synthetic features. We show that the consistency rate between the target and substitute models increases by only a logarithmic multiplicative factor of the normal active learning query complexity.

An active learning algorithm $(\mathcal{A}, \mathcal{T})$ refers to an interaction between two players - the oracle \mathcal{T} and the learner \mathcal{A} . \mathcal{A} chooses $x \in \mathcal{X}$ and sends it to \mathcal{T} , who responds with $y \in \mathcal{Y}$. (x, y) is then used by \mathcal{A} to obtain \hat{f} from a hypothesis class \mathcal{F} , which satisfies $\hat{f}(x) = y$. We assume that the oracle \mathcal{T} adopts $f^* \in \mathcal{F}$ to response x as $y := f^*(x)$. In this setting (Alabdulmohsin, Gao, and Zhang 2015; Chen, Hassani, and Karbasi 2017; Zhang, Shen, and Awasthi 2020), the different between \hat{f} and f^* is measured by

$$\text{error}(\hat{f}) = \Pr_{x \in \mathcal{X}} \{ \hat{f}(x) \neq f^*(x) \}. \quad (16)$$

Then $q(\epsilon, \delta)$ denotes the query complexity of \mathcal{A} as \mathcal{A} chooses $\hat{f} \in \mathcal{F}$ such that $\text{error}(\hat{f}) \leq \epsilon$ with probability at least $1 - \delta$.

Theorem. *Let \mathcal{F} be a hypothesis class and $(\mathcal{A}, \mathcal{T})$ refer to an active learning algorithm as described above with the query complexity of $q(\epsilon, \delta)$. Suppose that an adversary \mathcal{S} disguises as \mathcal{A} but he/she can only receive imperfect feedback as*

$$\phi(f^*, x) \stackrel{\text{def}}{=} \Pr\{y_g \neq f^*(x)\} > 0, \forall x \in \mathcal{X}, \quad (17)$$

where y_g is the random variable that represents the feedback of \mathcal{T} to the \mathcal{S} 's query x . If $\max_x \phi(f^*, x) < 1/2$, \mathcal{A} could obtain $\hat{f} \in \mathcal{F}$ such that

$$\Pr\{\text{error}(\hat{f}) \leq \epsilon\} \geq 1 - 2\delta, \quad (18)$$

with query complexity

$$Q = \frac{8}{(1 - 2 \cdot \max_x \phi(f^*, x))^2} \cdot q(\epsilon, \delta) \cdot \ln \frac{q(\epsilon, \delta)}{\delta}. \quad (19)$$

It is noteworthy that (17) is a common situation for model stealing adversary. On the one hand, the oracle only exists in an ideal setting and the MLaaS API could provide some wrong predictions with the given queries. On the other hand, the potential defense mechanism would perturb the predictions of the target model to prevent model stealing. We provide the proof details in the supplementary materials.

Experiments

Experiment Setup

Datasets and Model Architectures We conducted experiments on five well-known datasets: MNIST, FMNIST, SVHN, CIFAR10, CIFAR100. We unified the target models

Type	Method	Query	MNIST			FMNIST		
			FGSM	BIM	PGD	FGSM	BIM	PGD
Untar.	HEE	500K	40.09	45.00	47.35	73.65	67.46	69.41
	DaST	250K	12.40	4.14	5.68	67.46	29.26	46.68
	DFME	250K	27.12	14.24	15.85	70.46	47.16	57.16
	IDEAL	20K	30.38	14.35	16.87	75.76	64.59	70.88
	DFTA	20K	34.67	33.97	39.19	83.40	71.59	80.98
	Ours	5K	41.37	45.50	50.21	80.86	77.30	81.33
Tar.	HEE	500K	15.50	15.34	16.29	14.79	15.75	15.86
	DaST	250K	10.27	9.97	9.59	9.97	10.28	9.48
	DFME	250K	12.09	11.23	11.76	10.02	9.84	10.14
	IDEAL	20K	12.80	11.26	11.85	13.89	12.57	12.73
	DFTA	20K	14.34	13.15	14.15	16.11	17.91	21.33
	Ours	5K	15.85	16.62	17.58	14.19	21.42	21.60

Table 1: ASR(%) comparisons between our method and baselines. We highlight the highest and second-highest values in orange and blue, respectively. ‘‘Untar.’’ refers to non-targeted attacks, ‘‘Tar.’’ refers to targeted attacks.

of different methods, as well as the substitute models of different methods, noting that the target models and substitute models are different. We separately tested different model structures, including ResNet (He et al. 2016), VGG (Simonyan and Zisserman 2015), and our own designed CNN model. Furthermore, we adopted the architecture described in (Zhang et al. 2022) and employed the same generator as used in StyleGAN (Karras, Laine, and Aila 2019).

