

# Updates-Leak: Data Set Inference and Reconstruction Attacks in Online Learning

Ahmed Salem\*, Apratim Bhattacharya<sup>†</sup>, Michael Backes\*, Mario Fritz\*, Yang Zhang\*

\*CISPA Helmholtz Center for Information Security

<sup>†</sup>Max Planck Institute for Informatics

**Abstract**—Machine learning (ML) has progressed rapidly during the past decade and the major factor that drives such development is the unprecedented large-scale data. As data generation is a continuous process, this leads to ML service providers updating their models frequently with newly-collected data in an online learning scenario. In consequence, if an ML model is queried with the same set of data samples at two different points in time, it will provide different results.

In this paper, we investigate whether the change in the output of a black-box ML model before and after being updated can leak information of the dataset used to perform the update. This constitutes a new attack surface against black-box ML models and such information leakage severely damages the intellectual property and data privacy of the ML model owner/provider. In contrast to membership inference attacks, we use an encoder-decoder formulation that allows inferring diverse information ranging from detailed characteristics to full reconstruction of the dataset. Our new attacks are facilitated by state-of-the-art deep learning techniques. In particular, we propose a hybrid generative model (BM-GAN) that is based on generative adversarial networks (GANs) but includes a reconstructive loss that allows generating accurate samples. Our experiments show effective prediction of dataset characteristics and even full reconstruction in challenging conditions.

**Index Terms**—Machine Learning Security and Privacy, Dataset Reconstruction Attack

## I. INTRODUCTION

Machine learning (ML) has progressed rapidly during the past decade. Nowadays, it has become the core component in many industrial domains ranging from automotive manufacturing to financial services. Leading Internet companies, such as Google,<sup>1</sup> Amazon,<sup>2</sup> and Microsoft,<sup>3</sup> further provide Machine Learning as a Service (MLaaS) to simplify ML deployment. In this setting, an MLaaS provider trains a machine learning model at their backend and provides the trained model to public as a black-box API.

The major factor that drives the current ML development is the unprecedented large-scale data. In consequence, collecting high-quality data becomes essential for building advanced ML models. Data collection is a continuous process as enormous data is being generated at every second. This turns ML model training into a continuous process as well: Instead of training an ML model for once and keeping on using it afterwards, the model provider, such as an MLaaS provider, needs to keep on

updating the model with newly-collected data. In practice, this is also known as *online learning*. And we refer to the dataset used to perform model update as the *updating set*.

Regularly updating an ML model results in the model having different versions with respect to different model parameters. This indicates that if an ML model is queried with the same set of data samples at two different points in time, it will provide different outputs.

### A. Our Contributions

In this paper, our main research question is: *Can different outputs of an ML model's two versions queried with the same set of data samples leak information of the corresponding updating set?* This constitutes a new attack surface against machine learning models. Information leakage of the updating set can severely damage the intellectual property and data privacy of the model provider/owner.

We concentrate on the most common ML application – classification. More importantly, we target on black-box ML models – the most difficult attack setting where an adversary does not have access to her target model's parameters but can only query the model with her data samples and obtain the corresponding prediction results, i.e., *posteriors* in the case of classification.

In total, we propose four different attacks in this surface which can be categorized into two classes, namely, *single-sample attack class* and *multi-sample attack class*. The two attacks in the single-sample attack class concentrate on a simplified case when the target ML model is updated with one single data sample. We investigate this case to show whether an ML model's two versions' different outputs indeed constitute a valid attack surface. The two attacks in the multi-sample attack class tackle a more general and complex case when the updating set contains multiple data samples.

Among our four attacks, two (one for each attack class) aim at reconstructing the updating set which to our knowledge, are the first attempt in this direction. Compared to many previous attacks inferring certain properties of a target model's training set [6], [8], [13], dataset reconstruction attack leads to more severe consequences [20]. In theory, membership inference attacks [19], [29], [31] can also be leveraged to reconstruct the dataset from a black-box ML model. However, membership inference is not scalable in the real-world setting as the adversary needs to collect a large data sample which happens to include all the training set samples of the target

<sup>1</sup><https://cloud.google.com/ml-engine/>

<sup>2</sup><https://aws.amazon.com/machine-learning/>

<sup>3</sup><https://azure.microsoft.com/en-us/services/machine-learning-studio/>

model. Though our two reconstruction attacks are designed specifically for the online learning setting, we believe they can provide further insights on reconstructing a black-box ML model’s training set in other settings.

Extensive experiments show that indeed, the output difference of the same ML model’s two different versions can be exploited to infer information about the updating set. We detail our contributions as the following.

**General Attack Construction.** Our four attacks follow a general structure, which can be formulated into an encoder-decoder style. The encoder realized by a multilayer perceptron (MLP) takes the difference of the target ML model’s outputs, namely *posterior difference*, as its input while the decoder produces different types of information about the updating set with respect to different attacks.

To obtain the posterior difference, we randomly select a fixed set of data samples, referred to as the *probing set*, and probe the target model’s two different versions (the second-version model is obtained by updating the first-version model with an updating set). Then, we calculate the difference between the two sets of posteriors as the input for our attack’s encoder.

**Single-sample Attack Class.** The single-sample attack class contains two attacks: *Single-sample label inference attack* and *single-sample reconstruction attack*. The first attack predicts the label of the single sample used to update the target model. We realize the corresponding decoder for the attack by a two-layer MLP. Our evaluation shows that our attack is able to achieve a strong performance, e.g., 0.96 accuracy on the CIFAR-10 dataset.<sup>4</sup>

The single-sample reconstruction attack aims at reconstructing the updating sample. We rely on autoencoder (AE). In detail, we first train an AE on a different set of data samples. Then, we transfer the AE’s decoder into our attack model as its sample reconstructor. Experimental results show that we can construct the single sample with a mean squared error (MSE) of 0.06355 for the MNIST dataset<sup>5</sup> and 0.01352 for the CIFAR-10 dataset, respectively. Moreover, we show that our attack learns to generate the specific sample used in the updating set [20], [31] instead of a general representation of samples affiliated with the same label.

**Multi-sample Attack Class.** The multi-sample attack class includes *multi-sample label distribution estimation attack* and *multiple-sample reconstruction attack*. Multi-sample label distribution estimation attack estimates the label distribution of the updating set’s data samples. It is a generalization of the label inference attack in the single-sample attack class. We realize this attack by setting up the attack model’s decoder as a multilayer perceptron with a fully connected layer and a softmax layer. Kullback-Leibler divergence (KL-divergence) is adopted as the model’s loss function. Extensive experiments demonstrate the effectiveness of this attack. For the CIFAR-10

dataset, when the updating set’s cardinality is 100, our attack model achieves a 0.00376 KL-divergence which outperforms the baseline model by a factor of 3. Moreover, the accuracy of predicting the most frequent label is 0.32 which is also 3 times higher than the baseline model.

Our last attack, namely multiple-sample reconstruction attack, aims at generating all samples in the updating set. This is a much more complex attack than the previous ones. The decoder for this attack is assembled with two components. The first one learns the data distribution of the updating set samples. To this end, we propose a novel hybrid generative model, namely BM-GAN. Different from the standard generative adversarial networks (GANs), our BM-GAN introduces a “Best Match” loss which ensures that each sample in the updating set is reconstructed. The second component of our decoder relies on machine learning clustering to group the generated data samples by BM-GAN into clusters and take the central sample of each cluster as one final reconstructed sample. Our evaluation shows that we are able to reconstruct very similar samples as those in the original updating set on both MNIST and CIFAR-10 datasets.

In summary, we make the following contributions in this paper:

- We discover a new attack surface against black-box ML models, i.e., different outputs of the same ML model’s two versions queried with a same set of data samples.
- We propose four different attacks in this surface based on advanced machine learning techniques. Extensive experiments demonstrate that the updating set’s information can be effectively inferred.
- Two of our attacks aim at reconstructing the updating set itself. Though designed for our specific attack surface, we believe they can provide further insights on dataset reconstruction attacks in other settings.

