



# Machine unlearning: linear filtration for logit-based classifiers

Thomas Baumhauer<sup>1</sup> · Pascal Schöttle<sup>2</sup> · Matthias Zeppelzauer<sup>1</sup>

Received: 14 September 2020 / Revised: 13 September 2021 / Accepted: 11 April 2022 /  
Published online: 11 July 2022  
© The Author(s) 2022

## Abstract

Recently enacted legislation grants individuals certain rights to decide in what fashion their personal data may be used and in particular a “right to be forgotten”. This poses a challenge to machine learning: how to proceed when an individual retracts permission to use data which has been part of the training process of a model? From this question emerges the field of *machine unlearning*, which could be broadly described as the investigation of how to “delete training data from models”. Our work complements this direction of research for the specific setting of class-wide deletion requests for classification models (e.g. deep neural networks). As a first step, we propose *linear filtration* as an intuitive, computationally efficient sanitization method. Our experiments demonstrate benefits in an adversarial setting over naive deletion schemes.

**Keywords** Machine learning · Machine unlearning · Privacy

## 1 Introduction

Recently enacted legislation, such as the European Union’s General Data Protection Regulation (GDPR) (Council of European Union 2016), previously its “right to be forgotten” (Council of European Union 2012), and the California Consumer Privacy Act (State of California 2018) grant individuals certain rights to decide in what fashion their personal data may be used, and in particular the right to ask for personal data collected about them to be deleted.

---

Editor: Tijl De Bie.

---

✉ Pascal Schöttle  
pascal.schoettle@mci.edu  
Thomas Baumhauer  
thomas.baumhauer@fhstp.ac.at  
Matthias Zeppelzauer  
matthias.zeppelzauer@fhstp.ac.at

<sup>1</sup> St. Pölten University of Applied Sciences, St. Pölten, Austria

<sup>2</sup> Management Center Innsbruck, Innsbruck, Austria

At present the implementation of such rights in the context of machine learning models trained on personal data is largely an open problem. In Villaronga et al. (2018) the authors even conclude that “it may be impossible to fulfill the legal aims of the Right to be Forgotten in artificial intelligence environments”.

Indeed, machine learning models may unintentionally memorize (part of) their training data, leading to privacy issues in many applications, e.g. image classification (Yeom et al. 2018; Sablayrolles et al. 2019) or natural language processing (Carlini et al. 2019), and potentially enabling an adversary to extract sensitive information from a trained model by so-called *model inversion* (Veale et al. 2018).

Informally, deletion of part of the training data from a machine learning model can be understood as removal of its influence on the model parameters, in order to obtain a model that “looks as if it has never seen that part of the data”. We refer to this process as *unlearning*. Clearly, the problem of unlearning can be solved in a trivial fashion, by simply retraining the model without using the part of the data we wish to unlearn. For large, real-world models, retraining from scratch may incur significant computational costs, and may thus be practically infeasible, if deletion requests are expected to arrive at any time. We therefore wish to find more efficient unlearning algorithms, which is notoriously difficult, owing to the fact that for many popular learning algorithms every training data point potentially affects every model parameter.


First approaches towards directed unlearning were introduced in Cao and Yang (2015), and more recently in Ginart et al. (2019); Bourtole et al. (2019). In our work we consider the problem of unlearning in the setting of classification models for which, in contrast to previous work, we assume that single individuals own all training data associated with a particular class, as may be the case e.g. in biometric applications.

In this setting, we consider classifiers that predict logits, i.e. (rescaled) logarithmic probabilities that a data point belongs to certain classes. For such classifiers we propose a novel sanitization method that applies a linear transformation to these predictions. For an appropriate hypothesis class this transformation can be absorbed into the original classifier. The computation of the transformation requires barely more than computing predictions by the original classifier for a (small) number of data points per class. We call this method *linear filtration*. Figure 1 shows the results of this method when used as a defense against model inversion (Fredrikson et al. 2015).

In summary, the main contributions of our work are:

- (1) We develop *linear filtration*, a novel algorithm for the sanitization of classification models that predict logits, after class-wide deletion requests.
- (2) On the theoretical side, we add to the definition of unlearning in the sense of Ginart et al. (2019), by proposing a weakened, “black-box” variant of the definition, which may serve as a more realistic goal in practice.
- (3) As practical methodology, we suggest that the quality of an empirical unlearning operation may be evaluated in an adversarial setting, i.e. by testing how well it prevents certain privacy attacks on machine learning models.

The rest of this paper is organized as follows: Section 2 gives an overview of work related to machine unlearning, as well the adversarial methodology employed in our experiments. Section 3 formalizes the unlearning of training data from machine learning models. Section 4 establishes several variants of our method of linear filtration. Section 5



**Fig. 1** Results of a model inversion attack for a toy model trained on the AT&T Faces dataset with 4 classes. For one of the classes, from left to right: one of the training images, reconstruction of the class by model inversion, reconstruction after naive unlearning, reconstruction after unlearning by our proposed method of *normalizing linear filtration* (defined in Sect. 4). The reconstructions of the other classes remain visually unchanged by normalizing linear filtration, see Fig. 5

experimentally evaluates linear filtration, primarily in an adversarial context. Section 6 features discussion on the definition of unlearning.

## 2 Related work

### 2.1 Machine unlearning

The term *machine unlearning* first appears in Cao and Yang (2015). There the authors consider unlearning the framework of statistical query learning (Kearns 1998). This allows them to unlearn data points for learning algorithms where all queries to the data are decided upfront. However, many popular learning algorithms (such as gradient descent) query data adaptively. In the adaptive setting the approach of (Cao and Yang 2015) does not give any bounds and quickly falls apart if the number of queries is large, as is the case for neural networks.

Ginart et al. (2019) features a discussion of the problem of efficient unlearning of training data points from models, establishes several engineering principles, and on the practical side proposes unlearning algorithms for  $k$ -means clustering. In particular, they recognize that given the stochastic nature of many learning algorithms a probabilistic definition of unlearning (there “deletion”) is necessary. We adopt this view in our work.

Bourtoule et al. (2019) propose a framework they refer to as SISA (sharded, isolated, sliced, and aggregated training), which can be thought of as bookkeeping method seeking to limit and keep track of the influence of training data points on model parameters, thus reducing the amount of retraining necessary upon receiving a deletion request. This approach comes at the cost of a large storage overhead.

Guo et al. (2019) define  $\epsilon$ -certified removal, “a very strong theoretical guarantee that a model from which data is removed cannot be distinguished from a model that never observed the data to begin with”, a concept akin to that of differential privacy (Dwork et al. 2006). Combining this with the idea of influence functions (Koh and Liang 2017), they then develop a certified removal mechanism for linear classifiers.

Golatkar et al. (2020a, 2020b) adopt an information theoretic view of unlearning and develop unlearning operations based on linearized model dynamics (drawing inspiration from the *neural tangent kernel* (Jacot et al. 2018; Lee et al. 2019), a technique to describe the gradient descent dynamics of the training of neural networks using kernel methods).

The work done in Guo et al. (2019); Golatkar et al. (2020a, b) may be considered complementary to the method of *linear filtration* we are going to develop in this paper, in the sense that they use strong assumptions (in particular work with linear/linearized models) obtaining stronger guarantees, while linear filtration is an intuitive heuristic, largely agnostic to model architecture.

Sommer et al. (2020) propose a formal framework for verification of machine unlearning, based on machine learning backdoor attacks.

## 2.2 Membership inference

It is an open problem to find a suitable measure for the quality of unlearning, when employing a heuristic unlearning operation with no or weak theoretical guarantees, i.e. to quantify the remaining influence of “deleted” training data on a model’s parameters. In our experiments we thus take ideas from *membership inference*. The goal of membership inference is to determine whether a given data point has been used in the training process of a given model. A few recent works on membership inference include (Shokri et al. 2017; Yeom et al. 2018; Hayes et al. 2019; Sablayrolles et al. 2019). In particular, our adversarial setup in Sect. 5.1 draws a large amount of inspiration from Shokri et al. (2017), where a binary classifier is trained on the outputs of so-called *shadow models* to decide membership.

