

# High Accuracy and High Fidelity Extraction of Neural Networks

Matthew Jagielski<sup>†,\*</sup>, Nicholas Carlini<sup>\*</sup>, David Berthelot<sup>\*</sup>, Alex Kurakin<sup>\*</sup>, and Nicolas Papernot<sup>\*</sup>

<sup>†</sup>Northeastern University

<sup>\*</sup>Google Research

## Abstract

In a model extraction attack, an adversary steals a copy of a remotely deployed machine learning model, given oracle prediction access. We taxonomize model extraction attacks around two objectives: *accuracy*, i.e., performing well on the underlying learning task, and *fidelity*, i.e., matching the predictions of the remote victim classifier on any input.

To extract a high-accuracy model, we develop a learning-based attack exploiting the victim to supervise the training of an extracted model. Through analytical and empirical arguments, we then explain the inherent limitations that prevent any learning-based strategy from extracting a truly high-fidelity model—i.e., extracting a functionally-equivalent model whose predictions are identical to those of the victim model on all possible inputs. Addressing these limitations, we expand on prior work to develop the first practical functionally-equivalent extraction attack for direct extraction (i.e., without training) of a model’s weights.

We perform experiments both on academic datasets and a state-of-the-art image classifier trained with 1 billion proprietary images. In addition to broadening the scope of model extraction research, our work demonstrates the practicality of model extraction attacks against production-grade systems.

## 1 Introduction

Machine learning, and neural networks in particular, are widely deployed in industry settings. Models are often deployed as prediction services or otherwise exposed to potential adversaries. Despite this fact, the trained models themselves are often proprietary and are closely guarded.

There are two reasons models are often seen as sensitive. First, they are expensive to obtain. Not only is it expensive to train the final model [1] (e.g., Google recently trained a model with 340 million parameters on hardware costing 61,000 USD *per training run* [2]), performing the work to identify the optimal set of model architecture, training algorithm, and hyper-parameters often eclipses the cost of training the final

model. Further, training these models also requires investing in expensive collection process to obtain the training datasets necessary to obtain an accurate classifier [3–6]. Second, there are security [7, 8] and privacy [9, 10] concerns for revealing trained models to potential adversaries.

Concerningly, prior work found that an adversary with query access to a model can steal the model to obtain a copy that largely agrees with the remote victim models [8, 11–16]. These extraction attacks are therefore important to consider.

In this paper, we systematize the space of model extraction around two adversarial objectives: *accuracy* and *fidelity*. Accuracy measures the correctness of predictions made by the extracted model on the test distribution. Fidelity, in contrast, measures the general agreement between the extracted and victim models on any input. Both of these objectives are desirable, but they are in conflict for imperfect victim models: a high-fidelity extraction should replicate the errors of the victim, whereas a high-accuracy model should instead try to make an accurate prediction. At the high-fidelity limit is *functionally-equivalent* model extraction: the two models agree on all inputs, both on and off the underlying data distribution.

While most prior work considers accuracy [7, 11, 13], we argue that fidelity is often equally important. When using model extraction to mount black-box adversarial example attacks [7], fidelity ensures the attack is more effective because more adversarial examples transfer from the extracted model to the victim. Membership inference [9, 10] benefits from the extracted model closely replicating the confidence of predictions made by the victim. Finally, a functionally-equivalent extraction enables the adversary to inspect whether internal representations reveal unintended attributes of the input—that are statistically uncorrelated with the training objective, enabling the adversary to benefit from overlearning [17].

We design one attack for each objective. First, a *learning-based attack*, which uses the victim to generate labels for training the extracted model. While existing techniques already achieve high accuracy, our attacks are 16× more query-efficient and scale to larger models. We perform experiments

that surface inherent limitations of learning-based extraction attacks and argue that learning-based strategies are ill-suited to achieve high-fidelity extraction. Then, we develop the first practical *functionally-equivalent attack*, which directly recovers a two-layer neural network’s weights exactly given access to double-precision model inference. Compared to prior work, which required a high-precision power side-channel [18] or access to model gradients [19], our attack only requires input-output access to the model, while simultaneously scaling to larger networks than either of the prior methods.

We make the following contributions:

- We taxonomize the space of model extraction attacks by exploring the objective of *accuracy* and *fidelity*.
- We improve the query efficiency of learning attacks for accuracy extraction and make them practical for millions-of-parameter models trained on billions of images.
- We achieve high-fidelity extraction by developing the first practical functionally-equivalent model extraction.
- We mix the proposed methods to obtain a hybrid method which improves both accuracy and fidelity extraction.

## 2 Preliminaries

We consider classifiers with domain  $\mathcal{X} \subseteq \mathbb{R}^d$  and range  $\mathcal{Y} \subseteq \mathbb{R}^K$ ; the output of the classifier is a distribution over  $K$  class labels. The class assigned to an input  $x$  by a classifier  $f$  is  $\arg \max_{i \in [K]} f(x)_i$  (for  $n \in \mathbb{Z}$ , we write  $[n] = \{1, 2, \dots, n\}$ ). In order to satisfy the constraint that a classifier’s output is a distribution, a *softmax*  $\sigma(\cdot)$  is typically applied to the output of an arbitrary function  $f_L : \mathcal{X} \rightarrow \mathbb{R}^K$ :

$$\sigma(f_L(x))_i = \frac{\exp(f_L(x)_i)}{\sum_j \exp(f_L(x)_j)}.$$

We call the function  $f_L(\cdot)$  the *logit function* for a classifier  $f$ . To convert a class label into a probability vector, it is common to use *one-hot encoding*: for a value  $j \in [K]$ , the one-hot encoding  $OH(j; K)$  is a vector in  $\mathbb{R}^K$  with  $OH(j; K)_i = \mathbb{1}(i = j)$ —that is, it is 1 only at index  $j$ , and 0 elsewhere.

Model extraction concerns reproducing a victim model, or oracle, which we write  $O : \mathcal{X} \rightarrow \mathcal{Y}$ . The model extraction adversary will run an extraction algorithm  $\mathcal{A}(O)$ , which outputs the extracted model  $\hat{O}$ . We will sometimes parameterize the oracle (resp. extracted model) as  $O_\theta$  (resp.  $\hat{O}_\theta$ ) to denote that it has model parameters  $\theta$ —we will omit this when unnecessary or apparent from context.

In this work, we consider  $O$  and  $\hat{O}$  to both be neural networks. A neural network is a sequence of operations—alternatingly applying linear operations and non-linear operations—a pair of linear and non-linear operations is called a *layer*. Each linear operation projects onto some space  $\mathbb{R}^h$ —the dimensionality  $h$  of this space is referred to as the *width*

of the layer. The number of layers is the *depth* of the network. The non-linear operations are typically fixed, while the linear operations have parameters which are learned during training. The function computed by layer  $i$ ,  $f_i(a)$ , is therefore computed as  $f_i(a) = g_i(A^{(i)}a + B^{(i)})$ , where  $g_i$  is the  $i$ th non-linear function, and  $A^{(i)}, B^{(i)}$  are the parameters of layer  $i$  ( $A^{(i)}$  is the weights,  $B^{(i)}$  the biases). A common choice of activation is the rectified linear unit, or ReLU, which sets  $\text{ReLU}(x) = \max(0, x)$ . Introduced to improve the convergence of optimization when training neural networks, the ReLU activation has established itself as an effective default choice for practitioners [20]. Thus, we consider primarily ReLU networks in this work.

The network structure described here is called *fully connected* because each linear operation “connects” every input node to every output node. In many domains, such as computer vision, this is more structure than necessary. A neuron computing edge detection, for example, only needs to use information from a small region of the image. *Convolutional networks* were developed to combat this inefficiency—the linear functions become filters, which are still linear, but are only applied to a small (e.g.,  $3 \times 3$  or  $5 \times 5$ ) window of the input. They are applied to every window using the same weights, making convolutions require far fewer parameters than fully connected networks.

Neural networks are trained by *empirical risk minimization*. Given a dataset of  $n$  samples  $\mathcal{D} = \{x_i, y_i\}_{i=1}^n \subseteq \mathcal{X} \times \mathcal{Y}$ , training involves minimizing a loss function  $L$  on the dataset with respect to the parameters of the network  $f$ . A common loss function is the cross-entropy loss  $H$  for a sample  $(x, y)$ :  $H(y, f(x)) = -\sum_{k \in [K]} y_k \log(f(x)_k)$ , where  $y$  is the probability (or one-hot) vector for the true class. The cross-entropy loss on the full dataset is then

$$L(\mathcal{D}; f) = \frac{1}{n} \sum_{i=1}^n H(y_i, f(x_i)) = -\frac{1}{n} \sum_{i=1}^n \sum_{k \in [K]} y_k \log(f(x)_k).$$

The loss is minimized with some form of gradient descent, often stochastic gradient descent (SGD). In SGD, gradients of parameters  $\theta$  are computed over a randomly sampled batch  $B$ , averaged, and scaled by a learning rate  $\eta$ :

$$\theta_{t+1} = \theta_t - \frac{\eta}{|B|} \sum_{i \in B} \nabla_{\theta} H(y_i, f(x_i)).$$

Other optimizers [21–23] use gradient statistics to reduce the variance of updates which can result in better performance.

A less common setting, but one which is important for our work, is when the target values  $y$  which are used to train the network are not one-hot values, but are probability vectors output by a different model  $g(x)$ . When training using the dataset  $\mathcal{D}_g = \{x_i, g(x_i)^{1/T}\}_{i=1}^n$ , we say the trained model is *distilled* from  $g$  with *temperature*  $T$ , referring to the process of distillation introduced in Hinton *et al.* [24]. Note that the values of  $g(x_i)^{1/T}$  are always scaled to sum to 1.

### 3 Taxonomy of Threat Models

We now address the spectrum of adversaries interested in extracting neural networks. As illustrated in Table 1, we taxonomize the space of possible adversaries around two overarching goals—*theft* and *reconnaissance*. We detail why extraction is not always practically realizable by constructing models that are impossible to extract, or require a large number of queries to extract. We conclude our threat model with a discussion of how adversarial capabilities (e.g., prior knowledge of model architecture or information returned by queries) affect the strategies an adversary may consider.

