

# Forgetting Outside the Box: Scrubbing Deep Networks of Information Accessible from Input-Output Observations

Aditya Golatkar, Alessandro Achille, and Stefano Soatto

Department of Computer Science  
University of California, Los Angeles  
{aditya29, achille, soatto}@cs.ucla.edu

**Abstract.** We describe a procedure for removing dependency on a cohort of training data from a trained deep network that improves upon and generalizes previous methods to different readout functions, and can be extended to ensure forgetting in the final activations of the network. We introduce a new bound on how much information can be extracted per query about the forgotten cohort from a black-box network for which only the input-output behavior is observed. The proposed forgetting procedure has a deterministic part derived from the differential equations of a linearized version of the model, and a stochastic part that ensures information destruction by adding noise tailored to the geometry of the loss landscape. We exploit the connections between the final activations and weight dynamics of a DNN inspired by Neural Tangent Kernels to compute the information in the final activations.

**Keywords:** Forgetting, Data Removal, Neural Tangent Kernel, Information Theory

## 1 Introduction

We study the problem of removing information pertaining to a given set of data points from the weights of a trained network. After removal, a potential attacker should not be able to recover information about the forgotten cohort. We consider both the cases in which the attacker has full access to the weights of the trained model, and the less-studied case where the attacker can only query the model by observing some input data and the corresponding output, for instance through a web Application Programming Interface (API). We show that we can quantify the maximum amount of information that an attacker can extract from observing inputs and outputs (black-box attack), as well as from direct knowledge of the weights (white-box), and propose tailored procedures for removing such information from the trained model in *one shot*. That is, assuming the model has been obtained by fine-tuning a pre-trained generic backbone, we compute a single perturbation of the weights that, in one go, can erase information about a cohort to be forgotten in such a way that an attacker cannot access it.




Fig. 1: **Scrubby procedure:** PCA-projection of training paths on  $\mathcal{D}$  (blue),  $\mathcal{D}_r$  (orange) and the weights after scrubbing, using **(Left)** The Fisher method of [16], and **(Right)** the proposed scrubbing method. Our proposed scrubbing procedure (red cross) moves the model towards  $w(\mathcal{D}_r)$ , which reduces the amount of noise (point cloud) that needs to be added to achieve forgetting.

More formally, let a dataset  $\mathcal{D}$  be partitioned into a subset  $\mathcal{D}_f$  to be forgotten, and its complement  $\mathcal{D}_r$  to be retained,  $\mathcal{D} = \mathcal{D}_f \sqcup \mathcal{D}_r$ . A (possibly stochastic) training algorithm  $A$  takes  $\mathcal{D}_f$  and  $\mathcal{D}_r$  and outputs a weight vector  $w \in \mathbb{R}^p$ :

$$A(\mathcal{D}_f, \mathcal{D}_r) \rightarrow w.$$

Assuming an attacker knows the training algorithm  $A$  (e.g., stochastic gradient descent, or SGD), the weights  $w$ , and the retainable data  $\mathcal{D}_r$ , she can exploit their relationship to recover information about  $\mathcal{D}_f$ , at least for state-of-the-art deep neural networks (DNNs) using a “readout” function [16], that is, a function  $R(w)$  that an attacker can apply on the weights of a DNN to extract information about the dataset. For example, an attacker may measure the confidence of the classifier on an input, or measure the time that it takes to fine-tune the network on a given subset of samples to decide whether that input or set was used to train the classifier (membership attack). We discuss additional examples in Section 5.4. Ideally, the forgetting procedure should be robust to different choices of readout functions. Recent work [16, 17], introduces a “scrubbing procedure” (forgetting procedure or data deletion/removal)  $S_{\mathcal{D}_f}(w) : \mathbb{R}^p \rightarrow \mathbb{R}^p$ , that attempts to remove information about  $\mathcal{D}_f$  from the weights, i.e.,

$$A(\mathcal{D}_f, \mathcal{D}_r) \rightarrow w \rightarrow S_{\mathcal{D}_f}(w)$$

with an upper-bound on the amount of information about  $\mathcal{D}_f$  that can be extracted after the forgetting procedure, provided the attack has access to the scrubbed weights  $S_{\mathcal{D}_f}(w)$ , a process called “white-box attack.”

Bounding the information that can be extracted from a white-box attack is often complex and may be overly restrictive: Deep networks have large sets of equivalent solutions – a “null space” – that would give the same activations on all test samples. Changes in  $\mathcal{D}_f$  may change the position of the weights in the null space. Hence, the position of the weight in the null-space, even if irrelevant for the input-output behavior, may be exploited to recover information about  $\mathcal{D}_f$ .

This suggests that the study of forgetting should be approached from the perspective of the *final activations*, rather than the weights, since there could be infinitely many different models that produce the same input-output behavior, and we are interested in preventing attacks that affect the behavior of the network, rather than the specific solution to which the training process converged. More precisely, denote by  $f_w(x)$  the final activations of a network on a sample  $x$  (for example the softmax or pre-softmax vector). We assume that an attacker makes queries on  $n$  images  $\mathbf{x} = (x_1, \dots, x_n)$ , and obtains the activations  $f_w(\mathbf{x})$ . The pipeline then is described by the Markov Chain

$$A(\mathcal{D}_f, \mathcal{D}_r) \rightarrow w \rightarrow S_{\mathcal{D}_f}(w) \rightarrow f_{S_{\mathcal{D}_f}(w)}(\mathbf{x}).$$

The key question now is to determine how much information can an attacker recover about  $\mathcal{D}_f$ , starting from the final activations  $f_{S_{\mathcal{D}_f}(w)}(\mathbf{x})$ ? We provide a new set of bounds that quantifies the average information per query an attacker can extract from the model.

The forgetting procedure we propose is obtained using the Neural Tangent Kernel (NTK). We show that this forgetting procedure is able to handle the null space of the weights better than previous approaches when using over-parametrized models such as DNNs. In experiments, we confirm that it works uniformly better than previous proposals on all forgetting metrics introduced [16], *both in the white-box and black-box case* (Figure 1).