## B. Organization

The rest of the paper is organized as follows. We provide necessary preliminaries in Section II. In Section III, we introduce our general attack pipeline. Section IV and Section V present our attacks and their evaluation results in single-sample and multi-sample attack classes, respectively. In Section VI, we discuss possible defense mechanisms. Section VII presents the related work in the field and Section VIII concludes the paper.

## II. PRELIMINARIES

In this section, we start by introducing online learning, then present our threat model, and finally introduce the datasets used in our experiments.

### A. Online Learning

In this paper, we focus on the most common machine learning task – classification. An ML classifier  $\mathcal{M}$  is essentially a function that maps a data sample  $x \in \mathcal{X}$  to posterior probabilities  $y \in \mathcal{Y}$ .

$$\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Y}$$

<sup>4</sup><https://www.cs.toronto.edu/~kriz/cifar.html>

<sup>5</sup><http://yann.lecun.com/exdb/mnist/>

TABLE I: Notations

Notation	Description
$\mathcal{M}$	A machine learning model
$\mathcal{M}'$	The updated version of $\mathcal{M}$
$\mathcal{D}_{update}$	Updating set
$\mathcal{D}_{probe}$	Probing set
$\mathcal{D}_{shadow}$	Shadow set
$x$	A data sample
$x_{update}$	A data sample used to update the model
$\delta$	Posterior difference
$\mu$	Latent vector learned by attack model's encoder
$\ell$	Probability vector of $x_{update}$ affiliated with a certain label
$q$	Probability vector of $\mathcal{D}_{update}$ 's samples' label distribution

Here,  $y \in \mathcal{Y}$  is a vector with each entry indicating the probability of  $x$  being classified to a certain class or affiliated with a certain label. The sum of all values in  $y$  is 1. To train an ML model, we need a set of data samples, i.e., training set. The training process is performed by a certain optimization algorithm, such as ADAM, following a predefined loss function.

A trained ML model  $\mathcal{M}$  can be updated with an updating set denoted by  $\mathcal{D}_{update}$ . The model update is performed by further training the model with the updating set using the same optimization algorithm on the basis of the current model's parameters. More formally, given an updating set  $\mathcal{D}_{update}$  and a trained ML model  $\mathcal{M}$ , the updating process  $\mathcal{F}_{update}$  can be defined as the following:

$$\mathcal{F}_{update} : \mathcal{D}_{update}, \mathcal{M} \rightarrow \mathcal{M}'$$

where  $\mathcal{M}'$  is the updated version of  $\mathcal{M}$ . For presentation purposes, we summarize the notations used throughout the paper in [Table I](#).

### B. Threat Model

For all of our four attacks, we consider an adversary with black-box access to the target ML model. This means that the adversary can only query the model with a set of data samples, i.e., her probing set, and obtain the corresponding posteriors. This follows the typical setting of MLaaS, which is also the most difficult attack setting for the adversary [31]. We also assume that the adversary has a local dataset which comes from the same distribution as the target model's training set following previous works [8], [29], [31]. Moreover, we consider the adversary to be able to establish the same ML model as the target ML model with respect to model architecture. This can be achieved by using the same MLaaS which constructs the target model [31] or performing model hyperparameter stealing attacks [24], [38]. The adversary needs these two information to establish a shadow model which mimics the behavior of the target model to derive data for training her attack model (see [Section III-D](#)). Also, part of the adversary's local dataset will be used as her probing set.

### C. Datasets Description

For our experimental evaluation, we use two well-known image datasets including MNIST and CIFAR-10. Both of them

are benchmark datasets for various computer vision as well as machine learning security and privacy tasks. MNIST is a 10-class image dataset, it consists of 70,000  $28 \times 28$  grey-scale images. Each image contains in its center a handwritten digit. Images in MNIST are equally distributed over 10 classes, i.e., 7,000 images per class. CIFAR-10 contains 60,000  $32 \times 32$  color images. It also contains 10 classes including airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck. Similar to MNIST, CIFAR-10 is also a balanced dataset.

## III. GENERAL ATTACK PIPELINE

Our general attack pipeline can be divided into three phases. In the first phase, the adversary generates her attack input, i.e., posterior difference. In the second phase, our encoder transforms the posterior difference into a latent vector. In the last phase, the decoder decodes the latent vector to produce different information of the updating set with respect to different attacks. [Figure 1](#) provides a schematic view of our attack pipeline.

In this section, we provide a general introduction for each phase of our attack pipeline. In the end, we present our strategy of deriving data to train our attack models.

### A. Attack Input

Recall that we aim at investigating the information leaked from posterior difference of a model's two versions when queried with the same set of data samples. To create this posterior difference, the adversary first needs to pick a set of data samples as her probing set, denoted by  $\mathcal{D}_{probe}$ . In this work, the adversary picks a random sample of data points (from her local dataset) to form  $\mathcal{D}_{probe}$ . Choosing or crafting [24] a specific set of data samples as the probing set may further improve attack efficiency, we leave this as a future work. Next, the adversary queries the target ML model  $\mathcal{M}$  with all samples in  $\mathcal{D}_{probe}$  and concatenates the received outputs to form a vector  $y_{probe}$ . Then, she probes the updated model  $\mathcal{M}'$  with samples in  $\mathcal{D}_{probe}$  and creates a vector  $y'_{probe}$  accordingly. In the end, she sets the posterior difference, denoted by  $\delta$ , to the difference of both outputs:

$$\delta = y_{probe} - y'_{probe}$$

Note that the dimension of  $\delta$  is the product of  $\mathcal{D}_{probe}$ 's cardinality and the number of classes of the target dataset. For this paper, both CIFAR-10 and MNIST are 10-class datasets, and our probing set always contains 100 data samples, this indicates the dimension of  $\delta$  is 1,000.

### B. Encoder Design

All our attacks share the same encoder structure, we model it with a multilayer perceptron. The number of layers inside the encoder depends on the dimension of  $\delta$ : Longer  $\delta$  requires more layers in the encoder. As our  $\delta$  is always a 1,000-dimension vector, we use two fully connected layers in the encoder. The first layer transforms  $\delta$  to a 128-dimension vector and the second layer further reduces the dimension to 64. The concrete architecture of our encoder is as follows:




Fig. 1: A schematic view of the general attack pipeline.

*Encoder architecture:*

$$\delta \rightarrow \text{FullyConnected}(128) \\ \text{FullyConnected}(64) \rightarrow \mu$$

where  $\text{FullyConnected}(128)$  denotes a fully connected layer with 128 hidden units and  $\mu$  denotes the latent vector which serves as the input for our decoder. Furthermore, we use LeakyReLU as our encoder’s activation function and apply dropout on both layers for regularization.

*C. Decoder Structure*

Our four attacks aim at inferring different information of the updating set  $\mathcal{D}_{update}$ , ranging from sample labels to the updating set itself. Thus, we construct different decoders for different attacks with different techniques, such as multilayer perceptron, autoencoder, and generative adversarial networks. The details of these decoders will be presented in the following sections.

*D. Shadow Model*

Our encoder as well as decoder need to be trained jointly in a supervised manner. This indicates that we need ground truth data for model training. Due to our minimal assumptions, the adversary cannot get the ground truth from the target model. To solve this problem, we rely on shadow model following previous works [8], [29], [31]. A shadow model is designed to mimic the target model. By controlling the training process of the shadow model, the adversary can derive the ground truth data needed to train her attack models.

