

SCME: A Self-contrastive Method for Data-Free and Query-Limited Model Extraction Attack

Renyang Liu¹, Jinhong Zhang¹, Kwok-Yan Lam², Jun Zhao², and Wei Zhou¹(⊠)

¹ Yunnan University, Kunming, China {ryliu,jhnova}@mail.ynu.edu.cn, zwei@ynu.edu.cn
² Nanyang Technological University, Singapore, Singapore {kwokyan.lam,junzhao}@ntu.edu.sg

Abstract. Previous studies have revealed that artificial intelligence (AI) systems are vulnerable to adversarial attacks. Among them, model extraction attacks fool the target model by generating adversarial examples on a substitute model. The core of such an attack is training a substitute model as similar to the target model as possible, where the simulation process can be categorized in a data-dependent and data-free manner. Compared with the data-dependent method, the data-free one has been proven to be more practical in the real world since it trains the substitute model with synthesized data. However, the distribution of these fake data lacks diversity and cannot detect the decision boundary of the target model well, resulting in the dissatisfactory simulation effect. Besides, these data-free techniques need a vast number of queries to train the substitute model, increasing the time and computing consumption and the risk of exposure. To solve the aforementioned problems, in this paper, we propose a novel data-free model extraction method named SCME (Self-Contrastive Model Extraction), which considers both the inter- and intra-class diversity in synthesizing fake data. In addition, SCME introduces the Mixup operation to augment the fake data, which can explore the target model's decision boundary effectively and improve the simulating capacity. Extensive experiments show that the proposed method can yield diversified fake data. Moreover, our method has shown superiority in many different attack settings under the query-limited scenario, especially for untargeted attacks, the SCME outperforms SOTA methods by 11.43% on average for five baseline datasets.

Keywords: Adversarial Attacks \cdot Model Extraction Attacks \cdot Black-Box Attacks \cdot Model Robustness \cdot Information security

1 Introduction

Recently, Trusted AI, which contains fairness, trustworthiness and explainability, has received increasing attention and plays an essential role in the AI development process. The security of the AI models, however, is being doubted and has bought concerns in academia and industry. A lot of research has shown that AI models (including Machine Learning (ML) models and Deep Learning (DL) models) are vulnerable to adversarial examples [2], which are crafted by adding a virtually imperceptible perturbation to the benign input but can lead the well-trained AI model to make wrong decisions. For example, in the physical world, the attackers can maliciously alter traffic signs by sticking a small patch [12], changing the content style [4] and shooting a laser on it [5]. Although these modifications do not affect human senses, but can easily trick autonomous vehicles. Therefore, it is imperative to devise effective attack techniques to identify



Fig. 1. Synthetic example distribution, decision boundaries and whether the attack is successful. **Top left:** bad synthetic example distribution failed to fit the target model decision boundary. Bottom left: unfitting of the decision boundary leads to attack failure. **Top right:** good example distribution and decision boundary fit. Bottom right: good decision boundary.

the deficiencies of AI models beforehand in security-sensitive applications [13].

Existing adversarial attack methods on DL models can be categorized into white-box attacks and black-box attacks. In the white-box settings, the attackers can access the whole information of the target model, including weights, inner structures and gradients. In contrast, in the black-box one, the attackers have no permission to access the models' details but the final output [3,9]. With such rules, it is clear that black-box attacks are more challenging but practical in the physical world, where the attacker lacks details of the target models. To attack DL models in black-box settings more effectively, model extraction attacks have been proposed [17], which is implemented by training a substitute model and generating adversarial examples on such model to attack the target model successfully.

Most of the previous model extraction attacks [14, 15] concentrates on training the substitute model by querying the target model with real data, called data-dependent model extraction. However, it is infeasible to the physical world, where the adversary can not access the models' training data. As the counterpart, the data-free model extraction attack solved this problem by synthesizing fake data [20, 23]. In this scenario, the attackers use generators to synthesize the fake data to train substitute models. For launching attacks with a high success rate, as shown in Fig. 1, the decision boundary of the substitute model should maintain a very high similarity to the target black-box model. Besides, training the substitute model with synthetic data is challenging to the problem of *How* to generate valuable synthetic data for the substitute model training? Generally speaking, the synthetic examples should have the following two properties: 1) **inter-class diversity** and 2) **intra-class diversity**. The inter-class diversity means that the synthetic examples' categories classified by the target model should contain all the expected classes, while the intra-class diversity indicates that the examples should differ from each other, even they belong to the same category. However, existing methods [18,20,22,23] still suffer from the following two challenges: 1) They only consider inter-class diversity but ignore intra-class diversity, resulting in synthetic data not serving as well as real data. 2) The other is that the query data are generated in the substitute model's training process. However, once the substitute model is not well-trained, it can hardly provide the effective target model's decision information.

