

An Information Theoretic Approach to Machine Unlearning

Anonymous authors

Paper under double-blind review

Abstract

To comply with AI and data regulations, the need to forget private or copyrighted information from trained machine learning models is increasingly important. The key challenge in unlearning is forgetting the necessary data in a timely manner, while preserving model performance. In this work, we address the zero-shot unlearning scenario, whereby an unlearning algorithm must be able to remove data given only a trained model and the data to be forgotten. We explore unlearning from an information theoretic perspective, connecting the influence of a sample to the information gain a model receives by observing it. From this, we derive a simple but principled zero-shot unlearning method based on the geometry of the model. Our approach takes the form of minimising the gradient of a learned function with respect to a small neighbourhood around a target forget point. This induces a smoothing effect, causing forgetting by moving the boundary of the classifier. We explore the intuition behind why this approach can jointly unlearn forget samples while preserving general model performance through a series of low-dimensional experiments. We perform extensive empirical evaluation of our method over a range of contemporary benchmarks, verifying that our method is competitive with state-of-the-art performance under the strict constraints of zero-shot unlearning.

1 Introduction

Regulations such as the General Data Protection Regulation (GDPR) enshrine an individual’s data autonomy rights, including the right to be forgotten. While deleting an entry from a database is relatively straightforward, removing the influence of that data from a trained model is a challenging open problem. The process of removal, referred to as *unlearning*, is difficult for several reasons. It is known that neural networks memorise instance-level information (Arpit et al., 2017; Zhang et al., 2021; Feldman, 2020), and it is practically intractable to ascribe parameter changes to a specific training sample post-hoc (Kurmanji et al., 2023). At its core, unlearning is a multi-objective optimization problem with three key desiderata. An effective unlearning algorithm must remove the influence of the selected subset of data, maintain model performance on retained data, and minimise computational cost. These goals are antagonistic since inducing forgetting inevitably disrupts the model’s performance, thus balancing these objectives is key. Naïvely, one may achieve perfect forgetting by retraining a model on the training data sans the forget samples every time there is a forget request. However, this is prohibitively expensive thus violating the third desideratum.

Many of the unlearning methods proposed are effective, however they make strong assumptions about the problem setting that simplifies the task considerably. Primarily, existing methods typically assume access to all, or a subset of, the training data. This data is used in different ways, such as to fine-tune the model post-forgetting (Chundawat et al., 2023a; Graves et al., 2021), or to conduct parameter importance calculations after the initial training period (Foster et al., 2023). In reality, there are many reasons why this data could be unavailable, such as cost of storage, limited duration access to datasets, or an oversight in considering machine unlearning during model development. As such, Chundawat et al. (2023b) introduces a novel problem setting for unlearning, termed *zero-shot* (ZS) unlearning, whereby only the data to be forgotten and the trained model are available (Figure 1). This is extremely challenging, since the remaining data is not available to protect model performance, and thus more delicate methods are required. Insightful treatment of the ZS scenario can be found in Chen et al. (2023), where unlearning is formulated as reconstructing a

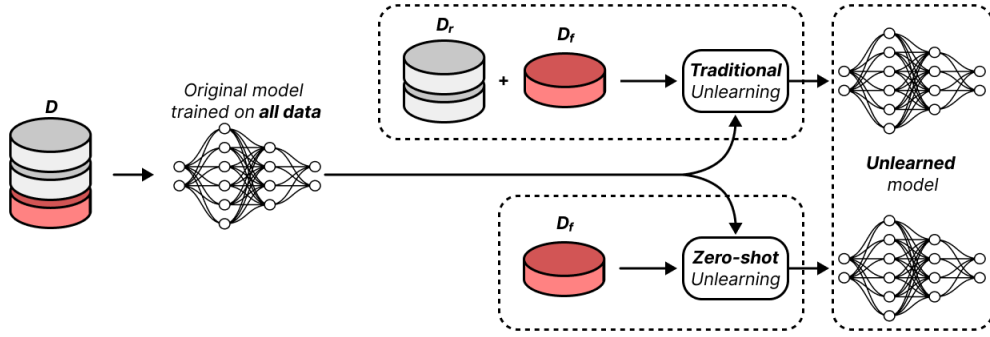


Figure 1: Visualization of the zero-shot unlearning scenario. Contrary to traditional unlearning there is no access to, or prior knowledge of, any data other than the forget set or the model at any point beyond its current state. These constraints make the problem considerably more challenging.

decision boundary that could be reasonably learnt by a model trained without the forget data, achieved through learning the nearest false label for each forget sample.