Note that one may think of forgetting in a black-box setting as just changing the activations (*e.g.*, adding noise or hiding one class output) so that less information can be extracted. This, however, is not proper forgetting as the model still contains information, it is just not visible outside. We refer to forgetting as removing information from the weights, but we provide bounds for how much information can be extracted after scrubbing in the black-box case, and show that they are order of magnitudes smaller than the corresponding bounds for white boxes for the same target accuracy.

**Key contributions:** To summarize, our contributions are as follow:

1. We introduce methods to scrub information from, and analyze the content of deep networks from their final activations (black-box attacks).
2. We introduce a “one-shot” forgetting algorithms that work better than the previous method [16] for both white-box and black-box attacks for DNNs.
3. This is possible thanks to an elegant connection between activations and weights dynamics inspired by the neural tangent kernel (NTK), which allows us to better deal with the null-space of the network weights. Unlike the NTK formalism, we do not need infinite-width DNNs.
4. We show that better bounds can be obtained against black-box attacks than white-box, which gives a better forgetting vs error trade-off curve.

## 2 Related work

**Differential privacy** [12] aims to learn the parameters of a model in such a way that no information about any particular training sample can be recovered. This

is a much stronger requirement than forgetting, where we only want to remove – *after* training is done – information about a given subset of samples. Given the stronger requirements, enforcing differential privacy is difficult for deep networks and often results in significant loss of accuracy [1, 10].

**Forgetting:** The term “machine unlearning” was introduced by [9], who shows an efficient forgetting algorithm in the restricted setting of statistical query learning, where the learning algorithm cannot access individual samples. [30] provided the first framework for instantaneous data summarization with data deletion using robust streaming submodular optimization. [14] formalizes the problem of efficient data elimination, and provides engineering principles for designing forgetting algorithms. However, they only provide a data deletion algorithms for k-means clustering. [6] propose a forgetting procedure based on sharding the dataset and training multiple models. Aside from the storage cost, they need to retrain subset of the models, while we aim for one-shot forgetting. [5] proposed a forgetting method for logit-based classification models by applying linear transformation to the output logits, but do not remove information from the weights. [17] formulates data removal mechanisms using differential privacy, and provides an algorithm for convex problems based on a second order Newton update. They suggest applying this method on top of the features learned by a differentially private DNN, but do not provide a method to remove information from a DNN itself. [22] provides a projective residual update based method using synthetic data points to delete data points for linear regression based models. [24, 15] provide a newton based method for computing the influence of a training point on the model predictions in the context of model interpretation and cross-validation, however, such an approach can also be used for data removal. Closer to us, [16] proposed a selective forgetting procedure for DNNs trained with SGD, using an information theoretic formulation and exploiting the stability of SGD [18]. They proposed a forgetting mechanism which involves a shift in weight space, and addition of noise to the weights to destroy information. They also provide an upper bound on the amount of remaining information in the weights of the network after applying the forgetting procedure. We extend this framework to activations, and show that using an NTK based scrubbing procedure uniformly improves the scrubbing procedure in all metrics that they consider.

**Membership Inference Attacks** [37, 21, 19, 31, 34, 35, 32] try to guess if a particular sample was used for training a model. Since a model has forgotten only if an attacker cannot guess at better than chance level, these attacks serve as a good metric for measuring the quality of forgetting. In Figure 2 we construct a black-box membership inference attack similar to the shadow model training approach in [34]. Such methods relate to model inversion methods [13] which aim to gain information about the training data from the model output.

**Neural Tangent Kernel:** [23, 27] show that the training dynamics of a linearized version of a Deep Network — which are described by the so called NTK matrix — approximate increasingly better the actual training dynamics as the network width goes to infinity. [4, 28] extend the framework to convolutional

networks. [33] compute information-theoretic quantities using the closed form expressions for various quantities that can be derived in this settings. While we do use the infinite width assumption, we show that the same linearization framework and solutions are a good approximation of the network dynamics during fine-tuning, and use them to compute an optimal scrubbing procedure.

### 3 Out of the box forgetting

In this section, we derive an upper-bound for how much information can be extracted by an attacker that has black-box access to the model, that is, they can query the model with an image, and obtain the corresponding output. We will then use this to design a forgetting procedure (Section 4).

While the problem itself may seem trivial — can the relation  $A(\mathcal{D}_f, \mathcal{D}_r) = w$  be inverted to extract  $\mathcal{D}_f$ ? — it is made more complex by the fact that the algorithm is stochastic, and that the map may not be invertible, but still partially invertible, that is, only a subset of information about  $\mathcal{D}_f$  can be recovered. Hence, we employ a more formal information-theoretic framework, inspired by [16] and that in turns generalizes Differential Privacy [12]. We provide a-posteriori bounds which provide tighter answers, and allow one to design and benchmark scrubbing procedures even for very complex models such as deep networks, for which a-priori bounds would be impossible or vacuous. In this work, we focus on a-posteriori bounds for Deep Networks and use them to design a scrubbing procedure.