To solve these challenges mentioned above, in this paper, we propose a novel data-free model extraction method, named **S**elf-**C**ontrastive **M**odel **E**xtraction (SCME for short). SCME introduces the idea of contrastive learning [1] and proposes a self-contrastive mechanism to guide the training of the generator. Specifically, we design a self-contrastive loss to enlarge the distance of the substitute model's latent representation. Benefiting from this, the generator will be encouraged to synthesize more diversified fake data. Furthermore, SCME introduces the Mixup operation to interpolate two random images into a single one to build the query examples, which can improve the efficiency of the substitute model in learning the target model's decision boundaries in a model-independent manner. Extensive experiments illustrate SCME can synthesize fake data with diversity and improve the attack performance. Our contributions are summarized as follows:

- We propose a novel data-free model extraction method, called SCME, to generate efficient fake data for the substitute model training under querylimited settings.
- We use a self-contrastive mechanism to guide the generator to synthesize the fake data with inter- and intra-class diversity to help the substitute model imitate the target model efficiently.
- We introduce the Mixup into SCME, which can build fake data in a modelindependent manner, to detect decision boundaries of the target model effectiveness and further help the imitating processing.
- Extensive empirical results show the SCME's superiority in the synthetic diversified fake data and the adversarial examples' attack performance in query-limited situations.

2 Related Work

Previous researches contend that the DL models are sensitive to adversarial attacks, which can be classified into white-box and black-box. In white-box settings, the attackers can generate adversarial examples with a nearly 100% attack success rate because they can access the target model. The black-box attack, however, is more threatening to the DL models in various realistic applications because they do not need the models' details. Among them, the model extraction

attack [21] has received much attention recently due to its high attack performance.

The success of model extraction attacks relies heavily on adversarial examples' transferability, which means the adversarial examples generate on model A can also attack model B successfully. To implement such an attack, the attacker first trains a surrogate local model by simulating the target models' output. When the surrogate model is well-trained, it will have the same decision boundary as the target model, i.e., output the same results for the same input; this imitation process is called model extraction. However, due to the data bias between the query data used for surrogate model training and the real data used for the target model training, creating a valid query dataset is the crucial point of model extraction attacks. Papernot et al. [15] first used adversarial examples to query the target model for model extraction. However, due to the surrogate model is not well-trained, the adversarial examples generated on it cannot perform well in the imitating process. Orekondy et al. [14] propose the Knockoff to try to find valid query examples in a huge dataset, e.g., ImageNet [8], and adopt an adaptive strategy in the extraction process. Zhou et al. [23] proposed DAST, which is the first work to use a generator-based data-free distillation technique in knowledge distillation for model extraction. Later, subsequent studies have improved this approach to achieve better results [18]. However, the generator-based approach cannot obtain sufficient supervised information as in white-box knowledge distillation, leading to a huge number of queries and low attack results.

Therefore, the block-box attack with adversarial examples' transferability poses the request to guarantee that the local model is highly similar to the target model. To achieve this goal, we know from previous studies that model extraction can steal the target model from a decision boundary perspective, even in a data-independent way. However, the previous data-free works can not guarantee the synthetic data's diversity and need a massive number of queries to the target model. Hence, we are well-motivated to develop a better model extraction strategy adapted to data-free settings for carrying out attacks with high performance. Besides, it can improve the diversity of the generated fake data to be suitable for query-limited settings.

3 Preliminary

3.1 Adversarial Attack

Given a classifier $\mathcal{F}(\cdot)$ and an input x with its corresponding label y, we have $\mathcal{F}(x) = y$. The adversarial attack aims to find a small perturbation δ added to x, so the generated input x' misleads the classifier's output. The perturbation δ is usually constrained by L_p -norm $(p = 1, 2, ..., \infty)$, i.e., $\|\delta\|_p \leq \epsilon$. Then, the definition of adversarial examples x' can be written as:

$$\mathcal{F}(x') \neq y_{true}, \quad s.t. \ \|x' - x\|_p \le \epsilon, \tag{1}$$

where ϵ is the noise budget, y_{true} is the ground-truth label of example x.



