

Data-Free Model Extraction

Jean-Baptiste Truong*

Worcester Polytechnic Institute

jtruong2@wpi.edu

Robert J. Walls

Worcester Polytechnic Institute

rjwalls@wpi.edu

Pratyush Maini*[†]

Indian Institute of Technology Delhi

pratyush.maini@gmail.com

Nicolas Papernot

University of Toronto and Vector Institute

nicolas.papernot@utoronto.ca

Abstract

Current model extraction attacks assume that the adversary has access to a surrogate dataset with characteristics similar to the proprietary data used to train the victim model. This requirement precludes the use of existing model extraction techniques on valuable models, such as those trained on rare or hard to acquire datasets. In contrast, we propose data-free model extraction methods that do not require a surrogate dataset. Our approach adapts techniques from the area of data-free knowledge transfer for model extraction. As part of our study, we identify that the choice of loss is critical to ensuring that the extracted model is an accurate replica of the victim model. Furthermore, we address difficulties arising from the adversary’s limited access to the victim model in a black-box setting. For example, we recover the model’s logits from its probability predictions to approximate gradients. We find that the proposed data-free model extraction approach achieves high-accuracy with reasonable query complexity – $0.99\times$ and $0.92\times$ the victim model accuracy on SVHN and CIFAR-10 datasets given 2M and 20M queries respectively.

1. Introduction

Machine learning (ML) and deep learning, in particular, often require large amounts of training data to achieve high performance on a particular task [39]. Curating such data necessitates significant time and monetary investment [17, 12]. Thus, the resulting ML model becomes valuable intellectual property, especially when considering the computing resources and human expertise required [5, 13]. Often to monetize these models, companies make them available as a service via APIs over the web (MLaaS). These models are also deployed to end-user devices, making their

predictions directly accessible to customers. However, the exposure of the model’s predictions represents a significant risk as an adversary can leverage this information to steal the model’s knowledge [26, 41, 7, 32, 31, 11, 28, 19]. The threat of such model extraction attacks is two-fold: adversaries may use the stolen model for monetary gains or as a reconnaissance step to mount further attacks [33, 37].

While model extraction is in many ways similar to model distillation, it differs in that the victim’s proprietary training set is not accessible to the adversary. To stage a model extraction attack, the adversary typically queries the victim using samples from a surrogate dataset with semantic or distributional similarity to the original training set [31]. In the classification setting, the victim’s response may be limited to the most-likely label [6] or include confidence values for different class labels [19]. The number of queries—i.e., the *query complexity*—is also an important consideration for the adversary. The greater the query complexity, the higher the cost of the attack—unless the victim model is available offline (e.g., deployed on-device).

In this work, we first demonstrate in Section 3 that the success of current established practices for model extraction, which often take the form of distillation, depends on the *closeness* of the surrogate distribution to the victim’s proprietary training distribution. This finding has important implications for the practicality of existing model extraction techniques.

To remedy this, we propose techniques for *data-free model extraction* (DFME). In short, we demonstrate the feasibility of extracting ML models without *any* knowledge of the distribution of the proprietary training data. In practice, gathering a surrogate dataset for the purpose of model extraction can be a very expensive process, both in terms of the time and money required to curate it. In particular, the most valuable models are often those for which it is most challenging to curate an appropriate surrogate dataset, i.e., when the victim model’s value arises from its proprietary

*equal contribution

[†]Work done while an intern at the Vector Institute.

dataset. Our work builds on recent advances in data-free knowledge distillation, which involve a generative model to synthesize queries that maximize disagreement between the student and teacher models [27, 14]. Here, the teacher is the victim model whereas the student is the stolen extracted model. We innovate on two fronts: the choice of loss to quantify student-teacher disagreement and an approach for training the generator without the ability to backpropagate through the teacher to compute its gradients (because we only have black-box access to the victim/teacher predictions in our setting). We observe that it is essential to ensure the stability of the loss computed, and find that the ℓ_1 norm loss is particularly conducive to data-free model extraction. We also demonstrate that using inexpensive gradient approximation (based on the victim model’s outputs) is sufficient to train a generative model that produces queries relevant to distill the knowledge of a victim to a student model. In summary, our main contributions are:

- We demonstrate in Section 3 that successful distillation-based model extraction attacks require the adversary to sample queries from a surrogate dataset whose distribution is *close* to the victim training data.
- In Section 4, we propose data-free model extraction (DFME) to extract ML models without knowledge of private training data, and only using the victim’s black-box predictions. As a by-product of DFME needing to approximate gradients of the victim, this leads us to present a method for recovering per-example logits out of the probability vector output by a ML model.
- We validate¹ our DFME technique in Section 5 on the SVHN and CIFAR10 datasets and successfully extract a model with 0.99x the victim accuracy with only 2M queries for SVHN, and 0.92x the victim accuracy with 20M queries for CIFAR10.
- An ablation study of our approach in Section 6 provides two key insights: (1) measuring disagreement between the victim and extracted models with the ℓ_1 norm achieves higher extraction accuracy than losses previously considered in the literature; (2) weak gradient estimates yield sufficient signal to train a generator despite only having access to the victim’s predictions.

2. Related Work

We covered the seminal results in model extraction based on surrogate datasets in the introduction. Here, we discuss data-free knowledge distillation—the technique that underlies our approach to data-free model extraction—as well as the rudiments of generative modeling and gradient approximation required to understand our method.