## 2.3 Model inversion

Broadly, *model inversion* may be defined as drawing inferences about private training data from the outputs of a model trained on this data. The term was introduced in Fredrikson et al. (2014). Fredrikson et al. (2015) reconstruct human-recognizable images of individuals from neural networks trained for face recognition, using gradient ascent on the input space. Recent approaches leverage generative adversarial networks (GANs) for model inversion (Zhang et al. 2020).

We remark that model inversion is at its core the result of a correlation between input and output space that is simply captured by the model (and may exist independently of the model), and thus does not necessarily constitute a privacy breach. McSherry (2016) features a highly recommended elaboration of this point in much detail.

Shokri et al. (2017) conclude their discussion of model inversion with the statement that “model inversion produces the average of the features that at best can characterize an entire output class.” Thus, model inversion is of some interest in the specific context of our paper, which focuses on class-wide unlearning (and hence implicitly makes the assumption that a single individual owns all data for an entire output class). We experiment with model inversion in Sect. 5.5, see also Fig. 1.

## 2.4 Differential privacy

Differential privacy (Dwork et al. 2006; Dwork and Roth 2014; Abadi et al. 2016) limits the influence of individual training points on a model in a precise probabilistic way. We briefly remark that (differential) privacy and data deletion may be considered orthogonal problems, in the sense that private models need not support efficient deletion, and models supporting efficient deletion need not be private. Ginart et al. (2019) discuss this point in additional detail.

### 3 Problem definition

In this section we formalize our notion of unlearning and the threat model we consider.

#### 3.1 Notation

For a vector  $v \in \mathbb{R}^k$  we use indices ranging from 0 to  $(k - 1)$  to denote its entries, i.e.  $v = (v_0, v_1, \dots, v_{k-1})^T$ . One-dimensional vector are always column vectors.

For a vector of logits  $\ell \in \mathbb{R}^k$  let

$$\sigma_{i^*}(\ell) = \frac{\exp(\ell_{i^*})}{\sum_{i < k} \exp(\ell_i)} \in [0, 1], \quad \text{for all } i^* < k$$

and  $\sigma(\ell) = (\sigma_0(\ell), \sigma_1(\ell), \dots, \sigma_{k-1}(\ell))$ . We call  $\sigma$  the *softmax function*.

We use uppercase, boldface letters for random variables. If  $\mathbf{R}$  is a random variable  $P(\mathbf{R})$  denotes its distribution.

#### 3.2 Classification

We consider a multiclass classification problem: Let  $\mathbf{X}$  be a random variable taking values in some input space  $\mathcal{X}$  (e.g.  $\mathcal{X} = \mathbb{R}^{28 \times 28}$ ), and let  $\mathbf{Y}$  be a random variable representing class labels taking values in  $\mathcal{Y} = \{0, 1, \dots, k - 1\}$  for some natural number  $k > 2$ , with some joint data generating distribution  $P(\mathbf{X}, \mathbf{Y})$ .

Given  $x \in \mathcal{X}$ , a classifier  $h : \mathcal{X} \rightarrow \mathbb{R}^k$  for this classification problem attempts to estimate logits  $\ell = h(x)$  such that

$$\sigma(\ell) \approx P(\mathbf{Y} \mid \mathbf{X} = x).$$

**Hypothesis class** In this paper we consider the class  $\mathcal{H}$  of all classifiers  $h$  of the form

$$h = \text{logistic regression} \circ \text{feature extraction},$$

i.e. classifiers that can be decomposed into a (potentially non-linear) feature extraction followed by a multinomial logistic regression.

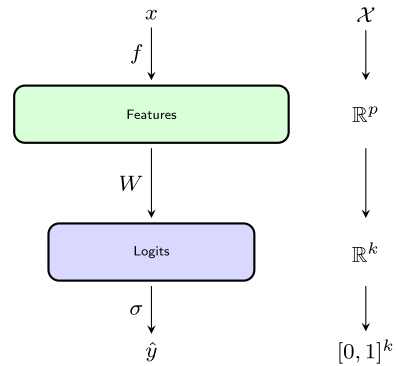
More formally any  $h \in \mathcal{H}$  can be expressed as

$$h : x \mapsto W \cdot f(x) \tag{1}$$

where  $f : \mathcal{X} \rightarrow \mathbb{R}^p$  denotes the feature extraction,  $p$  denotes the dimension of the feature space, and  $W$  is a  $(k \times p)$ -matrix representing a linear transformation  $\mathbb{R}^p \rightarrow \mathbb{R}^k$ . Figure 2 shows a schematic representation of the elements of  $\mathcal{H}$ , where  $\sigma$  denotes the softmax function. To simplify notation we do not consider affine transformations, i.e. classifiers of the form  $h : x \mapsto W \cdot f(x) + b$  for some  $b \in \mathbb{R}^k$ . However, we remark that our method is easily adapted to this case.

Observe that in particular all deep neural networks for which the output is a densely connected layer with softmax activations fit into the schema discussed in this section.

**Fig. 2** Schematic representation of  $\sigma \circ h$ , for a classifier  $h = W \circ f$  in the hypothesis class considered throughout this paper. Here  $f$  denotes a feature extraction,  $W$  is a linear transformation, and  $\sigma$  is the softmax function. In deep learning terminology “Logits” represents a fully connected layer with  $k$  units and weights  $W$



### 3.3 Learning algorithms

For a finite training set  $D \subseteq \mathcal{X} \times \mathcal{Y}$ , a learning algorithm  $\mathbf{A}$  calculates a classifier

$$h = \mathbf{A}(D) \in \mathcal{H}.$$

Note that if  $\mathbf{A}$  is non-deterministic  $\mathbf{A}(D)$  can be considered a random variable taking values in  $\mathcal{H}$ .

### 3.4 Unlearning of classes

Let  $\mathcal{C} \subsetneq \mathcal{Y}$  be a set of classes, which we want to unlearn. Consider the multiclass classification problem for the distribution  $P(\mathbf{X}, \mathbf{Y} \mid \mathbf{Y} \notin \mathcal{C})$ , i.e. the original problem with the classes  $\mathcal{C}$  removed. We define the hypothesis class  $\mathcal{H}_{-\mathcal{C}}$  for this problem similarly to  $\mathcal{H}$ , i.e.  $h \in \mathcal{H}_{-\mathcal{C}}$  is of form (1) with  $W$  a  $((k - |\mathcal{C}|) \times p)$ -Matrix. For  $D \subseteq \mathcal{X} \times \mathcal{Y}$  let  $D_{-\mathcal{C}} = \{(x, y) \in D : y \notin \mathcal{C}\}$  and let  $\mathbf{A}_{-\mathcal{C}}(D_{-\mathcal{C}}) \in \mathcal{H}_{-\mathcal{C}}$  denote a classifier calculated by some learning algorithm  $\mathbf{A}_{-\mathcal{C}}$ .

**Definition 1** (Unlearning) We say that a map

$$\mathfrak{D} : \mathcal{H} \rightarrow \mathcal{H}_{-\mathcal{C}}$$

“unlearns  $\mathcal{C}$  from  $\mathcal{H}$  with respect to  $\mathbf{A}, \mathbf{A}_{-\mathcal{C}}, D$ ” if the random variables  $\mathfrak{D}(\mathbf{A}(D))$  and  $\mathbf{A}_{-\mathcal{C}}(D_{-\mathcal{C}})$  have the same distribution over  $\mathcal{H}_{-\mathcal{C}}$ . We call  $\mathfrak{D}$  an *unlearning operation* (for  $\mathcal{C}$ ).