#### 3.1 Adversarial Motivations

Model extraction attacks target the *confidentiality* of a victim model deployed on a remote service. A model refers here to both the architecture and its parameters. Architectural details include the learning hypothesis (i.e., neural network in our case) and corresponding details (e.g., number of layers and activation functions for neural networks). Parameter values are the result of training.

First, we consider *theft* adversaries, motivated by economic incentives. Generally, the defender went through an expensive process to design the model’s architecture and train it to set parameter values. Here, the model can be viewed as intellectual property that the adversary is trying to steal. A line of work has in fact referred to this as “model stealing” [11].

In the latter class of attacks, the adversary is performing *reconnaissance* to later mount attacks targeting other security properties of the learning system: e.g., its integrity with adversarial examples [7], or privacy with training data membership inference [9, 10]. Model extraction enables an adversary previously operating in a black-box threat model to mount attacks against the extracted model in a white-box threat model. The adversary has—by design—access to the extracted model’s parameters. In the limit, this adversary would expect to extract an *exact* copy of the oracle.

The goal of **exact extraction** is to produce  $\hat{O}_\theta = O_\theta$ , so that the model’s architecture and all of its weights are identical to the oracle. This definition is purely a strawman—it is the strongest possible attack, but it is fundamentally impossible for many classes of neural networks, including ReLU networks, because any individual model belongs to a large equivalence class of networks which are indistinguishable from input-output behavior. For example, we can scale an arbitrary neuron’s input weights and biases by some  $c > 0$ , and scale its output weights and biases by  $c^{-1}$ ; the resulting model’s behavior is unchanged. Alternatively, in any intermediate layer of a ReLU network, we may also add a *dead neuron* which never contributes to the output, or might permute the (arbitrary) order of neurons internally. Given access to input-output behavior, the best we can do is identify the equivalence class the oracle belongs to.




Figure 1: Illustrating fidelity vs. accuracy. The solid blue line is the oracle; functionally equivalent extraction recovers this exactly. The green dash-dot line achieves high fidelity: it matches the oracle on all data points. The orange dashed line achieves perfect accuracy: it classifies all points correctly.

#### 3.2 Adversarial Goals

This perspective yields a natural spectrum of realistic adversarial goals characterizing decreasingly precise extractions.

**Functionally Equivalent Extraction** The goal of functionally equivalent extraction is to construct an  $\hat{O}$  such that  $\forall x \in \mathcal{X}, \hat{O}(x) = O(x)$ . This is a tractable weakening of the exact extraction definition from earlier—it is the hardest possible goal using only input-output pairs. The adversary obtains a member of the oracle’s equivalence class. This goal enables a number of downstream attacks, including those involving inspection of the model’s internal representations like over-learning [17], to operate in the white-box threat model.

**Fidelity Extraction** Given some target distribution  $\mathcal{D}_F$  over  $\mathcal{X}$ , and goal similarity function  $S(p_1, p_2)$ , the goal of fidelity extraction is to construct an  $\hat{O}$  that maximizes  $\Pr_{x \sim \mathcal{D}_F} [S(\hat{O}(x), O(x))]$ . In this work, we consider only *label agreement*, where  $S(p_1, p_2) = \mathbb{1}(\arg \max(p_1) = \arg \max(p_2))$ ; we leave exploration of other similarity functions to future work.

A natural distribution of interest  $\mathcal{D}_F$  is the data distribution itself—the adversary wants to make sure the mistakes and correct labels are the same between the two models. A reconnaissance attack for constructing adversarial examples would care about a perturbed data distribution; mistakes might be more important to the adversary in this setting. Membership inference would use the natural data distribution, including any outliers. These distributions tend to be concentrated on a low-dimension manifold of  $\mathcal{X}$ , making fidelity extraction significantly easier than functionally equivalent extraction.

Attack	Type	Model type	Goal	Query Output
Low & Meek [8]	Direct Recovery	LM	Functionally Equivalent	Labels
Tramer <i>et al.</i> [11]	(Active) Learning	LM, NN	Task Accuracy, Fidelity	Probabilities, labels
Tramer <i>et al.</i> [11]	Path finding	DT	Functionally Equivalent	Probabilities, labels
Milli <i>et al.</i> [19] (theoretical)	Direct Recovery	NN (2 layer)	Functionally Equivalent	Gradients, logits
Milli <i>et al.</i> [19]	Learning	LM, NN	Task Accuracy	Gradients
Pal <i>et al.</i> [15]	Active learning	NN	Fidelity	Probabilities, labels
Chandrasekharan <i>et al.</i> [13]	Active learning	LM	Functionally Equivalent	Labels
Copycat CNN [16]	Learning	CNN	Task Accuracy, Fidelity	Labels
Papernot <i>et al.</i> [7]	Active learning	NN	Fidelity	Labels
CSI NN [25]	Direct Recovery	NN	Functionally Equivalent	Power Side Channel
Knockoff Nets [12]	Learning	NN	Task Accuracy	Probabilities
Functionally equivalent (this work)	Direct Recovery	NN (2 layer)	Functionally Equivalent	Probabilities, logits
Efficient learning (this work)	Learning	NN	Task Accuracy, Fidelity	Probabilities

Table 1: Existing Model Extraction Attacks. Model types are abbreviated: LM = Linear Model, NN = Neural Network, DT = Decision Tree, CNN = Convolutional Neural Network.

Indeed, functionally equivalent extraction achieves a perfect fidelity of 1 on *all distributions* and *all similarity functions*.

**Task Accuracy Extraction** For the true task distribution  $\mathcal{D}_A$  over  $X \times \mathcal{Y}$ , the goal of task accuracy extraction is to construct an  $\hat{O}$  maximizing  $\Pr_{(x,y) \sim \mathcal{D}_A} [\arg \max(\hat{O}(x)) = y]$ . This goal is to match (or exceed) the accuracy of the target model, which is the easiest goal to consider in this taxonomy (because it doesn’t need to match the mistakes of  $O$ ).

**Existing Attacks** In Table 1, we fit previous model extraction work into this taxonomy, as well as discuss their techniques. Functionally equivalent extraction has been considered for linear models [8, 13], decision trees [11], both given probabilities, and neural networks [19, 25], given extra access. Task accuracy extraction has been considered for linear models [11] and neural networks [12, 16, 19], and fidelity extraction has also been considered for linear models [11] and neural networks [7, 15]. Notably, functionally equivalent attacks require model-specific techniques, while task accuracy and fidelity typically use generic learning-based approaches.

### 3.3 Model Extraction is Hard

Before we consider adversarial capabilities in Section 3.4 and potential corresponding approaches to model extraction, we must understand how successful we can hope to be. Here, we present arguments that will serve to bound our expectations. First, we will identify some limitations of functionally equivalent extraction by constructing networks which require arbitrarily many queries to extract. Second, we will present another class of networks that cannot be extracted with fidelity without querying a number of times exponential in its depth. We provide intuition in this section and later prove these statements in Appendix A.

**Exponential hardness of functionally equivalent attacks.** In order to show that functionally equivalent extraction is intractable in the worst case, we construct of a class of

neural networks that are hard to extract without making exponentially many queries in the network’s width.

**Theorem 1.** *There exists a class of width  $3k$  and depth  $2$  neural networks on domain  $[0, 1]^d$  (with precision  $p$  numbers) with  $d \geq k$  that require, given logit access to the networks,  $\Theta(p^k)$  queries to extract.*

The precision  $p$  is the number of possible values a feature can take from  $[0, 1]$ . In images with 8-bit pixels, we have  $p = 256$ . The intuition for this theorem is that a width  $3k$  network can implement a function that returns a non-zero value on at most a  $p^{-k}$  fraction of the space. In the worst case,  $p^k$  queries are necessary to find this fraction of the space.

Note that this result assumes the adversary can only observe the input-output behavior of the oracle. If this assumption is broken then functionally equivalent extraction becomes practical. For example, Batina *et al.* [25] perform functionally equivalent extraction by performing a side channel attack (specifically, differential power analysis [26]) on a microprocessor evaluating the neural network.

We also observe in Theorem 2 that, given white-box access to two neural networks, it is NP-hard in general to test if they are functionally equivalent. We do this by constructing two networks that differ only in coordinates satisfying a subset sum instance. Then testing functional equivalence for these networks is as hard as finding the satisfying subset.

**Theorem 2 (Informal).** *Given their weights, it is NP-hard to test whether two neural networks are functionally equivalent.*

Any attack which can claim to perform functionally equivalent extraction efficiently (both in number of queries used and in running time) must make some assumptions to avoid these pathologies. In Section 6, we will present and discuss the assumptions of a functionally equivalent extraction attack for two-layer neural network models.

**Learning approaches struggle with fidelity.** A final difficulty for model extraction comes from recent work in learnability [27]. Das *et al.* prove that, for deep random networks with input dimension  $d$  and depth  $h$ , model extraction approaches that can be written as Statistical Query (SQ) learning algorithms require  $\exp(O(h))$  samples for fidelity extraction. SQ algorithms are a restricted form of learning algorithm which only access the data with noisy aggregate statistics; many learning algorithms, such as (stochastic) gradient descent and PCA, are examples. As a result, most learning-based approaches to model extraction will inherit this inefficiency. A sample-efficient approach therefore must either make assumptions about the model to be extracted (to distinguish it from a deep random network), or must access its dataset without statistical queries.

**Theorem 3** (Informal [27]). *Random networks with domain  $\{0, 1\}^d$  and range  $\{0, 1\}$  and depth  $h$  require  $\exp(O(h))$  samples to learn in the SQ learning model.*

### 3.4 Adversarial Capabilities

We organize an adversary’s prior knowledge about the oracle and its training data into three categories—*domain knowledge*, *deployment knowledge*, and *model access*.

#### 3.4.1 Domain Knowledge

Domain knowledge describes what the adversary knows about the task the model is designed for. For example, if the model is an image classifier, then the model output should not change under standard image data augmentations, such as shifts, rotations, or crops. Usually, the adversary should be assumed to have as much domain knowledge as the oracle’s designer.

In some domains, it is reasonable to assume the adversary has access to public task-relevant pretrained models or datasets. This is often the case for learning-based model extraction, which we develop in Section 4. We consider an adversary using part of a public dataset of 1.3 million images [4] as unlabeled data to mount an attack against a model trained on a proprietary dataset of 1 billion labeled images [28].