¹Code and models for reproducing our work can be found at <https://github.com/cake-lab/datafree-model-extraction>

2.1. Data-Free Knowledge Distillation

Knowledge distillation aims to compress, i.e., transfer, the knowledge of a (larger) teacher model to a (smaller) student model [3, 18]. It was originally introduced to reduce the size of models deployed on devices with limited computational resources. Since then, this line of work has attracted a lot of attention [47, 16, 34, 45, 46]. While the model owner usually performs knowledge distillation, the original dataset used to train the teacher model may not be available during distillation [27], e.g., because the dataset is too large or confidential. Therefore, others have proposed distillation techniques that leverage a surrogate dataset with a similar feature space or distribution [25, 31]. Others proposed techniques that altogether remove the need for a surrogate dataset, i.e., data-free knowledge distillation [14, 27]. Techniques addressing data-free knowledge distillation have relied on training a generative model to synthesize the queries that the student makes to the teacher [10, 27].

The success of data-free knowledge distillation hints at the feasibility of data-free model extraction. Kariyappa et al. observe this as well in concurrent work [20]. They also tackle data-free model extraction through the synthesis of queries by a generative model. Key differences include our loss formulation and optimizer choice (see Section 4). We show in Sections 5 and 6 that our approach consistently outperforms theirs.

2.2. Generative Models

Model extraction through data-free distillation involves the generation of training data with which the *student* (i.e., adversary) queries the *teacher* (i.e., victim) model. Naively, one could generate these queries randomly [27, 14]. In this paper, we instead build on a min-max game between two adversaries that try to optimize opposite loss functions. This approach is analogous to the optimization performed in Generative Adversarial Networks (GANs) [15] to train the generator and discriminator. Here, we use GANs in a fashion analogous to their application to semi-supervised learning [36]: our student and teacher models, in conjunction, play the discriminator’s role. The key difference here is that GANs are generally trained to recover an underlying *fixed* data distribution. However, our generator chases a moving target: the distribution of data which is most indicative of the discrepancies between the decision surfaces of the current student model and its teacher model.

2.3. Black-box Gradient Approximation

Zeroth-order optimization is a common approach to approximating gradients [43, 29, 8, 24]. Such techniques have previously been used to mount attacks against ML models in a *black-box* setting, e.g., to craft adversarial examples [42, 8, 4]. Various gradient estimation methods solve




Figure 1. Dataset Interpolation with CIFAR10 as target. $\lambda = 0$ implies that the inputs are sampled from the target distribution, while $\lambda = 1$ implies sampling from the surrogate.

different trade-offs between query complexity and the quality of the gradient estimate [42, 8, 4]. We use the forward differences [44] method for its relatively low query utilization, and systematically study the impact of its main parameter (e.g. the number of random directions) in Section 4.3.

3. How Hard is it to Find a Surrogate Dataset?

To motivate the need for data-free approaches to model extraction, we evaluate if an adversary must ensure that the distribution of its surrogate dataset is *close* to that of the victim’s training dataset. We hypothesize that in the absence of this condition, distillation-based model extraction will return a poor approximation of the victim. We perform an analysis on the *closeness* of the distributions along three axes: (1) similarity in feature space, (2) marginal probability distribution of inputs, and (3) class-conditional probability distribution of the inputs. In our experiments, we attempt to steal ML models trained on CIFAR10 [21] and SVHN [30] using various surrogate datasets that align differently with the axes defined above. We study in details the experimental setting, optimization problem, surrogate datasets and hyperparameters in Appendix A.

Our experiments support our hypothesis. For instance, in case of CIFAR10, with a victim model of accuracy 95.5%, extracting it using CIFAR100 [21] as surrogate dataset results in extraction accuracy of 93.5%. This can be largely attributed to the fact that both the CIFAR10 and CIFAR100 datasets are subsets from the same TinyImages [40] dataset. However, on using SVHN as surrogate dataset, the model extraction performance dropped remarkably, attaining a maximum of 66.6% across all the hyperparameters tried. In the extreme scenario when querying the CIFAR10 teacher with MNIST [23]—a dataset with disjoint feature space both in terms of number of pixels, and number of channels—accuracy did not improve beyond 37.2%.

On the contrary, we notice that the victim trained on the SVHN dataset is much easier for the adversary to extract. Surprisingly, even when the victim is queried with completely random inputs, the extracted model attains an

accuracy of over 84% on the original SVHN test set. We hypothesize that this observation is linked to how the digit classification task, at the root of SVHN, is a simpler task for neural networks to solve, and the underlying representations (hence, not being as complex as for CIFAR10) can be learnt even when queried over random inputs.

While these correlations agree with our hypothesis, these experiments can not systematically quantify the *distance* between two distributions (viz. the surrogate and the target). To more systematically understand how the shift away from the target distribution affects extraction performance we interpolated inputs (x_{in}) from the surrogate (x_s) and target (x_t) datasets, s.t. $x_{in} = (1 - \lambda) \cdot x_t + \lambda \cdot x_s$. Figure 1 shows the decrease in extraction accuracy as the distribution diverges from target (CIFAR10) for two different surrogate datasets (SVHN and MNIST).