### 3.5 Weak unlearning of classes

A good choice of  $\mathfrak{D}$  will of course depend on the learning algorithm  $\mathbf{A}$ . We mostly concern ourselves with the case where  $\mathbf{A}$  trains a neural network with a densely connected output layer with softmax activations. Unfortunately it is difficult to understand how a neural network represents knowledge internally (e.g. (Achille et al. 2019)), hence unlearning as defined above may currently be out of reach. We therefore propose a weakening of the above definition.

**Definition 2** (Weak unlearning) As before let  $\mathfrak{D}$  be a map  $\mathcal{H} \rightarrow \mathcal{H}_C$  and for  $\mathbf{X}$  (taking values in the input space  $\mathcal{X}$  according to the data generating distribution) consider the random variables

$$\begin{aligned} \mathbf{L}_{\text{seen}} &= h_0(\mathbf{X}), \quad \text{where } h_0 = \mathfrak{D}(\mathbf{A}(D)) \\ \mathbf{L}_{\neg\text{seen}} &= h_1(\mathbf{X}), \quad \text{where } h_1 = \mathbf{A}_{-C}(D_{-C}) \end{aligned}$$

i.e. the logit outputs of the respective classifiers, taking values in  $\mathbb{R}^{k-|C|}$ . We say that “ $\mathfrak{D}$  weakly unlearns  $C$  from  $\mathcal{H}$  with respect to  $\mathbf{A}, \mathbf{A}_{-C}, D$ ” if  $\mathbf{L}_{\text{seen}}$  and  $\mathbf{L}_{\neg\text{seen}}$  have the same distribution over  $\mathbb{R}^{k-|C|}$ . We call  $\mathfrak{D}$  a *weak unlearning operation* (for  $C$ ).

**Fact 3** If  $\mathfrak{D}$  is an unlearning operation, then  $\mathfrak{D}$  is a weak unlearning operation. □

### 3.6 Discussion

Definition 1 demands that the distributions over the hypothesis class (i.e. in practical terms the parameter space) are the same, while 2 relaxes this to the distributions over the output space being the same. Intuitively speaking: 1 demands  $\mathfrak{D}$  make the model  $h$  “look as if  $h$  had never seen the data”, while 2 demands  $\mathfrak{D}$  make the **outputs** of  $h$  “look as if  $h$  had never seen the data” We may therefore consider unlearning (in the strong sense) to be a white-box variant and unlearning in the weak sense to be a black-box variant of the definition of unlearning. See section 6 for further discussion.

Abusing the terminology introduced in this section we refer to maps  $\mathfrak{D}$  that roughly fit definitions 1 and 2, respectively in the sense that they make the relevant distributions similar in an appropriate divergence measure, as “good”, “satisfactorily performing”, etc. unlearning operations.

#### 3.6.1 Threat model

The considerations about the white- and black-box setting also nicely fit into the common classification of adversaries against machine learning models, as used in adversarial machine learning (e.g., Biggio and Roli (2018); Papernot et al. (2016)). Here, it is common that a white-box adversary has full knowledge of all classifier properties, including training data and weights, while in the black-box scenario, an adversary can either only access the classifier’s final decision or has at most access to the logits. In line with this, considering a white-box adversary makes only sense when evaluating attacks against an unlearning operation in the strong sense (i. e., according to definition 1), while the appropriate mode to evaluate a weak unlearning operation (cf. definition 2) is to consider an adversary who can only access the logit output of the classifier. In general, the white-box setting adheres to Kerckhoffs’ principle (Kerckhoffs 1883) and should be adapted for rigorous theoretical security guarantees. The black-box setting on the other hand, is more akin to practice. The goal of an adversary in our setting is to distinguish between the distributions of the original and unlearned model, given that she has no previous outputs of samples of the deleted class.

## 4 Method

In this section we propose an intuitive weak unlearning operation  $\mathfrak{D}$  for classes, exploiting the special structure of elements of  $\mathcal{H}$ . In our experiments in section 5 we demonstrate that  $\mathfrak{D}$  performs satisfactorily for neural networks. We reiterate that the method introduced here is a heuristic, and not an unlearning operation in the strict sense of definition 2.

### 4.1 Definition of weak unlearning operation $\mathfrak{D}_z$

For simplicity of notation let  $\mathcal{C} = \{0\}$ , i.e. we are going to unlearn class 0. However our method easily generalizes to arbitrary  $\mathcal{C}$ . Let  $h = \text{Wof} \in \mathcal{H}$  be a classifier. For  $j < k$  let

$$a_j = \mathbb{E}[h(\mathbf{X}) \mid \mathbf{Y} = j] \in \mathbb{R}^k$$

be the expected prediction for class  $j$ , and let

$$A = (a_0 \mid a_1 \mid \dots \mid a_{k-1}) \in \mathbb{R}^{k \times k}.$$

In practice, we may estimate  $a_j$  from the training data. Next, define a map  $\pi$  such that for  $v = (v_0, v_1, \dots, v_{k-1})^\top \in \mathbb{R}^k$  we have  $\pi(v) = (v_1, v_2, \dots, v_{k-1})^\top$ . For arbitrary  $z \in \mathbb{R}^{(k-1) \times k}$  let

$$B_z = (z \mid \pi(a_1) \mid \pi(a_2) \mid \dots \mid \pi(a_{k-1})) \in \mathbb{R}^{(k-1) \times k}.$$

Let  $F_z = B_z A^{-1}$  and note that  $F_z$  represents the linear transformation which maps the  $j$ -th row of  $A$  to the  $j$ -th row of  $B_z$ . We call  $F_z$  a *filtration matrix*. Let

$$W_z = F_z W \in \mathbb{R}^{(k-1) \times p}.$$

Finally, we define a new classifier

$$h_z : x \mapsto W_z \cdot f(x).$$

Our unlearning operation is thus

$$\mathfrak{D}_z : \begin{cases} \mathcal{H} \rightarrow \mathcal{H}_{-\mathcal{C}} \\ h \mapsto h_z. \end{cases}$$

We call  $\mathfrak{D}_z$  a *linear filtration*.

Note that  $\mathfrak{D}_z$  replaces  $W$  by  $W_z = F_z W$ , hence after applying  $\mathfrak{D}_z$  we may delete  $W$ . This means that, even though our unlearning operation essentially filters the outputs of the original classifier, the linearity of the filtering operations allows us to absorb the filter into the classifier. This is an important feature in a situation were it may no longer be permissible to store the original classifier. Furthermore, as we consider weak unlearning and thus black-box adversaries with access to the logit output only, in our scenario, an adversary can solely access the final, i.e., filtered, output of the newly derived classifier.

So how do we choose  $z \in \mathbb{R}^{k-1 \times k}$ ?



### 4.1.1 Naive unlearning

$z = \pi(a_0)$ . This gives  $F_z = (0 \mid I_{k-1})$ , i.e.  $F_z = \pi$ . Intuitively, we may think of this choice as simply cutting the output unit associated with  $\mathcal{C}$  out off the classifier. We call this unlearning operation the *naive method* and will use it as a baseline to measure the improvements other methods provide.

### 4.1.2 Normalization

$$z = \pi(a_0) - \frac{1}{k-1} \sum_{1 \leq i < k} (a_0)_i + \frac{1}{(k-1)^2} \sum_{1 \leq i, j < k} (a_j)_i.$$

This means we shift  $\pi(a_0)$  such that its mean becomes the mean of the remaining rows of  $B_z$ . The intuition behind this choice of  $z$  is that we would like inputs of the class in  $\mathcal{C}$  to be misclassified in a “natural” way. We base this method on the assumption that the values in  $\pi(a_0)$  encode a natural distribution for predictions of the class in  $\mathcal{C}$ . However, we expect the values in  $\pi(a_0)$  to be unnaturally low (absolutely speaking), hence we shift them. We refer to this method as *normalization* or *normalizing filtration*. It is the main method we propose. Figure 3 agrees with the intuition described above: the bars for normalization line up nicely with the bars of the retrained models, while the bars of randomization do not. Proposition 1 summarizes what normalizing filtration achieves. For comparison we define three more methods.