#### Learning-based extraction is hard without natural data

In learning-based extraction, we assume that the adversary is able to collect public *unlabeled* data to mount their attack. This is a natural assumption for a theft-motivated adversary who wishes to steal the oracle for local use—the adversary has data they want to learn the labels of without querying the model! For other adversaries, progress in generative modeling is likely to offer ways to remove this assumption [29]. We leave this to future work because our overarching aim in this paper is to characterize the model extraction attacker space around the notions of accuracy and fidelity. All progress achieved by our approaches is complementary to possible progress in synthetic data generation.

#### 3.4.2 Deployment Knowledge

Deployment knowledge describes what the adversary knows about the oracle itself, including the model architecture, training procedure, and training dataset. The adversary may have access to public artifacts of the oracle—a distilled version of the oracle may be available (such as for OpenAI GPT [30]) or the oracle may be transfer learned from a public pretrained model (such as many image classifiers [31] or language models like BERT [32]).

In addition, the adversary may not even know the features (the exact inputs to the model) or the labels (the classes the model may output). While the latter can generally be inferred by interacting with the model (e.g., making queries and observing the labels predicted by the model), inferring the former is usually more difficult. Our preliminary investigations suggest that these are not limiting assumptions, but we leave proper treatment of these constraints to future work.

#### 3.4.3 Model Access

Model access describes the information the adversary obtains from the oracle, including bounds on how many queries the adversary may make as well as the oracle’s response:

- *label*: only the label of the most-likely class is revealed.
- *label and score*: in addition to the most-likely label, the confidence score of the model in its prediction for this label is revealed.
- *top-k scores*: the labels and confidence scores for the  $k$  classes whose confidence are highest are revealed.
- *scores*: confidence scores for all labels are revealed.
- *logits*: raw logit values for all labels are revealed.

In general, the more access an adversary is given, the more effective they should be in accomplishing their goal. We instantiate practical attacks under several of these assumptions. Limiting model access has also been discussed as a defensive measure, as we elaborate in Section 8.

## 4 Learning-based Model Extraction

We present our first attack strategy where the victim model serves as a labeling oracle for the adversary. While many attack variants exist [7, 11], they generally stage an iterative interaction between the adversary and the oracle, where the adversary collects labels for a set of points from the oracle and uses them as a training set for the extracted model. These algorithms are typically designed for accuracy extraction; in this section, we will demonstrate improved algorithms for accuracy extraction, using task-relevant unlabeled data.

We realistically simulate large-scale model extraction by considering an oracle that was trained on 1 billion Instagram

images [28] to obtain (at the time of the experiment) state-of-the-art performance on the standard image classification benchmark, ImageNet [4]. The oracle, with 193 million parameters, obtained 84.2% top-1 accuracy and 97.2% top-5 accuracy on the 1000-class benchmark—we refer to the model as the "WSL model", abbreviating the paper title. We give the adversary access to the public ImageNet dataset. The adversary’s goal is to use the WSL model as a labeling oracle to train an ImageNet classifier that performs better than if we trained the model directly on ImageNet. *The attack is successful if access to the WSL model—trained on 1 billion proprietary images inaccessible to the adversary—enables the adversary to extract a model that outperforms a baseline model trained directly with ImageNet labels.* This is accuracy extraction for the ImageNet distribution, given unlabeled ImageNet training data.

We consider two variants of the attack: one where the adversary selects 10% of the training set (i.e., about 130,000 points) and the other where the adversary keeps the entire training set (i.e., about 1.3 million points). To put this number in perspective, recall that each image has a dimension of 224x224 pixels and 3 color channels, giving us  $224 \cdot 224 \cdot 3 = 150,528$  total input features. Each image belongs to one of 1,000 classes. Although ImageNet data is labeled, we always treat it as unlabeled to simulate a realistic adversary.

## 4.1 Fully-supervised model extraction

The first attack is fully supervised, as proposed by prior work [11]. It serves to compare our subsequent attacks to prior work, and to validate our hypothesis that labels from the oracle are more informative than dataset labels.

The adversary needs to obtain a label for each of the points it intends to train the extracted model with. Then it queries the oracle to label its training points with the oracle’s predictions. The oracle reveals *labels and scores* (in the threat model from Section 3) when queried.

The adversary then trains its model to match these labels using the cross-entropy loss. We used a distillation temperature of  $T = 1.5$  in our experiments after a random search. Our experiments use two architectures known to perform well on image classification: ResNet-v2-50 and ResNet-v2-200.

**Results.** We present results in Table 2. For instance, the adversary is able to improve the accuracy of their model by 1.0% for ResNet2-50 and 1.9% for ResNet\_v2\_200 after having queried the oracle for 10% of the ImageNet data. Recall that the task has 1,000 labels, making these improvements significant. The gains we are able to achieve as an adversary are in line with progress that has been made by the computer vision community on the ImageNet benchmark over recent years, where the research community improved the state-of-the-art top-1 accuracy by about one percent point per year.<sup>1</sup>

<sup>1</sup><https://paperswithcode.com/sota/image-classification-on-imagenet>

## 4.2 Unlabeled data improves query efficiency

For adversaries interested in theft, a learning-based strategy should minimize the number of queries required to achieve a given level of accuracy. A natural approach towards this end is to take advantage of advances in label-efficient ML, including active learning [33] and semi-supervised learning [34].

Active learning allows a learner to query the labels of arbitrary points—the goal is to query the best set of points to learn a model with. Semi-supervised learning considers a learner with some labeled data, but much more unlabeled data—the learner seeks to leverage the unlabeled data (for example, by training on guessed labels) to improve classification performance. Active and semi-supervised learning are complementary techniques [35, 36]; it is possible to pick the best subset of data to train on, while also using the rest of the unlabeled data without labels.

The connection between label-efficient learning and learning-based model extraction attacks is not new [11, 13, 15], but has focused on active learning. *We show that, assuming access to unlabeled task-specific data, semi-supervised learning can be used to improve model extraction attacks.* This could potentially be improved further by leveraging active learning, as in prior work, but our improvements are overall complementary to approaches considered in prior work. We explore two semi-supervised learning techniques: rotation loss [37] and MixMatch [38].

**Rotation loss.** We leverage the current state-of-the-art semi-supervised learning approach on ImageNet, which augments the model with a *rotation loss* [37]. The model contains two linear classifiers from the second-to-last layer of the model: the classifier for the image classification task, and a rotation predictor. The goal of the rotation classifier is to predict the rotation applied to an input—each input is fed in four times per batch, rotated by  $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$ . The classifier should output one-hot encodings  $\{OH(0;4), OH(1;4), OH(2;4), OH(3;4)\}$ , respectively, for these rotated images. Then, the rotation loss is written:

$$L_R(X; f_\theta) = \frac{1}{4N} \sum_{i=0}^N \sum_{j=1}^r H(f_\theta(R_j(x_i)), j)$$

where  $R_j$  is the  $j$ th rotation,  $H$  is cross-entropy loss, and  $f_\theta$  is the model’s probability outputs for the rotation task. Inputs need not be labeled, hence we compute this loss on unlabeled data for which the adversary did not query the model. That is, we train the model on both unlabeled data (with rotation loss), and labeled data (with standard classification loss), and both contribute towards learning a good representation for all of the data, including the unlabeled data.

We compare the accuracy of models trained with the rotation loss on data labeled by the oracle and data with ImageNet labels. Our best performing extracted model, with an accuracy

Architecture	Data Fraction	ImageNet	WSL	WSL-5	ImageNet + Rot	WSL + Rot	WSL-5 + Rot
Resnet_v2_50	10%	(81.86/82.95)	(82.71/84.18)	(82.97/84.52)	(82.27/84.14)	(82.76/84.73)	(82.84/84.59)
Resnet_v2_200	10%	(83.50/84.96)	(84.81/86.36)	(85.00/86.67)	(85.10/86.29)	(86.17/88.16)	(86.11/87.54)
Resnet_v2_50	100%	(92.45/93.93)	(93.00/94.64)	(93.12/94.87)	N/A	N/A	N/A
Resnet_v2_200	100%	(93.70/95.11)	(94.26/96.24)	(94.21/95.85)	N/A	N/A	N/A

Table 2: Extraction attack (top-5 accuracy/top-5 fidelity) of the WSL model [28]. Each row contains an architecture and fraction of public ImageNet data used by the adversary. ImageNet is a baseline using only ImageNet labels. WSL is an oracle returning WSL model probabilities. WSL-5 is an oracle returning only the top 5 probabilities. Columns with (+ Rot) use rotation loss on unlabeled data (rotation loss was not run when all data is labeled). An adversary able to query WSL always improves over ImageNet labels, even when given only top 5 probabilities. Rotation loss does not significantly improve the performance on ResNet\_v2\_50, but provides a (1.36/1.80) improvement for ResNet\_v2\_200, comparable to the performance boost given by WSL labels on 10% data. In the high-data regime, where we observe a (0.56/1.13) improvement using WSL labels.

Dataset	Algorithm	250 Queries	1000 Queries	4000 Queries
SVHN	FS	(79.25/79.48)	(89.47/89.87)	(94.25/94.71)
SVHN	MM	(95.82/96.38)	(96.87/97.45)	(97.07/97.61)
CIFAR10	FS	(53.35/53.61)	(73.47/73.96)	(86.51/87.37)
CIFAR10	MM	(87.98/88.79)	(90.63/91.39)	(93.29/93.99)

Table 3: Performance (accuracy/fidelity) of fully supervised (FS) and MixMatch (MM) extraction on SVHN and CIFAR10. MixMatch with 4000 labels performs nearly as well as the oracle for both datasets, and MixMatch at 250 queries beats fully supervised training at 4000 queries for both datasets.

of 64.5%, is trained with the rotation loss on oracle labels whereas the baseline on ImageNet labels only achieves 62.5% accuracy with the rotation loss and 61.2% without the rotation loss. This demonstrates the cumulative benefit of adding a rotation loss to the objective and training on oracle labels for a theft-motivated adversary.

We expect that as semi-supervised learning techniques on ImageNet mature, further gains should be reflected in the performance of model extraction attacks.

**MixMatch.** To validate this hypothesis, we turn to smaller datasets where semi-supervised learning has made significant progress. We investigate a technique called MixMatch [38] on two datasets: SVHN [39] and CIFAR10 [40]. MixMatch uses a combination of techniques, including training on "guessed" labels, regularization, and image augmentations.

