

GAMIN: An Adversarial Approach to Black-Box Model Inversion

Ulrich Aïvodji

Université du Québec à Montréal
aivodji.ulrich@courrier.uqam.ca

Sébastien Gambs

Université du Québec à Montréal
gambs.sebastien@uqam.ca

Timon Ther

ISAE-Supaéro
timon.ther@isae-alumni.net

Abstract

Recent works have demonstrated that machine learning models are vulnerable to model inversion attacks, which lead to the exposure of sensitive information contained in their training dataset. While in the past some model inversion attacks have been developed in the black-box attack setting, in which the adversary does not have direct access to the structure of the model, few of these have been conducted so far against complex models such as deep neural networks. In this paper, we introduce GAMIN (for *Generative Adversarial Model INversion*), a new black-box model inversion attack framework achieving significant results even against deep models such as convolutional neural networks at a reasonable computing cost. GAMIN is based on the continuous training of a surrogate model for the target model under attack and a generator whose objective is to generate inputs resembling those used to train the target model. The attack was validated against various neural networks used as image classifiers. In particular, when attacking models trained on the MNIST dataset, GAMIN is able to extract recognizable digits for up to 60% of labels produced by the target. Attacks against skin classification models trained on the pilot parliament dataset also demonstrated the capacity to extract recognizable features from the targets.

Introduction

In recent years, the combination of increasing computing power, significant improvement in machine learning algorithms as well as the augmentation of storage capacity have led to the development of prediction services. Those services rely on exploiting a wide variety of data, the collection of which has become an integral part of the business of many actors. This push for personalized services has driven forward the use of sensitive data even in public-oriented applications such as evaluating the risk of recidivism of convicted inmates. In addition, data has also become a vital asset for providers, which often seek to ensure the exclusivity over it.

Unfortunately, data-driven algorithms such as machine learning models are not immune to information leakage (Barreno et al. 2006; Kasiviswanathan, Rudelson, and Smith 2013). In particular, if those models have been trained over personal information, an attack exploiting those leakages

could result in a privacy breach (e.g., by inferring a sensitive feature) for the individuals contained in the dataset used to train a learning algorithm. The potency of such an attack was first demonstrated in 2015 by Fredrikson, Jha and Ristenpart (Fredrikson, Jha, and Ristenpart 2015), extending as far as retrieving recognizable facial features or inferring with high accuracy a single sensitive attribute. In subsequent works, efforts have been made to quantify the vulnerability of machine learning models against inference attacks related to privacy. For instance, new attacks, such as membership inference attack (Shokri et al. 2017), as well as new metrics such as differential training privacy (Long, Bindschaedler, and Gunter 2017), have been developed.

Nonetheless many questions remain opened, especially with respect to the exploration of the possible attacks settings. Indeed, depending on the ability of the adversary to access the target model and his auxiliary knowledge of that model, the approach taken as well as the resulting efficiency can change significantly. Considering *model inversion attacks*, in which the attacker attempts to generate inputs resembling the original ones, Fredrikson and co-authors have left the design of effective black-box attacks as an open problem in their seminal paper (Fredrikson, Jha, and Ristenpart 2015). While a first significant step was proposed later (Tramèr et al. 2016), which consists in performing a black-box model extraction attack followed by the same white-box equation-solving attack used by Fredrikson and co-authors, this does not preclude the possibility that other approaches could be more efficient in terms of the reconstruction accuracy or the knowledge required on the attacked model. In particular, one limit of the equation-solving attack is that it requires detailed knowledge of the target’s architecture.

In this paper, we propose a novel approach that we coined as GAMIN (for *Generative Adversarial Model INversion*) to specifically address the problem of black-box model inversion. Our attack is agnostic in the sense that we make no assumptions on the design of the target model or the original data distribution, which could arguably be the most difficult situation for the adversary (Papernot et al. 2016). In a nutshell, our attack involves the simultaneous training of a *surrogate model* approximating the decision boundaries of the target and a *generator* that aims at producing the desired inputs, by mimicking the Generative Adversarial Networks (GANs) training process. To evaluate its cost and per-

formance and how those are impacted by the target’s architecture, our approach was tested against a collection of classifier models. To ensure the generality of this evaluation, we chose to consider targets of various complexity, the design of which has been observed and tested in previous works in the literature. In particular, we have been able to demonstrate the potency of this attack against image classification models for digit recognition or skin color identification.

The outline of the paper is as follows. First, we introduce the preliminaries necessary to the understanding of our work in particular with respect to machine learning models and their vulnerabilities with respect to various privacy attacks. Afterwards, we present the architecture of the GAMIN before describing in details how this attack performs against five models trained on the MNIST and the pilot parliament datasets. Finally, we briefly review the related work in before concluding and discussing future work.

Preliminaries

Generative adversarial networks. *Generative Adversarial Networks* (GANs) (Goodfellow et al. 2014) refers to a framework composed of two models competing against each other, hence the term “adversarial”. In this setting, one of the models has the role of the generator G while the other takes the role of the discriminator D . In a nutshell, D aims at distinguishing real data from fake one while G tries to fools D into taking a fake data it has generated for a real one. This process has been formally defined (Goodfellow et al. 2014) as a minimax game between G and D :

$$\min_G \max_D \mathbb{E}_{x \sim p_x(x)} [\log(D(x))] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))],$$

in which $p_x(x)$ is the real data distribution and $p_z(z)$ is usually a simple distribution (e.g., $\mathcal{N}(0, 1)$).

Adversarial setting. When attacking a machine learning model, an adversary can have different access and knowledge of the target model (Wu et al. 2016) that we briefly introduced hereafter.

White-box setting. In the white-box setting, the adversary knows the architecture of the target model as well as its parameters and internal states. Typically, this setting occurs when a malicious client obtains a complete trained model from a MLaaS (*Machine Learning as a Service*) provider rather than a dataset and wishes to learn more about the original proprietary data this model has been trained on.

Black-box setting. When only given black-box access, the adversary does not have a direct access to the model and can only query it with an input of his choice to receive the corresponding output (usually through an online API). For example, this scenario corresponds to the situation in which a company providing MLaaS is attacked by one of its users. However, in some cases the adversary may have partial knowledge of the training data (e.g., the characteristics of the distribution) or of the target model’s architecture.