As presented in Section II-B, our adversary knows (1) the architecture of the target model (by using the same MLaaS) and (2) a local dataset coming from the same distribution as the target dataset. To build a shadow model  $\mathcal{M}_{shadow}$ , the adversary first establishes an ML model with the same structure as the target model. Then, she gets a shadow dataset  $\mathcal{D}_{shadow}$  from her local dataset (the rest is used as  $\mathcal{D}_{probe}$ )

and splits it into two parts: Shadow training set  $\mathcal{D}_{shadow}^{train}$  and shadow updating set  $\mathcal{D}_{shadow}^{update}$ .  $\mathcal{D}_{shadow}^{train}$  is used to train the shadow model while  $\mathcal{D}_{shadow}^{update}$  is further split to  $m$  datasets:  $\mathcal{D}_{shadow}^{update^1} \dots \mathcal{D}_{shadow}^{update^m}$ . The number of samples in each of the  $m$  datasets depends on the attack. For instance, our two attacks in the single-sample attack class require each dataset containing a single sample. The adversary then generates  $m$  shadow updated models  $\mathcal{M}'^1_{shadow} \dots \mathcal{M}'^m_{shadow}$  by updating the shadow model  $\mathcal{M}_{shadow}$  with  $m$  shadow updating sets in parallel.

The adversary, in the end, probes the shadow and updated shadow models with her probing set  $\mathcal{D}_{probe}$ , and calculates the shadow posterior difference  $\delta^1_{shadow} \dots \delta^m_{shadow}$ . Together with the corresponding shadow updating set’s ground truth information (depending on the attack), the training data for her attack model is derived.

More generally, the training set for each of our attack models contains  $m$  samples corresponding to  $\mathcal{D}_{shadow}^{update^1} \dots \mathcal{D}_{shadow}^{update^m}$ . In all our experiments, we set  $m$  to 10,000. In addition, we create 1,000 updated models for the target model, this means the testing set for each attack model contains 1,000 samples.

IV. SINGLE-SAMPLE ATTACKS

In this section, we concentrate on the case when an ML model is updated with a single sample. This is a simplified attack scenario and we aim at examining the possibility of using posterior difference to infer information about the updating set. In Section V, we tackle a more general scenario when the updating set contains multiple samples.

We start by introducing the single-sample label inference attack, then, present the single-sample reconstruction attack.

*A. Single-sample Label Inference Attack*

**Attack Definition.** Our single-sample label inference attack takes the posterior difference as the input and outputs the label of the single updating sample. More formally, given a posterior

difference  $\delta$ , our single-sample label inference attack is defined as follows:

$$\mathcal{A}_{LI} : \delta \mapsto \ell$$

where  $\ell$  is a vector with each entry representing the probability of the updating sample affiliated with a certain label.

**Methodology.** To recap, the general construction of the attack model consists of an MLP-based encoder which takes the posterior difference as its input and outputs a latent vector  $\mu$ . For this attack, the adversary constructs her decoder also with an MLP which is assembled with a fully connected layer and a softmax layer to transform the latent vector to the corresponding updating sample’s label. Formally,  $\mathcal{A}_{LI}$ ’s decoder is realized with the following structure:

*$\mathcal{A}_{LI}$ ’s decoder architecture:*

$$\mu \rightarrow \text{FullyConnected}(n) \\ \text{softmax} \rightarrow \ell$$

where  $n$  is equal to the size of  $\ell$ , i.e.,  $n = |\ell|$ .

To obtain the data for training the attack model  $\mathcal{A}_{LI}$ , the adversary generates ground truth data by creating a shadow model as introduced in Section III-D while setting the shadow updating set’s cardinality to 1. Then, the adversary trains her attack model  $\mathcal{A}_{LI}$  with a cross-entropy loss. Our loss function in detail is,

$$\mathcal{L}_{CE} = \sum_i \ell_i \log(\hat{\ell}_i)$$

where  $\ell_i$  is the true probability of label  $i$  and  $\hat{\ell}_i$  is our predicted probability of label  $i$ . The optimization is performed by the ADAM optimizer.

To perform the label inference attack, the adversary constructs the posterior difference as introduced in Section III-A, then feeds it to the attack model  $\mathcal{A}_{LI}$  to obtain the label.

**Experimental Setup.** We evaluate the performance of our single-sample label inference attack using both MNIST and CIFAR-10 datasets. Firstly, we split each dataset into three disjoint datasets: The target dataset  $\mathcal{D}_{target}$ , the shadow dataset  $\mathcal{D}_{shadow}$ , and the probing dataset  $\mathcal{D}_{probe}$ . As mentioned before,  $\mathcal{D}_{probe}$  contains 100 data samples. We then split  $\mathcal{D}_{shadow}$  to  $\mathcal{D}_{shadow}^{train}$  and  $\mathcal{D}_{shadow}^{update}$  to train the shadow model as well as updating it (see Section III-D). The same process is applied to train and update the target model with  $\mathcal{D}_{target}$ . As mentioned in Section III-D, we build 10,000 and 1,000 updated models for target and shadow models, respectively. This means the training and testing sets for our attack model contain 10,000 and 1,000 samples, respectively.

We use convolutional neural network (CNN) to build shadow and target models for both CIFAR-10 and MNIST datasets. The CIFAR-10 model consists of two convolutional layers, one max pooling layer, three fully connected layers, and a softmax layer. The MNIST model consists of two convolutional layer, two fully connected layers, and a softmax layer.

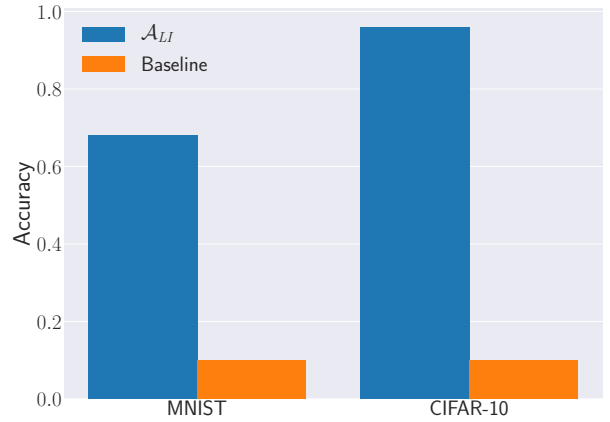


Fig. 2: [Higher is better] Performance of the single-sample label inference attack ( $\mathcal{A}_{LI}$ ) on both MNIST and CIFAR-10 datasets together with the baseline model. Accuracy is adopted as the evaluation metric.

All shadow and target models’ training sets contain 10,000 images. We train the CIFAR-10 and MNIST models for 50 and 25 epochs, respectively, with a batch size of 64. To create an updated ML model, we perform a single-epoch training.

Finally, we use balanced datasets in terms of label distribution over all classes to update the target model as well as the shadow model. Therefore, we adopt accuracy to measure the performance of the attack.

All of our experiments are implemented using Pytorch.<sup>6</sup> For reproducibility purposes, our code will be made available.

**Results.** The experimental results for our attack are depicted in Figure 2. As we can see,  $\mathcal{A}_{LI}$  achieves a strong performance with an accuracy of 0.96 on the CIFAR-10 dataset, and 0.68 on the MNIST dataset. Moreover, our attack significantly outperforms the baseline model which simply guesses a label over all possible labels. As both datasets contain 10 classes, the baseline model’s result is approximately 10%. Our evaluation shows that the different outputs of an ML model’s two versions indeed leak information of the corresponding updating set.

### B. Single-sample Reconstruction Attack

**Attack Definition.** Our next attack, i.e., single-sample reconstruction attack, takes one step further to construct the data sample used to update the model. Formally, given a posterior difference  $\delta$ , the attack, denoted by  $\mathcal{A}_{SSR}$ , is defined as follows:

$$\mathcal{A}_{SSR} : \delta \mapsto x_{update}$$

where  $x_{update}$  denotes the sample used to update the model ( $\mathcal{D}_{update} = \{x_{update}\}$ ).