In this work, we approach ZS unlearning from an information theoretic perspective. Golatkar et al. (2020) consider information leakage when observing model weights, whereas we consider the information gained by a model by training on a given sample. Data points offer a classifier different amounts of information gain when included in the training data (Lindley, 1956; Houlby et al., 2011). If a data point can be inferred from other training data, then it offers little information gain (Jeong and Qiu, 2018). In ZS unlearning, knowing the contribution a sample has made to a model is hard as we have access to only the model and the sample to be forgotten. We postulate that the information gain of a sample can serve as a good indicator for how much influence it has over a model. From this hypothesis, we derive a principled loss to induce forgetting based on how the geometry of the problem space changes with respect to the softmax classifier. If a sample offers little information gain then it may be inferred from other data, therefore the classifier’s output should change minimally over similar samples. This minimal change can be effectively measured via the curvature of the model should be low in the region surrounding such a point. In contrast, a high information gain sample cannot be inferred, and therefore one would expect they lie in high curvature regions of space. We present Just in Time (JiT) unlearning, a novel ZS unlearning algorithm based on minimising the gradients of a classifier with respect to local neighbourhood around each forget sample. We show in low dimensional experiments that removing training samples from low-curvature regions yields minimal changes to the learned decision boundary, whereas removing samples from high curvature regions has a more pronounced effect. We therefore propose that, with reference to Feldman (2020), samples with low information gain may be predicted with more generalised knowledge and thus do not infringe on privacy. In contrast, samples with large information gain have significant impact on the learned classifier, are more likely memorised, and do infringe on privacy. We demonstrate empirically that following these principles, JiT causes the removal of influence from the forget set while preserving generalisation performance across the wider space.

Our primary contributions are as follows:

- To the best of our knowledge, JiT is the first unlearning algorithm to be directly informed by the information gain of a sample.
- We provide extensive empirical analysis of the geometry of JiT unlearning in low dimensions.
- We show our method is competitive with existing SOTA in the zero-shot domain.

2 Related Work

Information theory is concerned with the transmission, quantification, and storage of information (Shannon, 1948), and has seen widespread use in machine learning. Most relevant here is its use in determining the

information gain of an experiment Lindley (1956). This notion has seen uses such as determining splits in decision trees (Quinlan, 1986), and active learning (Tong and Koller, 2001; Houlisby et al., 2011). Here we use this concept as a proxy to a training sample’s influence on a learned function.

Machine unlearning was first introduced in Cao and Yang (2015), and a probabilistic perspective of unlearning was explored in Ginart et al. (2019); Sekhari et al. (2021); Gupta et al. (2021); Neel et al. (2021); Triantafillou et al. (2023). Here we focus on post-hoc unlearning methods that operate on models that are already trained. Methods exist that alter the initial training scheme (Bourtoule et al., 2021; Mehta et al., 2022; Shah et al., 2023), but these are considered out of scope as they do not satisfy the ZS problem constraints.