Baselines We compare our approach with the following state-of-the-art (SOTA) methods. DaST (Zhou et al. 2020) and DFME (Truong et al. 2021) and HEE (Yuan et al. 2024) are three data-free model stealing methods that do not take into account the number of queries. DFTA (Zhang et al. 2022) and IDEAL (Zhang, Chen, and Lyu 2022) are query efficient and achieve improvements over the above methods.

Evaluation Metrics We employ three metrics for evaluation. Accuracy measures the precision with which the substitute model performs predictions on the actual test set. Motivated by theoretical analysis and (16), we also consider consistency metric. To further exclude the possibility that agreement may occur by chance between target model and substitute model. we adopt Cohen’s kappa (McHugh 2012). It measures the agreement between target model and substitute model. This metric is not affected by the accuracy of the target model. The calculation formula for Cohen’s kappa is as follows:

$$\kappa = \frac{p_o - p_e}{1 - p_e} \quad (20)$$

where p_e represents the relative consistency between the two models and is calculated by summing of the products of the marginal probabilities for each class. p_o represents the accidental consistency between the two models and is calculated by dividing the sum of the diagonal elements of the confusion matrix by the total number of samples. To further test the consistency between the substitute model and the target model, we also test the attack success rate. We employ three common attack methods to generate adversarial examples: FGSM (Goodfellow, Shlens, and Szegedy 2014), BIM

Method	MNIST			FMNIST			SVHN			CIFAR10			CIFAR100		
	Query	Acc	Con	Query	Acc	Con	Query	Acc	Con	Query	Acc	Con	Query	Acc	Con
HEE	500k	43.31	37.22	500k	27.07	19.26	1000k	39.75	31.36	1000k	34.82	28.69	2M	15.87	16.73
DaST	250k	11.62	1.14	250k	16.43	5.48	250k	24.08	13.62	250k	11.92	2.18	2M	17.26	19.38
DFME	250k	22.55	12.38	250k	10.12	1.38	250k	20.16	9.09	250k	12.35	2.94	2M	20.17	22.14
IDEAL	20k	58.82	54.26	20k	22.48	14.05	50k	60.71	54.93	50k	33.88	27.47	400k	18.21	19.85
DFTA	20k	64.42	65.61	20k	46.95	41.48	50k	57.42	51.68	50k	39.45	34.03	400k	20.08	21.93
Ours	5K	80.24	78.03	5k	54.95	50.86	30K	66.58	62.35	30K	43.86	39.45	200k	22.32	23.60

Table 2: Accuracy(%) and consistency(%) comparisons between our proposed method and baselines over five datasets. We highlight the highest and second-highest values in orange and blue, respectively.

Type	Method	Query	FGSM	SVHN BIM	PGD	Query	FGSM	CIFAR10 BIM	PGD	Query	FGSM	CIFAR100 BIM	PGD
UnTar.	HEE	1000K	72.48	86.07	86.17	1000K	72.12	74.20	72.36	2M	87.34	88.20	88.28
	DaST	250K	58.69	46.80	49.22	250K	58.61	43.48	47.24	2M	91.49	91.59	92.40
	DFME	250K	62.72	67.33	67.20	250K	59.76	48.99	50.29	2M	92.07	91.75	92.99
	IDEAL	50K	78.93	85.46	84.51	50K	68.24	70.58	70.86	400K	91.66	89.24	91.15
	DFTA	50K	72.14	86.95	87.30	50K	71.82	75.97	78.16	400K	92.20	90.17	91.16
	Ours	30K	73.90	89.31	87.83	30K	72.95	76.51	78.64	200K	91.94	91.84	93.07
Tar.	HEE	1000K	22.22	55.64	55.78	1000K	21.87	31.23	29.73	2M	2.04	12.03	11.45
	DaST	250K	11.40	22.24	21.85	1000K	9.62	10.22	10.56	2M	1.63	10.88	10.47
	DFME	250K	10.42	12.56	12.55	1000K	10.65	10.27	10.42	2M	1.61	11.26	10.34
	IDEAL	50K	20.35	62.26	58.20	50K	15.11	30.49	30.83	400K	1.83	9.04	8.81
	DFTA	50K	17.60	54.96	51.76	50K	14.95	29.96	27.93	400K	1.68	9.96	9.63
	Ours	30K	18.87	64.28	61.19	30K	15.46	31.88	31.33	200K	1.68	12.66	12.08

Table 3: ASR(%) comparisons between our method and baselines. We highlight the highest and second-highest values in orange and blue, respectively. “Untar.” refers to non-targeted attacks. “Tar.” refers to targeted attacks.