**Methodology.** Reconstructing a data sample is a much more complex task than predicting the sample’s label. To tackle this problem, we need an ML model which is able to generate a

<sup>6</sup><https://pytorch.org/>


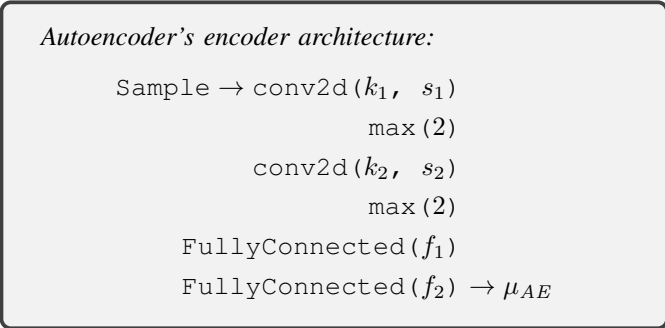


Fig. 3: Methodology of the single-sample reconstruction attack ( $\mathcal{A}_{SSR}$ ).

data sample in the complex space. To this end, we rely on autoencoder (AE).


Autoencoder is assembled with an encoder and a decoder. Different from our attacks, AE’s goal is to learn an efficient encoding for a data sample: Its encoder encodes a sample into a latent vector and its decoder tries to decode the latent vector to reconstruct the same sample. This indicates AE’s decoder itself is a data sample reconstructor. For our attack, we first train an AE, then transfer the AE’s decoder to our attack model as its pretrained decoder. Figure 3 provides an overview of the attack methodology.

Our AE’s encoder has the following architecture:



where  $\max(2)$  denotes a max-pooling layer with a  $2 \times 2$  kernel,  $\mu_{AE}$  is the latent vector output of the encoder. Moreover,  $k_i, s_i$ , and  $f_i$  represent the kernel size, number of filters, and number of units in the  $i$ th layer. Their concrete values depend on the target dataset. We adopt ReLU as the activation function for all layers and apply dropout after the first fully connected layer for regularization.

For AE’s decoder, we use the following model architecture:



Here,  $\text{ConvTranspose2d}(k', s')$  denotes a 2-dimension transposed convolution layer with kernel size  $k' \times k'$  and  $s'$  filters, and  $f'_i$  specifies the number of units in the  $i$ th fully connected layer. We again use ReLU as the activation function for all layers except for the last one where we adopt  $\tanh$ . We also apply dropout after the last fully connected layer for regularization.

After the autoencoder is trained, the adversary takes its decoder and appends it to her attack model’s encoder. To establish the link, the adversary adds an additional fully connected layer to its encoder which transforms the dimensions of the latent vector  $\mu$  to the same dimension as  $\mu_{AE}$ .

We divide the attack model training process into two phases. In the first phase, the adversary uses her shadow dataset to train an AE with the previously mentioned model architecture. In the second phase, she follows the same procedure for single-sample label inference attack to train her attack model. Note that the decoder from AE here serves as a pretrained decoder, this means it will be further trained together with the attack model’s encoder. To train both autoencoder and our attack model, we use mean squared error (MSE) as the loss function. Our objective in detail is,

$$\mathcal{L}_{MSE} = \|\hat{x}_{update} - x_{update}\|_2^2$$

where  $\hat{x}_{update}$  is our predicted data sample. We again adopt ADAM as the optimizer.

**Experimental Setup.** We use the same experimental setup as the previous attack (see Section IV-A) except for the evaluation metric. In detail, we adopt MSE to measure our attack’s performance instead of accuracy. We list the concrete values for the AE’s encoder and decoder hyperparameters in Table II in the Appendix.

We construct two baseline models, namely label-random and random. Both of these baseline models take a random data sample from the adversary’s shadow dataset. The difference is the label-random baseline picks a sample within the same class as the target updating sample, while the random baseline takes a random data sample from the whole shadow dataset of the adversary. The label-random baseline can be implemented by first performing our single-sample label inference attack to learn the label of the data sample and then picking a random sample affiliated with the same label.

**Results.** First of all, our single-sample reconstruction attack achieves very strong performance. As shown in Figure 4, our attack on the MNIST dataset outperforms the random




Fig. 4: [Lower is better] Performance of the single-sample reconstruction attack ( $\mathcal{A}_{SSR}$ ) together with autoencoder and two baseline models. Mean squared error is adopted as the evaluation metric. Autoencoder serves as an oracle as the adversary cannot use it for her attack.




Fig. 5: Visualization of some generated samples from the single-sample reconstruction attack ( $\mathcal{A}_{SSR}$ ) on the MNIST dataset. Samples are fair random draws, not cherry-picked. The first row shows the original samples. The second row shows the reconstructed samples by  $\mathcal{A}_{SSR}$ . The third shows the reconstructed samples by autoencoder, i.e., the upper bound of our reconstruction attack.

baseline by 36% and more importantly, outperforms the label-random baseline by 22%. Similarly, for the CIFAR-10 dataset, our attack achieves an MSE of 0.014 which is significantly better than the two baseline models, i.e., it outperforms the random and label-random baselines by a factor of 2.1 and 2.2, respectively. The difference between our attack’s performance gain over the baseline models on the MNIST and CIFAR-10 datasets is expected as images in CIFAR-10 are more complex than those in MNIST. In other words, the chance of picking a random image similar to the updating image is much higher in the MNIST dataset than in the CIFAR-10 dataset

Secondly, we compare our attack’s performance against the results of the autoencoder for sample reconstruction. Note that AE which takes the original data sample as input and outputs the reconstructed one is an oracle as the adversary does not have access to the original updating sample. Here, we just use AE’s result to show the best possible result for our attack. From Figure 4, we observe that AE achieves 0.042 and 0.0043 MSE for the MNIST and CIFAR-10 datasets, respectively, which indeed outperforms our attack. However, our attack still has a comparable performance.

Finally, Figure 5 visualizes some randomly sampled recon-

structed images by our attack on the MNIST dataset. The first row depicts the original images used to update the models and the second row shows the result of our attack. As we can see, our attack is able to reconstruct images that are visually similar to the original sample with respect to rotation and shape. We also show the result of AE in the third row in Figure 5 which as mentioned before, is the upper bound for our attack. The results from Figure 4 and Figure 5 demonstrate that our attack indeed learns to construct the specific updating data sample instead of a general representation of samples affiliated with the same label as the target updating sample. This follows the security and privacy definition presented in previous works [20], [31].

## V. MULTI-SAMPLE ATTACKS

After demonstrating the effectiveness of our attacks against the updating set with a single sample, we now focus on a more general attack scenario where the updating set contains multiple data samples. We introduce two attacks in the multi-sample attack class: Multi-sample label distribution estimation attack and multi-sample reconstruction attack.

### A. Multi-sample Label Distribution Estimation Attack

**Attack Definition.** Our first attack in the multi-label attack class aims at estimating the label distribution of the updating set’s samples. It can be considered as a generalization of the label inference attack in the single-sample attack class. Formally, the attack is defined as:

$$\mathcal{A}_{LDE} : \delta \mapsto q$$

where  $q$  as a vector denotes the distribution of labels over all classes for samples in the updating set.

**Methodology.** The adversary uses the same encoder structure as presented in Section III-B and the same decoder structure of the label inference attack (Section IV-A). However, since the label distribution estimation attack estimates a probability vector  $q$  instead of performing classification. We use Kullback-Leibler divergence (KL-divergence) as our objective function defined as follows,

$$\mathcal{L}_{KL} = \sum_i (\hat{q}_\ell)_i \log \frac{(\hat{q}_\ell)_i}{(q_\ell)_i}$$

where  $\hat{q}_\ell$  and  $q_\ell$  represent our attack’s estimated label distribution and the target label distribution, respectively, and  $(q_\ell)_i$  corresponds to the  $i$ th label.