Fig. 2. Framework of SCME, where \mathcal{L}_G consisting of inter- and intra-diverse loss.

3.2 Contrastive Learning

Contrastive learning models usually consist of two portions: self-supervised training in the upstream network and supervised fine-tuning in the downstream network. The upstream network $f(\cdot)$ aims to maximize the paired instance by augmenting the same data in different ways in the learned latent space while minimizing the agreement between different instances. Given a batch of examples $\{x_N\}$ without ground-truth labels, the random data transformation \mathcal{T} takes each example x in $\{x_N\}$ to a paired augmented data copies x_i and x_j , resulting in 2N augmented examples. The trained upstream network $f(\cdot)$ encodes the paired copies to latent vectors z_i and z_j . n SimCLR [1], the contrastive loss can be formulated as:

$$\ell_{i,j} = -\log \frac{\exp\left(\sin\left(\boldsymbol{z}_{i}, \boldsymbol{z}_{j}\right)/\tau\right)}{\sum_{k=1}^{2N} \mathbb{1}_{[k\neq i]} \exp\left(\sin\left(\boldsymbol{z}_{i}, \boldsymbol{z}_{k}\right)/\tau\right)},\tag{2}$$

where the z_i and z_j are the latent vectors of positive augmented examples, and z_k indicts the latent vector of negative examples from a different class. The $sim(\cdot)$ is a similarity function, such as cosine similarity loss, $\mathbb{1}$ is the indicator function, and τ is the temperature coefficient. A well-trained upstream network $f(\cdot)$ can extract effective features and use them in the downstream network, which usually is a simple MLP network, mapping the latent vectors to different classes through supervised learning.

4 Methodology

4.1 Overview

In this part, we illustrate the framework of our proposed data-free SCME in Fig. 2, which contains the following steps: 1) Synthesised Examples Generation and 2) Model Extraction. For step 1), we use a generator $\mathcal{G}(\cdot)$ to generate the



Fig. 3. The calculation process of intra-class diverse loss. C is the number of classes.

fake data \mathcal{X} . In step 2), we input the \mathcal{X} into both substitute model \mathcal{F}_{sub} and target model $\mathcal{F}_{tgt}(\cdot)$ to minimize the difference of their outputs. Notably, $\mathcal{F}_{sub}(\cdot)$ in SCME consists of upstream encoder network $\mathcal{F}_{up}(\cdot)$, a feature extraction network, projector network $\mathcal{F}_{down}(\cdot)$, and a classifier. Mathematically, the $\mathcal{F}_{sub}(\cdot)$ can be written as:

$$\mathcal{F}_{sub}(x) = F_{down}(F_{up}(x)),\tag{3}$$

where x is an arbitrary input example.

Based on the two steps mentioned above, the F_{sub} can imitate the \mathcal{F}_{tgt} in a data-free manner. Finally, we can generate adversarial examples by attacking \mathcal{F}_{sub} , and further attack the target model \mathcal{F}_{tgt} successfully.

4.2 Intra- and Inter-class Diverse

As mentioned above, \mathcal{X} should have both inter-class diversity and intra-class diversity to help the surrogate model training. Regarding this, as Fig. 3 show, we propose a self-contrastive loss to guide $\mathcal{G}(\cdot)$ in the \mathcal{X} generation. Inspired by the self-supervised loss in contrastive learning, we design a self-contrastive loss in SCME. Firstly, SCME uses the generator $\mathcal{G}(\cdot)$ to sample a batch of random noise $N = \{n_1, n_2, \dots, n_B\}$ to generate corresponding synthesize examples $\mathcal{X} = \{x_1, x_2, \dots, x_B\}$. SCME puts the \mathcal{X} into the feature extraction network $\mathcal{F}_{up}(\cdot)$ and gets the latent vectors z. Then, SCME calculates the self-contrastive intraclass diverse loss \mathcal{L}_{intra} by expanding the distance of each hidden vector z_i in z. The self-contrastive loss can be formulated as:

$$\mathcal{L}_{intra} = \log \sum_{i}^{B} \sum_{j}^{B} \mathbb{1}_{[i \neq j]} \cdot exp(sim(z_i, z_j)), \tag{4}$$

where B is the batch size, 1 is the indicator function and $sim(\cdot)$ is a similarity function.