We make two conclusions from our observations: (1) the success of distillation-based model extraction largely depends on the complexity of the task that the victim model aims to solve; and (2) similarity to source domain appears to be critical for extracting ML models that solve complex tasks. We posit that it may be nearly as expensive for the adversary to extract such models with a good surrogate dataset, as is training from scratch. A weaker or non-task specific dataset may have lesser costs, but has high accuracy trade-offs.

4. Data-Free Model Extraction

The goal of model extraction is to train a student model S to match the predictions of the victim \mathcal{V} on its private target domain \mathcal{D}_V . That is to say, find the student model’s parameters θ_S that minimize the probability of errors between the student and victim predictions $\mathcal{S}(x)$ and $\mathcal{V}(x) \forall x \in \mathcal{D}_V$:

$$\arg \min_{\theta_S} \mathcal{P}_{x \sim \mathcal{D}_V} \left(\arg \max_i \mathcal{V}_i(x) \neq \arg \max_i \mathcal{S}_i(x) \right) \quad (1)$$

Since the victim’s domain, \mathcal{D}_V , is not publicly available, the proposed data-free model extraction attack minimizes the student’s error on a synthesized dataset, \mathcal{D}_S . The error is minimized by optimizing a loss function, \mathcal{L} , which measures disagreement between the victim and student:

$$\arg \min_{\theta_S} \mathbb{E}_{x \sim \mathcal{D}_S} [\mathcal{L}(\mathcal{V}(x), \mathcal{S}(x))] \quad (2)$$

This section describes how we minimize the number of queries made to the victim model with a novel query generation process, and how we train the student model itself.

4.1. Overview

The overall attack setup is inspired by Generative Adversarial Networks [15]. A generator (\mathcal{G}) model is responsible for crafting some input images, and the student model S




Figure 2. Data-Free Model Extraction Attack Diagram

serves as a discriminator while trained to match the victim \mathcal{V} predictions on these images. In this setting, the two adversaries are \mathcal{S} and \mathcal{G} , which respectively try to minimize and maximize the disagreement between \mathcal{S} and \mathcal{V} .

The data flow is shown as a black arrow in Figure 2: a vector of random noise z is sampled from a standard normal distribution and fed into \mathcal{G} which produces an image x . Then the victim \mathcal{V} and student \mathcal{S} each perform inference on x to finally compute the loss function \mathcal{L} .

During the back-propagation phase (shown with red arrows) gradients from two different sources need to be computed: the gradients of \mathcal{L} with regards to the student’s parameters θ_S and the gradient of \mathcal{L} with regards to the generator’s parameters θ_G . Because the victim is only accessible as a black-box, it is not possible to propagate gradients through it. The dashed arrow indicates the need for gradient approximation (see Section 4.3).

Student. Prior work on knowledge distillation showed that a student model \mathcal{S} can learn from a teacher and reach high accuracy even though its architecture is smaller and different [9, 27]. Therefore, in the context of model extraction, the adversary only needs to select a model architecture which has sufficient capacity. This does not require knowledge of the victim architecture but rather generic knowledge of architectural choices made for the task solved by the victim (e.g., a convolutional neural network is appropriate for an object recognition task). In our work, we used a student with ResNet-18-8x architecture for model extraction.

The loss function \mathcal{L} is used to measure the disagreement between \mathcal{S} and \mathcal{V} . For this function, we use the ℓ_1 norm loss between victim and student logits (i.e. pre-softmax activations), $l_i(x)$ and $s_i(x)$ respectively. This requires us to recover the logits from the softmax outputs, since the adversary only has access to the later. We introduce an approach for doing so and further elaborate on the choice of \mathcal{L} is detailed in Subsection 4.2. It is important to note that the gradient of the loss with regard to the student’s weights θ_S does not require gradients of \mathcal{V} since the victim’s predictions don’t depend on the weights θ_S .

Algorithm 1: Data-Free Model Extraction

Input: Query budget Q , generator iterations n_G , student iterations n_S , learning rate η , random directions m , step size ϵ

Result: Trained \mathcal{S}

```

while  $Q > 0$  do
  for  $i = 1 \dots n_G$  do
     $z \sim \mathcal{N}(0, 1)$ 
     $x = \mathcal{G}(z; \theta_G)$ 
    approximate gradient  $\nabla_{\theta_G} \mathcal{L}(x)$ 
     $\theta_G = \theta_G - \eta \nabla_{\theta_G} \mathcal{L}(x)$ 
  end
  for  $i = 1 \dots n_S$  do
     $z \sim \mathcal{N}(0, 1)$ 
     $x = \mathcal{G}(z; \theta_G)$ 
    compute  $\mathcal{V}(x), \mathcal{S}(x), \mathcal{L}(x), \nabla_{\theta_S} \mathcal{L}(x)$ 
     $\theta_S = \theta_S - \eta \nabla_{\theta_S} \mathcal{L}(x)$ 
  end
  update remaining query budget  $Q$ 
end

```

Generator. The generator model \mathcal{G} is used to synthesize images that maximize the disagreement between \mathcal{S} and \mathcal{V} . The loss function used for \mathcal{G} is the same as for \mathcal{S} except that the goal is to maximize it. From this setting emerges an adversarial game in which \mathcal{S} and \mathcal{G} compete to respectively maximize and minimize the same function. In other words, the student is trained to match the victim’s predictions and the generator is trained to generate difficult examples for the student. The adversarial game can be written as:

$$\min_S \max_G \mathbb{E}_{z \sim \mathcal{N}(0,1)} [\mathcal{L}(\mathcal{V}(\mathcal{G}(z)), \mathcal{S}(\mathcal{G}(z)))] \quad (3)$$

As shown in Figure 2, computing the gradient of \mathcal{L} with regard to θ_G requires gradients of \mathcal{V} . As we only have access to \mathcal{V} as a black-box, gradient approximation techniques are required. These techniques are discussed in Section 4.3.