### 4.1.3 Randomization

$z \sim \mathcal{N}(0, I_{k-1})$ . We sample  $z$  from a multivariate normal distribution. We refer to this method as *randomization*.

### 4.1.4 Zeroing

$z = 0$ . We refer to this method as *zeroing*.

**Fig. 3** The probability distribution predicted for class “airplane” after its unlearning or randomization from models trained on CIFAR-10, compared to models retrained without class “airplane”. The bars are centered around the mean and have length of the standard deviation, over 100 models



### 4.1.5 Transfer

Instead of retraining the entire model, we only retrain the logistic regression head  $W$  with the same settings as the original model while holding  $f$  fixed. We refer to this method as *transfer unlearning*.

**Proposition 1** (Normalizing filtration.) *After applying normalizing filtration to delete class 0 we obtain a classifier  $h_z$  with the following property:*

$$\text{mean}(\mathbb{E}[h_z(\mathbf{X})|\mathbf{Y} = 0]) = \text{mean}(\text{mean}_{1 \leq j < k}(\pi(a_j))).$$

*I.e. the mean of the predictions of inputs belonging to the deleted class is equal to the mean of predictions of the remaining classes.*

$$\begin{aligned} \mathbb{E}[h_z(\mathbf{X})|\mathbf{Y} = 0] &= \mathbb{E}[W_z f(\mathbf{X})|\mathbf{Y} = 0] \\ &= \mathbb{E}[W_z f(\mathbf{X})|\mathbf{Y} = 0] \\ &= \mathbb{E}[F_z W f(\mathbf{X})|\mathbf{Y} = 0] \\ &= F_z \mathbb{E}[W f(\mathbf{X})|\mathbf{Y} = 0] \\ &= F_z a_0 = z. \end{aligned}$$

**Proof**

Thus we have:

$$\begin{aligned} \text{mean}(\mathbb{E}[h_z(\mathbf{X})|\mathbf{Y} = 0]) &= \text{mean}(z) \\ &= \text{mean}(\pi(a_0)) - \frac{1}{k-1} \sum_{1 \leq i < k} (a_0)_i + \frac{1}{(k-1)^2} \sum_{1 \leq i, j < k} (a_j)_i \\ &= \frac{1}{(k-1)^2} \sum_{1 \leq i, j < k} (a_j)_i \\ &= \text{mean}(\text{mean}_{1 \leq j < k}(\pi(a_j))). \end{aligned}$$

□

### 4.2 Computational complexity of $\mathfrak{D}_z$

To find  $\mathfrak{D}_z$  we need to compute the following: **1)**  $A$ , i.e. the expected predictions for all  $k$  classes; **2)**  $A^{-1}$ , the inversion of a  $(k \times k)$ -Matrix; **3)**  $z$ , in case  $z$  has a non-trivial definition (e.g. computing a sample of  $\mathcal{N}$ ); **4)**  $F_z$ , the multiplication of a  $((k-1) \times k)$  with a  $(k \times k)$ -Matrix; **5)**  $W_z$ , the multiplication of a  $((k-1) \times k)$  with a  $(k \times p)$ -Matrix.

Algorithm 1 lists this computation for the case of deleting class 0.

**Algorithm 1** Normalizing filtration deleting class 0Input:  $h = W \circ f$ Output:  $h_z = W_z \circ f$ 


---

```

1: for  $0 \leq j < k$  do
2:    $predictions \leftarrow []$ 
3:   for  $(x, j) \in D \cap \mathcal{X} \times \{j\}$  do
4:      $predictions.append(h(x))$ 
5:   end for
6:    $a_j \leftarrow \text{mean}(predictions)$ 
7: end for
8:  $A = (a_0|a_1|\dots|a_{k-1})$ 
9:  $z \leftarrow \pi(a_0) - \frac{1}{k-1} \sum_{1 \leq i < k} (a_0)_i + \frac{1}{(k-1)^2} \sum_{1 \leq i, j < k} (a_j)_i$ 
10:  $B_z = (z|\pi(a_1)|\dots|\pi(a_{k-1}))$ 
11:  $W_z \leftarrow B_z A^{-1} W$ 
12:  $h_z \leftarrow W_z \circ f$ 
13: return  $h_z$ 

```

---

In practice (1) will incur the majority of the computational costs, while (2)–(5) will be negligible. Thus, if we estimate  $A$  by predicting  $\ell$  samples per class the computational complexity of finding  $\mathfrak{D}_z(h)$  is  $\ell \cdot k$  times the complexity of computing a prediction of  $h$ . We find that  $\mathfrak{D}_z$  is robust in respect to the quality of the estimation of  $A$ , hence a small amount of sample points per class suffice, see Sect. 5.4.

Note that the costs incurred by linear filtration are virtually the same for several concurrent deletion requests: we just need to change  $B_z$  appropriately, e.g. if we want to delete class 3 in addition to class 0 we need to replace  $\pi(a_3)$  in column 3 with some  $z_3$ . The only additional costs incurred are the computation of  $z_3$  (which for our proposed methods are negligible).

## 5 Experiments

### 5.1 Evaluation method

By definition 2 the quality of a weak unlearning operation depends on the similarity of the resulting distributions  $P(\mathbf{L}_{\text{seen}})$  and  $P(\mathbf{L}_{\text{-seen}})$ . We begin by defining a divergence measure for distributions based on the Bayes error rate. This then motivates us to empirically evaluate the performance of the unlearning operations proposed in Sect. 4 by training binary classifiers on the pre-softmax outputs of our models.

#### 5.1.1 Classifier advantage

Let  $\mathbf{B}$  be a Bernoulli random variable uniformly taking values in  $\{0, 1\}$ . Let

$$\mathbf{U} = \mathbf{L}_{\text{seen}} \cdot (1 - \mathbf{B}) + \mathbf{L}_{\text{-seen}} \cdot \mathbf{B}$$

be the mixture of  $\mathbf{L}_{\text{seen}}$  and  $\mathbf{L}_{\text{-seen}}$ . Let

$$b : \text{dom}(\mathbf{U}) \rightarrow \{0, 1\}$$

be a binary classifier operating on the mixture  $\mathbf{U}$  and define

$$\alpha_b = 2 \left( \mathbb{E} [P(\mathbf{B} = b(\mathbf{U}) \mid \mathbf{U})] - \frac{1}{2} \right). \quad (2)$$

We call  $\alpha_b$  the *classifier advantage* of  $b$ . Intuitively speaking,  $\alpha_b$  is a measure for how good  $b$  is at telling the mixture  $\mathbf{U}$  apart. Let

$$b^* : u \mapsto \operatorname{argmax}_{i < 2} P(\mathbf{B} = i \mid \mathbf{U} = u)$$

be the Bayes optimal classifier for  $P(\mathbf{U}, \mathbf{B})$ . Then  $\alpha_{b^*} \in [0, 1]$  and it is a measure for the difference between  $P(\mathbf{L}_{\text{seen}})$  and  $P(\mathbf{L}_{\neg\text{seen}})$ , based on how much better the Bayes optimal classifier  $b^*$  performs than random guessing. A value of  $\alpha_{b^*}$  close to 0 indicates that  $P(\mathbf{L}_{\text{seen}})$  and  $P(\mathbf{L}_{\neg\text{seen}})$  are similar.

### 5.1.2 Experimental setup

Assume that  $b^*$  can be approximated sufficiently well by a classifier  $b$  derived via a state-of-the-art binary classification algorithm. Then  $\alpha_b$  approximates  $\alpha_{b^*}$ , hence we consider a low value of  $\alpha_b$  to be evidence for the similarity of the distributions of  $\mathbf{L}_{\text{seen}}$  and  $\mathbf{L}_{\neg\text{seen}}$ . Note that even if we do not believe that  $b$  approximates  $b^*$  well, we may still use  $\alpha_b$  as a relative performance measure for different unlearning operations.