To train the attack model  $\mathcal{A}_{LDE}$ , the adversary starts by generating her training data as mentioned in Section III-D. She then trains  $\mathcal{A}_{LDE}$  with the posterior difference  $\delta_{shadow}^1 \cdots \delta_{shadow}^m$  as the input and the normalized label distribution of their corresponding updating sets, i.e.,  $\mathcal{D}_{shadow}^{update^1} \cdots \mathcal{D}_{shadow}^{update^m}$ , as the output. We assume the adversary knows the cardinality of the updating set. We try to relax this assumption later in our evaluation.

**Experimental Setup.** We evaluate our label distribution estimation attack using updating set of sizes 10 and 100. For the




Fig. 6: [Lower is better for (a) and (b), higher is better for (c) and (d)] Performance of the multi-sample label distribution estimation attack ( $\mathcal{A}_{LDE}$ ) together with the baseline model and transfer attack. KL-divergence and accuracy are adopted as the evaluation metric. Accuracy here is used to measure the prediction of the most frequent label over samples in the updating set. (10) and (100) mean the target updating set contains 10 and 100 data samples, respectively. Transfer 10-100 means each of the training sample for the attack model corresponds to an updating set containing 10 data samples and each of the testing sample for the attack model corresponds to an updating set containing 100 data samples.

two different sizes, we build attack models as mentioned in the methodology. All data samples in each updating set for the shadow and target models are sampled uniformly, thus each sample (in both training and testing set) for the attack model has a uniform label distribution. We use a batch size of 64 when updating the models.

For evaluation metrics, we calculate KL-divergence for each testing sample (corresponding to an updating set on the target model) and report the average result over all testing samples (1,000 in total). Besides, we also measure the accuracy of predicting the most frequent label over samples in the updating set. We randomly sample a dataset with the same size as the updating set and use its samples' label distribution as the baseline.

**Results.** We report the result for our label distribution estimation attack in Figure 6. As shown,  $\mathcal{A}_{LDE}$  achieves a significantly better performance than the baseline model on

both datasets. For the updating set with 100 data samples on the CIFAR-10 dataset, our attack achieves 3 and 1.5 times better accuracy and KL-divergence, respectively, than the baseline model. Similarly, for the MNIST dataset, our attack also achieves 1.8 times better accuracy and 2 times better KL-divergence. Furthermore,  $\mathcal{A}_{LDE}$  achieves a similar improvement over the baseline model for the updating set of size 10.

Recall that the adversary is assumed to know the cardinality of the updating set in order to train her attack model, we further test whether we can relax this assumption. To this end, we first update the shadow model with 100 samples while updating the target model with 10 samples. As shown in Figure 6a and Figure 6c Transfer 100-10, our attack still has a similar performance as the original attack. However, when the adversary updates her shadow model with 10 data samples while the target model is updated with 100 data



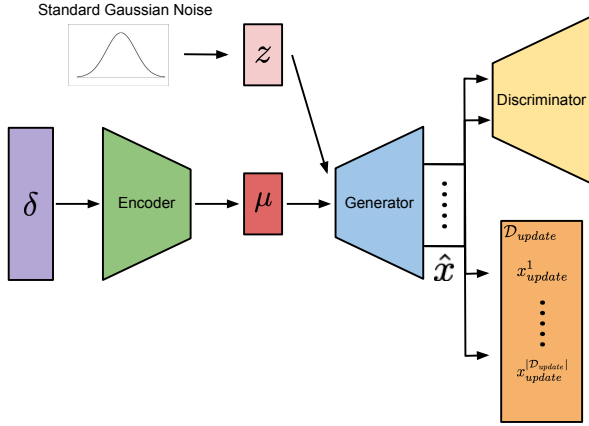


Fig. 7: Methodology of the multi-sample reconstruction attack ( $\mathcal{A}_{MSR}$ ).

samples (Figure 6b and Figure 6d Transfer 10-100), our attack performance drops significantly, in particular for KL-divergence on the CIFAR-10 dataset. This is expected as 10 samples does not provide enough information for the attack model to generalize to a larger updating set.

### B. Multi-sample Reconstruction Attack

**Attack Definition.** Our last attack, namely multi-sample reconstruction attack, aims at reconstructing the updating set. This attack can be considered as a generalization of the single-sample reconstruction attack, and a step towards the goal of reconstructing the training set of a black-box ML model.

Formally, the multi-sample reconstruction attack is defined as follows:

$$\mathcal{A}_{MSR} : \delta \mapsto \mathcal{D}_{update}$$

where  $\mathcal{D}_{update} = \{x_{update}^1, \dots, x_{update}^{|\mathcal{D}_{update}|}\}$  contains the samples used to update the model.

**Methodology.** The complexity of the task for reconstructing an updating set increases significantly when the updating set size grows from one to multiple. Our previous single-sample reconstruction attack (Section IV-B) uses autoencoder to reconstruct a single sample. However, AE cannot generate a set of samples.

*Generative Adversarial Networks.* Samples from a dataset, e.g.,  $\mathcal{D}_{update}$ , are essentially samples drawn from a complex data distribution. If an adversary is able to learn the data distribution of  $\mathcal{D}_{update}$ , then she is able to generate multiple samples with the distribution which is equivalent to reconstruct  $\mathcal{D}_{update}$ . In this work, we leverage the state-of-the-art generative model, namely generative adversarial networks (GANs) [9], the effectiveness of which has been demonstrated on learning image data distributions.

A GAN consists of a pair of ML models: a generator (G) and a discriminator (D). The generator G learns to transform a Gaussian noise vector  $z \sim \mathcal{N}(0, 1)$  to a data sample  $\hat{x}$ ,

$$G : z \mapsto \hat{x}$$

such that the generated sample  $\hat{x}$  is indistinguishable from a true data sample. This is enabled by the discriminator D which is jointly trained. The generator G tries to fool the discriminator, which is trained to distinguish between samples from the Generator (G) and true data samples. The objective function maximized by GAN’s discriminator D is,

$$\mathcal{L}_D = \mathbb{E}_{x \in \mathcal{D}_{update}} \log(D(x)) + \mathbb{E}_{\hat{x}} \log(1 - D(\hat{x})) \quad (1)$$

Thus, the GAN discriminator D is trained to output 1 (“true”) for real data and 0 (“false”) for fake data. On the other hand, the generator G maximizes:

$$\mathcal{L}_G = \mathbb{E}_{\hat{x}} \log(D(\hat{x}))$$

Thus, G is trained to produce samples  $\hat{x} = G(z)$  that are classified as “true” (real) by D.

However, the goal of our attack is to reconstruct  $\mathcal{D}_{update}$  for any given posterior difference  $\delta$ , which the standard GAN does not support. Therefore, we propose a novel hybrid generative model, referred to as BM-GAN, which conditions its generated samples  $\hat{x}$  on  $\delta$ .

*BM-GAN.* The decoder of our attack model is casted as our BM-GAN’s generator (G). To enable this, we concatenate the noise vector  $z$  and the latent vector  $\mu$  produced by our attack model’s encoder (with posterior difference as input), and use it as BM-GAN’s generator’s input, as in Conditional GANs [21]. This allows our decoder to map the posterior difference  $\delta$  to samples in  $\mathcal{D}_{update}$ .

However, Conditional GANs are severely prone to mode collapse [2], [39]. To deal with this, we introduce a reconstruction cost. This reconstruction cost forces our GAN to cover all the modes of the distribution (set) of data samples used to update the model. However, it is unclear, given a posterior difference  $\delta$  and a noise vector  $z$  pair, which point in the data distribution we should force BM-GAN to reconstruct. Therefore, we allow our GAN full flexibility in learning a mapping from posterior difference and noise vector  $z$  pairs to data samples – this means we allow it to choose the data sample to reconstruct. We realize this using a novel “Best Match” based objective in the BM-GAN formulation,

$$\mathcal{L}_{BM} = \sum_{x \in \mathcal{D}_{update}} \min_{\hat{x} \sim G} \|\hat{x} - x\|_2^2 + \sum_{\hat{x}} \log(D(\hat{x})) \quad (2)$$

where  $\hat{x} \sim G$  represents samples produced by our BM-GAN given a latent vector  $\mu$  and noise sample  $z$ . The first part of the  $\mathcal{L}_{BM}$  objective is based on the standard MSE reconstruction cost and forces our BM-GAN to reconstruct all samples in  $\mathcal{D}_{update}$  as the error is summed across  $x \in \mathcal{D}_{update}$ . However, unlike the standard MSE reconstruction cost, given a data sample  $x \in \mathcal{D}_{update}$ , the cost is based only on the generated sample  $\hat{x}$  which is closest to the data sample  $x \in \mathcal{D}_{update}$ . This allows BM-GAN to reconstruct samples in  $\mathcal{D}_{update}$  without having an explicit mapping from posterior difference and noise vector  $z$  pairs to data samples, as only the “Best Match” is penalized. Finally, the discriminator D ensures that the samples  $\hat{x}$  are indistinguishable from the “true” samples of  $\mathcal{D}_{update}$ .