(Kurakin, Goodfellow, and Bengio 2017), and PGD (Madry et al. 2018).

Implementation details We train the substitute model using an SGD optimizer with a batch size of 256, an initial learning rate of 0.01, a momentum of 0.9, and no weight decay. For the generator, we adopt a batch size of 256, utilizing an adam optimizer (Kingma and Ba 2015) with a fixed learning rate of 0.001. For both MNIST and FMNIST, the perturbation bound is respectively set to 0.3, 0.2, and 0.1 with a step size α of 0.01. For SVHN, CIFAR10 and CIFAR100, the perturbation bound is set to 32/255, 20/255, 8/255 with a step size α of 0.01. In untargeted attacks, we only attack images correctly classified by the target model. In targeted attacks, we only attack images correctly classified by the target model with labels different from their original labels. This is as the same as DFTA (Zhang et al. 2022).

Experiment Results

Attacks on the Microsoft Azure Online Model¹ We conducted experiments in real-world scenarios using Microsoft’s Azure API service. This service restricts users to uploading their own datasets, and we utilized the MNIST dataset for tests, consistent with (Zhang, Chen, and Lyu 2022). The architecture and parameters of the model provided by the service are not disclosed. For the substitute

¹<https://www.customvision.ai>




Figure 4: The accuracy of our method and other methods varies with the number of queries on Microsoft Azure.

model, we employed ResNet18. From the Figure 4, we can observe that our method achieves an accuracy comparable to DFTA’s accuracy at 100k queries when using only 20k queries. Moreover, with a sufficient number of queries, our method is capable of achieving even higher accuracy.

Experiments on MNIST and FMNIST We used ResNet18 as the substitute model and ResNet34 as the target model. We tested three evaluation metrics: consistency

rate, accuracy, and attack success rate. The perturbation bound was set to 0.3. The attack success rates for other perturbation bounds can be found in the appendix. We also tested results when both the target and substitute models were CNN structures. These results are also in the appendix. Table 2 shows the accuracy and consistency rate. Our method outperforms others by a significant margin in both accuracy and consistency rate, even with fewer queries. This is because our method generates more samples in the disputed region, enabling each query to effectively guide the substitute model’s update. In contrast, other methods generate overly confident samples for the substitute model. Few samples fall into the disputed area, resulting in significant query waste. We test the success rates of targeted and untargeted attacks using three attack methods. Table 1 shows that regardless of the attack method, our method achieves the highest attack success rates with the fewest queries. This indicates that the decision boundary learned by our substitute model is most similar to the target model’s.

Experiments on SVHN, CIFAR10 and CIFAR-100 Due to the simplicity of grayscale images, we extended our experiments to more complex datasets like SVHN and CIFAR10, CIFAR100 which consist of color images. We used the same model structure and evaluation metric. The perturbation bound was set to 32/255. The attack success rates for other perturbation bounds can be found in the appendix. We also tested results with the target model set to VGG19. These results are also in the appendix. We used more queries on both datasets because the decision boundaries of the target models are more complex. Table 2 shows that our method surpasses others in consistency rate and accuracy, requiring only 30k queries compared to 50k or even 250k needed by other methods. In addition, we also test the transfer attack success rate under different attack methods. From the table 3, it can be seen that our method achieves a higher transfer attack success rate in most cases with fewer query number. We also plotted line charts showing the relationship between accuracy, consistency, and the number of queries in the appendix. Figure shows that the accuracy and consistency of other methods fluctuate, while ours steadily increase.

Sensitivity Analysis We conducted a hyperparameter

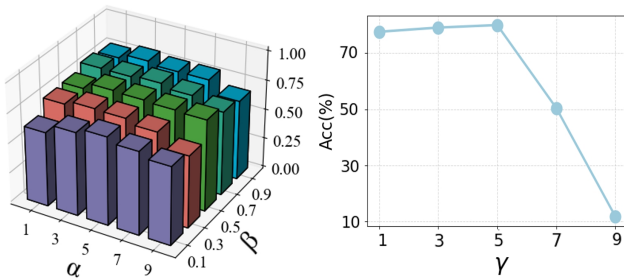


Figure 5: Sensitivity analysis about hyperparameters α and β and γ .

analysis on the MNIST dataset and found that optimal performance can only be achieved when the two loss functions are on the same order of magnitude. We set $\alpha = 5$ and

$\beta = 0.7$. Regarding the selection of γ , the experiments revealed that assigning a higher weight to contentious samples can improve the performance of our method. However, if the weight is too high, the performance will quickly decline. Here we set $\gamma = 5$.