Black-box agnostic setting. The arguably hardest setting for the adversary is when it has absolutely no information on the training data (including no prior knowledge of the data distribution) as well as the model architecture or any of its parameters besides the dimensions of inputs and outputs.

Basically, in this setting, which we refer to as the *agnostic black-box setting*, the adversary can only query the target model to make inferences about it.

Model inversion attacks. *Model inversion attacks* (Fredrikson et al. 2014; Fredrikson, Jha, and Ristenpart 2015) (also known as *attribute inference attacks* (Yeom et al. 2018; Jayaraman and Evans 2019)) aim to infer hidden sensitive attribute of instances that belong to a particular output class of a target machine learning model. Pioneering works in this field (Fredrikson et al. 2014; Fredrikson, Jha, and Ristenpart 2015) have proposed a generic framework to conduct model inversion attacks that can be described as follows. Given a target model T , an instance x — whose non-sensitive attributes $x_1, x_2, \dots, x_{x_d-1}$ and attributes’ prior probabilities are known to the adversary — and the prediction $y = T(x)$ of the target model T , the attack identifies the value of the sensitive attribute x_d maximizing the posterior probability $P(x_d | x_1, x_2, \dots, x_{x_d-1}, y)$.

Model inversion attacks have been successfully conducted against a variety of models in the white-box setting, including neural networks with equation-solving methods (Fredrikson, Jha, and Ristenpart 2015), through the reconstruction of recognizable portraits. The question of effective black-box attack with respect to this method was left as an open question. One of the proposed approach to realize this is to run a white-box inversion attack after a successful black-box model extraction attack (Tramèr et al. 2016), thus achieving similar results at a fraction of the cost. However, none of these attacks were performed against a convolutional neural network or other types of deep neural networks. In addition, equations-solving attacks require prior knowledge of the target model’s architecture, which severely undermines their feasibility in an agnostic black-box setting.

For image recognition tasks, the output of a model inversion attack is usually not a member of the training set, but rather an average of the features characterizing the class. Depending on the diversity of inputs yielding the targeted class, this average may lead to a more or less severe privacy breach. For instance, if a class correspond to a particular individual and the training instances corresponding to the latter have small variance (e.g., same location for the face or same shooting angle) then the result of the attack can help in re-identifying the individual. In contrast, if the output class characterizes highly diverse instances (e.g., car of very different types), then the output of the attack is more difficult to interpret. In this case, the output of the attack is very similar to that of a *property inference attack* (Ateniese et al. 2015).

Our approach: The GAMIN

First, we introduce the *Generative Adversarial Model INversion* (GAMIN) framework before detailing the attack.

Overview

Let us consider a target model T and a label y_t corresponding to one of its labels. The objective of our attack is to characterize the inputs x_t such that $T(x_t) = y_t$.

At a high level, GAMIN is composed of two deep neural networks, namely a *generator* $G : Z \rightarrow X$ that maps noise $z \sim \mathcal{N}(0, 1)$ to an input x_G and a *surrogate* model $S : X \rightarrow Y$ that outputs an estimation \hat{y} of the target model’s output. GAMIN allows to simultaneously train the surrogate model S and the generator G while performing a model inversion attack over S . Thus, the generator aims at learning the distribution of input x_t associated to label y_t .

Algorithm 1 GAMIN training

Require: (G, θ_G) generator model, (S, θ_S) shadow model, T target model, y_t target label to invert, k_0 initial boundary-equilibrium factor, λ_k, γ_k boundary-equilibrium update parameters
 $k \leftarrow k_0$
for n epochs **do**
 $Z_G \sim \mathcal{N}(0, 1)$ \triangleright Generate artificial inputs from noise
 $X_G \leftarrow G(Z_G)$
 $X_S \sim \mathcal{N}(0, 1)$ \triangleright Generate raw noise input
 $Y_G \leftarrow T(X_G)$ \triangleright Query from the target model
 $Y_S \leftarrow T(X_S)$
 \triangleright Compute boundary-equilibrium loss and train surrogate
 $L_S \leftarrow L_H(y_t, Y_S) - k * L_H(y_t, Y_G)$
 $\theta_S \leftarrow \text{train}(X_S, Y_S, L_S)$
 $k \leftarrow k + \lambda_k(\gamma_k L_H(y_t, Y_S) - L_H(y_t, Y_G))$ \triangleright Train generator
 $Z_G \sim \mathcal{N}(0, 1)$
 $\theta_G \leftarrow \text{train}(Z_G, y_t)$
end for

Algorithm 1 summarizes the training sequence of GAMIN. For each step of the algorithm, a batch Z_G is sampled from random noise $\mathcal{N}(0, 1)$. Then, the generator uses the latter to produce a batch X_G . Afterwards, the target model T is queried with (1) the output X_G of the generator and (2) a batch X_S sampled from random noise $\mathcal{N}(0, 1)$. Subsequently, the predictions of T on both X_G and X_S are used to compute the surrogate loss L_S , which is then used to inform the training of the surrogate S . Finally, the generator’s parameters are updated through the training of the combined model $S \circ G$, on a batch Z_G sampled from random noise $\mathcal{N}(0, 1)$. After convergence, the surrogate learns the target model’s decision boundaries and the generator learns to approximate x_t .

Given the black-box agnostic setting of the adversary, the architecture of GAMIN has to be as generic as possible. More precisely, the only constraint put on the architecture is that the dimensions and ranges of the outputs of the generator match the inputs of the target model T and GAMIN’s surrogate S . We demonstrate this by using the same architecture — for both GAMIN’s surrogate and the generator — for all the attacks conducted in this paper. In addition, we have deliberately chosen for the surrogate model an architecture different from that of the target models to emphasize the absence of prior knowledge by the adversary. Technical

details of GAMIN’s surrogate and generator can be found in appendices.

Loss, metrics and convergence

The adversarial nature of GAMIN requires the surrogate and the generator to be trained alternatively, each with its own losses and metrics. However, the training of one component impact the performances of the other. To overcome this challenge, we need to devise appropriate loss functions and metrics to have better control of the training process.