In the synthesised examples generation, the loss function \mathcal{L}_G of generator $\mathcal{G}(\cdot)$ contains inter-class diversity loss \mathcal{L}_{inter} and intra-class diversity loss \mathcal{L}_{intra} . To generate inter-class diversity examples, we use the inter-class information entropy to guide the generator $\mathcal{G}(\cdot)$. That is, SCME randomly sets a batch of target label y_{tgt} and reduces the entropy between y_{tgt} and the substitute model's output of the generated examples \mathcal{X} . Mathematically, the inter-class loss function is:

$$\mathcal{L}_{inter} = \sum_{i=1}^{B} \mathcal{F}_{sub}(\mathcal{X}_i) log[F_{sub}(\mathcal{X}_i)],$$
(5)

where B is the batch size.

4.3 Model-Independent Boundary Example

Although the synthesis examples have been generated, however, they are still challenging to detect the target model's decision boundary adequately for substitute model training, resulting in a low attack performance. To solve this problem, we further modify the boundary examples by the Mixup augmentation to improve the substitute model's training efficiency. Specifically, SCME randomly selects two synthesized examples \mathcal{X}_i and \mathcal{X}_j first and then uses the Mixup to fuse them together to get the new boundary examples $\hat{\mathcal{X}}$, the process can be written as follows:

$$\hat{\mathcal{X}} = \lambda \mathcal{X}_i + (1 - \lambda) \mathcal{X}_j, \tag{6}$$

where the $\lambda \in [0, 1]$ is the mix weight and randomly sampled from β distribution.

4.4 Objective Function

By combining the above inter-class diversity loss \mathcal{L}_{inter} and the intra-class diversity loss \mathcal{L}_{intra} , we obtain the generate loss \mathcal{L}_G as the objective function for training the generator:

$$\mathcal{L}_G = \mathcal{L}_{intra} + \alpha \mathcal{L}_{inter},\tag{7}$$

where the α is the hyperparameter to adjust the weight of each loss.

Once the intra-class and inter-class diverse examples are generated, we input them into both substitute model $\mathcal{F}_{sub}(\cdot)$ and the target model $\mathcal{F}_{tgt}(\cdot)$ to minimize the distance between their outputs. To craft more suitable examples for training the substitute model and make its decision boundary close to the target model in the training process, we first craft the generated examples by Mixup operation to get the boundary examples $\hat{\mathcal{X}}$. Mathematically, the objective loss function \mathcal{L}_{train} of training substitute model is:

$$\mathcal{L}_{train} = \sum_{i=1}^{B} d(\mathcal{F}_{sub}(\hat{\mathcal{X}}), \mathcal{F}_{tgt}(\hat{\mathcal{X}})),$$
(8)

where the distance function $d(\cdot)$ is the Cross-Entropy loss in the hard label scenario and is the Mean Square Error loss in the soft label scenario.

Once the surrogate model is well-trained, we are able to generate adversarial examples on the substitute model and further attack the target black-box model.

5 Experiments

5.1 Setup

Datasets: We consider five benchmark datasets, namely MNIST [11], Fashion-MNIST [19], CIFAR-10 [7], CIFAR-100 [7], Tiny-ImageNet [10] for comprehensive experiments.

Dataset	Methods	Targeted, Hard Label			Untargeted, Hard Label			Targeted, Soft Label			Untargeted, Soft Label		
		FGSM	BIM	PGD	FGSM	BIM	PGD	FGSM	BIM	PGD	\mathbf{FGSM}	BIM	PGD
MNIST	JPBA	3.89	6.89	5.31	18.14	23.56	20.18	4.29	7.02	5.49	18.98	25.14	21.98
	Knockoff	4.18	6.03	4.66	19.55	27.32	22.18	4.67	6.86	5.26	21.35	28.56	23.34
	DaST	4.33	6.49	5.17	20.15	27.45	27.13	4.57	6.41	5.34	25.36	29.56	29.14
	Del	6.45	9.14	6.13	22.13	25.69	23.18	6.97	9.67	6.24	24.56	25.35	25.28
	EBFA	14.45	28.71	9.86	39.73	57.54	52.73	16.99	36.82	14.55	36.45	58.48	48.46
	SCME	9.98	9.96	10.04	63.45	74.51	78.47	10.05	10.00	10.04	72.46	78.54	82.54
Fashion-MNIST	JPBA	6.45	8.46	7.57	24.22	30.56	30.11	6.89	8.56	7.56	26.23	31.35	31.11
	Knockoff	6.34	8.35	7.32	28.19	36.88	35.92	6.65	8.98	8.23	30.21	36.94	36.22
	DaST	5.38	7.18	6.53	30.45	36.17	34.23	5.33	7.46	7.84	32.14	37.34	34.91
	Del	3.89	8.19	7.47	28.14	34.14	32.45	3.23	8.59	8.11	31.43	36.26	33.87
	EBFA	30.08	76.46	32.42	84.85	80.93	89.30	29.11	66.02	43.56	75.19	79.94	79.30
	SCME	31.82	70.79	70.01	82.26	84.76	85.11	32.14	72.07	72.07	82.58	85.46	85.86