For both datasets, inputs are color images of 32x32 pixels belonging to one of 10 classes. The training set of SVHN contains 73257 images and the test set contains 26032 images. The training set of CIFAR10 contains 50000 images and the test set contains 10000 images. We train the oracle with a WideResNet-28-2 architecture on the labeled training set. The oracles achieve 97.36% accuracy on SVHN and 95.75% accuracy on CIFAR10.

The adversary is given access to the same training set but without knowledge of the labels. Our goal is to validate the effectiveness of semi-supervised learning by demonstrating

that the adversary only needs to query the oracle on a small subset of these training points to extract a model whose accuracy on the task is comparable to the oracle’s. To this end, we run 5 trials of fully supervised extraction (no use of unlabeled data), and 5 trials of MixMatch, reporting for each trial the median accuracy of the 20 latest checkpoints, as done in [38].

**Results.** In Table 3, we find that with only 250 queries (293x smaller label set than the SVHN oracle and 200x smaller for CIFAR10), MixMatch reaches 95.82% test accuracy on SVHN and 87.98% accuracy on CIFAR10. This is higher than fully supervised training that uses 4000 queries. With 4000 queries, MixMatch is within 0.29% of the accuracy of the oracle on SVHN, and 2.46% on CIFAR10. The variance of MixMatch is slightly higher than that of fully supervised training, but is much smaller than the performance gap. These gains come from the prior MixMatch is able to build using the unlabeled data, making it effective at exploiting few labels. We observe similar gains in test set fidelity.

## 5 Limitations of Learning-Based Extraction

Learning-based approaches have several sources of non-determinism: the random initializations of the model parameters, the order in which data is assembled to form batches for SGD, and even non-determinism in GPU instructions [41, 42]. Non-determinism impacts the model parameter values obtained from training. Therefore, even an adversary with full access to the oracle’s training data, hyperparameters, etc., would still need all of the learner’s non-determinism to achieve the *functionally equivalent* extraction goal described in Section 3. In this section, we will attempt to quantify this: for a strong adversary, with access to the exact details of the training setup, we will present an experiment to determine the limits of learning-based algorithms to achieving fidelity extraction.

We perform the following experiment. We query an oracle to obtain a labeled substitute dataset  $\mathcal{D}$ . We use  $\mathcal{D}$  for a learning-based extraction attack which produces a model  $f_{\theta}^1(x)$ . We run the learning-based attack a second time using  $\mathcal{D}$ , but with different sources of non-determinism to obtain

Query Set	Init & SGD	Same SGD	Same Init	Different
Test	93.7%	93.2%	93.1%	93.4%
Adv Ex	73.6%	65.4%	65.3%	67.1%
Uniform	65.7%	60.2%	59.0%	60.2%

Table 4: Impact of non-determinism on extraction fidelity. Even models extracted using the same SGD and initialization randomness as the oracle do not reach 100% fidelity.

a new set of parameters  $f_{\theta}^2(x)$ . If there are points  $x$  such that  $f_{\theta}^1(x) \neq f_{\theta}^2(x)$ , then the prediction on  $x$  is dependent not on the oracle, but on the non-determinism of the learning-based attack strategy—we are unable to guarantee fidelity.

We independently control the initialization randomness and batch randomness during training on Fashion-MNIST [43] with fully supervised SGD (we use Fashion-MNIST for training speed). We repeated each run 10 times and measure agreement between the ten obtained models on the test set, adversarial examples generated by running FGSM with  $\epsilon = 25/255$  with the oracle model and the test set, and uniformly random inputs. The oracle uses initialization seed 0 and SGD seed 0—we also use two different initialization and SGD seeds.

Even when both training and initialization randomness are fixed (so that only GPU non-determinism remains), fidelity peaks at 93.7% on the test set (see Table 4). With no randomness fixed, extraction achieves 93.4% fidelity on the test set. (Agreement on the test set should be considered in reference to the base test accuracy of 90%.) Hence, even an adversary who has the victim model’s *exact* training set will be unable to exceed  $\sim 93.4\%$  fidelity. Using prototypicality metrics, as investigated in Carlini et al. [44], we notice that test points where fidelity is easiest to achieve are also the most prototypical (i.e., more representative of the class it is labeled as). This connection is explored further in Appendix B. The experiment of this section is also related to uncertainty estimation using deep ensembles [42]; we believe a deeper connection may exist between the fidelity of learning-based approaches and uncertainty estimation. Also relevant is the work mentioned earlier in Section 3, that shows that random networks are hard for learning-based approaches to extract. Here, we find that learning-based approaches have limits even for trained networks, on some portion of the input space.

It follows from these arguments that non-determinism of both the victim and extracted model’s learning procedures potentially compound, limiting the effectiveness of using a learning-based approach to reaching high fidelity.

## 6 Functionally Equivalent Extraction

Having identified fundamental limitations that prevent learning-based approaches from perfectly matching the oracle’s mistakes, we now turn to a different approach where the adversary extracts the oracle’s weights directly, seeking

to achieve functionally-equivalent extraction.

This attack can be seen as an extension of two prior works.

- Milli *et al.* [19] introduce an attack to extract neural network weights under the assumption that the adversary is able to make *gradient queries*. That is, each query the adversary makes reveals not only the prediction of the neural network, but also the gradient of the neural network with respect to the query. To the best of our knowledge this is the only functionally-equivalent extraction attack on neural networks with one hidden layer, although it was not actually implemented in practice.
- Batina *et al.* [25], at USENIX Security 2019, develop a side-channel attack that extracts neural network weights through monitoring the power use of a microprocessor evaluating the neural network. This is a much more powerful threat model than made by any of the other model extraction papers. To the best of our knowledge this is the only practical direct model extraction result—they manage to extract essentially arbitrary depth networks.

In this section we introduce an attack which only requires standard queries (i.e., that return the model’s prediction instead of its gradients) and does not require any side-channel leakages, yet still manages to achieve *higher fidelity* extraction than the side-channel extraction work for two-layer networks, assuming double-precision inference.

**Attack Algorithm Intuition.** As in [19], our attack is tailored to work on neural networks with the ReLU activation function (the ReLU is an effective default choice of activation function [20]). This makes the neural network a piecewise linear function. Two samples are within the same linear region if all ReLU units have the same sign, illustrated in Figure 2.

By finding adjacent linear regions, and computing the difference between them, we force a single ReLU to change signs. Doing this, it is possible to almost completely determine the weight vector going into that ReLU unit. Repeating this attack for all ReLU units lets us recover the first weight matrix completely. (We say almost here, because we must do some work to recover the sign of the weight vector.) Once the first layer of the two-layer neural network has been determined, the second layer can be uniquely solved for algebraically through least squares. This attack is optimal up to a constant factor—the query complexity is discussed in Appendix D.

### 6.1 Notation and Assumptions

As in [19], we only aim to extract neural networks with one hidden layer using the ReLU activation function. We denote the model weights by  $A^{(0)} \in \mathbb{R}^{d \times h}$ ,  $A^{(1)} \in \mathbb{R}^{h \times K}$  and biases by  $B^{(0)} \in \mathbb{R}^h$ ,  $B^{(1)} \in \mathbb{R}^K$ . Here,  $d$ ,  $h$ , and  $K$  respectively refer



Symbol	Definition
$d$	Input dimensionality
$h$	Hidden layer dimensionality ( $h < d$ )
$K$	Number of classes
$A^{(0)} \in \mathbb{R}^{d \times h}$	Input layer weights
$B^{(0)} \in \mathbb{R}^h$	Input layer bias
$A^{(1)} \in \mathbb{R}^{h \times K}$	Logit layer weights
$B^{(1)} \in \mathbb{R}^K$	Logit layer bias

Table 5: Parameters for the functionally-equivalent attack.




Figure 2: 2-dimension intuition for the functionally equivalent extraction attack.

to the input dimensionality, the size of the hidden layer, and the number of classes. This is found in Table 6.1.

We say that  $\text{ReLU}(x)$  is at a critical point if  $x = 0$ ; this is the location at which the unit’s gradient changes from 0 to 1. We assume the adversary is able to observe the raw logit outputs as 64-bit floating point values. We will use the notation  $O_L$  to denote the logit oracle. Our attack implicitly assumes that the rows of  $A^{(0)}$  are linearly independent. Because the dimension of the input space is larger than the hidden space by at least 100, it is exceedingly unlikely for the rows to be linearly dependent (and we find this holds true in practice).

Note that our attack is not an SQ algorithm, which would only allow us to look at aggregate statistics of our dataset. Instead, our algorithm is very particular in its analysis of the network, computing the differences between linear regions, for example, cannot be done with aggregate statistics. This structure allows us to avoid the pathologies of Section 3.3.

## 6.2 Attack Overview

The algorithm is broken into four phases:

- **Critical point search** identifies inputs  $\{x_i\}_{i=1}^n$  to the neural network so that exactly one of the ReLU units is at a critical point (i.e., has input identically 0).
- **Weight recovery** takes an input  $x$  which causes the  $i$ th neuron to be at a critical point. We use this point  $x$  to compute the difference between the two adjacent linear regions induced by the critical point, and thus the weight vector row  $A_i^{(0)}$ . By repeating this process for each ReLU we obtain the complete matrix  $A^{(0)}$ . Due to technical reasons discussed below, we can only recover the row-vector up to sign.
- **Sign recovery** determines the sign of each row-vector  $A_j^{(0)}$  for all  $j$  using global information about  $A^{(0)}$ .
- **Final layer extraction** uses algebraic techniques (least squares) to solve for the second layer of the network.

## 6.3 Critical Point Search

For a two layer network, observe that the logit function is given by the equation  $O_L(x) = A^{(1)}\text{ReLU}(A^{(0)}x + B^{(0)}) + B^{(1)}$ . To find a critical point for every ReLU, we sample two random vectors  $u, v \in \mathbb{R}^d$ , and consider the function

$$L(t; u, v, O_L) = O_L(u + tv).$$

for  $t$  varying between a small and large appropriately selected value (discussed below). This amounts to drawing a line in the inputs of the network; passed through ReLUs, this line becomes the piecewise linear function  $L(\cdot)$ . The points  $t$  where  $L(t)$  is non-differentiable are exactly locations where some  $\text{ReLU}_i$  is changing signs (i.e., some ReLU is at a critical point). Figure 3 shows an example of what this sweep looks like on a trained MNIST model.