Algorithm. Each iteration alternates training the generator \mathcal{G} and student \mathcal{S} . To finely tune the balance between \mathcal{G} and \mathcal{S} training, each of these training phases is repeated n_G and n_S times, respectively, before moving on to the next epoch. While setting n_G higher allows \mathcal{G} to train faster and to produce more difficult examples for \mathcal{S} , it can also be wasteful if \mathcal{S} does not see enough examples. The trade-off between n_G and n_S is an additional hyperparameter that needs tuning. The additional hyperparameters m and ϵ are related to gradient approximation (see Section 4.3).

4.2. Loss function

Here we discuss different loss functions to measure the disagreement between \mathcal{V} and \mathcal{S} . These losses are commonly

used in the knowledge distillation literature given the similarity with the model extraction task [9, 14]. The choice of the loss function is key to the outcome of the attack since gradients computed through \mathcal{S} and \mathcal{V} can easily impede the convergence of optimizers, e.g., if they vanish because the wrong loss function is used.

Kullback–Leibler (KL) Divergence Most prior work in model distillation optimized over the KL divergence between the student and the teacher [9, 18, 22]. As a result, KL divergence between the outputs of \mathcal{S} and \mathcal{V} is a natural candidate for the loss function to train the student network. For a probability distribution over K classes indexed by i , the KL divergence loss for a single image x is defined as:

$$\mathcal{L}_{\text{KL}}(x) = \sum_{i=1}^K \mathcal{V}_i(x) \log \left(\frac{\mathcal{V}_i(x)}{\mathcal{S}_i(x)} \right) \quad (4)$$

However, as the student model matches more closely the victim model, the KL divergence loss tends to suffer from vanishing gradients [14]. Hypothesis 1 suggests that \mathcal{L}_{KL} can make it difficult to achieve convergence while training \mathcal{G} (refer to Appendix D for justification). Specifically, back-propagating such vanishing gradients through the generator can harm its learning. We confirm this through empirical evaluation as well in Section 6.1.

Hypothesis 1. *The gradients of the KL divergence loss with respect to the image x should be small compared to the gradients of the ℓ_1 norm loss when \mathcal{S} converges to \mathcal{V} :*

$$\|\nabla_x \mathcal{L}_{\text{KL}}(x)\| \ll_{\mathcal{S} \rightarrow \mathcal{V}} \|\nabla_x \mathcal{L}_{\ell_1}(x)\|$$

The ℓ_1 norm loss. To prevent gradients from vanishing, we use the ℓ_1 norm loss (\mathcal{L}_{ℓ_1}) computed with the victim and student logits v_i and s_i where $i \in \{1..K\}$ and K is the number of classes. This was previously found by Fang et al. to prevent gradients from vanishing in knowledge distillation [14]. Even though \mathcal{L}_{ℓ_1} is not differentiable everywhere, it does not suffer from the vanishing gradients issue and yields better results in practice (see Sec. 6.1). Lastly, the probabilities output by \mathcal{V} need to be transformed into logits to be used in \mathcal{L}_{ℓ_1} . We describe how to perform logit approximation in Appendix B, and evaluate in Sec. 6.3.

$$\mathcal{L}_{\ell_1}(x) = \sum_{i=1}^K |v_i - s_i| \quad (5)$$

4.3. Gradient Approximation

Because only black-box access is provided for \mathcal{V} , the optimizer aims at maximizing a function for which it only has an evaluation oracle. Yet, in order to train \mathcal{G} , gradients of the loss with regards to \mathcal{G} 's parameters $\nabla_{\theta_{\mathcal{G}}} \mathcal{L}$ must be computed. Thus, we approximate gradients by interacting with the oracle: we maximize \mathcal{L} with zeroth-order optimization.

4.3.1 Images as a Proxy

The number of parameters in \mathcal{G} is typically large (millions of parameters) and it would be very query-expensive for a zeroth-order optimizer to get accurate gradient estimations on this large space. Instead, one can approximate gradients with regards to the input images x , and then back-propagate this gradient through \mathcal{G} [20]. This way the dimensionality of gradients being approximated is much smaller, which yields more accurate zeroth-order approximations.

Additionally the oracle might only accept images that lie within a pre-defined input domain, for example $[-1, 1]^d$. To force \mathcal{G} to respect this constraint, we use a hyperbolic tangent activation at the end of the generator architecture. Furthermore, zeroth-order gradients approximation methods usually evaluate the function in the neighborhood of a given point, which can result in query images slightly outside the input domain. To avoid this, we approximate gradients with regard to the pre-activation images (i.e. just before the hyperbolic tangent function is applied).