Table 1. Attack performance on MNIST and Fashion-MNIST Datasets.

Table 2. Attack performance CIFAR-10 and CIFAR-100 Datasets.

Dataset	Methods	Targeted, Hard Label			Untargeted, Hard Label			Targeted, Soft Label			Untargeted, Soft Label		
		FGSM	BIM	PGD	FGSM	BIM	PGD	FGSM	BIM	PGD	FGSM	BIM	PGD
CIFAR-10	JPBA	6.32	7.70	7.92	27.82	33.23	31.70	7.28	8.56	7.64	28.77	33.38	31.96
	Knockoff	6.26	7.02	7.04	29.61	31.86	30.68	6.46	8.27	7.35	30.02	31.98	30.35
	DaST	6.54	7.81	7.41	27.61	34.43	26.99	8.15	8.40	8.26	27.58	34.75	27.47
	Del	7.14	7.44	6.95	25.33	30.45	30.34	7.86	8.29	7.17	26.38	31.53	31.47
	EBFA	14.57	16.95	12.27	86.13	87.02	84.32	31.54	13.93	69.14	83.89	87.68	85.11
	SCME	16.53	14.76	14.22	91.01	91.56	91.33	16.22	15.14	15.31	91.23	91.62	91.66
CIFAR-100	JPBA	4.35	6.20	6.17	33.58	38.54	37.08	5.73	7.50	6.41	34.21	39.12	37.31
	Knockoff	4.40	5.86	5.25	34.84	36.92	36.34	4.88	7.05	6.18	36.01	37.61	35.47
	DaST	4.97	6.19	5.92	33.57	39.86	32.71	6.38	7.04	7.01	32.80	40.34	32.78
	Del	5.38	5.72	5.69	30.80	35.63	36.15	6.30	6.53	5.23	31.64	36.63	37.44
	EBFA	16.64	16.88	12.77	78.61	91.31	91.21	7.91	16.15	12.54	83.69	94.53	94.14
	SCME	18.46	14.23	13.13	94.81	95.40	95.32	10.50	16.02	15.49	94.72	95.09	94.95

Models: For MNIST and Fashion-MNIST datasets, we use a simple network as the target model, which has four convolution layers and pooling layers and two fully-connected layers. For CIFAR-10 and CIFAR-100, we use the ResNet-18 [6] as the target model. For Tiny-ImageNet, we use the ResNet-50 [6] as the target model. The substitute model for all the datasets is the VGG-16 [16].

Baselines: To evaluate the performance of SCME, we compare it with the datadependent method, JPBA [15], Knockoff [14], and data-free methods, DAST [23], Del [18], EBFA [22].

Training Details: SCME and the baseline methods are trained with Adam optimizer with batch size 256. For the generator in SCME, we use an initial learning rate of 0.001 and a momentum of 0.9, and for the substitute model, we set the initial learning rate as 0.01 and momentum as 0.9. Furthermore, we set the maximal query times as 20 K, 100 K and 250 K for the MNIST dataset, CIFAR dataset and the Tiny-ImageNet dataset, respectively.