Furthermore, notice that given a pair  $u, v$ , there is exactly one value  $t$  for which each ReLU is at a critical point, and if  $t$  is allowed to grow arbitrarily large or small that every ReLU unit will switch sign exactly once. Intuitively, the reason this is true is that each ReLU’s input, (say  $w x + b$  for some  $w, b$ ), is a monotone function of  $t$  ( $w^T u t + w^T v + b$ ). Thus, by varying  $t$ , we can identify an input  $x_i$  that sets the  $i$ th ReLU to 0 for every relu  $i$  in the network. This assumes we are not moving parallel to any of the rows (where  $w^T u = 0$ ), and that we vary  $t$  within a sufficiently large interval (so the  $w^T u t$  term may overpower the constant term). The analysis of [19] suggests that these concerns can be resolved with high probability by varying  $t \in [-h^2, h^2]$ .

While in theory it would be possible to sweep all values of  $t$  to identify the critical points, this would require a large number of queries. Thus, to efficiently search for the locations




Figure 3: An example sweep for critical point search. Here we plot the partial derivative across  $t$  and see that  $O_L(u+tv)$  is piecewise linear, enabling a binary search.




Figure 4: Efficient and accurate 2-linear testing subroutine in Algorithm 1. Left shows a successful case where the algorithm succeeds; right shows a potential failure case, where there are multiple nonlinearities. We detect this by observing the expected value of  $O(x)$  is not the observed (queried) value.

of critical points, we introduce a refined search algorithm which improves on the binary search as used in [19]. Standard binary search requires  $O(n)$  model queries to obtain  $n$  bits of precision. Therefore, we propose a refined technique which does not have this restriction and requires just  $O(1)$  queries to obtain high (20+ bits) precision. The key observation we make is that if we are searching between two values  $[t_1, t_2]$  and there is exactly one discontinuity in this range, we can precisely identify the location of that discontinuity efficiently.

An intuitive diagram for this algorithm can be found in Figure 4 and the algorithm can be found in Algorithm 1. The property this leverages is that the function is piecewise linear—if we know the range is composed of two linear segments, we can identify the linear segments and compute their intersection. In Algorithm 1, lines 1-3 describe computing the two linear regions' slopes and intercepts. Lines 4 and 5 compute the intersection of the two lines (also shown in the red dotted line of Figure 4). The remainder of the algorithm performs

---

**Algorithm 1** Algorithm for 2-linearity testing. Computes the location of the only critical point in a given range or rejects if there is more than one.

---

```

Function  $f$ , range  $[t_1, t_2]$ ,  $\epsilon$ 
 $m_1 = \frac{f(t_1+\epsilon) - f(t_1)}{\epsilon}$  ▷ Gradient at  $t_1$ 
 $m_2 = \frac{f(t_2) - f(t_2-\epsilon)}{\epsilon}$  ▷ Gradient at  $t_2$ 
 $y_1 = f(a), y_2 = f(b)$ 
 $x = a + \frac{y_2 - y_1 - (b-a)m_2}{m_1 - m_2}$  ▷ Candidate critical point
 $\hat{y} = y_1 + m_1 \frac{y_2 - y_1 - (b-a)m_2}{m_1 - m_2}$  ▷ Expected value at candidate
 $y = f(x)$  ▷ True value at candidate
if  $\hat{y} = y$  then return  $x$ 
else return "More than one critical point"
end if

```

---

the correctness check, also illustrated in Figure 4; if there are more than 2 linear components, it is unlikely that the true function value will match the function value computed in line 5, and we can detect that the algorithm has failed.

## 6.4 Weight Recovery

After running critical point search we obtain a set  $\{x_i\}_{i=1}^h$ , where each critical point corresponds to a point where a single ReLU flips sign. In order to use this information to learn the weight matrix  $A^{(0)}$  we measure the second derivative of  $O_L$  in each input direction at the points  $x_i$ . Taking the second derivative here corresponds to measuring the difference between the linear regions on either side of the ReLU. Recall that prior work assumed direct access to gradient queries, and thus did not require any of the analysis in this section.

### 6.4.1 Absolute Value Recovery

To formalize the intuition of comparing adjacent hyperplanes, observe that for the oracle  $O_L$  and for a critical point  $x_i$  (corresponding to  $\text{ReLU}_i$  being zero) and for a random input-space direction  $e_j$  we have

$$\begin{aligned}
\left. \frac{\partial^2 O_L}{\partial e_j^2} \right|_{x_i} &= \left. \frac{\partial O_L}{\partial e_j} \right|_{x_i + c \cdot e_j} - \left. \frac{\partial O_L}{\partial e_j} \right|_{x_i - c \cdot e_j} \\
&= \sum_k A_k^{(1)} \mathbb{1}(A_k^{(0)}(x_i + c \cdot e_j) + B_k^{(0)} > 0) A_{kj}^{(0)} \\
&\quad - \sum_k A_k^{(1)} \mathbb{1}(A_k^{(0)}(x_i - c \cdot e_j) + B_k^{(0)} > 0) A_{kj}^{(0)} \\
&= A_i^{(1)} \left( \mathbb{1}(A_i^{(0)} \cdot e_j > 0) - \mathbb{1}(-A_i^{(0)} \cdot e_j > 0) \right) A_{ji}^{(0)} \\
&= \pm (A_{ji}^{(0)} A_i^{(1)})
\end{aligned}$$

for a  $c > 0$  small enough so that  $x_i \pm c \cdot e_j$  does not flip any other ReLU. Because  $x_i$  is a critical point and  $c$  is small, the sums in the second line differ only in the contribution of

ReLU<sub>*i*</sub>. However at this point we only have a product involving both weight matrices. We now show this information is useful.

If we compute  $|A_{1i}^{(0)} A^{(1)}|$  and  $|A_{2i}^{(0)} A^{(1)}|$  by querying along directions  $e_1$  and  $e_2$ , we can divide these quantities to obtain the value  $|A_{1i}^{(0)} / A_{2i}^{(0)}|$ , the ratio of the two weights. By repeating the above process for each input direction we can, for all  $k$ , obtain the pairwise ratios  $|A_{1i}^{(0)} / A_{ki}^{(0)}|$ .

Recall from Section 3 that obtaining the ratios of weights is the theoretically optimal result we could hope to achieve. It is always possible to multiply all of the weights *into* a ReLU by a constant  $c > 0$  and then multiply all of the weights *out* of the ReLU by  $c^{-1}$ . Thus, without loss of generality, we can assign  $A_{1i}^{(0)} = 1$  and scale the remaining entries accordingly. Unfortunately, we have lost a small amount of information here. We have only learned the absolute value of the ratio, and not the value itself.

### 6.4.2 Weight Sign Recovery

Once we reconstruct the values  $|A_{ji}^{(0)} / A_{1i}^{(0)}|$  for all  $j$  we need to recover the sign of these values. To do this we consider the following quantity:

$$\left. \frac{\partial^2 O_L}{\partial(e_j + e_k)^2} \right|_{x_i} = \pm(A_{ji}^{(0)} A_i^{(1)} \pm A_{ki}^{(0)} A_i^{(1)}).$$

That is, we consider what would happen if we take the second partial derivative in the direction  $(e_j + e_k)$ . Their contributions to the gradient will either cancel out, indicating  $A_{ji}^{(0)}$  and  $A_{ki}^{(0)}$  are of opposite sign, or they will compound on each other, indicating they have the same sign. Thus, to recover signs, we can perform this comparison along each direction  $(e_1 + e_j)$ .

Here we encounter one final difficulty. There are a total of  $n$  signs we need to recover, but because we compute the signs by comparing ratios along different directions, we can only obtain  $n - 1$  relations. That is, we now know the correct signed value of  $A_i^{(0)}$  up to a single sign for the entire row.

It turns out this is to be expected. What we have computed is the normal direction to the hyperplane, but because any given hyperplane can be described by an infinite number of normal vectors differing by a constant scalar, we can not hope to use local information to recover this final sign bit.

Put differently, while it is possible to push a constant  $c > 0$  through from the first layer to the second layer, it is not possible to do this for negative constants, because the ReLU function is not symmetric. Therefore, it is necessary to learn the sign of this row.

### 6.5 Global Sign Recovery

Once we have recovered the input vector’s weights, we still don’t know the sign for the given inputs—we only measure the difference between linear functions at each critical point, but

do not know which side is the positive side of the ReLU [19]. Now, we need to leverage global information in order to reconcile all of inputs’ signs.

Notice that recovering  $\hat{A}_i^{(0)}$  allows us to obtain  $B_i^{(0)}$  by using the fact that  $A_i^{(0)} \cdot x_i + B_i^{(0)} = 0$ . Then we can compute  $\hat{B}_i^{(0)}$  up to the same global sign as is applied to  $\hat{A}_i^{(0)}$ .

Now, to begin recovering sign, we search for a vector  $z$  that is in the null space of  $\hat{A}^{(0)}$ , that is,  $\hat{A}^{(0)} z = \vec{0}$ . Because the neural network has  $h < d$ , the null-space is non-zero, and we can find many such vectors using least squares. Then, for each ReLU<sub>*i*</sub>, we search for a vector  $v_i$  such that  $v_i A^{(0)} = e_i$  where here  $e_i$  is the *i*th basis vector in the hidden space. That is, moving along the  $v_i$  direction only changes ReLU<sub>*i*</sub>’s input value. Again we can search for this through least squares.

Given  $z$  and these  $v_i$  we query the neural network for the values of  $O_L(z)$ ,  $O_L(z + v_i)$ , and  $O_L(z - v_i)$ . On each of these three queries, all hidden units are 0 except for ReLU<sub>*i*</sub> which receives as input either 0, 1, or  $-1$  by the construction of  $v_i$ . However, notice that the output of ReLU<sub>*i*</sub> can only be either 0 or 1, and the two  $\{-1, 0\}$  cases collapse to just output 0. Therefore, if  $O_L(z + v_i) = O_L(z)$ , we know that  $A_i^{(0)} \cdot v_i < 0$ . Otherwise, we will find  $O_L(z - v_i) = O_L(z)$  and  $A_i^{(0)} \cdot v_i > 0$ . This allows us to recover the sign bit for ReLU<sub>*i*</sub>.