Current SOTA methods rely on accessing all or a subset of the original dataset that is *not* to be forgotten (i.e. the retain set), thus violating the ZS constraints. Bad Teacher unlearning (Chundawat et al., 2023a) and SCRUB (Kurmanji et al., 2023) leverage a student-teacher framework while Amnesiac unlearning (Graves et al., 2021) trains with randomised labels for forget data before fine-tuning on the retained data to repair the model. UNSIR (Tarun et al., 2023) learns an error-maximising noise to induce forgetting of the necessary data, before also employing a finetuning step. Warnecke et al. (2021) minimise the divergence in model output over a sample and its noisy perturbations and then finetune. A key limitation of these methods is that protecting model performance necessitates access to retain-set data for the entire duration of the model’s lifetime. To address this, Golatkar et al. (2020) and Foster et al. (2023) propose methods that do not require fine-tuning or repair steps. Golatkar et al. (2020) derives an unlearning algorithm that minimises information gained about the training data when observing model weights. However, this scales quadratically with dataset size and often performs considerably worse than state-of-the-art (Tarun et al., 2023; Foster et al., 2023). Selective Synaptic Dampening is a scalable retrain-free approach, based on inducing forgetting by selectively dampening parameters that are disproportionately important to the forget-set (Foster et al., 2023). This requires access to the whole dataset at least once, to calculate the importance over the retained data. Chundawat et al. (2023b) introduces two methods to address ZS unlearning. The first method, an extension of Tarun et al. (2023), replaces the repair step with an error-minimising noise. The second utilises a generator network and an attention loss to distil knowledge from an expert teacher, with a band-pass filter preventing the flow of knowledge for specific classes. Both methods are slow, do not scale well to large problem spaces, and can only forget entire classes. Chen et al. (2023) present boundary shrinking and boundary expanding. Shrinking causes unlearning by training over the nearest false label for forget samples, found via a fast gradient sign attack (Goodfellow et al., 2014). While performant, shrinking scales poorly with model and input size. Boundary expanding is faster but less performant, remapping forget samples by training them to fit a new output neuron, before removing the neuron leaving the forget samples in high entropy states.

Membership inference attacks (MIA) are a way of measuring information leakage of a machine learning model (Shokri et al., 2017). An auxiliary model (e.g. a logistic regression) is trained to infer whether a given data point was included in a model’s training data. MIAs are used as typical evaluation metrics in machine unlearning; if a MIA cannot recognise a forgotten sample as an element from the train set, then this presents empirical evidence that the sample has indeed been forgotten.

3 Preliminaries

We introduce the notation for machine unlearning in a supervised classification task, consistent with the approach outlined in Chen et al. (2023). Consider some input space $\mathcal{X} \subset \mathbb{R}^d$ and some output label space $\mathcal{Y} \subset \mathbb{R}^c$, where d is the dimensionality of the input and c is the number of classes. We define a training dataset $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N \subseteq \mathcal{X} \times \mathcal{Y}$, where \mathbf{x}_i is a training input sample with the label \mathbf{y}_i . We denote a subset $\mathcal{D}_f \subseteq \mathcal{D}$ as the forget set, and $\mathcal{D}_r = \mathcal{D} \setminus \mathcal{D}_f$ as the retain set.

Let $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ be a neural network, with parameters θ . We assume that f_θ is well trained and generalises to in distribution test samples well. The objective of ZS unlearning is, given only the model f_θ trained on \mathcal{D} , to remove the influence of \mathcal{D}_f from the learned model such that the unlearned model $f_{\theta'}$ is approximately equivalent to a model retrained on only \mathcal{D}_r which we define as the optimal solution, f_{θ^*} . Since direct access to f_{θ^*} is, by definition of the unlearning problem, impossible, existing works construct approximate heuristics to induce forgetting and use a membership inference attack to measure forgetting. These attacks typically evaluate the difference in the output distributions of the model over train and test samples.