Metrics: We utilize three classical attack methods, which include FGSM, BIM and PGD, to generate adversarial examples for the surrogate model. For MNIST and Fashion-MNIST, we set perturbation budget $\epsilon = 32/255$. And for CIFAR-10, CIFAR-100 and Tiny-ImageNet, we set $\epsilon = 8/255$. In the untargeted attack

	Hard L	abel		Soft Label				
Methods	FGSM	BIM	PGD	FGSM	BIM	PGD		
JPBA	15.37	25.16	14.23	26.54	28.91	26.83		
Knockoff	22.33	21.39	11.26	29.99	27.64	26.17		
DaST	16.23	18.26	15.86	28.81	29.37	26.51		
Del	28.31	32.54	29.73	34.28	38.49	36.72		
EBFA	78.29	81.12	78.23	80.26	85.32	78.29		
SCME	90.16	90.25	89.72	96.44	96.29	96.32		

 Table 3. Attack performance on Tiny-ImageNet Dataset.



Fig. 4. Synthesised examples without (left) and with (right) data augmentation.

scenario, we only generate adversarial examples for the images which can be classified correctly by the victim model, while in targeted attacks, we only generate adversarial examples for the images which are not classified to the specific wrong labels. The **attack success rate** (ASR) is calculated by:

$$ASR = \begin{cases} \frac{1}{N} \sum_{i=1}^{N} [f(x_i^{adv}) \neq y_i], \text{ for untargeted} \\ \frac{1}{N} \sum_{i=1}^{N} [f(x_i^{adv}) = y_t], \text{ for targeted} \end{cases}$$
(9)

where N is the total number of generated adversarial examples.

Besides, for given a batch of query examples \mathcal{X} , we input them to the target model \mathcal{F}_{tgt} to get the output of each example and calculate its **Boundary Values** (BV) to verify whether the query samples are close to the decision boundary of the target model or not. The proposed BV can be calculated as follows:

$$BV = \sum_{i=1}^{B} (p(\mathcal{F}_{tgt}(\mathcal{X}_i))_{top_1} - p(\mathcal{F}_{tgt}(\mathcal{X}_i))_{top_2}),$$
(10)

where p is the Soft-max function, the top_1 and top_2 are the maximum value and sub-maximal value in the output probability vector, and the B is the total example counts.

5.2 Attack Performance

Experiments on MNIST and Fashion-MNIST: We report the ASR under targeted and untargeted attacks for both label-only and probability-only scenarios. As shown in Table 1, the ASR of SCME is much higher than the SOTA baselines on MNIST and Fashion-MNIST datasets. Obviously, our method can obtain higher ASR than other baselines in most cases with a small number of queries (here is 20K). This phenomenon shows that the proposed method is more applicable to the real world than the baselines.



Fig. 5. The T-SNE of original CFIAR-10 data (left), synthetic data by EBFA (middle) and synthetic data by SCME (right).

Experiments on CIFAR-10, CIFAR-100 and Tiny-ImageNet: We further investigate the performance of our method on complex datasets. From the results shown in Tables 2 and 3, our method achieves the best attack performance over probability-only and label-only scenarios under all datasets. In addition, compared to the strong baselines EBFA, our method still outperforms it significantly. Although the number of categories directly affects the training of the substitute model, our method still achieves a very high ASR on the CIFAR-100 and Tiny-ImageNet datasets, which have 100 categories and 200 categories, respectively. On the Tiny-ImageNet dataset, our method even achieves the highest ASR of 96.44% in the soft label setting. These improvements effectively demonstrate the superiority of the proposed SCME.



Table 4. Boundary value of EBFA and SCME.

Fig. 6. The ablation results of model accuracy (left) and model ASR (right), where "- L_sc " means without self-contrastive loss, and "- Mixup" means without Mixup operation.

5.3 Evaluation on Data Diversity

To evaluate the generated data's diversity of the strong baseline EBFA and the proposed SCME, we generated 10,000 examples and fed them into the same model trained on the CIFAR-10 dataset to get the predicted labels. The results in Fig. 4 show the data with data augmentation or not. The results show most of the examples generated by EBFA were classified as "deer", while synthetic examples by SCME have preferable inter-class diversity. Further, we plot the T-SNE for real data and synthetic examples generated by EBFA and SCME, respectively, in Fig. 5. The results illustrated that our method generates examples similar to the real data, i.e., with more intra-class diversity. These phenomena strongly support that our method can generate data with high inter- and intra- diversity.

5.4 Evaluation on Boundary Value

To verify whether the query examples are closer to the decision boundary of the target model, we compared the BV of 10,000 examples generated by EBFA and SCME. The results in Table 4 show although EBFA achieves smaller BV without data augmentation, SCME can achieve substantially lower BV with data augmentation. This further demonstrates the effectiveness of the Mixup operation in SCME for generating query examples close to the decision boundary.