Drawing i.i.d. samples from  $P(\mathbf{L}_{\text{seen}})$  and  $P(\mathbf{L}_{\neg\text{seen}})$  is computationally expensive as it requires us to run the algorithms  $\mathbf{A}(D)$  respectively  $\mathbf{A}_{\neg\mathcal{C}}(D_{\neg\mathcal{C}})$  for every sample point. In our experiments we thus take a pragmatic approach. We train 100 models which get to see the full training data and 100 models which get to see the training data with  $\mathcal{C}$  removed, i.e. models that unlearned by retraining from scratch. We then apply our unlearning operation to each of the models that got to see the full training data. For every single model we then calculate the predictions for the full test data by that model, without applying the softmax activations of the output layer. We label the predictions made by models that originally got to see the full training data with “**seen**” and predictions made by models which never got to see the training data for  $\mathcal{C}$  as “**not seen**”.

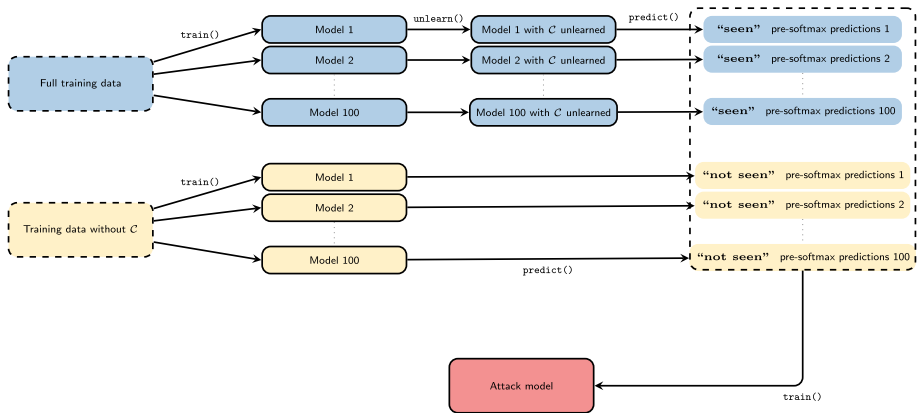
Finally, we train a binary classifier  $b$  that given a prediction attempts to predict its label “**seen**” or “**not seen**”. For this purpose we use the predictions of 70 models of either category as training data and the predictions of the remaining 30 models of either category as test data. Figure 4 shows a schematic representation of our setup. In practice, we train a separate binary classifier on the predictions of each class.

Note the subtle difference of our setup to the shadow model setup of (Shokri et al. 2017), where the authors ask their binary classifier (there “attack model”): “Do these outputs look like they come from a member of the training set?”, and hope for an accurate answer, such that their membership inference attack succeeds. We ask our binary classifier: “Do these outputs look like they come from a model that has seen  $\mathcal{C}$ ?”, and hope for an inaccurate answer, as we hope that our unlearning operation prevents the attack.

## 5.2 Data

### 5.2.1 MNIST

The MNIST dataset (Lecun et al. 1998) contains 70, 000 28×28 images of handwritten digits in 10 classes, with 7, 000 images per class, split into 60, 000 training and 10, 000 test images.



**Fig. 4** Experimental setup: On the full training data we train 100 models by  $\mathbf{A} = \text{train}()$ . To these models we then apply an unlearning operation  $\mathfrak{D} = \text{unlearn}()$ . We then  $\text{predict}()$  our test data for each of these models and label these predictions “seen”. On the training data with  $C$  removed we train 100 models by  $\mathbf{A}_{-C} = \text{train}()$ . We then  $\text{predict}()$  our test data for each of these models and label these predictions “not seen”. Finally, we use all labeled predictions as the training/test data of a binary classifier  $b$ , which we employ as our “attack model”. We interpret low test accuracy of  $b$  as evidence for good performance of a weak unlearning operation

## 5.2.2 CIFAR-10

The CIFAR-10 dataset (Krizhevsky 2009) contains 60,000  $32 \times 32 \times 3$  images in 10 classes, with 6,000 images per class, split into 50,000 training and 10,000 test images.

## 5.2.3 AT&T Faces

The AT&T Laboratories Cambridge Database of Faces (AT&T Laboratories Cambridge 1994) contains 400  $92 \times 112$  images of 40 subjects, with 10 images per subject.

## 5.3 Network architectures

### 5.3.1 MLP

For the MNIST dataset, we evaluate our unlearning method on a multilayer perceptron (MLP), with one hidden layer of 50 units.

### 5.3.2 CNN

For the CIFAR-10 dataset, we evaluate our unlearning method on convolutional neural networks (CNNs). Our networks consist of two convolutional (with 16 respectively 32  $3 \times 3$  filters) and two max-pooling layers, followed by a fully connected layer with  $p$  units with rectified linear activations and a softmax output layer. We experiment with  $p \in \{64, 256, 1024\}$ . Note that our unlearning operation works by

**Table 1** Classifier advantage for 1 unlearned and 9 remaining classes, for MLPs trained on MNIST

Unlearned	NN	RF	AB
Naive	0.593	0.609	0.641
Normalization	0.327	0.362	0.438
Transfer	0.462	0.552	0.550
Remaining	NN	RF	AB
Naive	0.048	0.090	0.098
Normalization	0.041	0.089	0.095
Transfer	0.315	0.383	0.428

**Table 2** Classifier advantage for 1 unlearned and 9 remaining classes, for CNNs trained on CIFAR-10, with  $p = 256$ 

Unlearned	NN	RF	AB
Naive	0.609	0.457	0.590
Normalization	0.146	0.110	0.093
Randomization	0.634	0.579	0.582
Zeroing	0.642	0.566	0.575
Transfer	0.951	0.840	0.952
Remaining	NN	RF	AB
Naive	0.115	0.080	0.109
Normalization	0.148	0.097	0.118
Randomization	0.416	0.279	0.230
Zeroing	0.421	0.276	0.219
Transfer	0.949	0.819	0.948

manipulating  $W \in \mathbb{R}^{k \times p}$ , thus we want to investigate how the size of  $W$  affects unlearning performance.

### 5.3.3 ResNet

For the AT&T faces dataset we evaluate our unlearning method on a residual neural network architecture (He et al. 2016). We use a convolutional layer (with 8  $5 \times 5$  filters), followed by 5 downsampling residual blocks (with  $2^{i+3} 3 \times 3$  filters in the  $i$ -th block), followed by global max-pooling and a softmax output layer.

## 5.4 Results

We experiment with the following classification algorithms: nearest neighbors (NN), random forests (RF), and AdaBoost (AB). For the models trained on MNIST and CIFAR-10 we unlearn the first class (of 10) in the dataset (the digit “zero” and the class “airplane”, respectively). For the models trained on AT&T Faces we unlearn the first 4 individuals (of 40) in the dataset.

Table 1 shows classifier advantages (remember equation (2)) for the MLPs trained on MNIST. We observe a significant decrease of advantage when unlearning by normalization compared to the naive method for the unlearned class, and similar advantage for the remaining classes. The classifier advantage for the transfer learning approach is

**Table 3** Classifier advantage for 4 unlearned and 36 remaining classes, for residual networks trained on AT&T faces

Unlearned	NN	RF	AB
Naive	0.467	0.573	0.574
Normalization	0.381	0.454	0.462
Transfer	0.985	0.977	0.978
Remaining	NN	RF	AB
Naive	0.149	0.266	0.245
Normalization	0.152	0.263	0.246
Transfer	0.998	0.996	0.996

**Table 4** Classifier advantage for CNNs trained on CIFAR-10, for different sample size  $s$  per class, with  $p = 256$ 

Unlearned	$s$	NN	RF	AB
Naive	–	0.609	0.457	0.590
Normalization	10	0.193	0.114	0.122
Normalization	100	0.156	0.095	0.111
Normalization	1000	0.146	0.110	0.093
Remaining	$s$	NN	RF	AB
Naive	–	0.115	0.080	0.109
Normalization	10	0.205	0.142	0.171
Normalization	100	0.169	0.108	0.136
Normalization	1000	0.148	0.097	0.118

between the values for naive unlearning and normalization for the unlearned class, and considerably higher for the remaining class. Table 2 shows classifier advantages for the CNNs trained on CIFAR-10 (with  $p = 256$ ). We observe a vast decrease of advantage when unlearning by normalization compared to the naive method for the unlearned class, while randomization and zeroing do not provide such benefits. For the transfer learning approach, the classifier advantages are high for both unlearning and remaining classes. Table 3 shows classifier advantages for the residual networks trained on AT&T Faces. We observe a slight decrease of advantage when unlearning by normalization compared to the naive method for the unlearned classes, and similar advantage for the remaining classes. Again, for the transfer learning approach, the classifier advantages are high for both unlearning and remaining classes.