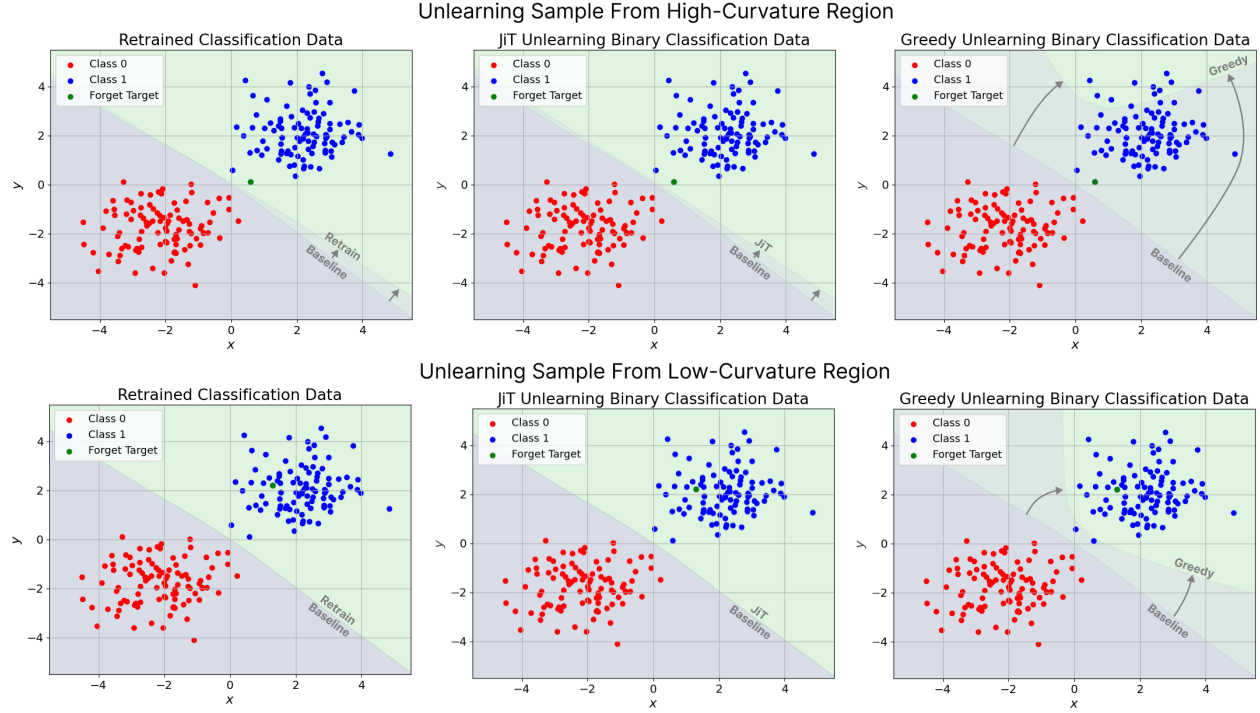


Figure 2: Demonstration of how the boundary of a classifier moves during unlearning. Retrained model is the gold standard. Removing a sample from a low-curvature region has almost no effect on the retrained model, whereas removing a sample from high curvature space has significant impact. In this low-dimensional setting, JiT successfully reconstructs the retrained boundary, whereas naively training to mislabel the forget sample completely destroys the trained model.

4 Proposed Method

In this section, we introduce our JiT unlearning method and provide intuition for its effectiveness by examining the geometry of a learned classifier and analyzing how the location of a forget sample in the input space can impact the decision boundary of a model retrained without it. This section considers the case where $|\mathcal{D}_f| = 1$, noting that larger subsets have an additive effect.

Consider the hypothetical of removing a single image of a black cat from a dataset comprised of 1 million black cats and 1 million white dogs. This likely has minimal effects on the learning process, since a well trained classifier should generalise and infer the class of the sample easily. If a sample may be removed from the training dataset without significant changes in the resultant model, we posit that an unlearning algorithm should also have minimal effect on the model when unlearning such a sample. As such, it is logical to design an unlearning algorithm that accounts for the information gain of a sample. However, directly measuring this quantity is difficult, especially in a ZS setting where there is no access to other data points. We therefore seek to derive a heuristic that can approximate how much information gain a sample may have offered the model, based on only that sample and the model itself. We begin by formally introducing the notion of a neighbourhood of a target sample:

Definition 4.1 (Neighbourhood of a sample). For a data point $\mathbf{x} \in \mathcal{X}$, let its neighbourhood $\mathcal{B}_r(\mathbf{x})$ be the bounded subspace of \mathcal{X} containing \mathbf{x} , such that $\forall \hat{\mathbf{x}} \in \mathcal{B}_r(\mathbf{x})$ the $\|\hat{\mathbf{x}} - \mathbf{x}\|_2 \leq r$ for some bound $r \in \mathbb{R}$.

From this definition, we introduce the concept of the amount of information contained within a target sample depending upon its neighbourhood.