Figure 7 presents a schematic view of our multi-sample reconstruction attack’s methodology. The concrete architecture of BM-GAN’s generator and discriminator for the two datasets used in this paper are listed as the following.

*BM-GAN’s generator architecture for MNIST:*

```

 $\mu, z \rightarrow$  FullyConnected(2048)
FullyConnected(2048)
FullyConnected(2048)
FullyConnected(784)  $\rightarrow$  Sample

```

*BM-GAN’s discriminator architecture for MNIST:*

```

 $\mu, z \rightarrow$  FullyConnected(1024)
FullyConnected(512)
FullyConnected(256)
FullyConnected(1)
Sigmoid  $\rightarrow$  {1, 0}

```

*BM-GAN’s generator architecture for CIFAR-10:*

```

 $\mu, z \rightarrow$  conv2d(2, 512)
conv2d(4, 256)
conv2d(4, 128)
conv2d(4, 64)
conv2d(4, 3)  $\rightarrow$  Sample

```

*BM-GAN’s discriminator architecture for CIFAR-10:*

```

 $\mu, z \rightarrow$  conv2d(2, 64)
conv2d(4, 128)
conv2d(4, 256)
conv2d(4, 512)
conv2d(4, 1)
Sigmoid  $\rightarrow$  {1, 0}

```

Here, for both generators and discriminators, Sigmoid is the Sigmoid function, batch normalization is applied on the output of each layer except the last layer, and LeakyReLU is used as the activation function for all layers except the last one, which uses tanh.

*Training of BM-GAN.* The training of the attack model  $\mathcal{A}_{MSR}$  is more complicated than previous attacks, hence we provide more details here. Similar to the previous attacks, the adversary starts the training by generating the training data as mentioned in Section III-D. She then jointly trains her encoder and BM-GAN with the posterior difference  $\delta_{shadow}^1 \cdots \delta_{shadow}^m$  as the inputs and samples inside their corresponding updating sets, i.e.,  $\mathcal{D}_{shadow}^{update^1} \cdots \mathcal{D}_{shadow}^{update^m}$  as the output. More concretely,

for each posterior difference  $\delta_{shadow}^i$ , she updates her attack model  $\mathcal{A}_{MSR}$  as follows:

- 1) The adversary sends the posterior difference  $\delta_{shadow}^i$  to her encoder to get the latent vector  $\mu_i$ .
- 2) She then generates  $|\mathcal{D}_{shadow}^{update^i}|$  noise vectors.
- 3) To create the BM-GAN’s generator input, she concatenates each of the noise vectors with the latent vector  $\mu_i$ .
- 4) On the input of the concatenated vectors, the BM-GAN generates  $|\mathcal{D}_{shadow}^{update^i}|$  samples, i.e., each vector corresponds to each sample.
- 5) The adversary then calculates the generator loss as introduced by Equation 2, and uses it to update the generator and the encoder.
- 6) Finally, she calculates and updates the BM-GAN’s discriminator according to the loss function introduced in Equation 1.

As before, all the optimization is performed by the ADAM optimizer.

*Clustering.* BM-GAN only provides a generator which learns the distribution of the samples in the updating set. However, to reconstruct the exact data samples in  $\mathcal{D}_{update}$ , we need a final step assisted by machine learning clustering. In detail, we assume the adversary knows the cardinality of  $\mathcal{D}_{update}$  as in Section V-A. After BM-GAN is trained, the adversary utilizes BM-GAN’s generator to generate a large number of samples. She then clusters the generated samples into  $|\mathcal{D}_{update}|$  clusters. Here, the K-means algorithm is adopted to perform clustering where we set K to  $|\mathcal{D}_{update}|$ . In the end, for each cluster, the adversary calculates its centroid, and takes the nearest sample to the centroid as one reconstructed sample.

**Experimental Setup.** We evaluate the multi-sample reconstruction attack on the updating set of size 100 and generate 20,000 images for each updating set reconstruction with BM-GAN. For the rest of the experimental settings, we follow the one mentioned in Section V-A except for evaluation metrics and baseline.

We use MSE between the updating and reconstructed data samples to measure the multi-sample reconstruction attack’s performance. For the baseline model of this attack, we perform K-means clustering on the adversary’s shadow dataset. More concretely, we cluster the adversary’s shadow dataset into 100 clusters and take the nearest sample to the centroid of each cluster as one reconstructed sample.

**Results.** In Figure 8, we first present some visualization of the intermediate result of our attack, i.e., the BM-GAN’s output before clustering, on the CIFAR-10 dataset. For each randomly sampled image in the updating set, we show the 5 nearest reconstructed images with respect to MSE generated by BM-GAN. As we can see, our attack model generates images with similar characteristics to the original images. For instance, the 5 reconstructed images for the airplane image in Figure 8b all show a blue background and a blurry version of the airplane itself. The similar result can be observed from the



Fig. 8: Visualization of some generated samples from the multi-sample reconstruction attack ( $\mathcal{A}_{MSR}$ ) before clustering on the CIFAR-10 dataset. Samples are fair random draws, not cherry-picked. The left column shows the original samples and the next 5 columns show the 5 nearest reconstructed samples with respect to mean squared error.

boat image in Figure 8a, the car image in Figure 8c, and the boat image in Figure 8d. It is also interesting to see that BM-GAN provides different samples for the two different horse images in Figure 8b. The blurriness in the results is expected, due to the complex nature of the CIFAR-10 dataset and the weak assumptions for our adversary, i.e., access to black-box ML model.

We also quantitatively measure the performance of our intermediate results, by calculating the MSE between each image in the updating set and its nearest reconstructed sample. We refer to this as one-to-one match. Figure 10 shows for the CIFAR-10 dataset, we achieve 0.0283 MSE, and for the MNIST dataset, we achieve 0.043. It is important to note that the adversary cannot perform one-to-one match as she does not have access to ground truth samples in the updating set, i.e., one-to-one match is an oracle.

Figure 10 shows the mean squared error of our full attack with clustering for both CIFAR-10 and MNIST datasets. To match each of our reconstructed samples to a sample in  $\mathcal{D}_{update}$ , we rely on the Hungarian algorithm [16]. This guarantees that each reconstructed sample is only matched with one ground truth sample in  $\mathcal{D}_{update}$  and vice versa. As we can see, our attack outperforms the baseline model on both datasets by 20% and 22%, respectively. In detail, our attack achieves an MSE of 0.036 on the CIFAR-10 dataset and 0.051 on the MNIST dataset. As expected, the result of our final attack is lower than one-to-one match, i.e., the above mentioned intermediate results.

We further visualize our full attack’s result on the MNIST dataset. Figure 9 shows a sample of a full MNIST updating set reconstruction, i.e., the BM-GAN’s reconstructed images for the 100 original images in an updating set. We observe



Fig. 9: Visualization of a full MNIST updating set together with the output of the multi-sample reconstruction attack ( $\mathcal{A}_{MSR}$ ) after clustering. Samples are fair random draws, not cherry-picked. The left column shows the original samples and the right column shows the reconstructed samples. The match between the original and reconstructed samples is performed by the Hungarian algorithm.