Boundary-equilibrium loss for surrogate model. Initial experiments with this architecture proved its difficulty to assess convergence and to compare the performance performances of both the surrogate and the generator. Since these issues are common in the GAN literature, we adapted the *Boundary-Equilibrium Adaptive* loss from BEGAN (Berthelot, Schumm, and Metz 2017) to GAMIN. The main idea is to have the loss function updating itself to reflect the trade-off made by the surrogate loss on samples according to their nature. More precisely, the distinction operated by the surrogate will be between noise inputs and generated samples. Further differences with the original BEGAN include adapting the setting from an auto-encoder perspective to our GAMIN setting. The loss for the surrogate model L_S is defined as:

$$L_S = L_H(y_t, Y_S) - k_t * L_H(y_t, Y_G)$$

$$k_{t+1} = k_t + \lambda_k(\gamma_k L_H(y_t, Y_S) - L_H(y_t, Y_G)),$$

in which L_H is the cross-entropy loss function, Y_S and Y_G are respectively the predictions of the target model T on X_S (*i.e.*, noise inputs for the surrogate S) and X_G (*i.e.*, inputs crafted by the generator G), λ_k is a “learning rate” for the parameter k_t and γ_k is an equilibrium ratio objective.

By allowing to balance the training of the surrogate and the generator, this loss helps in avoiding the issue of *the unforgiving teacher*. This phenomenon is observed when a GAN’s generator does not improve because its outputs are so far from expectations that it receives uniformly negative output (Salimans et al. 2016). In addition, the adaptive loss helps in determining and improving the component that is limiting the overall performance. Finally, it reduces the occurrences of *catastrophic forgetting*, which plague GAN-based architectures and occurs when the learning of a new skill damages a previously learned one (French 1999).

Generator model training through combined networks.

The update of the generator is achieved by merging both the generator and surrogate into a virtual “combined” neural network $S \circ G$. Then, the parameters of the combined network are updated through a standard training procedure. During this training, all the parameters related to the surrogate are kept frozen. Recall that the objective of the “combined” model is to map noise z_G to the estimation of the target outputs upon receiving the generated inputs $S(G(z_G))$, to achieve $S(G(z_G)) \approx y_t$ through training. Thus, the “combined” model is trained as performing a simple classification task with a cross-entropy loss function.

Relevance metrics. In this section, we discuss the metrics used to monitor and control the training and convergence of GAMIN.

Surrogate fidelity. One of the fundamental measure for our approach is the *surrogate fidelity*, which reflects the capacity of the surrogate to imitate the target model’s behaviour. With the same input x for the target model T and the surrogate S , let $y = T(x)$ and $\hat{y} = S(x)$ be the outputs of these models. The fidelity F_S of the surrogate can be expressed as $1 - \frac{\|y - \hat{y}\|}{\|y\|}$. In practice, we compute the mean of this fidelity, which is then $1 - mae(y, \hat{y})$, mae being the mean average error, and we test the fidelity over a batch of 64 samples of uniform random noise. Note that fidelity differs from precision in that it needs not be evaluated over testing samples, which means that it can be measured during a black-box attack with no knowledge of the original data. A surrogate with high fidelity (up to 1) estimates with good precision the decision boundaries of the target model.

Combined accuracy. The generator’s ability to craft inputs that are classified in the desired way by the surrogate is measured by the *categorical accuracy* $A_{S \circ G}$ of the “combined” model. When assessed in a batch, the categorical accuracy is the proportion of inputs categorized in the desired class.

BEGAN M-global score. The BEGAN-inspired adaptive loss comes with another useful metric from the same study: the *global convergence score* M_{global} (Berthelot, Schumm, and Metz 2017). This metric is defined as :

$$M_{global} = L_S + |\gamma_k L_S - L_{S \circ G}|,$$

in which L_S is the loss of the surrogate on noise inputs and $L_{S \circ G}$ is the loss of the surrogate on inputs generated by the GAMIN generator.

The M-global convergence score provides an insight on the convergence of both networks, which is helpful to assess the overall convergence since the adversarial nature of the GAMIN architecture means an improvement of the surrogate may impact the generator (and vice-versa). One of the strength of this score is that it can be computed in a black-box setting as it only relies on the discrimination between different types of inputs used to train the surrogate. This score and the adaptive loss are updated at the same time.

Attack protocol and budget

To perform an attack, GAMIN is given a query budget, which represents the number of queries that can be sent to the target model T . GAMIN is then trained against the target model, saving both the surrogate and the generator whenever the M-global score improves while tracking the remaining query budget. Upon exhausting the query budget, the generator with the lowest M-global score is used as the result of the attack to generate the targeted inputs.

This approach provides two main benefits. First, it helps to determine the actual number of queries needed to achieve the best result. Second, it preserves the best model observed during training, which means that catastrophic forgetting or a decrease in performance will have a lesser impact on the final performance of the attack.

A post-processing routine has been added, which has considerably enhanced the visual quality of generated images by

removing a good part of the noise cluttering the individual output (see Figure 1 for an example).




Figure 1: Raw GAMIN output example (left), result of the post-processing on the same attack (center), and mean of the attacked class in original dataset (right).

Further post-processing such as thresholds or posterization could still be added to improve the quality of images, but those are left to explore as future work.

Black-box model inversion attacks

In this section, we describe the results obtained against several target models. First, we present the datasets used to train the target models as well as the models themselves before reporting on the results obtained. Afterwards, we discuss on how to evaluate the success of a GAMIN attack as well as the associated cost before finally reporting how models trained with differential privacy as a possible countermeasure are affected by the attack. The source code of GAMIN is available on *github*¹.

Targets and evaluation setting

MNIST dataset. The MNIST Dataset of 70,000 handwritten digits was used to train target models for a image recognition task. More precisely, in this case the recognition task consists in associating an image passed as input to a class corresponding to the digit associated to the image. In this dataset, images are of low-sizes (square of size 28 by 28 pixels) and black-and-white, which provides a relatively low-cost task. The target models were trained on a proportion of the dataset accounting for 60,000 images, the rest being used for testing.