#### 5.4.1 Effect of sample size

We investigate the effect of sample size for estimating the matrix of mean predictions  $A$  (recall Sect. 4.1). Table 4 shows the results for CNNs trained on CIFAR-10 (with  $p = 256$ ). Unlearning by normalization, compared to the naive method, we observe a strong decrease of advantage for the unlearned class for a sample size of 10 per class (1% of the test data), and for a sample size of 100 per class (10% of the test data) we observe performance comparable to the estimation based on the full test data.

**Table 5** Classifier advantage, for CNNs trained on CIFAR-10, for different numbers  $p$  of units in the fully connected layer

Unlearned	$p$	NN	RF	AB
Naive	64	0.290	0.321	0.380
Normalization	64	0.086	0.106	0.119
Naive	256	0.609	0.457	0.590
Normalization	256	0.146	0.110	0.093
Naive	1024	0.627	0.485	0.603
Normalization	1024	0.174	0.138	0.138
Remaining	$p$	NN	RF	AB
Naive	64	0.040	0.043	0.044
Normalization	64	0.044	0.040	0.045
Naive	256	0.115	0.080	0.109
Normalization	256	0.148	0.097	0.118
Naive	1024	0.129	0.090	0.115
Normalization	1024	0.174	0.104	0.129

### 5.4.2 Effect of $p$

For the CNNs trained on CIFAR-10 we investigate the effect of the number  $p$  of units in the fully connected layer. Table 5 shows that unlearning by normalization significantly decreases classifier advantage for all tested values of  $p$  on the unlearned class. We observe a slight increase of advantage for the remaining classes, that appears to get slightly worse with  $p$  increasing.

### 5.4.3 Effect on classification accuracy

While implicit in definition 2, an important consideration for the design of an unlearning operation is that we do not want to decrease the performance of a classifier on the remaining classes. Table 6 reports the test images for which the most likely label predicted was changed by one of our unlearning methods, compared to the naive method. Normalization did not change any labels (and thus in particular non of the correct ones).

**Table 6** The amount of labels changed in percent, when compared to the naive method, for 100 models trained on CIFAR-10, and  $p = 256$ . The “all”-column reports the value for predictions of the entire test set, the “unl.”-column for predictions of images of the unlearned class, and the “cor.”-column for predictions of images of the remaining classes that are correctly predicted by the naive method. The “acc.”-column shows the mean classification accuracy

Deletion method	All	Unl.	Cor.	Acc
Naive	–	–	–	69.2
Normalization	0.0	0.0	0.0	69.2
Randomization	20.1	48.8	11.8	64.8
Zeroing	18.3	45.4	10.3	65.7



**Table 7**  $KS_{\phi}$  for  $\Phi$  consisting of 1000 unit vectors pointing in uniform random directions. (Lower values are better.)

Unlearned	CIFAR-10	MNIST
Naive	0.115	0.184
Normalization	0.038	0.1
Baseline	0.015	0.017
Remaining	CIFAR-10	MNIST
Naive	0.037	0.034
Normalization	0.038	0.034
Baseline	0.013	0.023

#### 5.4.4 Summary of experimental results

Normalization (1) decreased classifier advantage on the unlearned classes in all three experiments; (2) showed robustness with regards to sample size for parameter estimation; (3) performed well for different dimensions of the feature space, and (4) did not negatively affect classification accuracy. The transfer learning approach performed poorly for all three datasets.

#### 5.5 Model inversion

Figures 1 (on page 2) and 5 show the results of a “model inversion attack”, i.e. gradient ascend on the input space for a model trained on AT&T Database of Faces. We use a neural network consisting of two fully connected layers with 1000 and 300 units respectively with sigmoid activations and a softmax output layer.

Visually, naive unlearning barely affects the quality of the reconstruction for any of the classes (in particular not the unlearned class). On the other hand the normalization method greatly disturbs the reconstruction of the unlearned class, while barely affecting

**Fig. 5** Model inversion for a toy model trained on the AT&T Faces dataset with 4 classes. The top row shows one training image of each class, the second row reconstructions of classes by model inversion, the third row reconstructions after naive unlearning of the class in the first column, the bottom row reconstructions after unlearning the class in the first column by *normalizing filtration*. See Fig. 6 for the remaining classes



the remaining classes. In accordance with our discussion of model inversion in Sect. 2, we thus interpret Fig. 1 as visual evidence suggesting a desirable effect of our normalization method on the correlation between input and output space represented by our model. We leave a more detailed investigation of this phenomenon for future work.

## 5.6 Random direction Kolmogorov–Smirnov statistics

Given  $\Phi \subseteq \mathbb{R}^{k-|C|}$ , and having drawn samples  $S_1, S_2$  from  $P(\mathbf{L}_{\text{seen}})$  and  $P(\mathbf{L}_{\text{-seen}})$  respectively in the fashion described in Sect. 5.1, we compute the statistic

$$\text{KS}_{\Phi}(S_1, S_2) = \frac{1}{|\Phi|} \sum_{\phi \in \Phi} \text{KS}_{\text{two sample}}(\phi \cdot S_1, \phi \cdot S_2). \quad (3)$$

Here  $\text{KS}_{\text{two sample}}$  denotes the two-sample Kolmogorov–Smirnov statistic and  $\phi \cdot S$  is shorthand for  $\{\phi \cdot s : s \in S\} \subseteq \mathbb{R}$ .

Table 7 reports the  $\text{KS}_{\Phi}$  statistic for  $\Phi$  consisting of 1000 unit vectors pointing in uniform random directions. We compare naive unlearning to normalizing filtration for the CIFAR-10 experiment employing CNNs ( $p = 256$ ), and the MNIST experiment employing MLPs. As noted in Sect. 5.1 the samples we draw are not i.i.d., hence we also report a baseline  $\text{KS}_{\Phi}$  statistic computed from two independent batches of 100 models trained without the class we unlearn.

The results paint a picture similar to Sect. 5.4. In both experiments normalizing filtration decreases  $\text{KS}_{\Phi}$  for the unlearned class, with somewhat better performance in the CIFAR-10 experiment, while the  $\text{KS}_{\Phi}$  remains unchanged for the remaining classes.

## 6 Discussion

In Sect. 3 we made a didactic choice to introduce the special hypothesis class  $\mathcal{H}$  (which permitted absorption of the filtration operation) before defining unlearning. It should however be clear how definitions 1 and 2 are applicable to any class of classifiers, and in the somewhat more common setting of deletion requests of single data points. We would like to emphasize our belief that in the light of the non-deterministic nature of many learning algorithms a probabilistic definition of unlearning, such as chosen in (Ginart et al. 2019) (there “deletion”) and our work, is necessary.

For contrast let us consider the definition of unlearning in (Bourtoule et al. 2019) that asks to find a model that “could have been obtained” without looking at the deleted training data. Since all models on discrete digital systems necessarily have a finite parameter space we very much could obtain any model without looking at any data by guessing its parameters. What happened here?