5.5 Ablation Study

To investigate the contribution of Self-Contrastive loss \mathcal{L}_G (described in Sect. 4.2 and Mixup operation, we plot the model classification accuracy (ACC) and the model ASR in the model training process. The results in Fig. 6 shows that using both \mathcal{L}_G and Mixup augmentation performs best on both ACC and ASR, besides the model training convergence faster. For instance, the standard SCME is close to convergence with 6K queries, and the ASR is also beyond 80%.

6 Conclusion

In this paper, we proposed a novel data-free model extraction attack, namely SCME, to boost the attack performance under query-limited settings. Specifically, we first design a self-contrastive loss to guide the generator to synthesize the query data with high inter- and intra-class diversity. Besides, we introduce the Mixup augmentation to combine two generated query samples as the final query input to obtain effective decision boundaries and further help the simulation process of the substitute model. Extensive empirical results show that the proposed SCME framework can achieve SOTA attack performance.

Acknowledgements. This work is supported in part by Yunnan Province Education Department Foundation under Grant No.2022j0008, in part by the National Natural Science Foundation of China under Grant 62162067 and 62101480, Research and Application of Object Detection based on Artificial Intelligence, in part by the Yunnan Province expert workstations under Grant 202205AF150145.

References

- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. In: ICML, vol. 119, pp. 1597–1607 (2020)
- Demuynck, K., Triefenbach, F.: Porting concepts from dnns back to gmms. In: ASRU, pp. 356–361 (2013)
- Dong, Y., Pang, T., Su, H., Zhu, J.: Evading defenses to transferable adversarial examples by translation-invariant attacks. In: CVPR, pp. 4312–4321 (2019)
- Duan, R., Ma, X., Wang, Y., Bailey, J., Qin, A.K., Yang, Y.: Adversarial camouflage: Hiding physical-world attacks with natural styles. In: CVPR, pp. 997–1005 (2020)
- Duan, R., et al.: Adversarial laser beam: Effective physical-world attack to dnns in a blink. In: CVPR, pp. 16062–16071 (2021)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR, pp. 770–778 (2016)
- Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
- Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS, pp. 1106–1114 (2012)
- Kurakin, A., Goodfellow, I.J., Bengio, S.: Adversarial examples in the physical world. In: ICLR (2017)

- 10. Le, Y., Yang, X.S.: Tiny imagenet visual recognition challenge (2015)
- 11. LeCun, Y., Cortes, C., Burges, C.: Mnist handwritten digit database (1998)
- Liu, A., et al.: Perceptual-sensitive GAN for generating adversarial patches. In: AAAI, pp. 1028–1035 (2019)
- Liu, J., Park, J.: Seeing is not always believing: detecting perception error attacks against autonomous vehicles. IEEE Trans. Dependable Sec. Comput. 18(5), 2209– 2223 (2021)
- Orekondy, T., Schiele, B., Fritz, M.: Knockoff nets: stealing functionality of blackbox models. In: CVPR, pp. 4954–4963 (2019)
- Papernot, N., McDaniel, P.D., Goodfellow, I.J.: Practical black-box attacks against machine learning. In: Asia@CCS, pp. 506–519 (2017)
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)
- 17. Tramèr, F., Zhang, F., Juels, A., Reiter, M.K., Ristenpart, T.: Stealing machine learning models via prediction apis. In: USENIX Security, pp. 601–618 (2016)
- Wang, W., et al.: Delving into data: effectively substitute training for black-box attack. In: CVPR, pp. 4761–4770 (2021)
- Xiao, H., Rasul, K., Vollgraf, R.: Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms (2017)
- Yu, M., Sun, S.: Fe-dast: fast and effective data-free substitute training for blackbox adversarial attacks. Comput. Sec. 113, 102555 (2022)
- Yuan, X., Ding, L., Zhang, L., Li, X., Wu, D.O.: ES attack: model stealing against deep neural networks without data hurdles. IEEE Trans. Emerging Topics Comput. Intell. 6(5), 1258–1270 (2022)
- Zhang, J., et al.: Towards efficient data free blackbox adversarial attack. In: CVPR, pp. 15094–15104 (2022)
- Zhou, M., Wu, J., Liu, Y., Liu, S., Zhu, C.: Dast: data-free substitute training for adversarial attacks. In: CVPR, pp. 231–240 (2020)