Types of models. Targets of various architectures were trained in order to see if a model’s complexity impact the performance of the GAMIN attack. Overall, five models were considered as targets, which are named by order of increasing complexity: one *softmax regression* (*Soft*), two *multilayer perceptron networks* (*MLP1* and *MLP2*), and two *convolutional neural networks* (*CNN1* and *CNN2*). The architectures of these models are detailed in the appendix. These architectures were chosen because of their prevalence in privacy and security-oriented literature. More precisely, *Soft* and *MLP1* models were attacked in (Fredrikson, Jha, and Ristenpart 2015) and (Tramèr et al. 2016), while *MLP2*, *CNN1* and *CNN2* were targeted in (Papernot et al. 2017).

Training setup. The surrogate has been trained with an Adam optimizer (Kingma and Ba 2015) regarding to this adaptive loss with the following parameters for the optimizer: learning rate of 10^{-5} , $\beta_1 = 0.99$, $\beta_2 = 0.99$, $\epsilon = 10^{-8}$ and for the adaptive loss: $\lambda_k = 0.01$, $\gamma_k = 0.5$

¹<https://github.com/definitively-not-a-lab-rat/gamin>

and $k_0 = 0.001$. The generator has also been trained with an Adam optimizer configured similarly to the surrogate’s with the following parameters for the optimizer: learning rate of 10^{-5} , $\beta_1 = 0.99$, $\beta_2 = 0.99$ and $\epsilon = 10^{-8}$. These parameters are fairly standard and yield good results for the problems the GAMIN was confronted with. Remark that an extensive hyperparameter search is a non-trivial process in the black-box setting.

Attack against MNIST-trained targets

Evaluation survey. To measure how well GAMIN solves the model inversion problem from a privacy point of view, we cannot rely directly on pixel-wise errors or distances between distributions. Indeed, since the objective is to extract sensitive information from a model, we have to rely on human interpretation of the results to evaluate whether they successfully disclose information about the training data or not. In order to measure this information leakage on MNIST, we set up a survey in which we showed to 13 participants the results of the model inversion of the GAMIN and asked them to guess which specific digit they thought was represented.

Results. The results of the survey are summarized in Figure 2. In this figure, *average* displays the percentage of correct inference of digits averaged across all 10 digits while *majority* depicts the percentage of classes (*i.e.*, digits) for which more than half the volunteers guessed the correct digit from the pictures shown. In our experiments, a *correct* guess is awarded only when the participant picked the right digit. In opposition, giving a wrong answer or selecting the option “unable to decide” answers were counted as negative results.




Figure 2: Results of the survey for MNIST-based models attack. For each target model, the average (over all the target digits) of the proportion of correct identification is given as well as the proportion of target digits with at least 50% of correct identification.

In the best case, when attacking the soft model, volunteers were able to pick the attacked digit in 55% of cases, assessing the reconstruction of 6 out of 10 digits. However, the ability of the GAMIN to invert models seems to be dramatically affected by the complexity of the target model, with both average number of correct answers and number of classes successfully inverted dropping to respectively 20% and 1 out of 10 for the CNN2 architecture.

Convolutional models, which were not targeted by model inversion attacks in previous works, proved to be more resistant to these attacks. In particular, the drop in reconstruction accuracy that occurs when attacking deeper models could be explained by their complexity, as the link between decision boundaries and classification becomes less and less direct.




Figure 3: Reconstruction of the same digit (*i.e.*, eight) issued from attacking different models. Raw outputs samples are located on the top while postprocessed outputs can be found on the bottom).

By taking a look at the GAMIN output for the same digit, we can observe the captured decision boundaries (see Figure 3), which confirm that the reconstruction given when attacking the soft architecture has usually sharp and identifiable features that are directly linked to the value of a region, whereas models using convolutions (*i.e.*, CNN1 and CNN2) use patterns and filters before selecting the class resulting in blurred and more abstract images.




Figure 4: Results of the survey for MNIST-based models attack. For each target model and for each target digit, the proportion of correct identification is given.

A closer look at the results breakdown between digits (displayed on Figure 4) shows a strong disparity among digits. Interestingly, some digits (*e.g.*, 2 and 3) are consistently inverted with success, whereas others such as for instance 1,

display the opposite behaviour as the GAMIN attack fails to reconstruct these most of the time. A plausible explanation of the difference in accuracy when attacking those different classes could be the amount of recognizable features each of these class have. While this might considered a subjective notion, recall that the mean of the GAMIN outputs tends to converge to the mean of the targeted class.

Related work

Prior works related to privacy attacks against discriminative models in the supervised learning setting include membership inference attacks (Shokri et al. 2017; Melis et al. 2019), property inference attacks (Ateniese et al. 2015; Ganju et al. 2018), model inversion attacks (Fredrikson, Jha, and Ristenpart 2015; Hitaj, Ateniese, and Perez-Cruz 2017) and model extraction attacks (Tramèr et al. 2016; Wang and Gong 2018; Milli et al. 2019). In addition to regularization techniques, differential privacy is the main approach used to address most of the inference attacks against machine learning models.

Inference attacks against machine learning models

Membership attacks against machine learning models have been introduced by Shokri, Stronati, Song and Shmatikov (Shokri et al. 2017). Given a data record d and a trained model M trained over a training dataset D_{train}^M , a membership inference attack consist in trying to evaluate if $d \in D_{train}^M$. For instance, the authors demonstrated in 2017 the possibility for an adversary to assess the presence of a given individual in hospital datasets in a true black-box setting, highlighting the potential privacy damage this type of attack can cause. This type of attack exploits the fact that machine learning models may be subject to overfitting. Membership attacks have also been studied by Melis, Song, de Cristofaro and Shmatikov (Melis et al. 2019) in the context of collaborative learning, in which the authors showed that the interactive nature of the collaboration can be exploited by a participant to conduct a membership attack on other participants' training sets. In addition, Hayes, Melis, Danezis and de Cristofaro have demonstrated in the context of generative models (Hayes et al. 2019) that GANs (Goodfellow et al. 2014) can be used to infer the presence of a particular individual in the training set.

Property inference attacks against machine learning models have been introduced by Ateniese and co-authors (Ateniese et al. 2015). This type of attack involves training a meta-classifier to detect if the target model has a given property P . The authors have used this attack to learn that the training set of a speech recognition system have been produced by people speaking a particular dialect. Remark that the shadow training technique developed here is the same that has inspired the membership attacks (Shokri et al. 2017). More recently, property inference attacks have been studied (Melis et al. 2019) in the context of collaborative learning, in which the authors demonstrated that the collaborative gradient update that occurs in this type of learning can be exploited by the adversary to infer properties that are true for a subset of the training set of other participants. Ganju,

Wang, Yang, Gunther and Borisov (Ganju et al. 2018) have further demonstrated the effectiveness of property inference attacks on more complex machine learning models such as fully connected neural networks.