### 6.6 Last Layer Extraction

Given the completely extracted first layer, the logit function of the network is just a linear transformation which we can recover with least squares, through making  $h$  queries where each ReLU is active at least once. In practice, we use the critical points discovered in the previous section so that we do not need to make additional neural network queries.

### 6.7 Results

**Setup.** We train several one-layer fully-connected neural networks with between 16 and 512 hidden units (for 12,000 and 100,000 trainable parameters, respectively) on the MNIST [45] and CIFAR-10 datasets [40]. We train the models with the Adam [23] optimizer for 20 epochs at batch size 128 until they converge. We train five networks of each size to obtain higher statistical significance. Accuracies of these networks can be found in the supplement in Appendix C. In Section 4, we used  $140,000 \approx 2^{17}$  queries for ImageNet model extraction. This is comparable to the number of queries used to extract the smallest MNIST model in this section, highlighting the advantages of both approaches.

**MNIST Extraction.** We implement the functionally-equivalent extraction attack in JAX [46] and run it on each trained oracle. We measure the fidelity of the extracted model, comparing predicted labels, on the MNIST test set.

Results are summarized in Table 6. For smaller networks, we achieve 100% fidelity on the test set: every single one

of the 10,000 test examples is predicted the same. As the network size increases, low-probability errors we encounter become more common, but the extracted neural network still disagrees with the oracle on only 2 of the 10,000 examples.

Inspecting the weight matrix that we extract and comparing it to the weight matrix of the oracle classifier, we find that we manage to reconstruct the first weight matrix to an average precision of 23 bits—we provide more results in Appendix C.

**CIFAR-10 Extraction.** Because this attack is data-independent, the underlying task is unimportant for how well the attack works; only the number of parameters matter. The results for CIFAR-10 are thus identical to MNIST when controlling for model size: we achieve 100% test set agreement on models with fewer than 200,000 parameters and and greater than 99% test set agreement on larger models.

**Comparison to Prior Work.** To the best of our knowledge, this is by orders of magnitude the highest fidelity extraction of neural network weights.

The only fully-implemented neural network extraction attack we are aware of is the work of Batina *et al.* [25], who uses an electromagnetic side channels and differential power analysis to recover an MNIST neural network with neural network weights with an average error of 0.0025. In comparison, we are able to achieve an average error in the first weight matrix for a similarly sized neural network of just 0.0000009—over two thousand times more precise. To the best of our knowledge no functionally-equivalent CIFAR-10 models have been extracted in the past.

We are unable to make a comparison between the fidelity of our extraction attack and the fidelity of the attack presented in Batina *et al.* because they do not report on this number: they only report the accuracy of the extracted model and show it is similar to the original model. We believe this strengthens our observation that comparing across accuracy and fidelity is not currently widely accepted as best practice.

**Investigating Errors.** We observe that as the number of parameters that must be extracted increases, the fidelity of the model decreases. We investigate why this happens and discovered that a small fraction of the time (roughly 1 in 10,000) the gradient estimation procedure obtains an incorrect estimate of the gradient and therefore one of the extracted weights  $\hat{A}_{ij}^{(0)}$  is incorrect by a non-insignificant margin.

Introducing an error into just one of the weights of the first matrix  $\hat{A}^{(0)}$  should not induce significant further errors. However, because of this error, when we solve for the bias vector, the extracted bias  $\hat{B}_i^{(0)}$  will have error proportional to the error of  $\hat{A}_{ij}^{(0)}$ . And when the bias is wrong, it impacts *every* calculation, even those where this edge is not in use.

Resolving this issue completely either requires reducing the failure rate of gradient estimation from 1 in 10,000 to practically 0, or would require a complex error-recovery procedure. Instead, we will introduce in the following section an improvement which almost completely solves this issue.

# of Parameters	12,500	25,000	50,000	100,000
Fidelity	100%	100%	100%	99.98%
Queries	$2^{17.2}$	$2^{18.2}$	$2^{19.2}$	$2^{20.2}$

Table 6: Fidelity of the functionally-equivalent extraction attack across different test distributions on an MNIST victim model. Results are averaged over five extraction attacks. For small models, we achieve perfect fidelity extraction; larger models have near-perfect fidelity on the test data distribution, but begins to lose accuracy at 100,000 parameters.

**Difficulties Extending the Attack.** The attack is specific to two layer neural networks; deeper networks pose multiple difficulties. In deep networks, the critical point search step of Section 6.3 will result in critical points from many different layers, and determining which layer a critical point is on is nontrivial. Without knowing which layer a critical point is on, we cannot control inputs to the neuron, which we need to do to recover the weights in Section 6.4. Even given knowledge of what layer a critical point is on, the inputs of any neuron past layer 1 are the outputs of other neurons, so we only have indirect control over their inputs. Finally, even with the ability to recover these weights, small numerical errors occur in the first layer extraction. These cause errors in every finite differences computation in further layers, causing the second layer to have even larger numerical errors than the first (and so on). Therefore, extending the attack to deeper networks will require at least solving each of the following: producing critical points belonging to a specific layer, recovering weights for those neurons without direct control of their inputs, and significantly reducing numerical errors in these algorithms.

## 7 Hybrid Strategies

Until now the strategies we have developed for extraction have been pure and focused entirely on learning or entirely on direct extraction. We now show that there is a continuous spectrum from which we can draw attack strategies, and these *hybrid* strategies can leverage both the query efficiency of learning extraction, and the fidelity of direct extraction.

### 7.1 Learning-Based Extraction with Gradient Matching

Milli *et al.* demonstrate that *gradient matching* helps extraction by optimizing the objective function

$$\sum_{i=1}^n H(O(x_i), f(x_i)) + \alpha \|\nabla_x O(x_i) - \nabla_x f(x_i)\|_2^2,$$

assuming the adversary can query the model for  $\nabla_x O(x)$ . This is more model access than we permit our adversary, but is an

# of Parameters	50,000	100,000	200,000	400,000
<b>Fidelity</b>	100%	100%	99.95%	99.31%
<b>Queries</b>	$2^{19.2}$	$2^{20.2}$	$2^{21.2}$	$2^{22.2}$

Table 7: Fidelity of extracted MNIST model is improved with the hybrid strategy. Note when comparing to Table 6 the model sizes are  $4\times$  larger.

example of using intuition from direct recovery to improve extraction. We found in preliminary experiments that this technique can improve fidelity on small datasets (increasing fidelity from 95% to 96.5% on Fashion-MNIST), but we leave scaling and removing the model access assumption of this technique to future work. Next, we will show another combination of learning and direct recovery, using learning to alleviate some of the limitations of the previous functionally-equivalent extraction attack.

## 7.2 Error Recovery through Learning

Recall from earlier that the functionally-equivalent extraction attack fidelity degrades as the model size increases. This is a result of low-probability errors in the first weight matrix inducing incorrect biases on the first layer, which in turn propagates and causes worse errors in the second layer.

We now introduce a method for performing a learning-based error recovery routine. While performing a fully-learning-based attack leaves too many free variables so that functionally-equivalent extraction is not possible, if we fix many of the variables to the values extracted through the direct recovery attack, we now show it is possible to learn the remainder of the variables.

Formally, let  $\hat{A}^{(0)}$  be the extracted weight matrix for the first layer and  $\hat{B}^{(0)}$  be the extracted bias vector for the first layer. Previously, we used least squares to directly solve for  $\hat{A}^{(1)}$  and  $\hat{B}^{(1)}$  assuming we had extracted the first layer perfectly. Here, we relax this assumption. Instead, we perform gradient descent optimizing for parameters  $W_{0,2}$  that minimize

$$\mathbb{E}_{x \in \mathcal{D}} \|f_{\theta}(x) - W_1 \text{ReLU}(\hat{A}^{(0)}x + \hat{B}^{(0)} + W_0) + W_2\|$$

That is, we use a single trainable parameter to adjust the bias term of the first layer, and then solve (via gradient descent with training data) for the remaining weights accordingly.

This hybrid strategy increases the fidelity of the extracted model substantially, detailed in Table 8. In the worst-performing example from earlier (with only direct extraction) the extracted 128-neuron network had 80% fidelity agreement with the victim model. When performing learning-based recovery, the fidelity agreement jumps all the way to 99.75%.

# of Parameters	50,000	100,000	200,000	400,000
<b>Transferability</b>	100%	100%	100%	100%

Table 8: Transferability rate of adversarial examples using the extracted neural network from our Section 7 attack.

### 7.2.1 Transferability

Adversarial examples *transfer*: an adversarial example [47] generated on one model often fools different models, too. Transferability is higher when the models are more similar [7].

We should therefore expect that we can generate adversarial examples on our extracted model, and that these will fool the remote oracle nearly always. In order to measure transferability, we run 20 iterations of PGD [48] with  $\ell_{\infty}$  distortion set to the value most often used in the literature: for MNIST: 0.1, and for CIFAR-10: 0.03.

The attack achieves functionally equivalent extraction (modulo floating point precision errors in the extracted weights), so we expect it to have high adversarial example transferability. Indeed, we find we achieve a 100% transferability success rate for all extracted models.

## 8 Related Work

Defenses for model extraction have fallen into two camps: limiting the information gained per query, and differentiating extraction adversaries from benign users. Approaches to limiting information include perturbing the probabilities returned by the model [11, 13, 49], removing the probabilities for some of the model’s classes [11], or returning only the class output [11, 13]. Another proposal has considered sampling from a distribution over model parameters [13, 50]. The other camp, differentiating benign from malicious users, has focused on analyzing query patterns [51, 52]. Non-adaptive attacks (such as supervised or MixMatch extraction) bypass query pattern-based detection, and are weakened by information limiting. We demonstrate the impact of removing complete access to probability values by considering only access to top 5 probabilities from WSL in Table 2. Our functionally-equivalent attack is broken by all of these measures. We leave consideration of defense-aware attacks to future work.

Queries to a model can also reveal hyperparameters [53] or architectural information [14]. Adversaries can use side channel attacks to do the same [18, 25]. These are orthogonal to, but compatible with, our work—information about a model, such as assumptions made in Section 6, empowers extraction.

Watermarking neural networks has been proposed [54, 55] to identify extracted models. Model extraction calls into question the utility of cryptographic protocols used to protect model weights. One unrealized approach is obfuscation [56], where an equivalent program could be released and queried as

many times as desired. A practical approach is secure multi-party computation, where each query is computed by running a protocol between the model owner and querier [57].