Definition 4.2. Let X be a random variable corresponding to a sample \mathbf{x} belonging to class C . Define $\mathcal{B}_r(\mathbf{x})$ as the neighbourhood of $\mathbf{x} \in \mathcal{X}$ and let Y be a random variable corresponding to a sample $\mathcal{B}_r(\mathbf{x})$ belonging to class C . Finally, let $H(X | Y)$ denote the conditional entropy of X given Y . Then we say a sample \mathbf{x}

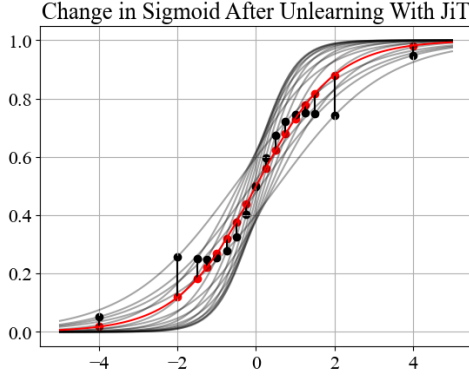


Figure 3: Change in sigmoid after unlearning with JiT. Red dots are unlearned samples, black dots are the location on the new sigmoid post-JiT.




Figure 4: Entropy, $\mathcal{H}(x)$, of the \mathcal{D}_f output distributions for full-class unlearning on CIFAR-10, showing JiT exhibits performance similar to the retrained model.

has **low information** if $H(X | Y) \approx 0$, meaning that \mathbf{x} can be well inferred from its neighbourhood $\mathcal{B}_r(\mathbf{x})$. Conversely, we say a sample is **high information** if $H(X | Y) \gg 0$.

Plainly, we can say that a data point may be said to be low information if it can be inferred from its neighbourhood; and high information if it can not. Consider a low information training sample x_l , from definitions 4.1 and 4.2 we can expect that for some bound r , $f_\theta(\mathbf{x}_l) \approx f_\theta(\hat{\mathbf{x}}_l) \forall \hat{\mathbf{x}}_l \in \mathcal{B}_r(\mathbf{x}_l)$. In other words, a model’s predictions over a low information sample and a set of similar data should have similar output distributions. As such, the curvature of the model in this space will be low. However, for a high information gain sample, this would not necessarily hold. From this we can describe an unlearning objective; if the classifier is smooth with respect to a forget sample, then the model’s prediction over this sample can be viewed as being interpolated or inferred from other data. Hence, we present a method based on minimising the gradient of the classifier with respect to the forget set. Since taking the gradient of the model with respect to the input is extremely expensive for larger problems, we instead construct a first order approximation to the gradient at the target via considering noisy perturbations within its neighbourhood. Formally, $\forall \mathbf{x} \in \mathcal{D}_f$, we seek to minimise the loss given below:

$$\ell = \mathbb{E} \left(\frac{\|f_\theta(\mathbf{x}) - f_\theta(\mathbf{x} + \boldsymbol{\xi})\|_2}{\|\mathbf{x} - (\mathbf{x} + \boldsymbol{\xi})\|_2} \right) \approx \frac{1}{N} \sum_{j=1}^N \left(\frac{\|f_\theta(\mathbf{x}) - f_\theta(\mathbf{x} + \boldsymbol{\xi}_j)\|_2}{\|\boldsymbol{\xi}_j\|_2} \right). \quad (1)$$

Where $\boldsymbol{\xi}$ is a noise vector of equivalent dimensionality to \mathbf{x} , and each component ξ_i of $\boldsymbol{\xi}$ is independently drawn from a Gaussian distribution such that $\xi_i \sim \mathcal{N}(0, \sigma^2)$. For samples that are highly influential, minimising this loss will smooth the local region and remove its influence from the model. For low-information samples that are generalised knowledge, the neighbourhood will be rather smooth resulting in minimal changes. A full algorithm for JiT is given in the appendix 10.1.

4.1 A Geometric Interpretation of JiT

We now present JiT from a geometric perspective, providing insight into why it causes forgetting and how it protects the wider function. Consider a simple 2D classification task, as visualised in figure 2. We pose a simple question: does it make sense to forget all regions of space equally? Naturally, the answer is no. Completely forgetting a sample from within a low-curvature region of space would necessitate the forgetting of almost the entire class, even if they are not part of the forget set. Furthermore, removing a sample from this region would not have significant ramifications on a model retrained from scratch, nor is it likely infringing on the privacy of an individual. Unlearning in this instance often requires minimal alterations to the model. In contrast, a sample that lies in a high-curvature region may not only have significant influence over the position of the learned boundary, but may have been misclassified had it not been included in the training data. Figure 2 shows that this intuition holds in low dimensions; when the forget sample is within