Guessing a model’s parameters can be considered drawing a sample from a uniform distribution over the parameter space. On the other hand a probabilistic definition such as 1 requires the distribution over the parameter space to be the same as if we had run the original learning algorithm without using the deleted data, which for reasonable learning algorithms is certainly not uniform. If we would like to stick to an informal definition we should therefore say that a model “could have been obtained, with reasonable likelihood”.

We further conclude that it is an important consideration whether a definition of unlearning makes sense when not read in a benevolent way (e.g. by a party whose interest

in unlearning stems from of legal obligation). In fact our definition of weak unlearning 2, i.e. unlearning in a black-box sense, suffers from a similar issue. A malicious way to define a weak unlearning operations  $\mathfrak{D}$  is the following: for any classifier  $h$ , simply train a new classifier  $h'$  without using the deleted data, then define  $\mathfrak{D}(h) = h'' = (h, h')$ , where  $h''(x) = h'(x)$ . The outputs of  $\mathfrak{D}(h)$  look exactly like those of  $h'$ , thus  $\mathfrak{D}$  is indeed an unlearning operation in the black-box sense, however we never actually deleted  $h$ . Let us thus emphasize that the weak definition of unlearning is only applicable when acting in good faith.

## 7 Conclusion

We considered the problem of unlearning in a class-wide setting, for classifiers predicting logits. We developed *normalizing filtration* as an unlearning method, with compelling visual results (figure 6). These are backed up by good experimental results with regards to our proposed definition of *weak unlearning* and our metrics of *classifier*



**Fig. 6** Model inversion for a toy model trained on the AT&T Faces dataset with 4 classes. On the left side we show results for naive unlearning, on the right side for *normalizing filtration*. On either side: The top row shows one training image of each class, the second row reconstructions of classes by model inversion, the  $(i + 2)$ -th row shows the reconstructions after unlearning the class in the  $i$ -th column by the respective unlearning operation

*advantage* (2) and the  $KS_{\phi}$  statistic (3). We emphasize once again the black-box nature of weak unlearning, entailing the black-box nature of our proposed unlearning operation. While linear filtration can be absorbed into the final layer of a classifier (for the hypothesis class considered), our approach remains somewhat limited with regards to its shallowness. In future work we hope to find methods that allow for deeper absorption, thus hopefully leading to stronger privacy guarantees. On the other hand our method's intuitive simplicity nicely complements concurrent approaches such as (Guo et al. 2019; Golatkar et al. 2020a, b). Another promising direction may be to enhance our method by some form of shrinkage of the logits, in case the unlearned class constitutes a large part of the misclassifications of one of the remaining classes and plan to explore the connection to the field of private prediction van der Maaten and Hannun (2020). Finally it should be noted that, while the “right to be forgotten” inspired our research, whether our approach is adequate in this context is for legal scholars to decide.

## Appendix

### A more complete versions of tables from section 5.4

#### MNIST – MLP experiment

See Tables 8, 9.

**Table 8** Accuracy (in percent) and cross-entropy loss, mean  $\pm$  standard deviation, for 100 MLPs trained on MNIST.  $p$  is the number of units in the fully connected layer, i.e. the dimension of the feature space.  $s$  is the sample size per class, used to estimate the mean predictions

Unlearned	NN	RF	AB
Naive	0.593	0.609	0.641
Normalization	0.327	0.362	0.438
Randomization	0.910	0.869	0.902
Zeroing	0.919	0.886	0.909
Remaining	NN	RF	AB
Naive	0.048	0.090	0.098
Normalization	0.041	0.089	0.095
Randomization	0.116	0.132	0.133
Zeroing	0.119	0.129	0.132

**Table 9** Classifier advantage for 1 unlearned and 9 remaining classes, for MLPs trained on MNIST.  $p$  is the number of units in the fully connected layer, i.e. the dimension of the feature space.  $s$  is the sample size per class, used to estimate the mean predictions

	Accuracy	Loss
Before unlearning	97.1 $\pm$ 0.19	0.10 $\pm$ 0.01
Baseline model	97.1 $\pm$ 0.18	0.09 $\pm$ 0.01
Naive	97.1 $\pm$ 0.20	0.10 $\pm$ 0.01
Normalization	97.1 $\pm$ 0.20	0.10 $\pm$ 0.01
Randomization	96.7 $\pm$ 0.33	0.11 $\pm$ 0.01
Zeroing	96.7 $\pm$ 0.30	0.11 $\pm$ 0.01

## CIFAR10 – CNN experiment

Tables 10, 11.

**Table 10** Classifier advantage for 1 unlearned and 9 remaining classes, for CNNs trained on CIFAR-10.  $p$  is the number of units in the fully connected layer, i.e. the dimension of the feature space.  $s$  is the sample size per class, used to estimate the mean predictions

Unlearned	$p$	$s$	NN	RF	AB
Naive	64	–	0.290	0.331	0.380
Normalization	64	10	0.113	0.135	0.137
Normalization	64	100	0.089	0.105	0.129
Normalization	64	–	0.086	0.106	0.119
Randomization	64	–	0.547	0.487	0.528
Zeroing	64	–	0.599	0.534	0.560
Naive	256	–	0.609	0.457	0.590
Normalization	256	10	0.193	0.114	0.122
Normalization	256	100	0.156	0.095	0.111
Normalization	256	–	0.146	0.110	0.093
Randomization	256	–	0.634	0.579	0.582
Zeroing	256	–	0.642	0.566	0.575
Naive	1024	–	0.627	0.485	0.603
Normalization	1024	10	0.230	0.164	0.179
Normalization	1024	100	0.216	0.157	0.172
Normalization	1024	–	0.174	0.138	0.138
Randomization	1024	–	0.741	0.637	0.682
Zeroing	1024	–	0.746	0.642	0.680
<i>Remaining</i>					
Naive	64	–	0.040	0.043	0.044
Normalization	64	10	0.072	0.075	0.080
Normalization	64	100	0.047	0.044	0.047
Normalization	64	–	0.044	0.040	0.045
Randomization	64	–	0.240	0.203	0.174
Zeroing	64	–	0.246	0.204	0.163
Naive	256	–	0.115	0.080	0.109
Normalization	256	10	0.205	0.142	0.171
Normalization	256	100	0.169	0.108	0.136
Normalization	256	–	0.148	0.097	0.118
Randomization	256	–	0.416	0.279	0.230
Zeroing	256	–	0.421	0.276	0.219
Naive	1024	–	0.129	0.090	0.115
Normalization	1024	10	0.242	0.149	0.182
Normalization	1024	100	0.217	0.128	0.146
Normalization	1024	–	0.174	0.097	0.118
Randomization	1024	–	0.506	0.309	0.249
Zeroing	1024	–	0.508	0.310	0.251

**Table 11** Accuracy (in percent) and cross-entropy loss, mean  $\pm$  standard deviation, for 100 CNNs trained on CIFAR10.  $p$  is the number of units in the fully connected layer, i.e. the dimension of the feature space.  $s$  is the sample size per class, used to estimate the mean predictions