4.3.2 Forward Differences Method

The Forward Differences method approximates gradients by computing directional derivatives $\mathcal{D}_{\mathbf{u}_i} f(x)$ of a function f at a point x along m random directions \mathbf{u}_i . The directional derivatives are computed by measuring the variation of f a small step of size ϵ in the direction \mathbf{u}_i . They are then averaged to form an estimator of the gradient $\nabla_{\text{FWD}} f(x)$. In a way, each directional derivative brings some amount of information about true gradient. The estimator being more accurate as the number of random directions increases.

$$\nabla_{\text{FWD}} f(x) = \frac{1}{m} \sum_{i=1}^m d \frac{f(x + \epsilon \mathbf{u}_i) - f(x)}{\epsilon} \mathbf{u}_i \quad (6)$$

The main advantage of this method is that the number of query directions m may be chosen independently of the input space dimensionality, offering a trade-off between query utilization and gradient accuracy. This makes it an appealing candidate for DFME [20]. The influence of the number of query directions m is further described in Section 6.2.

Finite differences, an alternative gradient approximation method used when crafting adversarial examples [4], requires too many queries per gradient estimate to be viable for data-free model extraction.

5. Experimental Validation

We evaluate data-free model extraction (DFME) against victim models trained SVHN and CIFAR-10. We show that the resulting student models can reach high accuracy (e.g., 95.2% on SVHN) even when the generator only has access to inaccurate gradient estimates. Later in Section 6, we perform an ablation study and evaluate the impact of each at-

tack component on the final student model accuracy and on the query budget Q .

5.1. Datasets and Architectures

We evaluate the effectiveness of the proposed DFME method on two datasets: SVHN and CIFAR-10. For each dataset, the victim model architecture is a ResNet-34-8x. These victim models were trained during 50 epochs for SVHN and 200 for CIFAR-10 with SGD at an initial learning rate of 0.1, decayed by a factor of 10 at 50% of training.

We use ResNet-18-8x as the architecture for our student model. This is inspired by previous works in knowledge distillation [14] that show how a smaller student is sufficient to distill the knowledge of a larger teacher. The network was trained with a batch size of 256 with SGD, with an initial learning rate of 0.1, a weight-decay of $5 \cdot 10^{-4}$, and a learning rate scheduler that multiplies the learning rate by a factor 0.3 at $0.1 \times$, $0.3 \times$, and $0.5 \times$ the total training epochs. The default query budget Q is 2M for SVHN, and 20M for CIFAR-10 in our experiments.

The generator used three convolutional layers, interleaved with linear up-sampling layers, batch normalization layers, and ReLU activations for all layers except the last one. The final activation function was the hyperbolic tangent function to output values in the range $[-1, 1]$ (see Section 4.3). It was also trained with a batch size of 256, but using an Adam optimizer with an initial learning rate of $5 \cdot 10^{-4}$ which is decayed by a factor 0.3 at 10%, 30%, and 50% of the training.

For gradient approximation we sample $m = 1$ random directions and a step size $\epsilon = 10^{-3}$.

5.2. Results

We compare the performance of different extraction attacks in Table 1. We measure the ratio between the student’s accuracy and the victim’s accuracy on the victim’s test set. This helps compare the performance of DFME across different datasets. The student model’s normalized accuracy is reported for each dataset and extraction method evaluated—our approach (DFME), our approach with KL divergence loss (DFME-KL), our approach without logit correction (Log-Probabilities), and concurrent work [20] (MAZE). Further, we perform DFME with a range of query budgets and reported the accuracy in Figure 3.

Without any knowledge of the original training distribution, the proposed DFME method achieved as high as 88.1% (0.92x target) of accuracy with $Q = 20M$ and 89.9% (0.94x target) with $Q = 30M$. The accuracy of the extracted model exceeds that reported in concurrent work which we refer to as MAZE in our results [20]. However, in our best-efforts at reproducing their results with the details in the paper,² we were unsuccessful in achieving the same re-




Figure 3. Test accuracy wrt query budget, for SVHN and CIFAR10

ported accuracy, and were only able to achieve an accuracy of 45.6% (0.48x target) at best on the student model. In addition, MAZE reports that they were unable to learn when using extremely few directions (such as $m = 1$) for the gradient approximation with CIFAR10, whereas we find that weak gradient approximations are beneficial to reduce the overall query budget of successful attacks.

We observed similar results for the SVHN victim model: reaching as high as 95.2% (0.99x target) accuracy with only 2M queries. The task for SVHN is much simpler than CIFAR10 given that a model with 84% (0.87x target) accuracy can be extracted from just random noise. The proposed method allows one to achieve far higher accuracy.

One limitation of our study is that the reported query budgets do not include the cost of hyperparameter tuning. This is an important direction for future work as preliminary experiments suggest that extraction accuracy can be sensitive to the choice of hyperparameters.

6. Ablation Studies

Our work systematically transitions from a data-free knowledge distillation paradigm [14, 27] to a data-free model extraction scenario. The main challenges in this transition were (1) to surpass the need for true gradients for training the student; (2) the lack of access to true victim logits; and (3) the need to restrict the query complexity of the attacks (to reduce the cost of stealing). With this goal, we made specific choices with regards to (a) the loss function; (b) gradient approximation; and (c) logit access. In this section, we detail the impact of each of these choices to the final performance of our proposed DFME method.