## 9 Conclusion

This paper characterizes and explores the space of model extraction attacks on neural networks. We focus this paper specifically around the objectives of *accuracy*, to measure the success of a theft-motivated adversary, and *fidelity*, an often-overlooked measure which compares the agreement between models to reflect the success of a recon-motivated adversary.

Our learning-based methods can effectively attack a model with several millions of parameters trained on a billion images, and allows the attacker to reduce the error rate of their model by 10%. This attack does not match perfect fidelity with the victim model due to what we show are inherent limitations of learning-based approaches: nondeterminism (including only the nondeterminism on the GPU) prohibits training identical models. In contrast, our direct functionally-equivalent extraction returns a neural network agreeing with the victim model on 100% of the test samples and having 100% fidelity on transferred adversarial examples.

We then propose a hybrid method which unifies these two attacks, using learning-based approaches to recover from numerical instability errors when performing the functionally-equivalent extraction attack.

Our work highlights many remaining open problems in model extraction, such as reducing the capabilities required by our attacks and scaling functionally-equivalent extraction.

## Acknowledgements

We would like to thank Ilya Mironov for lengthy and fruitful discussions regarding the functionally equivalent extraction attack. We also thank Úlfar Erlingsson for helpful discussions on positioning the work, and Florian Tramèr for his comments on an early draft of this paper.

## References

- [1] E. Strubell, A. Ganesh, and A. McCallum, “Energy and policy considerations for deep learning in nlp,” *arXiv preprint arXiv:1906.02243*, 2019.
- [2] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, “Xlnet: Generalized autoregressive pretraining for language understanding,” in *Advances in neural information processing systems*, 2019, pp. 5754–5764.
- [3] A. Halevy, P. Norvig, and F. Pereira, “The unreasonable effectiveness of data,” 2009.
- [4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [5] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” in *Neural information processing systems*, 2014, pp. 3104–3112.
- [6] A. Van Den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio.” *SSW*, vol. 125, 2016.
- [7] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical black-box attacks against machine learning,” in *Proceedings of the 2017 ACM on Asia conference on computer and communications security*. ACM, 2017, pp. 506–519.
- [8] D. Lowd and C. Meek, “Adversarial learning,” in *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. ACM, 2005, pp. 641–647.
- [9] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, “Membership inference attacks against machine learning models,” in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 3–18.
- [10] A. Salem, Y. Zhang, M. Humbert, P. Berrang, M. Fritz, and M. Backes, “MI-leaks: Model and data independent membership inference attacks and defenses on machine learning models,” *arXiv preprint arXiv:1806.01246*, 2018.
- [11] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, “Stealing machine learning models via prediction apis,” in *25th {USENIX} Security Symposium ({USENIX} Security 16)*, 2016, pp. 601–618.
- [12] T. Orekondy, B. Schiele, and M. Fritz, “Knockoff nets: Stealing functionality of black-box models,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4954–4963.
- [13] V. Chandrasekaran, K. Chaudhuri, I. Giacomelli, S. Jha, and S. Yan, “Model extraction and active learning,” *CoRR*, vol. abs/1811.02054, 2018. [Online]. Available: <http://arxiv.org/abs/1811.02054>
- [14] S. J. Oh, M. Augustin, B. Schiele, and M. Fritz, “Towards reverse-engineering black-box neural networks,” *arXiv preprint arXiv:1711.01768*, 2017.
- [15] S. Pal, Y. Gupta, A. Shukla, A. Kanade, S. K. Shevade, and V. Ganapathy, “A framework for the extraction of deep neural networks by leveraging public data,”

- CoRR*, vol. abs/1905.09165, 2019. [Online]. Available: <http://arxiv.org/abs/1905.09165>
- [16] J. R. Correia-Silva, R. F. Berriel, C. Badue, A. F. de Souza, and T. Oliveira-Santos, “Copycat cnn: Stealing knowledge by persuading confession with random non-labeled data,” in *2018 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2018.
- [17] C. Song and V. Shmatikov, “Overlearning reveals sensitive attributes,” *arXiv preprint arXiv:1905.11742*, 2019.
- [18] S. Hong, M. Davinroy, Y. Kaya, S. N. Locke, I. Rackow, K. Kulda, D. Dachman-Soled, and T. Dumitras, “Security analysis of deep neural networks operating in the presence of cache side-channel attacks,” *arXiv preprint arXiv:1810.03487*, 2018.
- [19] S. Milli, L. Schmidt, A. D. Dragan, and M. Hardt, “Model reconstruction from model explanations,” *arXiv preprint arXiv:1807.05185*, 2018.
- [20] V. Nair and G. E. Hinton, “Rectified linear units improve restricted boltzmann machines,” in *Proceedings of the 27th international conference on machine learning (ICML-10)*, 2010, pp. 807–814.
- [21] Y. E. Nesterov, “A method for solving the convex programming problem with convergence rate  $o(1/k^2)$ ,” in *Dokl. akad. nauk Sssr*, vol. 269, 1983, pp. 543–547.
- [22] J. Duchi, E. Hazan, and Y. Singer, “Adaptive subgradient methods for online learning and stochastic optimization,” *Journal of Machine Learning Research*, vol. 12, no. Jul, pp. 2121–2159, 2011.
- [23] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [24] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [25] L. Batina, S. Bhasin, D. Jap, and S. Picek, “Csi neural network: Using side-channels to recover your artificial neural network information,” *arXiv preprint arXiv:1810.09076*, 2018.
- [26] P. Kocher, J. Jaffe, and B. Jun, “Differential power analysis,” in *Annual International Cryptology Conference*. Springer, 1999, pp. 388–397.
- [27] A. Das, S. Gollapudi, R. Kumar, and R. Panigrahy, “On the learnability of deep random networks,” *CoRR*, vol. abs/1904.03866, 2019.
- [28] D. Mahajan, R. Girshick, V. Ramanathan, K. He, M. Paluri, Y. Li, A. Bharambe, and L. van der Maaten, “Exploring the limits of weakly supervised pretraining,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 181–196.
- [29] P. Micaelli and A. Storkey, “Zero-shot knowledge transfer via adversarial belief matching,” *arXiv preprint arXiv:1905.09768*, 2019.
- [30] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, “Language models are unsupervised multi-task learners,” *OpenAI Blog*, vol. 1, no. 8, 2019.
- [31] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, “Cnn features off-the-shelf: an astounding baseline for recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2014, pp. 806–813.
- [32] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [33] D. Angluin, “Queries and concept learning,” *Machine learning*, vol. 2, no. 4, pp. 319–342, 1988.
- [34] A. Blum and T. Mitchell, “Combining labeled and unlabeled data with co-training,” in *Proceedings of the eleventh annual conference on Computational learning theory*. Citeseer, 1998, pp. 92–100.
- [35] S. Song, D. Berthelot, and A. Rostamizadeh, “Combining mixmatch and active learning for better accuracy with fewer labels,” 2020. [Online]. Available: <https://openreview.net/forum?id=HJxW10NKPB>
- [36] O. Siméoni, M. Budnik, Y. Avrithis, and G. Gravier, “Rethinking deep active learning: Using unlabeled data at model training,” 2020. [Online]. Available: <https://openreview.net/forum?id=rJehlIrtDS>
- [37] X. Zhai, A. Oliver, A. Kolesnikov, and L. Beyer, “S4l: Self-supervised semi-supervised learning,” *arXiv preprint arXiv:1905.03670*, 2019.
- [38] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. Raffel, “Mixmatch: A holistic approach to semi-supervised learning,” *arXiv preprint arXiv:1905.02249*, 2019.
- [39] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, “Reading digits in natural images with unsupervised feature learning,” 2011.
- [40] A. Krizhevsky *et al.*, “Learning multiple layers of features from tiny images,” Citeseer, Tech. Rep., 2009.
- [41] D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, and

- D. Dennison, “Hidden technical debt in machine learning systems,” in *Advances in neural information processing systems*, 2015, pp. 2503–2511.
- [42] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” in *Advances in Neural Information Processing Systems*, 2017, pp. 6402–6413.
- [43] H. Xiao, K. Rasul, and R. Vollgraf. (2017) Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms.
- [44] N. Carlini, U. Erlingsson, and N. Papernot, “Prototypical examples in deep learning: Metrics, characteristics, and utility,” 2019. [Online]. Available: <https://openreview.net/forum?id=r1xyx3R9tQ>
- [45] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner *et al.*, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [46] Google, “Jax,” <https://github.com/google/jax>, 2019.
- [47] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” *arXiv preprint arXiv:1312.6199*, 2013.
- [48] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” *arXiv preprint arXiv:1706.06083*, 2017.
- [49] T. Lee, B. Edwards, I. Molloy, and D. Su, “Defending against model stealing attacks using deceptive perturbations,” *arXiv preprint arXiv:1806.00054*, 2018.
- [50] I. M. Alabdulmohsin, X. Gao, and X. Zhang, “Adding robustness to support vector machines against adversarial reverse engineering,” in *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. ACM, 2014, pp. 231–240.
- [51] M. Juuti, S. Szyller, A. Dmitrenko, S. Marchal, and N. Asokan, “Prada: protecting against dnn model stealing attacks,” *arXiv preprint arXiv:1805.02628*, 2018.
- [52] M. Kesarwani, B. Mukhoty, V. Arya, and S. Mehta, “Model extraction warning in mlaas paradigm,” in *Proceedings of the 34th Annual Computer Security Applications Conference*. ACM, 2018, pp. 371–380.
- [53] B. Wang and N. Z. Gong, “Stealing hyperparameters in machine learning,” in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 36–52.
- [54] J. Zhang, Z. Gu, J. Jang, H. Wu, M. P. Stoecklin, H. Huang, and I. Molloy, “Protecting intellectual property of deep neural networks with watermarking,” in *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*. ACM, 2018, pp. 159–172.
- [55] Y. Uchida, Y. Nagai, S. Sakazawa, and S. Satoh, “Embedding watermarks into deep neural networks,” in *Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval*. ACM, 2017, pp. 269–277.
- [56] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang, “On the (im) possibility of obfuscating programs,” in *Annual international cryptology conference*. Springer, 2001, pp. 1–18.
- [57] M. Barni, C. Orlandi, and A. Piva, “A privacy-preserving protocol for neural-network-based computation,” in *Proceedings of the 8th workshop on Multimedia and security*. ACM, 2006, pp. 146–151.
- [58] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, “Reluplex: An efficient smt solver for verifying deep neural networks,” in *International Conference on Computer Aided Verification*. Springer, 2017, pp. 97–117.