Fig. 10: [Lower is better] Performance of the multi-sample reconstruction attack ( $\mathcal{A}_{MSR}$ ) together with one-to-one match and the baseline model. Mean squared error (MSE) is adopted as the evaluation metric. The match between the original and reconstructed samples is performed by the Hungarian algorithm for both  $\mathcal{A}_{MSR}$  and the baseline model. One-to-one match serves as an oracle as the adversary cannot use it for her attack.

that our attack model reconstructs diverse digits of each class that for most of the cases match the actual ground truth data very well. This suggests BM-GAN is able to capture all modes in a data distribution well. Only in limited cases, our attack does not provide a perfect reconstruction, e.g., the digit “3” in the first image of Figure 9 is reconstructed to “8”. One

limitation of our attack is that BM-GAN’s sample generation and clustering are performed separately. In the future, we plan to combine them to perform an end-to-end training which may further boost our attack’s performance.

From all these results, we show that our attack in the most difficult setting still does not generate a general representation of data samples affiliated with the same label, but reconstructs the specific images inside the updating set.

## VI. POSSIBLE DEFENSES

Our attacks make the first attempt to explore a new attack surface against black-box ML models. Extensive experimental results have demonstrated their effectiveness. In this section, we discuss two possible defense mechanisms in order to mitigate the threat.

**Adding Noise to Posteriors.** All our attacks leverage posterior difference as the input. Therefore, to reduce our attacks’ performance, one could sanitize posterior difference. However, the model owner cannot directly manipulate the posterior difference, as she does not know with what or when the adversary probes her model. Therefore, she has to add noise to the posterior for each queried sample independently. We have tried adding noise sampled from a uniform distribution to the posteriors. Experimental results show that the performance for some of our attacks indeed drops to a certain degree. For instance, the single-sample label inference attack on the CIFAR-10 dataset drops by 17% in accuracy. However, the performance of our multi-sample reconstruction attack stays stable. One reason might be the noise vector  $z$  is part of BM-

GAN’s input which makes the attack model more robust to the noise input.

**Differential Privacy.** Another possible defense mechanism against our attacks is differentially private learning. Differential privacy [5] can help an ML model learn its main tasks while reducing its memory on the training data. If differentially private learning schemes [1], [4], [30] are used when updating the target ML model, this by design will reduce the performance of our attacks. However, it is also important to mention that depending on the privacy budget for differential privacy, the utility of the model can drop significantly.

We leave an in-depth exploration of effective defense mechanisms against our attacks as a future work.

## VII. RELATED WORKS

**Membership Inference.** Membership inference aims at determining whether a data sample is inside a dataset. It has been successfully performed in various settings such as biomedical data [11], [14] and location data [27], [28]. Shokri et al. [31] propose the first membership inference attack against machine learning models. In this attack, an adversary’s goal is to determine whether a data sample is in the training set of a black-box ML model. To mount this attack, the adversary relies on a binary machine learning classifier which is trained with the data derived from shadow models (similar to our attacks). More recently, multiple membership inference attacks have been proposed with new attacking techniques or targeting on different types of ML models [12], [18], [19], [22], [23], [29], [33], [40].

In theory, membership inference attack can be used to reconstruct the dataset, similar to our reconstruction attacks. However, it is not scalable in the real-world setting as the adversary needs to obtain a large-scale dataset which includes all samples in the target model’s training set. Though our two reconstruction attacks are designed specifically for the online learning setting, we believe the underlying techniques we propose, i.e., pretrained decoder from a standard autoencoder and BM-GAN, can be further extended to reconstruct datasets from black-box ML models in other settings.

**Model Inversion.** Fredrikson et al. [7] propose model inversion attack first on biomedical data. The goal of model inversion is to infer some missing attributes of an input feature vector based on the interaction with a trained ML model. Later, other works generalize the model inversion attack to other settings, for instance, to reconstruct recognizable human faces [6], [13]. As pointed out by other works [20], [31], model inversion attack reconstructs a general representation of data samples affiliated with certain labels, while our reconstruction attacks target on specific data samples used in the updating set.

**Model Stealing.** Another related line of work is model stealing. Tramèr et al. [36] are among the first to introduce the model stealing attack against black-box ML models. In this attack, an adversary tries to learn the target ML model’s

parameters. Tramèr et al. propose various attacking techniques including equation-solving and decision tree path-finding. The former has been demonstrated to be effective on simple ML models, such as logistic regression, while the latter is designed specifically for decision trees, a class of machine learning classifiers. Moreover, relying on an active learning based retraining strategy, the authors show that it is possible to steal an ML model even if the model only provides the label instead of posteriors as the output. More recently, Orekondy et al. [25] propose a more advanced attack on stealing the target model’s functionality and show that their attack is able to replicate a mature commercial machine learning API. In addition to model parameters, several works concentrate on stealing ML models’ hyperparameters [24], [38].

Besides the above, there exist a wide range of other attacks and defenses on machine learning models [1], [3], [4], [8], [10], [15], [17], [26], [32], [34], [35], [37].

## VIII. CONCLUSION

Large-scale data being generated at every second turns ML model training into a continuous process. In consequence, a machine learning model queried with a same set of data samples at two different time points will provide different results. In this paper, we investigate whether these different model outputs can constitute a new attack surface for an adversary to infer information of the dataset used to perform model update.

We propose four different attacks in this surface all of which follow a general encoder-decoder structure. The encoder encodes the difference in the target model’s output before and after being updated, and the decoder generates different types of information regarding the updating set.

We start by exploring a simplified case when a black-box ML model is only updated with one single data sample. We propose two different attacks for this setting, The first attack shows that the label of the single updating sample can be effectively inferred (0.96 accuracy on the CIFAR-10 dataset). The second attack utilizes an autoencoder’s decoder as the attack model’s pretrained decoder and accurately reconstructs the single updating sample.

We then generalize our attacks to the case when the updating set contains multiple samples. Our multi-sample label distribution estimation attack trained following a KL-divergence loss is able to infer the label distribution of the updating set’s data samples effectively. For our last and the most complex attack – the multiple-sample reconstruction attack, we propose a novel hybrid generative model, namely BM-GAN, which uses a “Best Matching” loss in its objective function. The “Best Matching” loss directs BM-GAN’s generator to reconstruct each sample in the updating set. Quantitative and qualitative results show that we are able to reconstruct the updating dataset in challenging scenarios. To the best of our knowledge, this constitutes the first attack of this type, which is able to infer very detailed information on the dataset and even lends itself to full reconstruction of the data.