6.1. Choice of Loss Function

The choice of loss is of paramount importance to a successful extraction. In our DFME approach, the choice of loss involves similar factors to those outlined in research on GANs: multiple works have discussed the problem of vanishing gradients as the discriminator becomes strong in case of GAN training [1, 2]. For our DFME approach, we minimize the ℓ_1 distance between the output logits of the student and the teacher. We find that this significantly improves convergence and stability over other possible losses, such as the KL divergence chosen in [20].

²The authors declined to share their code upon request.

Dataset (budget)	Victim accuracy	DFME	DFME-KL	MAZE* [20]	Log-Probabilities
CIFAR10 (20M)	95.5%	88.1% (0.92 \times)	76.7% (0.80 \times)	45.6% (0.48 \times)	73.2% (0.77 \times)
SVHN (2M)	96.2%	95.2% (0.99 \times)	84.7% (0.88 \times)	91.1% (0.95 \times)	94.4% (0.98 \times)

Table 1. Accuracy and normalized accuracy of data-free model extraction methods. Results for ‘MAZE’ reflect our best-effort reproduction.

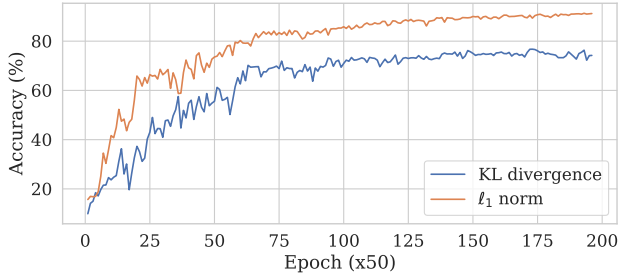


Figure 4. Test accuracy as training progresses for ℓ_1 and KL divergence losses.

We perform DFME in the same setting to evaluate the difference between the KL divergence and ℓ_1 losses. Below, we draw comparisons based on two metrics: (1) Final accuracy attained by the student at the end of a fixed number of queries as well as the learning curves of the student; and (2) The norm of gradients of the loss with respect to the input image as the training progresses.

Test Accuracy. The key metric of interest for this comparison is the normalized accuracy of the student model at the end of a designated query budget Q of 20M queries for CIFAR10 and 2M queries for SVHN. Table 1 shows that using the ℓ_1 loss achieves significantly better test accuracies compared to the KL divergence loss. For instance, on CIFAR10 the accuracy improves from 76.7% to 88.1% when switching from KL divergence to the ℓ_1 loss. We also visualize a learning curve for CIFAR10 in Figure 4: the KL divergence objective slows converge and tappers off earlier, even when the student has yet to plateau.

Gradient Vanishing. The KL divergence loss suffers from vanishing gradients, as explained in Section 4.2. In DFME, these gradients are used to update the generator’s parameters and are thus essential to synthesize queries which extract more information from the victim. In Figure 5 we empirically demonstrate that as the student accuracy approaches that of the victim model, the gradients of the KL divergence loss with respect to the input image reduce significantly in norm. The same decay is slower and less significant in case of the ℓ_1 loss. We hypothesize that these vanishing gradients are the cause for degraded accuracy when using the KL divergence loss.

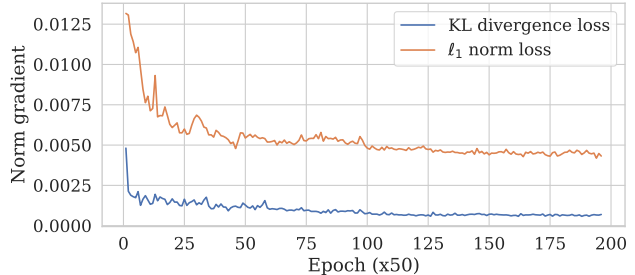


Figure 5. Norm of gradients with respect to the input image, for the KL divergence and ℓ_1 norm losses.

m	1	3	5	8	10
No. of Queries	10.04	10.02	16.33	13.80	20.00

Table 2. Minimum queries (in millions) to reach 85% accuracy on CIFAR10, for different number of gradient approximation steps.

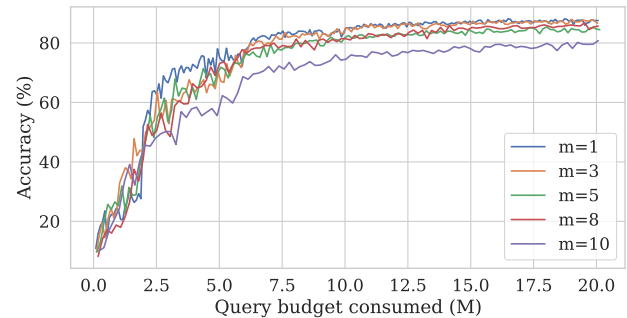


Figure 6. Accuracy of the model during training for different number of gradient approximation steps, m .

Model	MTL ($\pm 2.3e-6$)	MC ($\pm 2.3e-6$)	LP (± 1.3)
Resnet-34-8x	-1.24e-5	1.24e-5	4.98
Densenet 121	5.53e-7	1.88e-6	3.88
VGG 16	1.97e-5	1.97e-5	3.93

Table 3. Mean of true logits (MTL) for 3 victim architectures; reconstruction error (MAE) between approximate and true logits when using mean correction (MC) and log-probabilities (LP).