	$p$	$s$	Accuracy	Loss
Before unlearning	64	–	65.1 $\pm$ 1.01	1.01 $\pm$ 0.03
Retraining	64	–	66.0 $\pm$ 1.06	0.96 $\pm$ 0.03
Naive	64	–	66.4 $\pm$ 1.04	0.95 $\pm$ 0.03
Normalization	64	10	66.4 $\pm$ 1.04	0.95 $\pm$ 0.03
Normalization	64	100	66.4 $\pm$ 1.04	0.95 $\pm$ 0.03
Normalization	64	–	66.4 $\pm$ 1.04	0.95 $\pm$ 0.03
Randomization	64	–	61.1 $\pm$ 2.08	1.17 $\pm$ 0.10
Zeroing	64	–	62.3 $\pm$ 1.76	1.12 $\pm$ 0.09
Before unlearning	256	–	68.1 $\pm$ 0.85	0.93 $\pm$ 0.02
Retraining	256	–	69.2 $\pm$ 0.80	0.89 $\pm$ 0.02
Naive	256	–	69.2 $\pm$ 0.83	0.88 $\pm$ 0.02
Normalization	256	10	69.2 $\pm$ 0.83	0.88 $\pm$ 0.02
Normalization	256	100	69.2 $\pm$ 0.83	0.88 $\pm$ 0.02
Normalization	256	–	69.2 $\pm$ 0.83	0.88 $\pm$ 0.02
Randomization	256	–	64.9 $\pm$ 1.76	1.07 $\pm$ 0.08
Zeroing	256	–	65.7 $\pm$ 1.40	1.03 $\pm$ 0.06
Before unlearning	1024	–	69.7 $\pm$ 0.81	0.93 $\pm$ 0.03
Retraining	1024	–	70.5 $\pm$ 0.82	0.88 $\pm$ 0.03
Naive	1024	–	70.7 $\pm$ 0.82	0.88 $\pm$ 0.03
Normalization	1024	10	70.7 $\pm$ 0.82	0.88 $\pm$ 0.03
Normalization	1024	100	70.7 $\pm$ 0.82	0.88 $\pm$ 0.03
Normalization	1024	–	70.7 $\pm$ 0.82	0.88 $\pm$ 0.03
Randomization	1024	–	67.1 $\pm$ 1.55	1.06 $\pm$ 0.08
Zeroing	1024	–	67.6 $\pm$ 1.16	1.03 $\pm$ 0.05

**Acknowledgments** The authors thank Valentin Muhr for early discussion on Machine Unlearning.

**Author contributions** TB was the lead author of the paper, developed the method for machine unlearning and performed the experiments. PS supported the evaluation, supervised the method development, contributed to paper writing and performed the preliminary research\* on the investigated topic together with MZ. MZ raised the funding for the research, initiated the idea for the paper, supervised the method development and evaluation and contributed to the writing of the paper.

**Funding** Open access funding provided by MCI Management Center Innsbruck – Internationale Hochschule GmbH. The first and third author were funded by the Austrian Research Promotion Agency (FFG) COIN project 866880 “Big Data Analytics”. The second author was supported by the Austrian Science Fund (FWF) under Grant No. I 4057-N31 (“Game Over Eva(sion)”).

**Data availability** We only use data that is openly available.

**Code availability** We will make the code available at Github upon acceptance of the manuscript.

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

**Ethics approval** Not Applicable.

**Consent to participate** Not Applicable.

**Consent for publication** Not Applicable.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Abadi, M., Chu, A., Goodfellow, I., McMahan, B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In: 23rd ACM Conference on Computer and Communications Security (ACM CCS)
- Achille, A., Paolini, G., & Soatto, S. (2019). Where is the information in a deep neural network? arXiv 1905.12213
- AT&T Laboratories Cambridge (1994) The database of faces
- Biggio, B., & Roli, F. (2018). Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition*, 84, 317–331.
- Bourtole, L., Chandrasekaran, V., Choquette-Choo, C., Jia, H., Travers, A., Zhang, B., Lie, D., & Papernot, N. (2019). Machine unlearning. arXiv preprint [arXiv:1912.03817](https://arxiv.org/abs/1912.03817)
- Cao, Y., & Yang, J. (2015). Towards making systems forget with machine unlearning. In: 2015 IEEE Symposium on Security and Privacy
- Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., & Song, D. (2019). The secret sharer: Evaluating and testing unintended memorization in neural networks. In: 28th USENIX Security Symposium (USENIX Security 19)
- Council of European Union (2012) Regulation (eu) 2012/0011. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52012PC0011>, accessed: 2022/04/28 09:23:04
- Council of European Union (2016) Regulation (eu) 2016/679. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>, accessed: 2022/04/28 09:23:04
- Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations Trends in Theoretical Computer Science*, 4(3–4):211–407
- Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In: Halevi S, Rabin T (eds) Theory of Cryptography
- Fredrikson, M., Lantz, E., Jha, S., Lin, S., Page, D., & Ristenpart, T. (2014). Privacy in pharmacogenetics: An end-to-end case study of personalized warfarin dosing. In: Proceedings of the 23rd USENIX Conference on Security Symposium
- Fredrikson, M., Jha, S., & 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
- Ginart, A., Guan, M., Valiant, G., & Zou, JY. (2019). Making AI forget you: Data deletion in machine learning. In: Advances in Neural Information Processing Systems, pp 3518–3531
- Golatkar, A., Achille, A., & Soatto, S. (2020a). Eternal sunshine of the spotless net: Selective forgetting in deep networks. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)
- Golatkar, A., Achille, A., & Soatto, S. (2020b). Forgetting outside the box: Scrubbing deep networks of information accessible from input-output observations. arXiv 200302960
- Guo, C., Goldstein, T., Hannun, A., & van der Maaten, L. (2019). Certified data removal from machine learning models. arXiv 191103030
- Hayes, J., Melis, L., Danezis, G., & De Cristofaro, E. (2019). LOGAN: Membership inference attacks against generative models. PoPETS
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)

- Jacot, A., Gabriel, F., & Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. *Advances in Neural Information Processing Systems*
- Kearns, M. (1998). Efficient noise-tolerant learning from statistical queries. *Journal of the ACM*, 45(6), 983–1006.
- Kerckhoffs, A. (1883). La cryptographie militaire. *Journal des sciences militaires*, IX(1), 5–38.
- Koh, P.W., & Liang, P. (2017). Understanding black-box predictions via influence functions. *Proceedings of the 34th International Conference on Machine Learning*
- Krizhevsky, A. (2009). *Learning multiple layers of features from tiny images*. Tech. rep.: University of Toronto.
- Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*
- Lee, J., Xiao, L., Schoenholz, S., Bahri, Y., Sohl-Dickstein, J., & Pennington, J. (2019). Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in Neural Information Processing Systems*
- McSherry, F. (2016). Statistical inference considered harmful. <https://github.com/frankmsherry/blog/blob/master/posts/2016-06-14.md>, accessed: 2022/04/28 09:23:04
- Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, ZB., & Swami, A. (2016). The limitations of deep learning in adversarial settings. In: 2016 IEEE European symposium on security and privacy (EuroS&P), IEEE, pp 372–387
- Sablayrolles, A., Douze, M., Schmid, C., Ollivier, Y., & Jegou, H. (2019). White-box vs black-box: Bayes optimal strategies for membership inference. In: *International Conference on Machine Learning*
- Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In: 2017 IEEE Symposium on Security and Privacy (SP)
- Sommer, DM., Song, L., Wagh, S., & Mittal, P. (2020). Towards probabilistic verification of machine unlearning. arXiv 200304247
- State of California (2018) California Consumer Privacy Act. [https://leginfo.ca.gov/faces/billTextClient.xhtml?bill\\_id=201720180AB375](https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201720180AB375), accessed: 2022/04/28 09:23:04
- van der Maaten, L., & Hannun, A.Y. (2020). The trade-offs of private prediction. CoRR abs/2007.05089, <https://arxiv.org/abs/2007.05089>, 2007.05089
- Veale, M., Binns, R., & Edwards, L. (2018). *Algorithms that remember: Model inversion attacks and data protection law*. Mathematical, Physical and Engineering Sciences: Philosophical Transactions of the Royal Society A.
- Villaronga, EF., Kieseberg, P., & Li, TJ. (2018). Humans forget, machines remember: Artificial intelligence and the right to be forgotten. *Computer Law & Security Review*
- Yeom, S., Giacomelli, I., Fredrikson, M., & Jha, S. (2018). Privacy risk in machine learning: Analyzing the connection to overfitting. 2018 IEEE 31st Computer Security Foundations Symposium (CSF)
- Zhang, Y., Jia, R., Pei, H., Wang, W., Li, B., & Song, D. (2020). The secret revealer: generative model-inversion attacks against deep neural networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp 253–261

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.