## REFERENCES

- [1] M. Abadi, A. Chu, I. Goodfellow, B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep Learning with Differential Privacy," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2016, pp. 308–318. [13](#)
- [2] A. Brock, J. Donahue, and K. Simonyan, "Large Scale GAN Training for High Fidelity Natural Image Synthesis," in *Proceedings of the 2-19 International Conference on Learning Representations (ICLR)*, 2-19. [9](#)
- [3] N. Carlini and D. Wagner, "Towards Evaluating the Robustness of Neural Networks," in *Proceedings of the 2017 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2017, pp. 39–57. [13](#)
- [4] K. Chaudhuri and C. Monteleoni, "Privacy-preserving Logistic Regression," in *Proceedings of the 2009 Annual Conference on Neural Information Processing Systems (NIPS)*. NIPS, 2009, pp. 289–296. [13](#)
- [5] C. Dwork and A. Roth, "The Algorithmic Foundations of Differential Privacy," *Foundations and Trends in Theoretical Computer Science*, 2014. [12](#)
- [6] M. Fredrikson, S. Jha, and T. Ristenpart, "Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures," in *Proceedings of the 2015 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2015, pp. 1322–1333. [1](#), [13](#)
- [7] M. Fredrikson, E. Lantz, S. Jha, S. Lin, D. Page, and T. Ristenpart, "Privacy in Pharmacogenetics: An End-to-End Case Study of Personalized Warfarin Dosing," in *Proceedings of the 2014 USENIX Security Symposium (USENIX Security)*. USENIX, 2014, pp. 17–32. [13](#)
- [8] K. Ganju, Q. Wang, W. Yang, C. A. Gunter, and N. Borisov, "Property Inference Attacks on Fully Connected Neural Networks using Permutation Invariant Representations," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2018, pp. 619–633. [1](#), [3](#), [4](#), [13](#)
- [9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Nets," in *Proceedings of the 2014 Annual Conference on Neural Information Processing Systems (NIPS)*. NIPS, 2014. [9](#)
- [10] I. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and Harnessing Adversarial Examples," in *Proceedings of the 2015 International Conference on Learning Representations (ICLR)*, 2015. [13](#)
- [11] I. Hagestedt, Y. Zhang, M. Humbert, P. Berrang, H. Tang, X. Wang, and M. Backes, "MBeacon: Privacy-Preserving Beacons for DNA Methylation Data," in *Proceedings of the 2019 Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2019. [13](#)
- [12] J. Hayes, L. Melis, G. Danezis, and E. D. Cristofaro, "LOGAN: Evaluating Privacy Leakage of Generative Models Using Generative Adversarial Networks," *CoRR abs/1705.07663*, 2017. [13](#)
- [13] B. Hitaj, G. Ateniese, and F. Perez-Cruz, "Deep Models Under the GAN: Information Leakage from Collaborative Deep Learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2017, pp. 603–618. [1](#), [13](#)
- [14] N. Homer, S. Szlinger, M. Redman, D. Duggan, W. Tembe, J. Muehling, J. V. Pearson, D. A. Stephan, S. F. Nelson, and D. W. Craig, "Resolving Individuals Contributing Trace Amounts of DNA to Highly Complex Mixtures Using High-Density SNP Genotyping Microarrays," *PLOS Genetics*, 2008. [13](#)
- [15] M. Jagielski, A. Oprea, B. Biggio, C. Liu, C. Nita-Rotaru, and B. Li, "Manipulating Machine Learning: Poisoning Attacks and Countermeasures for Regression Learning," in *Proceedings of the 2018 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2018. [13](#)
- [16] H. W. Kuhn, "The Hungarian Method for the Assignment Problem," *Naval Research Logistics Quarterly*, 1955. [11](#)
- [17] B. Li and Y. Vorobeychik, "Scalable Optimization of Randomized Operational Decisions in Adversarial Classification Settings," in *Proceedings of the 2015 International Conference on Artificial Intelligence and Statistics (AISTATS)*. PMLR, 2015, pp. 599–607. [13](#)
- [18] Y. Long, V. Bindschaedler, and C. A. Gunter, "Towards Measuring Membership Privacy," *CoRR abs/1712.09136*, 2017. [13](#)
- [19] Y. Long, V. Bindschaedler, L. Wang, D. Bu, X. Wang, H. Tang, C. A. Gunter, and K. Chen, "Understanding Membership Inferences on Well-Generalized Learning Models," *CoRR abs/1802.04889*, 2018. [1](#), [13](#)
- [20] L. Melis, C. Song, E. D. Cristofaro, and V. Shmatikov, "Exploiting Unintended Feature Leakage in Collaborative Learning," in *Proceedings of the 2019 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2019. [1](#), [2](#), [7](#), [13](#)
- [21] M. Mirza and S. Osindero, "Conditional Generative Adversarial Nets," *CoRR abs/1411.1784*, 2014. [9](#)
- [22] M. Nasr, R. Shokri, and A. Houmansadr, "Machine Learning with Membership Privacy using Adversarial Regularization," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2018. [13](#)
- [23] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive Privacy Analysis of Deep Learning: Passive and Active White-box Inference Attacks against Centralized and Federated Learning," in *Proceedings of the 2019 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2019. [13](#)
- [24] S. J. Oh, M. Augustin, B. Schiele, and M. Fritz, "Towards Reverse-Engineering Black-Box Neural Networks," in *Proceedings of the 2018 International Conference on Learning Representations (ICLR)*, 2018. [3](#), [13](#)
- [25] T. Orekondy, B. Schiele, and M. Fritz, "Knockoff Nets: Stealing Functionality of Black-Box Models," in *Proceedings of the 2019 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2019. [13](#)
- [26] N. Papernot, P. D. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, "Practical Black-Box Attacks Against Machine Learning," in *Proceedings of the 2017 ACM Asia Conference on Computer and Communications Security (ASIACCS)*. ACM, 2017, pp. 506–519. [13](#)
- [27] A. Pyrgelis, C. Troncoso, and E. D. Cristofaro, "Knock Knock, Who's There? Membership Inference on Aggregate Location Data," in *Proceedings of the 2018 Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2018. [13](#)
- [28] A. Pyrgelis, C. Troncoso, and E. D. Cristofaro, "Under the Hood of Membership Inference Attacks on Aggregate Location Time-Series," *CoRR abs/1902.07456*, 2019. [13](#)
- [29] A. Salem, Y. Zhang, M. Humbert, P. Berrang, M. Fritz, and M. Backes, "ML-Leaks: Model and Data Independent Membership Inference Attacks and Defenses on Machine Learning Models," in *Proceedings of the 2019 Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2019. [1](#), [3](#), [4](#), [13](#)
- [30] R. Shokri and V. Shmatikov, "Privacy-Preserving Deep Learning," in *Proceedings of the 2015 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2015, pp. 1310–1321. [13](#)
- [31] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership Inference Attacks Against Machine Learning Models," in *Proceedings of the 2017 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2017, pp. 3–18. [1](#), [2](#), [3](#), [4](#), [7](#), [13](#)
- [32] C. Song, T. Ristenpart, and V. Shmatikov, "Machine Learning Models that Remember Too Much," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2017, pp. 587–601. [13](#)
- [33] C. Song and V. Shmatikov, "The Natural Auditor: How To Tell If Someone Used Your Words To Train Their Model," *CoRR abs/1811.00513*, 2018. [13](#)
- [34] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing Properties of Neural Networks," *CoRR abs/1312.6199*, 2013. [13](#)
- [35] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel, "Ensemble Adversarial Training: Attacks and Defenses," in *Proceedings of the 2017 International Conference on Learning Representations (ICLR)*, 2017. [13](#)
- [36] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Stealing Machine Learning Models via Prediction APIs," in *Proceedings of the 2016 USENIX Security Symposium (USENIX Security)*. USENIX, 2016, pp. 601–618. [13](#)
- [37] Y. Vorobeychik and B. Li, "Optimal Randomized Classification in Adversarial Settings," in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 2014, pp. 485–492. [13](#)
- [38] B. Wang and N. Z. Gong, "Stealing Hyperparameters in Machine Learning," in *Proceedings of the 2018 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2018. [3](#), [13](#)
- [39] D. Yang, S. Hong, Y. Jang, T. Zhao, and H. Lee, "Diversity-Sensitive Conditional Generative Adversarial Networks," in *Proceedings of the 2019 International Conference on Learning Representations (ICLR)*, 2019. [9](#)
- [40] S. Yeom, I. Giacomelli, M. Fredrikson, and S. Jha, "Privacy Risk in Machine Learning: Analyzing the Connection to Overfitting," in *Pro-*

APPENDIX

TABLE II: Hyperparameters for AE's encoder and decoder.

Variable	MNIST	CIFAR-10
$k_1$	3	3
$s_1$	16	32
$k_2$	3	3
$s_2$	8	16
$f_1$	15	50
$f_2$	10	30
$f'_1$	15	50
$f'_2$	32	64
$k'_1$	3	3
$s'_1$	16	32
$k'_2$	5	5
$s'_2$	8	16
$k'_3$	2	4
$s'_3$	1	3