6.2. Gradient Approximation

Recall that the ℓ_1 loss cannot be back-propagated through the victim since the adversary only has access to it as a black-box. Recent work on data-free model distillation [14] has claimed that the gradient information from the teacher is ‘indispensable at the beginning of adversarial training’ because the student alone can not provide useful

signal to the generator when randomly initialized. Below we consider: (1) the quality of approximation required; and (2) the overall impact on query complexity. In particular, we compare two approaches for improving the training of our DFME generator: using an increased number of queries to compute more accurate gradient estimates or training generator for longer using poorer gradient estimates.

Number of gradient approximation steps. When gradients used to update the generator are approximated with the forward differences method, a larger number of random directions m allows one to compute more accurate gradients. However, in a model extraction setting, each additional gradient approximation step comes at the cost of increased query complexity. In practice, with a fixed query budget Q , changing the number of random directions directly impacts the proportion of queries used to train each network. This ratio of queries r used to train the student is given by:

$$r = \frac{n_S}{n_S + (m + 1)n_G} \quad (7)$$

In our setting (i.e. $n_G = 1$, $n_S = 5$) choosing m equal to 1 or 10 respectively results in 71% and 31% of the query budget being used to directly train the student, while the remainder is spent to get better gradient estimates to train the generator. Despite using a majority of queries to train the generator, the setting with $m = 10$ achieves comparable accuracy. In Table 2, we observe how choosing lower values of m achieves 85% test accuracy in much fewer queries.

Amortizing the cost. We hypothesize that since early in the training the discriminator (or student) provides only little signal, it is beneficial for the generator to initially rely on weak signals of gradient approximation. Effectively, this helps amortize the cost of gradient approximation over multiple epochs, and effectively pushes the expense to a later stage when the discriminator (or student) provides stronger signal. Figure 6 shows that relative to the query budget utilization, different values of m perform similarly.

This also suggests a hybrid strategy where the adversary first extracts a (somewhat poor) student model through distillation from a surrogate dataset. Indeed, we show in Section 3 that surrogate datasets drawn from a different distribution (than the victim model’s training data) enable distillation-based model extraction—albeit to a lower accuracy than in-distribution surrogate datasets. This poor initial student can then be improved by synthesizing queries with the data-free model extraction’s generator to bring the student model closer to the victim model’s performance.

Case $m = 1$. In the extreme case where the number of gradient approximation steps for forward differences (m) is set to 1, the approximated gradient is colinear to the random

direction sampled for the approximation, but always points in the direction that helps maximize the loss (i.e. its projection onto the true gradient is positive). The cosine similarity with the true gradient is, thus, very small. To validate this effect, we additionally experiment with $m = 1$ where the approximate gradient was randomly flipped to the wrong direction with half probability. As hypothesized, the student accuracy did not improve beyond 20% in our experiments on CIFAR10. This suggests that computing gradients that are extremely inaccurate makes it possible to train the student as long as these gradients are in the correct direction.

6.3. Impact of Logits Correction

A model extraction attack should be applicable to the nature of predictions offered by MLaaS APIs. Most APIs provide per-class probability distribution rather than the true logits, since probabilities are more easily interpreted by the end user. To perform model extraction successfully we thus need to recover the logits from the victim’s prediction probabilities. We show that it is possible to do so and recover approximate logits whose Mean Average Error (MAE) with the true logits is low, on three different victim architecture.

The MAE reported in Table 3 are negligible compared to true logits which take values in the order of magnitude of 1. Therefore, the adversary can use these approximate logits in lieu of the true logits. In comparison, approximating true logits with plain log-probabilities resulted in a MAE in the order of magnitude of the true logits themselves. Using the log-probabilities with such a large error makes the student training harder—it did not yield accuracy above 75%.

This method is effective because the mean of the true logits is nearly 0 (see Table 3). Therefore, subtracting the mean from the log-probabilities is equivalent to subtracting the additive constant $C(x)$ itself.

7. Conclusions

In this paper, we demonstrate that data-free model extraction is not only practical but also yields accurate copies of the victim model. This means that model extraction attacks is a credible threat to the intellectual property of models released intentionally or not to the public. We believe an interesting direction for future work is to detect such queries without decreasing the model’s utility to legitimate users.

Acknowledgements.

We thank the reviewers and members of CleverHans Lab for their insightful feedback. This work was supported by a Canada CIFAR AI Chair, NSERC, a gift from Microsoft, and sponsors of the Vector Institute.