Model inversion attacks aim at inferring, given a target model f and an output class y , sensitive hidden features of inputs corresponding to the class y . As a result, the adversary will learn the average of the features of inputs that belong to the class y . Model inversion attacks have been used originally (Fredrikson, Jha, and Ristenpart 2015) to infer if participants to a survey admitted of having cheated on their significant other by inverting decision trees, and to reconstruct people's faces by inverting a facial recognition system trained using three different classifiers, namely a softmax classifier, a multilayer perceptron network and a denoising autoencoder. Recently, Hitaj, Ateniese and Perez-Cruz (Hitaj, Ateniese, and Perez-Cruz 2017) have demonstrated that model inversion can be more powerful in collaborative machine learning context in which the adversary can train a generative adversarial network during the update phase to create prototypical examples of its target's training set. Hidano, Murakami, Katsumata, Kiyomoto and Hanaoka (Hidano et al. 2018) have introduced a model inversion attack that does not leverage on the knowledge of non-sensitive attributes of the input, under the assumption that the target model is operating in the online setting. In particular, a poisoning attack (Biggio et al. 2013) is conducted on the target model to turn the regression coefficients of the non-sensitive attributes to zero. The resulting target model then becomes easy to invert, using the same approach proposed in (Fredrikson et al. 2014; Fredrikson, Jha, and Ristenpart 2015), without requiring additional knowledge of the non-sensitive attributes.

Finally, model extraction attacks aim at inferring, given a target model f and its predictions (or explanations) for a chosen set of inputs, the parameters and/or hyper-parameters of f . Tramèr and co-authors (Tramèr et al. 2016) have demonstrated the effectiveness of such attacks by reconstructing several machine learning models after querying online machine learning as a service platforms. Wang and Cong (Wang and Gong 2018) have also proposed attacks to steal hyper-parameters of several machine learning models. More recently, Milli, Schmidt, Dragan and Hardt (Milli et al. 2019) have shown how attacks exploiting gradient-based explanations can allow to extract models with significant less query budget compared to attacks based on predictions. Batina, Bashin, Jap and Picek (Batina et al. 2018) have proposed an attack against neural networks that leverages on a side-channel attack. In particular, it monitors the power use of the microprocessor on which the target model is evaluated to extract its weights. Jagielski, Carlini, Berthelot, Kurakin and Papernot (Jagielski et al. 2019) have leveraged on Mix-Match (Berthelot et al. 2019), a semi-supervised learning technique, to produce a query-efficient learning-based extraction attack, which trains a surrogate on inputs labelled by the target model. The authors also introduced a functionally-equivalent extraction attack, which produces a surrogate that belongs to the equivalence class of the target model, displaying a better fidelity. Finally, they proposed a hybrid approach

that combines functionally-equivalent extraction attacks and learning-based extraction attacks to improve the fidelity of the surrogate model. Remark that while model extraction is not *per se* a privacy attack, it can be used to build white-box surrogates of the target model, thus potentially improving the efficiency of other inference attacks.

The main difference between GAMIN and the existing attacks is that it operates in a black-box agnostic setting, which allows the attack to be easily deployed against any machine learning model. In addition, the computational cost of the attack (about 6 hours on a MLP) is acceptable. To compare, the estimated cost of a black-box attack — based on numeric gradient approximation (Fredrikson, Jha, and Ristenpart 2015) — against a MLP is about 50 to 80 days.

Countermeasures against inference attacks

Adding noise to individual records while maintaining global distribution and correlations is at the core of the concept of differential privacy, a privacy model developed by Dwork and collaborators (Dwork, Roth, and others 2014) to balance the risk of privacy loss and the degradation of accuracy of the results of an analysis run on a dataset of private records. Recently, the application of differential private methods to machine learning has drawn a lot of attention and stemmed multiple approaches for privacy protection. One of these approaches relied on the training of deep networks over data distributed among multiple users (Shokri and Shmatikov 2015). Afterwards, new attacks were designed to address this specific setting (Hitaj, Ateniese, and Perez-Cruz 2017; Melis et al. 2019). Differential privacy was also integrated in the design of a stochastic gradient descent algorithm (Abadi et al. 2016). However, the inherent trade-off between the privacy level and the performance of models trained with this method has already been pointed out in previous works (Rahman et al. 2018; Jayaraman and Evans 2019).

To measure the resistance of a model against membership inference attacks, Long, Bindschaedler and Gunter introduced the concept of differential training privacy (Long, Bindschaedler, and Gunter 2017) and showed that countermeasures such as *distillation* may not be sufficient to counter this type of attack. In their study of model inversion attacks (Fredrikson, Jha, and Ristenpart 2015), Fredrikson, Jha and Ristenpart suggest to round confidence scores to reduce the vulnerability of models to equation-solving attacks in the black-box setting. However, this can decrease the precision on the gradient retrieved from target model as intended. For instance against a black-box GAMIN attack, this would reduce the capacity of the surrogate to correctly approximate the target’s decision boundaries. While this may result in a lesser quality of the reconstruction, the efficiency of this defense could be also differ greatly as boundaries retrieved by the surrogate could appear less strict, which could ultimately facilitate the training of the generator.

Conclusion and future works

In this paper, we have demonstrated how adversarial approaches can be adapted to achieve model inversion attacks

against machine learning models in an agnostic black-box setting. More precisely, we have shown the potential threat of privacy leakage posed by these attacks by testing our architecture against models trained on a dataset of handwritten digits and another of face portraits for which we were able to reconstruct recognizable features. In addition our attack works in multiple scenarios, including the situation in which the adversary has absolutely no prior knowledge of the target models and without specific parameter optimization.

We are planning to extend our approach in several directions. We believe that the GAMIN can still be improved to further improve accuracy or to decrease the cost of this attack in terms of queries. GAMIN could benefit from being tested against targets trained on different types of data, either as inputs (*i.e.*, sequential data for instance) or as outputs (*i.e.*, regressors). Further development of this architecture could allow for instance a single GAMIN to invert multiple outputs.