## A Formal Statements for Section 3.3

Here, we give the formal arguments for the difficulty of model extraction to support informal statements from Section 3.3.

**Theorem 1.** *There exists a class of width  $3k$  and depth  $2$  neural networks on domain  $[0, 1]^d$  (with precision  $p$  numbers) with  $d \geq k$  that require, given logit access to the networks,  $\Theta(p^k)$  queries to extract.*

In order to prove Theorem 1, we introduce a family of functions we call  $k$ -rectangle bounded functions, which we will show satisfies this property.

**Definition A.1.** *A function  $f$  on domain  $[0, 1]^d$  with range  $\mathbb{R}$  is a rectangle bounded function if there exists two vectors  $a, b$  such that  $f(x) \neq 0 \implies a \preceq x \preceq b$ , where  $\preceq$  denotes element-wise comparison. The function  $f$  is a  $k$ -rectangle bounded function if there are  $k$  indices  $i$  such that  $a_i \neq 0$  or  $b_i \neq 1$ .*

Intuitively, a  $k$ -rectangle function only outputs a non-zero value on a multidimensional rectangle that is constrained in only  $k$  coordinates. We begin by showing that we can implement  $k$ -rectangle functions for any  $a, b$  using a ReLU network of width  $k$  and depth 2.

**Lemma 1.** *For any  $a, b$  with  $k$  indices  $i$  such that  $a_i \neq 0$  or  $b_i \neq 1$ , we can construct a  $k$ -rectangle bounded function for  $a, b$  with a ReLU network of width  $3k$  and depth 2.*



*Proof.* We will start by constructing a 3-ReLU gadget with output  $\geq 1$  only when  $a_i \leq x_i \leq b_i$ . We will then show how to compose  $k$  of these gadgets, one for each index of the  $k$ -rectangle, to construct the  $k$ -rectangle bounded function.

The 3-ReLU gadget only depends on  $x_i$ , so weights for all other ReLUs will be set to 0. Observe that the function  $T_i(x; a, b) = \text{ReLU}(x - a) + \text{ReLU}(x_i - b_i) - 2\text{ReLU}(x_i - (a_i + b_i)/2)$  is nonzero only on the interval  $(a_i, b_i)$ . This is easier to see when it is written as

$$\begin{aligned} & \text{ReLU}(x_i - a_i) - \text{ReLU}(x_i - (a_i + b_i)/2) \\ & - (\text{ReLU}(x_i - (a_i + b_i)/2) - \text{ReLU}(x_i - b_i)). \end{aligned}$$

The function  $\text{ReLU}(x - x_1) - \text{ReLU}(x - x_2)$  with  $x_1 < x_2$  looks like a sigmoid, and has the following form:

$$\text{ReLU}(x - x_1) - \text{ReLU}(x - x_2) = \begin{cases} 0 & x \leq x_1 \\ x - x_1 & x_1 \leq x \leq x_2 \\ x_2 - x_1 & x \geq x_2 \end{cases}$$

Now,  $T_i(x; a_i, b_i) \cdot 1/(b_i - a_i)$  has range  $[0, 1]$  for any value of  $a_i, b_i$ . Then the function

$$f_{a,b}(x) = \text{ReLU}\left(\sum_i (T_i(x; a_i, b_i)/(b_i - a_i)) - (k - 1)\right)$$

is  $k$ -rectangle bounded for vectors  $a, b$ . To see why, we need that no input  $x$  not satisfying  $a \preceq x \preceq b$  has  $\sum_i (T_i(x; a_i, b_i)/(b_i - a_i)) > k - 1$ . This is simply because each term  $T_i(x; a_i, b_i) \leq 1$ , so unless all  $k$  such terms are  $> 0$ , the inequality cannot hold.  $\square$

Now that we know how to construct a  $k$ -rectangle bounded function, we will introduce a set of  $p^k$  disjoint  $k$ -rectangle bounded functions, and then show that any one requires  $p^k$  queries to extract when the others are also possible functions.

**Lemma 2.** *There exists a family of  $k$ -rectangle bounded functions  $\mathcal{F}$  such that extracting an element of  $\mathcal{F}$  requires  $p^k$  queries in the worst case.*

Here,  $p$  is the feature precision; images with 8-bit pixels have  $p = 256$ .

*Proof.* We begin by constructing  $\mathcal{F}$ . The following  $p$  ranges are clearly pairwise disjoint:  $\{(\frac{i-1}{p}, \frac{i}{p})\}_{i=1}^p$ . Then pick any  $k$  indices, and we can construct  $p^k$  distinct  $k$ -rectangle bounded functions - one for each element in the Cartesian product of each index's set of ranges. Call this set  $\mathcal{F}$ .

The set of inputs with non-zero output is distinct for each function, because their rectangles are distinct. Now consider the information gained from any query. If the query returns a non-zero value, the function is learned. If not, at most one function from  $\mathcal{F}$  is ruled out - the function whose rectangle was queried. Then any sequence of  $n$  queries to an oracle can rule out at most  $n$  of the functions of  $\mathcal{F}$ , so that at least  $|\mathcal{F}| = p^k$  queries are required in the worst case.  $\square$




Figure 5: Fidelity is easier on more prototypical examples.

Putting Lemma 1 and 2 together gives us Theorem 1.

**Theorem 2.** *Checking whether two networks with domains  $\{0, 1\}^d$  are functionally equivalent is NP-hard.*

*Proof.* We prove this by reduction to subset sum. A similar reduction (reducing to 3-SAT instead of Subset Sum) for a different statement appears in [58].

Suppose we receive a subset sum instance  $T, p, [v_1, v_2, \dots, v_d]$  - the set is  $v$ , the target sum is  $T$ , and the problem's precision is  $p$ . We will construct networks  $f_1$  and  $f_2$  such that checking if  $f_1$  and  $f_2$  are functionally equivalent is equivalent to solving the subset sum instance. We start by setting  $f_1 = 0$  - it never returns a non-zero value. We now construct a network  $f_2$  that has nonzero output only if the subset sum instance can be solved (and finding an input with nonzero output reveals the satisfying subset).

The network  $f_2$  has three hidden units in the first layer with incoming weight for the  $i$ th feature equal to  $v_i$ . This means the dot product of the input  $x$  with weights will be the sum of the subset  $\{i|x_i = 1\}$ . We want to force this to accept iff there is an input where this sum is  $T$ . To do so, we use the same 3-ReLU gadget as in the proof of Theorem 1:

$$\begin{aligned} f_2(x; T, p, v) &= \text{ReLU}(x \cdot v - (T - p/2)) \\ &+ \text{ReLU}(x \cdot v - (T + p/2)) - 2\text{ReLU}(x \cdot v - T). \end{aligned}$$

As before, this will only be nonzero in the range  $[T - p/2, T + p/2]$ , and we are done.  $\square$

## B Prototypicality and Fidelity

We know from Section 5 that learning strategies struggle to achieve perfect fidelity due to non-determinism inherent in learning. What remains to be understood is whether some

samples are more difficult than others to achieve fidelity on. We investigate using recent work on identifying prototypical data points. Using each metric developed in Carlini *et al.* [44], we can rank the Fashion-MNIST test set in order of increasing prototypicality. Binning the prototypicality ranking into percentiles, we can measure how many of the 90 models we trained for Section 5 agree with the oracle’s prediction. The intuition here is that more prototypical examples should be more consistently learnable, whereas more outlying points may be harder to consistently classify. Indeed, we find that this is the case - all metrics find a correlation between prototypicality and model agreement (fidelity), as seen in Figure 5. Interestingly, the metrics which do not use ensembles of models (adversarial distance and holdout-retraining) have the best correlation with the model agreement metric—roughly the top 50% of prototypical examples by these metrics are classified the same by nearly all 90 models.

### C Supplement for Section 6

Accuracies for the oracles in Section 6 are found in Table 9.

MNIST		CIFAR-10	
Parameters	Accuracy	Parameters	Accuracy
12,500	94.3%	49,000	29.2%
25,000	95.6%	98,000	34.2%
50,000	97.2%	196,000	40.3%
100,000	97.7%	393,000	42.6%
200,000	98.0%	786,000	43.1%
400,000	98.3%	1,572,000	45.9%

Table 9: Statistics for the oracle models we train to extract.

Figure 6 shows a distribution over the bits of precision in the difference between the logits (i.e., pre-softmax prediction) of the 16 neuron oracle neural network and the extracted network. Formally, we measure the magnitude of the gap  $|f_{\theta}(x) - f_{\hat{\theta}}(x)|$ . Notice that this is a different (and typically stronger) measure of fidelity than used elsewhere in the paper.

### D Query Complexity of Functionally Equivalent Extraction

In this section, we briefly analyze the query complexity of the attack from Section 6. We assume that a simulated partial derivative requires  $O(1)$  queries using finite differences.

1. **Critical Point Search.** This step is the most nontrivial to analyze, but fortunately this was addressed in [19]. They found this step requires  $O(h \log(h))$  gradient queries, which we simulate with  $O(h \log(h))$  model queries.
2. **Weight Recovery.** This piece is significantly complicated by not having access to gradient queries. For each




Figure 6: For a 16-neuron MNIST model the attack works. Plotted here is number of bits of precision on the logits normalized by the value of the lot as done in the prior figure.

ReLU, absolute value recovery requires  $O(d)$  queries and weight sign recovery requires an additional  $O(d)$ , making this step take  $O(dh)$  queries total.

3. **Global Sign Recovery.** For each ReLU, we require only three queries. Then this step is  $O(h)$ .
4. **Last Layer Extraction.** This step requires  $h$  queries to make the system of linear equations full rank (although in practice we reuse previous queries here, making this step require 0 queries).

Overall, the algorithm requires  $O(h \log(h) + dh + h) = O(dh)$  queries. Extraction requires  $\Omega(dh)$  queries without auxiliary information, as there are  $dh$  parameters in the model. Then the algorithm is query-optimal up to a constant factor, removing logarithmic factors from Milli *et al.* [19].