References

- [1] Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks, 2017. **6**
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan, 2017. **6**
- [3] Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *Advances in neural information processing systems*, pages 2654–2662, 2014. **2**
- [4] Arjun Nitin Bhagoji, Warren He, Bo Li, and Dawn Song. Practical black-box attacks on deep neural networks using efficient query mechanisms. In *European Conference on Computer Vision*, pages 158–174. Springer, 2018. **2, 3, 5, 12**
- [5] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020. **1**
- [6] V. Chandrasekaran, K. Chaudhuri, I. Giacomelli, S. Jha, and Songbai Yan. Model extraction and active learning. *ArXiv*, abs/1811.02054, 2018. **1**
- [7] Varun Chandrasekaran, Kamalika Chaudhuri, Irene Giacomelli, Somesh Jha, and Songbai Yan. Exploring connections between active learning and model extraction, 2019. **1**
- [8] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pages 15–26, 2017. **2, 3, 12**
- [9] Jang Hyun Cho and Bharath Hariharan. On the efficacy of knowledge distillation, 2019. **4, 5, 11**
- [10] Yoojin Choi, Jihwan Choi, Mostafa El-Khamy, and Jungwon Lee. Data-free network quantization with adversarial knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 710–711, 2020. **2**
- [11] Jacson Rodrigues Correia-Silva, Rodrigo F. Berriel, Claudine Badue, Alberto F. de Souza, and Thiago Oliveira-Santos. Copycat cnn: Stealing knowledge by persuading confession with random non-labeled data. *2018 International Joint Conference on Neural Networks (IJCNN)*, Jul 2018. **1**
- [12] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. **1**
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2020. **1**
- [14] Gongfan Fang, Jie Song, Chengchao Shen, Xinchao Wang, Da Chen, and Mingli Song. Data-free adversarial distillation. *arXiv preprint arXiv:1912.11006*, 2019. **2, 5, 6, 7**
- [15] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. **2, 3**
- [16] Jianping Gou, Baosheng Yu, Stephen John Maybank, and Dacheng Tao. Knowledge distillation: A survey, 2020. **2**
- [17] Alon Halevy, Peter Norvig, and Fernando Pereira. The unreasonable effectiveness of data. *IEEE Intelligent Systems*, 24:8–12, 2009. **1**
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. **2, 5, 11**
- [19] Matthew Jagielski, Nicholas Carlini, David Berthelot, Alex Kurakin, and Nicolas Papernot. High accuracy and high fidelity extraction of neural networks. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*, 2020. **1**
- [20] Sanjay Kariyappa, Atul Prakash, and Moinuddin Qureshi. Maze: Data-free model stealing attack using zeroth-order gradient estimation, 2020. **2, 5, 6, 7**
- [21] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009. **3, 11**
- [22] Xu Lan, Xiatian Zhu, and Shaogang Gong. Knowledge distillation by on-the-fly native ensemble, 2018. **5, 11**
- [23] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. **3, 11**
- [24] Sijia Liu, Pin-Yu Chen, Bhavya Kailkhura, Gaoyuan Zhang, Alfred Hero, and Pramod K. Varshney. A primer on zeroth-order optimization in signal processing and machine learning, 2020. **2**
- [25] Raphael Gontijo Lopes, Stefano Fenu, and Thad Starner. Data-free knowledge distillation for deep neural networks, 2017. **2**
- [26] Daniel Lowd and Christopher Meek. Adversarial learning. In *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, KDD '05, page 641–647, New York, NY, USA, 2005. Association for Computing Machinery. **1**
- [27] Paul Micaelli and Amos J Storkey. Zero-shot knowledge transfer via adversarial belief matching. In *Advances in Neural Information Processing Systems*, pages 9551–9561, 2019. **2, 4, 6**
- [28] Smitha Milli, Ludwig Schmidt, Anca D. Dragan, and Moritz Hardt. Model reconstruction from model explanations, 2018. **1**
- [29] Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566, 2017. **2**
- [30] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011. **3, 11**

- [31] Tribhuvanesh Orekondy, Bernt Schiele, and Mario Fritz. Knockoff nets: Stealing functionality of black-box models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 2
- [32] Soham Pal, Yash Gupta, Aditya Shukla, Aditya Kanade, Shirish Shevade, and Vinod Ganapathy. A framework for the extraction of deep neural networks by leveraging public data, 2019. 1
- [33] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security, ASIA CCS '17*, page 506–519, New York, NY, USA, 2017. Association for Computing Machinery. 1
- [34] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets, 2015. 2
- [35] Sebastian Ruder. Transfer Learning - Machine Learning's Next Frontier. <http://ruder.io/transfer-learning/>, 2017. 11
- [36] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016. 2
- [37] Ilia Shumailov, Yiren Zhao, Daniel Bates, Nicolas Papernot, Robert Mullins, and Ross Anderson. Sponge examples: Energy-latency attacks on neural networks, 2020. 1
- [38] Leslie N. Smith. Cyclical learning rates for training neural networks, 2017. 11
- [39] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in nlp, 2019. 1
- [40] Antonio Torralba, Rob Fergus, and William T. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(11):1958–1970, Nov. 2008. 3, 12
- [41] Florian Tramèr, Fan Zhang, Ari Juels, Michael K Reiter, and Thomas Ristenpart. Stealing machine learning models via prediction apis. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*, pages 601–618, 2016. 1
- [42] Chun-Chen Tu, Paishun Ting, Pin-Yu Chen, Sijia Liu, Huan Zhang, Jinfeng Yi, Cho-Jui Hsieh, and Shin-Ming Cheng. Autozoom: Autoencoder-based zeroth order optimization method for attacking black-box neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 742–749, 2019. 2, 3
- [43] Yining Wang, Simon Du, Sivaraman Balakrishnan, and Aarti Singh. Stochastic zeroth-order optimization in high dimensions. In *International Conference on Artificial Intelligence and Statistics*, pages 1356–1365, 2018. 2
- [44] Andre Wibisono, Martin J Wainwright, Michael Jordan, and John C Duchi. Finite sample convergence rates of zero-order stochastic optimization methods. *Advances in Neural Information Processing Systems*, 25:1439–1447, 2012. 3
- [45] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer, 2017. 2
- [46] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation, 2019. 2
- [47] Ying Zhang, Tao Xiang, Timothy M. Hospedales, and Huchuan Lu. Deep mutual learning, 2017. 2