#### 3.1 Information Theoretic formalism

We start by modifying the framework of [16], developed for the weights of a network, to the final activations. We expect an adversary to use a readout function applied to the final activations. Given a set of images  $\mathbf{x} = (x_0, \dots, x_n)$ , we denote by  $f_w(\mathbf{x}) = (f(x_0), \dots, f(x_m))$  the concatenation of their respective final activations. Let  $\mathcal{D}_f$  be the set of training data to forget, and let  $y$  be some function of  $\mathcal{D}_f$  that an attacker wants to reconstruct (*i.e.*,  $y$  is some piece of information regarding the samples). To keep the notation uncluttered, we write  $S_{\mathcal{D}_f}(w) = S(w)$  for the scrubbing procedure to forget  $\mathcal{D}_f$ . We then have the following Markov chain

$$y \leftarrow \mathcal{D}_f \longrightarrow w \longrightarrow S(w) \longrightarrow f_{S(w)}(\mathbf{x})$$

connecting all quantities. Using the Data Processing Inequality [11] we have the following inequalities:

$$\underbrace{I(y; f_{S(w)}(\mathbf{x}))}_{\text{Recovered information}} \leq \underbrace{I(\mathcal{D}_f; f_{S(w)}(\mathbf{x}))}_{\text{Black-box upper bound}} \leq \underbrace{I(\mathcal{D}_f; S(w))}_{\text{White box upper-bound}} \quad (1)$$

where  $I(x; y)$  denotes the Shannon Mutual Information between random variables  $x$  and  $y$ . Bounding the last term — which is a general bound on how

much information an attacker with full access to the weights could extract — is the focus of [16]. In this work, we also consider the case where the attacker can only access the final activations, and hence focus on the central term. As we will show, if the number of queries is bounded then the central term provides a sharper bound compared to the black-box case.

### 3.2 Bound for activations

The mutual information in the central term is difficult to compute,<sup>1</sup> but in our case has a simple upper-bound:

**Lemma 1 (Computable bound on mutual information).** *We have the following upper bound:*

$$I(\mathcal{D}_f; f_{S(w)}(\mathbf{x})) \leq \mathbb{E}_{\mathcal{D}_f} [\text{KL}(p(f_{S(w)}(\mathbf{x})|\mathcal{D}_f \cup \mathcal{D}_r) \| p(f_{S_0(w)}(\mathbf{x})|\mathcal{D}_r))], \quad (2)$$

where  $p(f_{S(w)}(\mathbf{x})|\mathcal{D} = \mathcal{D}_f \cup \mathcal{D}_r)$  is the distribution of activations after training on the complete dataset  $\mathcal{D}_f \sqcup \mathcal{D}_r$  and scrubbing. Similarly,  $p(f_{S_0(w)}(\mathbf{x})|\mathcal{D} = \mathcal{D}_r)$  is the distribution of possible activations after training only on the data to retain  $\mathcal{D}_r$  and applying a function  $S_0$  (that does not depend on  $\mathcal{D}_f$ ) to the weights.

The lemma introduces the important notion that we can estimate how much information we erased by comparing the activations of our model with the activations of a reference model that was trained in the same setting, but without  $\mathcal{D}_f$ . Clearly, if the activations after scrubbing are identical to the activations of a model that has never seen  $\mathcal{D}_f$ , they cannot contain information about  $\mathcal{D}_f$ .

We now want to convert this bound in a more practical expected information gain per query. This is not yet trivial due to them stochastic dependency of  $w$  on  $\mathcal{D}$ : based on the random seed  $\epsilon$  used to train, we may obtain very different weights for the same dataset. Reasoning in a way similar to that used to obtain the “local forgetting bound” of [16], we have:

**Lemma 2.** *Write a stochastic training algorithm  $A$  as  $A(\mathcal{D}, \epsilon)$ , where  $\epsilon$  is the random seed and  $A(\mathcal{D}, \epsilon)$  is a deterministic function. Then,*

$$I(\mathcal{D}_f; f_{S(w)}(\mathbf{x})) \leq \mathbb{E}_{\mathcal{D}_f, \epsilon} [\text{KL}(p(f_{S(w_{\mathcal{D}})}(\mathbf{x})) \| p(f_{S_0(w_{\mathcal{D}_r}})}(\mathbf{x})))] \quad (3)$$

where we call  $w_{\mathcal{D}} = A(\mathcal{D}, \epsilon)$  the deterministic result of training on the dataset  $\mathcal{D}$  using random seed  $\epsilon$ . The probability distribution inside the KL accounts only for the stochasticity of the scrubbing map  $S(w_{\mathcal{D}})$  and the baseline  $S_0(w_{\mathcal{D}_r})$ .

The expression above is general. To gain some insight, it is useful to write it for a special case where scrubbing is performed by adding Gaussian noise.

<sup>1</sup> Indeed, it is even difficult to define, as the quantities at play,  $\mathcal{D}_f$  and  $\mathbf{x}$ , are fixed values, but that problem has been addressed by [3, 2] and we do not consider it here.

### 3.3 Close form bound for Gaussian scrubbing

We start by considering a particular class of scrubbing functions  $S(w)$  in the form

$$S(w) = h(w) + n, \quad n \sim N(0, \Sigma(w, \mathcal{D})) \quad (4)$$

where  $h(w)$  is a deterministic shift (that depends on  $\mathcal{D}$  and  $w$ ) and  $n$  is Gaussian noise with a given covariance (which may also depend on  $w$  and  $\mathcal{D}$ ). We consider a baseline  $S_0(w) = w + n'$  in a similar form, where  $n' \sim N(0, \Sigma_0(w, \mathcal{D}_r))$ .

Assuming that the covariance of the noise is relatively small, so that  $f_{h(w)}(\mathbf{x})$  is approximately linear in  $w$  in a neighborhood of  $h(w)$  (we drop  $\mathcal{D}$  without loss of generality), we can easily derive the following approximation for the distribution of the final activations after scrubbing for a given random seed  $\epsilon$ :

$$f_{S(w_{\mathcal{D}})}(\mathbf{x}) \sim N(f_{h(w_{\mathcal{D}})}(\mathbf{x}), \nabla_w f_{h(w_{\mathcal{D}})}(\mathbf{x}) \Sigma \nabla_w f_{h(w_{\mathcal{D}})}(\mathbf{x})^T) \quad (5)$$

where  $\nabla_w f_{h(w_{\mathcal{D}})}(\mathbf{x})$  is the matrix whose row is the gradient of the activations with respect to the weights, for each sample in  $\mathbf{x}$ . Having an explicit (Gaussian) distribution for the activations, we can plug it in Lemma 2 and obtain:

**Proposition 1.** *For a Gaussian scrubbing procedure, we have the bounds:*

$$I(\mathcal{D}_f; S(w)) \leq \mathbb{E}_{\mathcal{D}_f, \epsilon} [\Delta w^T \Sigma_0^{-1} \Delta w + d(\Sigma, \Sigma_0)] \quad (\text{white-box}) \quad (6)$$

$$I(\mathcal{D}_f; f_{S(w)}(\mathbf{x})) \leq \mathbb{E}_{\mathcal{D}_f, \epsilon} [\Delta f^T \Sigma'_{\mathbf{x}}{}^{-1} \Delta f + d(\Sigma_{\mathbf{x}}, \Sigma'_{\mathbf{x}})] \quad (\text{black-box}) \quad (7)$$

where for a  $k \times k$  matrix  $\Sigma$  we set  $d(\Sigma, \Sigma_0) := \text{tr}(\Sigma \Sigma_0^{-1}) + \log |\Sigma \Sigma_0^{-1}| - k$ , and

$$\begin{aligned} \Delta w &:= h(w_{\mathcal{D}}) - w_{\mathcal{D}_r}, & \Delta f &:= f_{h(w_{\mathcal{D}})}(\mathbf{x}) - f_{w_{\mathcal{D}_r}}(\mathbf{x}) \\ \Sigma_{\mathbf{x}} &:= \nabla_w f_{h(w_{\mathcal{D}})}(\mathbf{x}) \Sigma \nabla_w f_{h(w_{\mathcal{D}})}(\mathbf{x})^T, & \Sigma'_{\mathbf{x}} &:= \nabla_w f_{w_{\mathcal{D}_r}}(\mathbf{x}) \Sigma_0 \nabla_w f_{w_{\mathcal{D}_r}}(\mathbf{x})^T. \end{aligned}$$

**Avoiding curse of dimensionality:** Proposition 1. Comparing eq. (6) and eq. (7), we see that the bound in eq. (6) involves variables of the same dimension as the number of weights, while eq. (7) scales with the number of query points. Hence, for highly overparametrized models such as DNNs, we expect that the black-box bound in eq. (7) will be much smaller if the number of queries is bounded, which indeed is what we observe in the experiments (Figure 3).

**Blessing of the null-space:** The white-box bound depends on the difference  $\Delta w$  in weight space between the scrubbed model and the reference model  $w_{\mathcal{D}_r}$ , while the black-box bound depends on the distance  $\Delta f$  in the activations. As we mentioned in Section 1, over-parametrized models such as deep networks have a large null-space of weights with similar final activations. It may hence happen that even if  $\Delta w$  is large,  $\Delta f$  may still be small (and hence the bound in eq. (7) tighter) as long as  $\Delta w$  lives in the null-space. Indeed, we often observe this to be the case in our experiments.

**Adversarial queries:** Finally, this should not lead us to think that whenever the activations are similar, little information can be extracted. Notice that the

relevant quantity for the black box bound is  $\Delta f(J_{\mathbf{x}}\Sigma_0J_{\mathbf{x}}^T)^{-1}\Delta f^T$ , which involves also the gradient  $J_{\mathbf{x}} = \nabla_w f_{w_{\mathcal{D}_r}}(\mathbf{x})$ . Hence, if an attacker crafts an adversarial query  $\mathbf{x}$  such that its gradient  $J_{\mathbf{x}}$  is small, they may be able to extract a large amount of information even if the activations are close to each other. In particular, this happens if the gradient of the samples lives in the null-space of the reference model, but not in that of the scrubbed model. In Figure 3 (right), we show that indeed different images can extract different amount of information.

## 4 An NTK-inspired forgetting procedure

We now introduce a new scrubbing procedure, which aims to minimize both the white-box and black-box bounds of Proposition 1. It relates to the one introduced in [16, 17], but it enjoys better numerical properties and can be computed without approximations (Section 4.1). In Section 5 we show that it gives better results under all commonly used metrics.

The main intuition we exploit is that most networks commonly used are fine-tuned from pre-trained networks (*e.g.*, on ImageNet), and that the weights do not move much during fine-tuning on  $\mathcal{D} = \mathcal{D}_r \cup \mathcal{D}_f$  will remain close to the pre-trained values. In this regime, the network activations may be approximated as a linear function of the weights. This is inspired by a growing literature on the so called Neural Tangent Kernel, which posits that large networks during training evolve in the same way as their linear approximation [27]. Using the linearized model we can derive an analytical expression for the optimal forgetting function, which we validate empirically. However, we observe this to be misaligned with weights actually learned by SGD, and introduce a very simple “isocoles trapezium” trick to realign the solutions (see Supplementary Material).

Using the same notation as [27], we linearize the final activations around the pre-trained weights  $\theta_0$  as:

$$f_t^{\text{lin}}(x) \equiv f_0(x) + \nabla_{\theta} f_0(x)|_{\theta=\theta_0} w_t$$

where  $w_t = \theta_t - \theta_0$  and gives the following expected training dynamics for respectively the weights and the final activations:

$$\dot{w}_t = -\eta \nabla_{\theta} f_0(\mathcal{D})^T \nabla_{f_t^{\text{lin}}(\mathcal{D})} \mathcal{L} \quad (8)$$

$$\dot{f}_t^{\text{lin}}(x) = -\eta \Theta_0(x, \mathcal{D}) \nabla_{f_t^{\text{lin}}(\mathcal{D})} \mathcal{L} \quad (9)$$

The matrix  $\Theta_0 = \nabla_{\theta} f(\mathcal{D}) \nabla_{\theta} f(\mathcal{D})^T$  of size  $c|\mathcal{D}| \times c|\mathcal{D}|$ , where  $c$  the number of classes, is called the Neural Tangent Kernel (NTK) matrix [27, 23]. Using this dynamics, we can approximate in closed form the final training point when training with  $\mathcal{D}$  and  $\mathcal{D}_r$ , and compute the optimal “one-shot forgetting” vector to jump from the weights  $w_{\mathcal{D}}$  that have been obtained by training on  $\mathcal{D}$  to the weights  $w_{\mathcal{D}_r}$  that would have been obtained training on  $\mathcal{D}_r$  alone:



**Proposition 2.** *Assuming an  $L_2$  regression loss,<sup>2</sup> the optimal scrubbing procedure under the NTK approximation is given by*

$$h_{NTK}(w) = w + P \nabla f_0(\mathcal{D}_f)^T M V \quad (10)$$

where  $\nabla f_0(\mathcal{D}_f)^T$  is the matrix whose columns are the gradients of the sample to forget, computed at  $\theta_0$  and  $w = A(\mathcal{D}_r \cup \mathcal{D}_f)$ .  $P = I - \nabla f_0(\mathcal{D}_r)^T \Theta_{rr}^{-1} \nabla f_0(\mathcal{D}_r)$  is a projection matrix, that projects the gradients of the samples to forget  $\nabla f_0(\mathcal{D}_f)$  onto the orthogonal space to the space spanned by the gradients of all samples to retain. The terms  $M = [\Theta_{ff} - \Theta_{rf}^T \Theta_{rr}^{-1} \Theta_{rf}]^{-1}$  and  $V = [(Y_f - f_0(\mathcal{D}_f)) + \Theta_{rf}^T \Theta_{rr}^{-1} (Y_r - f_0(\mathcal{D}_r))]$  re-weight each direction before summing them together.

Given this result, our proposed scrubbing procedure is:

$$S_{NTK}(w) = h_{NTK}(w) + n \quad (11)$$

where  $h_{NTK}(w)$  is as in eq. (10), and we use noise  $n \sim N(0, \lambda F^{-1})$  where  $F = F(w)$  is the Fisher Information Matrix computed at  $h_{NTK}(w)$  using  $\mathcal{D}_r$ . The noise model is as in [16], and is designed to increase robustness to mistakes due to the linear approximation.

#### 4.1 Relation between NTK and Fisher forgetting

In [16] and [17], a different forgetting approach is suggested based on either the Hessian or the Fisher Matrix at the final point: assuming that the solutions  $w(\mathcal{D}_r)$  and  $w(\mathcal{D})$  of training with and without the data to forget are close and that they are both minima of their respective loss, one may compute the shift to to jump from one minimum to the other of a slightly perturbed loss landscape. The resulting “scrubbing shift”  $w \mapsto h(w)$  relates to the newton update:

$$h(w) = w - H(\mathcal{D}_r)^{-1} \nabla_w L_{\mathcal{D}_r}(w). \quad (12)$$

In the case of an  $L_2$  loss, and using the NTK model, the Hessian is given by  $H(w) = \nabla f_0(\mathcal{D}_r)^T \nabla f_0(\mathcal{D}_r)$  which in this case also coincides with the Fisher Matrix [29]. To see how this relates to the NTK matrix, consider determining the convergence point of the linearized NTK model, that for an  $L_2$  regression is given by  $w^* = w_0 + \nabla f_0(\mathcal{D}_r)^+ \mathcal{Y}$ , where  $\nabla f_0(\mathcal{D}_r)^+$  denotes the matrix pseudo-inverse, and  $\mathcal{Y}$  denotes the regression targets. If  $\nabla f_0(\mathcal{D}_r)$  is a tall matrix (more samples in the dataset than parameters in the network), then the pseudo-inverse is  $\nabla f_0(\mathcal{D}_r)^+ = H^{-1} \nabla f_0(\mathcal{D}_r)^T$ , recovering the scrubbing procedure considered by [16, 17]. However, if the matrix is wide (more parameters than samples in the network, as is often the case in Deep Learning), the Hessian is not invertible, and the pseudo-inverse is instead given by  $\nabla f_0(\mathcal{D}_r)^+ = \nabla f_0(\mathcal{D}_r)^T \Theta^{-1}$ , leading

<sup>2</sup> This assumption is to keep the expression simple, in the Supplementary Material we show the corresponding expression for a softmax classification loss.

to our proposed procedure. In general, when the model is over-parametrized there is a large null-space of weights that do not change the final activations or the loss. The degenerate Hessian is not informative of where the network will converge in this null-space, while the NTK matrix gives the exact point.

## 5 Experiments

### 5.1 Datasets

We report experiments on smaller versions of CIFAR-10 [26] and Lacuna-10 [16], a dataset derived from the VGG-Faces [8] dataset. We obtain the small datasets using the following procedure: we randomly sample 500 images (100 images from each of the first 5 classes) from the training/test set of CIFAR-10 and Lacuna-10 to obtain the small-training/test respectively. We also sample 125 images from the training set (5 classes  $\times$  25 images) to get the validation set. So, in short, we have 500 ( $5 \times 100$ ) examples for training and testing respectively, and 125 ( $5 \times 25$ ) examples for validation. On both the datasets we choose to forget 25 random samples (5% of the dataset). Without loss of generality we choose to forget samples from class 0.

### 5.2 Models and Training

We use All-CNN [36] (to which we add batch-normalization before non-linearity) and ResNet-18 [20] as the deep neural networks for our experiments. We pre-train the models on CIFAR-100/Lacuna-100 and then fine-tune them (all the weights) on CIFAR-10/Lacuna-10. We pre-train using SGD for 30 epochs with a learning rate of 0.1, momentum 0.9 and weight decay 0.0005. Pre-training helps in improving the stability of SGD while fine-tuning. For fine-tuning we use a learning rate of 0.01 and weight decay 0.1. While applying weight decay we bias the weights with respect to the initialization. During training we always use a batch-size of 128 and fine-tune the models till zero training error. Also, during fine-tuning we do not update the running mean and variance of the batch-normalization parameters to simplify the training dynamics. We perform each experiment 3 times and report the mean and standard deviation.