References

- Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318. ACM.
- Abadi, M.; Erlingsson, U.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Papernot, N.; Talwar, K.; and Zhang, L. 2017. On the protection of private information in machine learning systems: Two recent approaches. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, 1–6. IEEE.
- Ateniese, G.; Mancini, L. V.; Spognardi, A.; Villani, A.; Vitali, D.; and Felici, G. 2015. Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers. *International Journal of Security* 10(3):137–150.
- Barreno, M.; Nelson, B.; Sears, R.; Joseph, A. D.; and Tygar, J. D. 2006. Can machine learning be secure? In *Proceedings of the 2006 ACM Symposium on Information, computer and communications security*, 16–25. ACM.
- Batina, L.; Bhasin, S.; Jap, D.; and Picek, S. 2018. Csi neural network: Using side-channels to recover your artificial neural network information. *arXiv preprint arXiv:1810.09076*.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. 2019. Mixmatch: A holistic approach to semi-supervised learning. *arXiv preprint arXiv:1905.02249*.
- Berthelot, D.; Schumm, T.; and Metz, L. 2017. Began: Boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717*.
- Biggio, B.; Corona, I.; Maiorca, D.; Nelson, B.; Šrndić, N.; Laskov, P.; Giacinto, G.; and Roli, F. 2013. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*, 387–402. Springer.
- Buolamwini, J., and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classifi-

- cation. In *Conference on fairness, accountability and transparency*, 77–91.
- Dwork, C.; Roth, A.; et al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* 9(3–4):211–407.
- Fredrikson, M.; Lantz, E.; Jha, S.; Lin, S.; Page, D.; and Ristenpart, T. 2014. Privacy in pharmacogenetics: An end-to-end case study of personalized warfarin dosing. In *23rd {USENIX} Security Symposium ({USENIX} Security 14)*, 17–32.
- Fredrikson, M.; Jha, S.; and Ristenpart, T. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 1322–1333. ACM.
- French, R. M. 1999. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences* 3(4):128–135.
- Ganju, K.; Wang, Q.; Yang, W.; Gunter, C. A.; and Borisov, N. 2018. 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*, 619–633. ACM.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Hayes, J.; Melis, L.; Danezis, G.; and De Cristofaro, E. 2019. Logan: Membership inference attacks against generative models. *Proceedings on Privacy Enhancing Technologies* 2019(1):133–152.
- Hidano, S.; Murakami, T.; Katsumata, S.; Kiyomoto, S.; and Hanaoka, G. 2018. Model inversion attacks for online prediction systems: Without knowledge of non-sensitive attributes. *IEICE Transactions on Information and Systems* 101(11):2665–2676.
- Hitaj, B.; Ateniese, G.; and Perez-Cruz, F. 2017. Deep models under the gan: information leakage from collaborative deep learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 603–618. ACM.
- Jagielski, M.; Carlini, N.; Berthelot, D.; Kurakin, A.; and Papernot, N. 2019. High-fidelity extraction of neural network models. *arXiv preprint arXiv:1909.01838*.
- Jayaraman, B., and Evans, D. 2019. Evaluating differentially private machine learning in practice. In *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association.
- Kasiviswanathan, S. P.; Rudelson, M.; and Smith, A. 2013. The power of linear reconstruction attacks. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, 1415–1433. SIAM.
- Kingma, D. P., and Ba, J. 2015. Adam: A method for stochastic optimization. In *ICLR (Poster)*.
- Long, Y.; Bindschaedler, V.; and Gunter, C. A. 2017. Towards measuring membership privacy. *arXiv preprint arXiv:1712.09136*.
- Melis, L.; Song, C.; De Cristofaro, E.; and Shmatikov, V. 2019. Exploiting unintended feature leakage in collaborative learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, 691–706. IEEE.
- Milli, S.; Schmidt, L.; Dragan, A. D.; and Hardt, M. 2019. Model reconstruction from model explanations. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 1–9. ACM.
- Papernot, N.; McDaniel, P.; Sinha, A.; and Wellman, M. 2016. Towards the science of security and privacy in machine learning. *arXiv preprint arXiv:1611.03814*.
- Papernot, N.; McDaniel, P.; Goodfellow, I.; Jha, S.; Celik, Z. B.; and Swami, A. 2017. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, 506–519. ACM.
- Rahman, M. A.; Rahman, T.; Laganière, R.; Mohammed, N.; and Wang, Y. 2018. Membership inference attack against differentially private deep learning model. *Transactions on Data Privacy* 11(1):61–79.
- Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. In *Advances in neural information processing systems*, 2234–2242.
- Samaria, F. S., and Harter, A. C. 1994. Parameterisation of a stochastic model for human face identification. In *Proceedings of 1994 IEEE Workshop on Applications of Computer Vision*, 138–142. IEEE.
- Shokri, R., and Shmatikov, V. 2015. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, 1310–1321. ACM.
- Shokri, R.; Stronati, M.; Song, C.; and Shmatikov, V. 2017. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, 3–18. IEEE.
- Song, S.; Chaudhuri, K.; and Sarwate, A. D. 2013. Stochastic gradient descent with differentially private updates. In *2013 IEEE Global Conference on Signal and Information Processing*, 245–248. IEEE.
- Tramèr, F.; Zhang, F.; Juels, A.; Reiter, M. K.; and Ristenpart, T. 2016. Stealing machine learning models via prediction apis. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*, 601–618.
- Wang, B., and Gong, N. Z. 2018. Stealing hyperparameters in machine learning. In *2018 IEEE Symposium on Security and Privacy (SP)*, 36–52. IEEE.
- Wu, X.; Fredrikson, M.; Jha, S.; and Naughton, J. F. 2016. A methodology for formalizing model-inversion attacks. In *2016 IEEE 29th Computer Security Foundations Symposium (CSF)*, 355–370. IEEE.
- Yeom, S.; Giacomelli, I.; Fredrikson, M.; and Jha, S. 2018. Privacy risk in machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, 268–282. IEEE.