Ablation Study

QA	$\mathcal{L}_{\text{harm}}$	\mathcal{L}_{div}	CIFAR10		
			ACC	Con	ASR
✓	✓	✗	31.2	25.7	69.3
✓	✗	✓	41.9	37.3	71.6
✗	✓	✓	38.8	33.4	70.5
✓	✓	✓	43.8	39.4	72.9

Table 4: Ablation study on the proposed components. The highest score on each column is shown in **bold**.

In this section, we conducted comprehensive ablation experiments on CIFAR10 to assess the significance of each component of our method, with the results presented in Table 4. Our approach primarily consists of three improvement components: query-free augmentation (QA), $\mathcal{L}_{\text{harm}}$, and \mathcal{L}_{div} . We tested the experimental effects of adding or removing each module separately, resulting in a total of six sets of experiments. Under otherwise identical conditions, using \mathcal{L}_{div} yielded higher performance compared to using QA and $\mathcal{L}_{\text{harm}}$, and we achieved the best performance when all three components are used simultaneously. Because $\mathcal{L}_{\text{harm}}$ prefers the centers of decision surfaces with multiple classes, if only $\mathcal{L}_{\text{harm}}$ appears without \mathcal{L}_{div} , it leads to negative optimization of the loss function. This is evidenced by the experiments, where the results with only $\mathcal{L}_{\text{harm}}$ and without \mathcal{L}_{div} are the worst. The role of QA is to increase the diversity of samples without increasing the number of queries. Experimental results also indicate that having QA can significantly improve the model’s performance.

Conclusion and Limitations

To reduce the number of queries required in model stealing, we dissect the existing synthetic process and reconstruct the principal of the generator to ensure that the synthetic samples are close to and uniformly distributed along the decision boundary of the substitute model. The generated samples should obtain different predictions of the target and substitute models and cover as large a decision boundary surface as possible. Furthermore, a query-free sample augmentation is proposed to make full use of the supervisory information obtained by interacting with the target model. Then we establish the query complexity of the proposed active adversary with imperfect feedback of the target model. We further introduce the consistency rate metric, which more accurately evaluates the similarity between the substitute and target models. Our method outperformed the current state-of-the-art competitors on the real **MLaaS** scenario and five datasets. However, both our method and existing approaches generate samples that lack any semantic information.

Acknowledgements

This work was supported in part by the National Science and Technology Major Project 2022ZD0119204, in part by National Natural Science Foundation of China: 62236008, U21B2038, U23B2051, 62122075 and 62376257, in part by Youth Innovation Promotion Association CAS, in part by the Strategic Priority Research Program of the Chinese Academy of Sciences, Grant No. XDB0680000, in part by the Innovation Funding of ICT, CAS under Grant No.E000000, in part by the Tencent Marketing Solution Rhino-Bird Focused Research Program, and in part by Xiaomi Young Talents Program.