### 5.3 Baselines

We consider three baselines for comparison: (i) **Fine-tune**, we fine-tune  $w(\mathcal{D})$  on  $\mathcal{D}_r$  (similar to catastrophic forgetting) and (ii) **Fisher forgetting** [16], we scrubs the weights by adding Gaussian noise using the inverse of the Fisher Information Matrix as covariance matrix, (iii) **Original** corresponds to the original model trained on the complete dataset ( $w(\mathcal{D})$ ) without any forgetting. We compare those, and our proposal, with optimal reference the model  $w(\mathcal{D}_r)$  trained from scratch on the retain set, that is, without using  $\mathcal{D}_f$  in the first place. Values read from this reference model corresponds to the **green** region in Figure 2 and represent the gold standard for forgetting: In those plots, an optimal algorithm should lie inside the green area.

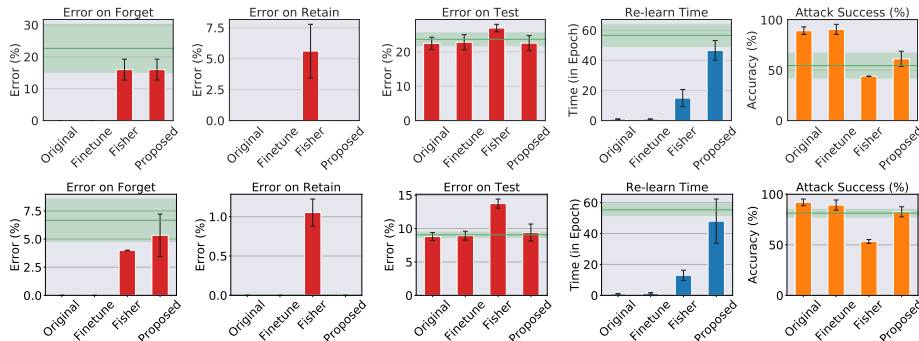


Fig. 2: Comparison of different models baselines (original, finetune) and forgetting methods (Fisher [16] and our NTK proposed method), using several readout functions ((**Top**) CIFAR and (**Bottom**) Lacuna). We benchmark them against a model that has never seen the data (the gold reference for forgetting): values (mean and standard deviation) measured from this models corresponds to the green region. Optimal scrubbing procedure should lie in the green region, or they will leak information about  $\mathcal{D}_f$ . We compute three read-out functions: (a) Error on forget set  $\mathcal{D}_f$ , (b) Error on retain set  $\mathcal{D}_r$ , (c) Error on test set  $\mathcal{D}_{\text{test}}$ . (d) Black-box membership inference attack: We construct a simple yet effective membership attack using the entropy of the output probabilities. We measures how often the attack model (using the activations of the scrubbed network) classify a sample belonging  $\mathcal{D}_f$  as a training sample rather than being fooled by the scrubbing. (e) Re-learn time for different scrubbing methods: How fast a scrubbed model learns the forgotten cohort when fine-tuned on the complete dataset. We measure the re-learn time as the first epoch when the loss on  $\mathcal{D}_f$  goes below a certain threshold.

#### 5.4 Readout Functions

We use multiple readout functions similar to [16]: (i) **Error on residual** (should be small), (ii) **Error on cohort to forget** (should be similar to the model re-trained from scratch on  $\mathcal{D}_r$ ), (iii) **Error on test set** (should be small), (iv) **Re-learn time**, measures how quickly a scrubbed model learns the cohort to forget, when fine-tuned on the complete data. Re-learn time (measured in epochs) is the first epoch when the loss during fine-tuning (the scrubbed model) falls below a certain threshold (loss of the original model (model trained on  $\mathcal{D}$ ) on  $\mathcal{D}_f$ ). (v) **Blackbox membership inference attack**: We construct a simple yet effective blackbox membership inference attack using the entropy of the output probabilities of the scrubbed model. Similar to the method in [34], we formulate the attack as a binary classification problem (class 1 - belongs to training set and class 0 - belongs to test set). For training the attack model (Support Vector Classifier with Radial Basis Function Kernel) we use the retain set ( $\mathcal{D}_r$ ) as class 1 and the test set as class 0. We test the success of the attack on the cohort to forget ( $\mathcal{D}_f$ ). Ideally, the attack accuracy for an optimally scrubbed model




Fig. 3: **Error-forgetting trade-off** Using the proposed scrubbing procedure, by changing the variance of the noise, we can reduce the remaining information in the weights (**white-box bound, left**) and activations (**black-box bound, center**). However, it comes at the cost of increasing the test error. Notice that the bound on activation is much sharper than the bound on error at the same accuracy. (**Right**) **Different samples leak different information.** An attacker querying samples from  $\mathcal{D}_f$  can gain much more information than querying unrelated images. This suggest that adversarial samples may be created to leak even more information.

should be the same as a model re-trained from scratch on  $\mathcal{D}_f$ , having a higher value implies incorrect (or no) scrubbing, while a lower value may result in Streisand Effect, (vi) **Remaining information in the weights** [16] and (vii) **Remaining information in the activations:** We compute an upper bound on the information the activations contain about the cohort to forget ( $\mathcal{D}_f$ ) (after scrubbing) when queried with images from different subsets of the data ( $\mathcal{D}_r, \mathcal{D}_r$ ).

## 5.5 Results

**Error readouts:** In Figure 2 (a-c), we compare error based readout functions for different forgetting methods. Our proposed method outperforms Fisher forgetting which incurs high error on the retain ( $\mathcal{D}_r$ ) and test ( $\mathcal{D}_{\text{Test}}$ ) set to attain the same level of forgetting. This is due the large distance between  $w(\mathcal{D})$  and  $w(\mathcal{D}_r)$  in weight space, which forces it to add too much noise to erase information about  $\mathcal{D}_f$ , and ends up also erasing information about the retain set  $\mathcal{D}_r$  (high error on  $\mathcal{D}_r$  in Figure 2). Instead, our proposed method first moves  $w(\mathcal{D})$  in the direction of  $w(\mathcal{D}_r)$ , thus minimizing the amount of noise to be added (Figure 1). Fine-tuning the model on  $\mathcal{D}_r$  (catastrophic forgetting) does not actually remove information from the weights and performs poorly on all the readout functions.