Post-processing

Early experiments lead us to observe that when we attack datasets in which diverse inputs correspond to the same output (*i.e.*, the same class), the individual outputs of the generator do not converge to one of these specific inputs but rather that *the mean of those generated inputs converged to the mean of the original inputs*. Despite this limitation, we are nonetheless able to exploit the results of the GAMIN training by aggregating its results using the following procedure.

Signal post-processing pipeline. Once the generator has been trained, the following process is used to retrieve the best version of the attacked class’ mean. This process is applied separately for each channel when dealing with color images, each channel containing one color distribution.

Batch generation of outputs. The first step consists in using our model to generate a large batch of outputs from noise. Our experiments have shown that 1000 samples are sufficient to achieve a satisfactory sample of the variation between individual pictures.

Low-pass common frequency filter. Afterwards, each picture is transposed into the frequency domain via a Fast-Fourier Transform (FTT). The frequencies are then counted pixel-wise on all FTT images. In addition, a filter is applied to remove frequencies that do not appear in at least 90% of the FTT batch samples. Due to the frequency distribution on our experiment, this is equivalent to a low-pass filter.

Reconstruction. The samples are mapped back to the spatial domain before applying a Gaussian blur to each of them, followed by an edge detection filter. Finally, the “final” output is obtained from the pixel-wise median of the transformed batch.

Metrics and relation to success of the attack

The average metrics defined previously for each target model trained on MNIST are displayed in Table 1. For each GAMIN attack, the average over all 10 digits is denoted as F_S for the fidelity of the surrogate model, $A_{S \circ G}$ for the categorical accuracy of the combined model and M_{global} for the best global score of convergence. The ratio of correct re-identification by the respondents of the survey is also provided for comparison.



Figure 5: Averages for digits 1 and 3 on the MNIST dataset.

First, we can observe that there is an absence of correlations between the accuracy $A_{S \circ G}$ of the combined model and the success rate of the re-identification by human volunteers. For instance, by comparing the mean of a digit with frequent success (the 3) to a digit with rare success (such as the 1) displayed on Figure 5, we observe that the classification task of a representation of an image’s mean might not be sufficient to evaluate completely the success of the

model inversion attack. Nonetheless, the relatively high fidelity scores yielded by most attacks demonstrate that the surrogate is able to capture, at least partially, the decision boundaries of the target model.

Target model	F_S	$A_{S \circ G}$	M_{global}	%correct
Soft	0.91	0.98	0.06	0.553
MLP1	0.73	1.0	0.08	0.492
MLP2	0.71	1.0	0.18	0.284
CNN1	0.83	0.74	0.19	0.269
CNN2	0.56	1.0	0.16	0.207

Table 1: Average performances across all digits of inversion attacks against the target models. F_S , $A_{S \circ G}$, M_{global} and %correct refer resp. to the fidelity of the surrogate, the categorical accuracy of the combined model, the global convergence score and the accuracy as evaluated by the survey.

To further investigate the link between F_S and M_{global} as well as the attack success, we provide an extensive review of the final score when attacking each digit against the MLP1 architecture in Table 2. These results show that a surrogate model achieving a very high fidelity or very low global convergence score might not be sufficient to achieve a good model inversion. In particular, this table provides examples with digits 1 and 9 of attacks that failed to reconstruct recognizable digits, yet achieve M_{global} scores among the lowest. Furthermore, attacks against digit 1 also achieved fidelity scores comparable to that of any successful attack.

Target digit	F_S	M_{global}	%correct
0	0.83	0.06	0.31
1	0.75	0.07	0.15
2	0.78	0.07	0.92
3	0.74	0.09	1.00
4	0.73	0.10	0.23
5	0.75	0.11	0.69
6	0.77	0.07	0.23
7	0.86	0.09	0.62
8	0.77	0.08	0.69
9	0.36	0.08	0.08

Table 2: Per-digit performances of inversion attacks against the MLP1 model. F_S , M_{global} , and %correct are respectively the fidelity of the surrogate, the global convergence score, and the accuracy as evaluated with the survey

The results demonstrate that the GAMIN surrogate needs to achieve a high fidelity F_S to succeed, which is consistent with using fidelity as a measure of how well the decision boundaries are captured. However, this is a necessary but not a sufficient condition to ensure the success of the attack.

Cost of the attack

The budget allowed for attacking each model was 1,280,000 queries (or equivalently 20,000 batches of 64 queries). However, the attack protocol enables to easily re-

trieve the required number of batches to achieve convergence in terms of the M-global score. Depending on the attacked digits, the GAMIN achieves convergence towards “final” results after 2,000 to 20,000 batches of training, with an average of 2,500 batches. In particular, each epoch requires 128 queries to the target model, giving an average of 320,000 queries per attack.

Compared to other black-box attacks targeting even larger datasets ($\sim 41,000$ online queries in (Tramèr et al. 2016), and 20,600 in (Fredrikson, Jha, and Ristenpart 2015) for attacking a model trained on the AT&T Faces dataset of larger dimension (Samaria and Harter 1994)), the GAMIN is significantly more expensive in terms of queries. This increased cost could be attributed to the higher complexity of target models and is balanced by the flexibility of GAMIN.

Despite the higher complexity in terms of queries, if we estimate that it takes 70 ms in terms of latency for querying the target, as in (Fredrikson, Jha, and Ristenpart 2015; Papernot et al. 2017) and not accounting for the speed increase offered by batch queries, a *GAMIN attack could be completed in six hours*. In addition assuming the adversary benefits from GPU acceleration, the networks training takes negligible time compared to querying the oracle.

Attacks against models trained with differential privacy

To evaluate how well our attack fare against protection mechanisms, we evaluate the effect of differential privacy on the success on an inversion attack against a MLP model. For this experiment, we focus on the MLP1 architecture described previously. To train models with this architecture with differential privacy, we use the noisy batched stochastic gradient descent algorithm (Song, Chaudhuri, and Sarwate 2013), with the implementation provided in *tensorflow/privacy*². In particular, we set the privacy parameter δ to $1e^{-5}$ and train a model achieving a test accuracy of 0.89 for a privacy parameter $\epsilon = 2.9$.

Target model	F_S	$A_{S \circ G}$	M_{global}
MLP1 without DP	0.73	1.0	0.08
MLP1 with DP	0.74	0.96	0.16

Table 3: Average performances across all digits after model inversion attacks against MLP models trained with (and without) differential privacy. F_S , $A_{S \circ G}$ and M_{global} are resp. the fidelity of the surrogate, the categorical accuracy of the combined model and the global convergence score.