References

- Alabdulmohsin, I. M.; Gao, X.; and Zhang, X. 2015. Efficient Active Learning of Halfspaces via Query Synthesis. In *AAAI Conference on Artificial Intelligence*, 2483–2489.
- Andriushchenko, M.; Croce, F.; Flammarion, N.; and Hein, M. 2020. Square attack: a query-efficient black-box adversarial attack via random search. In *European Conference on Computer Vision*, 484–501.
- Binici, K.; Aggarwal, S.; Pham, N. T.; Leman, K.; and Mitra, T. 2022. Robust and resource-efficient data-free knowledge distillation by generative pseudo replay. In *AAAI Conference on Artificial Intelligence*, 6089–6096.
- Borra, P. 2024. A Survey of Google Cloud Platform (GCP): Features, Services, and Applications. *International Journal of Advanced Research in Science, Communication and Technology*, 191–199.
- Brendel, W.; Rauber, J.; and Bethge, M. 2018. Decision-Based Adversarial Attacks: Reliable Attacks Against Black-Box Machine Learning Models. In *International Conference on Learning Representations*.
- Chandrasekaran, V.; Chaudhuri, K.; Giacomelli, I.; Jha, S.; and Yan, S. 2020. Exploring Connections between Active Learning and Model Extraction. In *USENIX Security Symposium*, 1309–1326.
- Chen, L.; Hassani, S. H.; and Karbasi, A. 2017. Near-Optimal Active Learning of Halfspaces via Query Synthesis in the Noisy Setting. In *AAAI Conference on Artificial Intelligence*, 1798–1804.
- Chen, P.-Y.; Zhang, H.; Sharma, Y.; Yi, J.; and Hsieh, C.-J. 2017. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *ACM Workshop on Artificial Intelligence and Security*, 15–26.
- Chen, S.; Huang, Z.; Tao, Q.; Wu, Y.; Xie, C.; and Huang, X. 2022. Adversarial Attack on Attackers: Post-Process to Mitigate Black-Box Score-Based Query Attacks. In *Annual Conference on Neural Information Processing Systems*, 14929–14943.
- Do, K.; Le, T. H.; Nguyen, D.; Nguyen, D.; Harikumar, H.; Tran, T.; Rana, S.; and Venkatesh, S. 2022. Momentum adversarial distillation: Handling large distribution shifts in data-free knowledge distillation. In *Annual Conference on Neural Information Processing Systems*, 10055–10067.
- Gao, S.; Jia, X.; Ren, X.; Tsang, I.; and Guo, Q. 2025. Boosting transferability in vision-language attacks via diversification along the intersection region of adversarial trajectory. In *European Conference on Computer Vision*, 442–460. Springer.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2014. Explaining and harnessing adversarial examples. *arXiv*.
- Gu, J.; Jia, X.; de Jorge, P.; Yu, W.; Liu, X.; Ma, A.; Xun, Y.; Hu, A.; Khakzar, A.; Li, Z.; et al. 2023. A survey on transferability of adversarial examples across deep neural networks. *arXiv preprint arXiv:2310.17626*.
- Guo, C.; Gardner, J.; You, Y.; Wilson, A. G.; and Weinberger, K. 2019. Simple black-box adversarial attacks. In *International Conference on Machine Learning*, 2484–2493.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 770–778.
- He, Z.; Rakin, A. S.; and Fan, D. 2019. Parametric Noise Injection: Trainable Randomness to Improve Deep Neural Network Robustness Against Adversarial Attack. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Jia, X.; Gu, J.; Huang, Y.; Qin, S.; Guo, Q.; Liu, Y.; and Cao, X. 2023. TranSegPGD: Improving Transferability of Adversarial Examples on Semantic Segmentation. *arXiv preprint arXiv:2312.02207*.
- Jia, X.; Zhang, Y.; Wu, B.; Ma, K.; Wang, J.; and Cao, X. 2022. LAS-AT: adversarial training with learnable attack strategy. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13398–13408.
- Juuti, M.; Szyller, S.; Marchal, S.; and Asokan, N. 2019. PRADA: protecting against DNN model stealing attacks. In *European Symposium on Security and Privacy*, 512–527.
- Kang, M. S.; Park, S. Y.; Chung, M.-A.; and Han, D.-h. 2024. Microsoft Azure. In *NO-CODE AI: Concepts and Applications in Machine Learning, Visualization, and Cloud Platforms*, 255–262. World Scientific.
- Kariyappa, S.; Prakash, A.; and Qureshi, M. K. 2021. Maze: Data-free model stealing attack using zeroth-order gradient estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13814–13823.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4401–4410.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*.
- Kurakin, A.; Goodfellow, I. J.; and Bengio, S. 2017. Adversarial Machine Learning at Scale. In *International Conference on Learning Representations*.