**Relearn time:** In Figure 2 (d), we compare the re-learn time for different methods. Re-learn time can be considered as a proxy for the information remaining in the weights about the cohort to forget ( $\mathcal{D}_f$ ) after scrubbing. We observe that the proposed method outperforms all the baselines which is in accordance with the previous observations (in Figure 2(a-c)).




Fig. 4: **Scrubbing brings activations closer to the target.** We plot the  $L_1$  norm of the difference between the final activations (post-softmax) of the target model trained only on  $\mathcal{D}_r$ , and models sampled along the line joining the original model  $w(D)$  ( $\alpha = 0$ ) and the proposed scrubbed model ( $\alpha = 1$ ). The distance between the activations decreases as we move along the scrubbing direction. The  $L_1$  distance is already low on the retain set ( $\mathcal{D}_r$ ) (red) as it corresponds to the data common to  $w(D)$  and  $w(\mathcal{D}_r)$ . However, the two models differ on the forget set ( $\mathcal{D}_f$ ) (blue) and we observe that the  $L_1$  distance decreases as move along the proposed scrubbing direction.

**Membership attacks:** In Figure 2 (e), we compare the robustness of different scrubbed models against blackbox membership inference attacks (attack aims to identify if the scrubbed model was ever trained on  $\mathcal{D}_f$ ). This can be considered as a proxy for the remaining information (about  $\mathcal{D}_f$ ) in the activations. We observe that attack accuracy for the proposed method lies in the optimal region (green), while Fine-tune does not forget ( $\mathcal{D}_f$ ), Fisher forgetting may result in Streisand effect which is undesirable.

**Closeness of activations:** In Figure 4, we show that the proposed scrubbing method brings the activations of the scrubbed model closer to retrain model (model retrained from scratch on  $\mathcal{D}_r$ ). We measure the closeness by computing:  $\mathbb{E}_{x \sim \mathcal{D}_f / \mathcal{D}_r} [\|f_{\text{scrubbed}}(x) - f_{\text{retrain}}(x)\|_1]$  along the scrubbing direction, where  $f_{\text{scrubbed}}(x)$ ,  $f_{\text{retrain}}(x)$  are the activations (post soft-max) of the proposed scrubbed and retrain model respectively. The distance between the activations on the cohort to forget ( $\mathcal{D}_f$ ) decreases as we move along the scrubbing direction and achieves a minimum value at the scrubbed model, while it almost remains constant on the retain set. Thus, the activations of the scrubbed model shows desirable behaviour on both  $\mathcal{D}_r$  and  $\mathcal{D}_f$ .

**Error-forgetting trade-off** In Figure 3, we plot the trade-off between the test error and the remaining information in the weights and activations respectively by changing the scale of the variance of Fisher noise. We can reduce the remaining information but this comes at the cost of increasing the test error. We observe that the black-box bound on the information accessible with one query is much tighter than the white box bound at the same accuracy (compare left and center plot x-axes). Finally, in (right), we show that query samples belonging to the cohort to be forgotten ( $\mathcal{D}_f$ ) leaks more information about the  $\mathcal{D}_f$  rather

than the retain/test set, proving that indeed carefully selected samples are more informative to an attacker than random samples.

## 6 Discussion

Recent work [3, 16, 17] has started providing insights on both the amount of information that can be extracted from the weights about a particular cohort of the data used for training, as well as give constructive algorithms to “selectively forget.” Note that forgetting alone could be trivially obtained by replacing the model with a random vector generator, obviously to the detriment of performance, or by retraining the model from scratch, to the detriment of (training time) complexity. In some cases, the data to be retained may no longer be available, so the latter may not even be an option.

We introduce a scrubbing procedure based on the NTK linearization which is designed to minimize both a white-box bound (which assumes the attacker has the weights), and a newly introduced black-box bound. The latter is a bound on the information that can be obtained about a cohort using only the observed input-output behavior of the network. This is relevant when the attacker performs a bounded number of queries. If the attacker is allowed infinitely many observations, the matter of whether the black-box and white-box attack are equivalent remains open: Can an attacker always craft sufficiently exciting inputs so that the exact values of all the weights can be inferred? An answer would be akin to a generalized “Kalman Decomposition” for deep networks. This alone is an interesting open problem, as it has been pointed out recently that good “clone models” can be created by copying the response of a black-box model on a relatively small set of exciting inputs, at least in the restricted cases where every model is fine-tuned from a common pre-trained models [25].

While our bounds are tighter than others proposed in the literature, our model has limitations. Chiefly, computational complexity. The main source of computational complexity is computing and storing the matrix  $P$  in eq. (10), which naively would require  $O(c^2|\mathcal{D}_r|^2)$  memory and  $O(|w| \cdot c^2|\mathcal{D}_r|^2)$  time. For this reason, we conduct experiments on a relatively small scale, sufficient to validate the theoretical results, but our method is not yet scalable to production level. However, we notice that  $P$  is the projection matrix on the orthogonal of the subspace spanned by the training samples, and there is a long history [7] of numerical methods to incrementally compute such operator incrementally and without storing it fully in memory. We leave these promising option to scale the method as subject of future investigation.

## Acknowledgement

We would like to thank the anonymous reviewers for their feedback and suggestions. This work is supported by ARO W911NF-17-1-0304 and ONR N00014-19-1-2229.