Table 3 and Figure 6 summarize the results obtained. These results seem to indicate that the use of differential privacy does not prevent the convergence of the training of GAMIN and the success of the attack. This observation is consistent with findings from previous works that have analyzed the impact of differential privacy on inversion attacks (Hitaj, Ateniese, and Perez-Cruz 2017; Abadi et al. 2017; Jayaraman and Evans 2019). Intuitively this can be

²<https://github.com/tensorflow/privacy>




Figure 6: Results of the survey for attacks against MLP1 model with (respectively without) differential privacy. For each target model and for each target digit, the proportion of correct identification is given.

explained by the fact that differentially private learning techniques are essentially designed to protect against membership attack. In contrast protecting against model inversion attacks is similar in spirit as protecting the membership of all items that belong to a particular class, which require to apply some notion of group privacy. However, by doing so, the utility of the model will be negatively impacted.

Additional experiments

Attacks against skin color classifiers

Pilot parliament dataset. Compiled by Joy Buolamwini (Buolamwini and Gebru 2018) from official portraits of parliament representatives across the globe, this dataset of 1270 color images of different sizes is classified according to sex and skin color (across the 6 Fitzpatrick skin categories), which are inherently sensitive attributes. In contrast to MNIST, this dataset is composed of colored pictures. Note that for our experiment we have reshape all images to the same size, which implies that some faces are distorted compared to the original versions. The target models were trained on 1079 portraits, the remainder being kept for the testing phase. For this experiment, we test our attack against a Softmax and a CNN1 skin color classifiers.

Results. Subtle, yet recognizable facial shapes can be observed on the images generated after attacking a skin color classifier trained on the pilot parliament dataset (see Figures 8 and 9). Interestingly enough, even the central area of the image does not give a hint of the skin color of individuals. Instead, the decision boundaries captured by the surrogate model are based on contours. While the outputs of the GAMIN attack against the CNN1 architecture remain blurry and noisy enough to disclose few features of the members of each class, the attack against the soft architecture gives away more precise contours. However, the attack against the Softmax skin classifier yields results which, unlike the ones obtained on MNIST classifier, did not seem to converge to the attacked class mean (e.g, compare Figure 9 to Figure 10).

In this case, it could be argue that even if the faces are undetermined, the clothes alone might be enough to constitute a privacy breach. The privacy leak depicted here is more



Figure 7: Pilot parliament dataset samples.

similar to a property inference attack (Ateniese et al. 2015; Ganju et al. 2018), which aim at inferring global properties of the training set of the target model. In contrast to previous implementations of such attacks (Ateniese et al. 2015; Ganju et al. 2018) that work in the white-box setting, GAMIN aims to infer properties of the target model’s training set in the black-box agnostic setting.




Figure 8: Outputs of attack against a CNN1 skin color classifier.

While Fredrikson and co-authors have already obtained recognizable faces against a model of similar architecture in the white-box setting (Fredrikson, Jha, and Ristenpart 2015) and Tramèr and colleagues proposed a more efficient black-box attack based on successive extraction and inversion of the model (Tramèr et al. 2016), the GAMIN provides another black-box alternative that does not require the adversary to have information about the target’s architecture.



Figure 9: Outputs of attack against a soft skin color classifier.




Figure 10: Mean of original attacked class.

Gallery of models

This appendix lists the models used in this study, both as parts of the GAMIN architecture and as targets models.

Image classification

IN	OUT	CM	CM	RL	D	RL	D	S
784	10	32	64	128	0.5	32	0.5	10

Table 4: GAMIN surrogate for MNIST image classification model attacks. IN denotes the dimension of inputs, OUT the dimension of outputs, CM is a 2D convolution of kernel size (3x3) with ReLU activation followed by a max pooling over (2x2), RL is a Rectified Linear Unit fully-connected layer, D is the proportion of dropout between these layers and S a softmax (output) layer.

GAMIN surrogate model. Table 4 describes the model used to perform the GAMIN attack against models trained over the MNIST dataset. Note that since the convolutions used in this model have a different filter size than those used for the targets, the surrogate has very little in common with the target beyond the dimensions of its inputs and outputs.

Target models. Table 5 describes all models used to classify MNIST digits.

Name	IN	OUT	CM	CM	RL	RL	S
Soft	784	10					10
MLP1	784	10			100		10
MLP2	784	10			200	200	10
CNN1	784	10	32		200		10
CNN2	784	10	32	64	200		10

Table 5: Target models for MNIST image classification. IN denotes the dimension of inputs, OUT the dimension of outputs, CM is a 2D convolution of kernel size (2x2) with ReLU activation followed by a max pooling over (2x2), RL us a Rectified Linear Unit fully-connected layer and S is a soft-max (output) layer.

GAMIN generator mode.

IN	OUT	L	CM	D	CM	D	CM	TAN
10	784	784	128	0.5	128	0.5	64	784

Table 6: GAMIN Generator for MNIST image classification model attacks. IN denotes the dimension of inputs, OUT the dimension of outputs, CM is a 2D convolution of kernel size (3x3) with ReLU activation followed by a max pooling over (2x2), D is the proportion of dropout between these layers and TAN is a fully-connected layer with \tanh activation (output layer).

Table 6 displays the details of the GAMIN Generator architecture for reconstructing handwritten digits.

Target performance breakdown

All the accuracies of target models are listed in Tables 7 and 8 for models trained respectively, on the MNIST and on the pilot parliament datasets. We denote by A_{train} & A_{test} the accuracy respectively over the training and testing datasets.

MNIST target digit	A_{train}	A_{test}
Soft	0.85	0.83
MLP1	0.93	0.91
MLP2	0.97	0.94
CNN1	0.99	0.99
CNN2	0.98	0.99
Random baseline	0.1	0.1

Table 7: Target accuracy on MNIST Dataset

Pilot Parliament: skin color	A_{train}	A_{test}
Soft	0.43	0.33
CNN1	1.0	0.99
CNN2	1.0	1.0
Random baseline	0.17	0.17

Table 8: Target accuracy on pilot parliament dataset.