- Liang, S.; Li, L.; Fan, Y.; Jia, X.; Li, J.; Wu, B.; and Cao, X. 2022a. A large-scale multiple-objective method for black-box attack against object detection. In *European Conference on Computer Vision*, 619–636. Springer.
- Liang, S.; Liu, A.; Liang, J.; Li, L.; Bai, Y.; and Cao, X. 2022b. Imitated detectors: Stealing knowledge of black-box object detectors. In *30th ACM International Conference on Multimedia*, 4839–4847.
- Liang, S.; Wu, B.; Fan, Y.; Wei, X.; and Cao, X. 2021. Parallel Rectangle Flip Attack: A Query-based Black-box Attack against Object Detection. In *2021 IEEE/CVF International Conference on Computer Vision*, 7677–7687. IEEE Computer Society.
- Liang, S.; Zhu, M.; Liu, A.; Wu, B.; Cao, X.; and Chang, E.-C. 2024. Badclip: Dual-embedding guided backdoor attack on multimodal contrastive learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24645–24654.
- Liu, X.; Jia, X.; Gu, J.; Xun, Y.; Liang, S.; and Cao, X. 2024. Does few-shot learning suffer from backdoor attacks? In *AAAI Conference on Artificial Intelligence*, 19893–19901.
- Ma, K.; Xu, Q.; Zeng, J.; Cao, X.; and Huang, Q. 2021. Poisoning attack against estimating from pairwise comparisons. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10): 6393–6408.
- Ma, K.; Xu, Q.; Zeng, J.; Li, G.; Cao, X.; and Huang, Q. 2022. A tale of hodgerank and spectral method: Target attack against rank aggregation is the fixed point of adversarial game. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4): 4090–4108.
- Ma, K.; Xu, Q.; Zeng, J.; Liu, W.; Cao, X.; Sun, Y.; and Huang, Q. 2024. Sequential manipulation against rank aggregation: theory and algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*.
- McHugh, M. L. 2012. Interrater reliability: the kappa statistic. *Biochemia Medica*, 22(3): 276–282.
- Pal, S.; Gupta, Y.; Shukla, A.; Kanade, A.; Shevade, S.; and Ganapathy, V. 2020. Activethief: Model extraction using active learning and unannotated public data. In *AAAI Conference on Artificial Intelligence*, 865–872.
- Ribeiro, M.; Grolinger, K.; and Capretz, M. A. M. 2015. MLaaS: Machine Learning as a Service. In *IEEE International Conference on Machine Learning and Applications*, 896–902.
- Sanyal, S.; Addepalli, S.; and Babu, R. V. 2022. Towards data-free model stealing in a hard label setting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15284–15293.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In Bengio, Y.; and LeCun, Y., eds., *International Conference on Learning Representations*.
- Sun, X.; Cheng, G.; Li, H.; Pei, L.; and Han, J. 2022. Exploring effective data for surrogate training towards black-box attack. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15355–15364.
- Tramèr, F.; Zhang, F.; Juels, A.; Reiter, M. K.; and Ristenpart, T. 2016. Stealing Machine Learning Models via Prediction APIs. In *USENIX Security Symposium*, 601–618.
- Truong, J.-B.; Maini, P.; Walls, R. J.; and Papernot, N. 2021. Data-free model extraction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4771–4780.
- Wang, W.; Yin, B.; Yao, T.; Zhang, L.; Fu, Y.; Ding, S.; Li, J.; Huang, F.; and Xue, X. 2021a. Delving into data: Effectively substitute training for black-box attack. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4761–4770.
- Wang, X.; and He, K. 2021. Enhancing the transferability of adversarial attacks through variance tuning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1924–1933.
- Wang, X.; He, X.; Wang, J.; and He, K. 2021b. Admix: Enhancing the transferability of adversarial attacks. In *IEEE/CVF International Conference on Computer Vision*, 16158–16167.
- Wittig, A.; and Wittig, M. 2023. *Amazon Web Services in Action: An in-depth guide to AWS*. Simon and Schuster.
- Xie, C.; Zhang, Z.; Zhou, Y.; Bai, S.; Wang, J.; Ren, Z.; and Yuille, A. L. 2019. Improving transferability of adversarial examples with input diversity. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2730–2739.
- Yuan, X.; Chen, K.; Huang, W.; Zhang, J.; Zhang, W.; and Yu, N. 2024. Data-free hard-label robustness stealing attack. In *AAAI Conference on Artificial Intelligence*, 6853–6861.
- Zhang, C.; Shen, J.; and Awasthi, P. 2020. Efficient active learning of sparse halfspaces with arbitrary bounded noise. In *Annual Conference on Neural Information Processing Systems*, 7184–7197.
- Zhang, J.; Chen, C.; and Lyu, L. 2022. Ideal: Query-efficient data-free learning from black-box models. In *International Conference on Learning Representations*.
- Zhang, J.; Li, B.; Xu, J.; Wu, S.; Ding, S.; Zhang, L.; and Wu, C. 2022. Towards efficient data free black-box adversarial attack. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15115–15125.
- Zhou, M.; Wu, J.; Liu, Y.; Liu, S.; and Zhu, C. 2020. Dast: Data-free substitute training for adversarial attacks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 234–243.