## References

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., Zhang, L.: Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 308–318. ACM (2016)
2. Achille, A., Soatto, S.: Emergence of invariance and disentanglement in deep representations. *The Journal of Machine Learning Research* **19**(1), 1947–1980 (2018)
3. Achille, A., Soatto, S.: Where is the Information in a Deep Neural Network? arXiv e-prints arXiv:1905.12213 (May 2019)
4. Arora, S., Du, S.S., Hu, W., Li, Z., Salakhutdinov, R.R., Wang, R.: On exact computation with an infinitely wide neural net. In: Advances in Neural Information Processing Systems. pp. 8139–8148 (2019)
5. Baumhauer, T., Schöttle, P., Zeppelzauer, M.: Machine unlearning: Linear filtration for logit-based classifiers. arXiv preprint arXiv:2002.02730 (2020)
6. Bourtole, L., Chandrasekaran, V., Choquette-Choo, C., Jia, H., Travers, A., Zhang, B., Lie, D., Papernot, N.: Machine unlearning. arXiv preprint arXiv:1912.03817 (2019)
7. Bunch, J.R., Nielsen, C.P., Sorensen, D.C.: Rank-one modification of the symmetric eigenproblem. *Numerische Mathematik* **31**(1), 31–48 (1978)
8. Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: Vggface2: A dataset for recognising faces across pose and age. In: International Conference on Automatic Face and Gesture Recognition (2018)
9. Cao, Y., Yang, J.: Towards making systems forget with machine unlearning. In: 2015 IEEE Symposium on Security and Privacy. pp. 463–480. IEEE (2015)
10. Chaudhuri, K., Monteleoni, C., Sarwate, A.D.: Differentially private empirical risk minimization. *Journal of Machine Learning Research* **12**(Mar), 1069–1109 (2011)
11. Cover, T.M., Thomas, J.A.: Elements of information theory. John Wiley & Sons (2012)
12. Dwork, C., Roth, A., et al.: The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* **9**(3–4), 211–407 (2014)
13. Fredrikson, M., Jha, S., Ristenpart, T.: Model inversion attacks that exploit confidence information and basic countermeasures. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. pp. 1322–1333. ACM (2015)
14. Ginart, A., Guan, M., Valiant, G., Zou, J.Y.: Making ai forget you: Data deletion in machine learning. In: Advances in Neural Information Processing Systems. pp. 3513–3526 (2019)
15. Giordano, R., Stephenson, W., Liu, R., Jordan, M., Broderick, T.: A swiss army infinitesimal jackknife. In: The 22nd International Conference on Artificial Intelligence and Statistics. pp. 1139–1147 (2019)
16. Golatkar, A., Achille, A., Soatto, S.: Eternal sunshine of the spotless net: Selective forgetting in deep networks (2019)
17. Guo, C., Goldstein, T., Hannun, A., van der Maaten, L.: Certified data removal from machine learning models. arXiv preprint arXiv:1911.03030 (2019)
18. Hardt, M., Recht, B., Singer, Y.: Train faster, generalize better: Stability of stochastic gradient descent. arXiv preprint arXiv:1509.01240 (2015)
19. Hayes, J., Melis, L., Danezis, G., De Cristofaro, E.: Logan: Membership inference attacks against generative models. *Proceedings on Privacy Enhancing Technologies* **2019**(1), 133–152 (2019)

20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
21. Hitaj, B., Ateniese, G., Perez-Cruz, F.: Deep models under the gan: information leakage from collaborative deep learning. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. pp. 603–618. ACM (2017)
22. Izzo, Z., Smart, M.A., Chaudhuri, K., Zou, J.: Approximate data deletion from machine learning models: Algorithms and evaluations. arXiv preprint arXiv:2002.10077 (2020)
23. Jacot, A., Gabriel, F., Hongler, C.: Neural tangent kernel: Convergence and generalization in neural networks. In: Advances in neural information processing systems. pp. 8571–8580 (2018)
24. Koh, P.W., Liang, P.: Understanding black-box predictions via influence functions. arXiv preprint arXiv:1703.04730 (2017)
25. Krishna, K., Tomar, G.S., Parikh, A.P., Papernot, N., Iyyer, M.: Thieves on sesame street! model extraction of bert-based apis (2019)
26. Krizhevsky, A., et al.: Learning multiple layers of features from tiny images. Tech. rep., Citeseer (2009)
27. Lee, J., Xiao, L., Schoenholz, S., Bahri, Y., Novak, R., Sohl-Dickstein, J., Pennington, J.: Wide neural networks of any depth evolve as linear models under gradient descent. In: Advances in neural information processing systems. pp. 8570–8581 (2019)
28. Li, Z., Wang, R., Yu, D., Du, S.S., Hu, W., Salakhutdinov, R., Arora, S.: Enhanced convolutional neural tangent kernels. arXiv preprint arXiv:1911.00809 (2019)
29. Martens, J.: New insights and perspectives on the natural gradient method. arXiv preprint arXiv:1412.1193 (2014)
30. Mirzasoleiman, B., Karbasi, A., Krause, A.: Deletion-robust submodular maximization: Data summarization with “the right to be forgotten”. In: International Conference on Machine Learning. pp. 2449–2458 (2017)
31. Pyrgelis, A., Troncoso, C., De Cristofaro, E.: Knock knock, who’s there? membership inference on aggregate location data. arXiv preprint arXiv:1708.06145 (2017)
32. Sablayrolles, A., Douze, M., Ollivier, Y., Schmid, C., Jégou, H.: White-box vs black-box: Bayes optimal strategies for membership inference. arXiv preprint arXiv:1908.11229 (2019)
33. Schwartz-Ziv, R., Alemi, A.A.: Information in infinite ensembles of infinitely-wide neural networks. arXiv preprint arXiv:1911.09189 (2019)
34. Shokri, R., Shmatikov, V.: Privacy-preserving deep learning. In: Proceedings of the 22nd ACM SIGSAC conference on computer and communications security. pp. 1310–1321. ACM (2015)
35. Song, C., Ristenpart, T., Shmatikov, V.: Machine learning models that remember too much. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. pp. 587–601. ACM (2017)
36. Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M.: Striving for simplicity: The all convolutional net. arXiv preprint arXiv:1412.6806 (2014)
37. Truex, S., Liu, L., Gursoy, M.E., Yu, L., Wei, W.: Demystifying membership inference attacks in machine learning as a service. IEEE Transactions on Services Computing (2019)