the low-curvature region, a model retrained on \mathcal{D}_r exhibits almost no change, whereas when the sample in \mathcal{D}_f is in a high-curvature region, the boundary is shifted considerably. Enshrining such behaviour into a ZS unlearning algorithm is tantamount; as protection through fine-tuning or regularization is not possible, a ZS algorithm must be surgical in its forgetting methodology. Figure 2 shows that unlearning using JiT yields a classifier almost identical to the retrained model in this low-dimensional setting whereas greedily training over \mathcal{D}_f with a false label causes complete destruction of the model.

The heuristic behind JiT’s performance is based on the gradient field of the classifier. The crux of this rests upon the inherent non-linearity of neural networks. By definition, the model will experience a large rate of change at the decision boundary. As such, given two unit noise vectors ξ_i, ξ_j , where ξ_i points towards the decision boundary and ξ_j points away, the gradient of the classifier between \mathbf{x} and $\mathbf{x} + \xi_i$ will be larger than for $\mathbf{x} + \xi_j$. As such, minimising equation 1 for samples near a boundary will be biased towards moving the boundary *towards* \mathbf{x} . This has the consequence of increasing the uncertainty of the prediction and potentially changing the samples’ predicted class. To further highlight this phenomena, figure 3 shows how unlearning forget samples (red dots) from a learned sigmoid function (red line) changes the learned function. Two things should be observed here: first, samples that lie in low curvature regions have relatively small changes and secondly, the updates to the function have the effect of pulling the the forget sample towards the centre of the sigmoid, which is the decision boundary. Unlearning in this way increases model uncertainty over forgotten samples, without destroying the wider function.

4.2 Entropy Similarity

In the previous section, we demonstrated how in low dimensions JiT can induce forgetting of a single sample in a way similar to retraining the model from scratch. Now, we demonstrate that the same loss can be used to forget arbitrary subsets \mathcal{D} in higher dimensions, including full classes. Since visualising decision boundaries in high dimensions is challenging, we instead evaluate the entropy of the model output. We train a 2-layer CNN on the CIFAR-10 dataset, focusing specifically on the task of forgetting class 0. We compare the entropy of the unlearned model over class 0 with that of the original model, and a model retrained from scratch without class 0. Intuitively, low-entropy predictions indicate higher model confidence, and therefore we expect that the entropy of the model after JiT unlearning is applied will be higher, aligning closely with that of the retrained model.

Figure 4 shows the entropy of the forget-set output distributions for a CNN trained on CIFAR-10. Our unlearning approach increases the entropy over the forget set, reducing the divergence between it and that of a model retrained from scratch on \mathcal{D}_r . In fact, under a Wilcoxon signed-rank test (Woolson, 2007), we find there is **no statistically significant difference** between the model unlearned with JiT and the retrained model for $p = 0.10$. JiT and retraining both increase the entropy over the forget set, suggesting the resultant models behave in a similar way, possessing less knowledge of the forget samples compared to the baseline model. Alongside matching the entropy over the forget set, JiT preserve model performance, as the unlearned model drops only 2% accuracy (From 99% to 97%) on \mathcal{D}_r . Our algorithm demonstrates promising characteristics that are indicative of an effective unlearning algorithm.

5 Experimental Setup

5.1 Benchmarks

We implement the same benchmarks from Foster et al. (2023), which are similar to that of Chundawat et al. (2023a), Golatkar et al. (2020) and Kurmanji et al. (2023). We run experiments 10 times, reporting the mean and standard deviation of these performances. Where classes or sub-classes are forgotten, we show performance over the same class/sub-class as in Foster et al. (2023); performance on additional classes can be found in the appendix (10.3). We perform a hyper-parameter search across a single forget class/sub-class, then use these parameters for all classes. This is more realistic, as it cannot be known *a priori* what future forget sets may be presented to the method. The reported \mathcal{D}_r accuracy refers to accuracy over a test set of samples from the classes in \mathcal{D}_r .

Unlearning scenarios: Typically the three unlearning scenarios are: i) Full-class forgetting, where a full

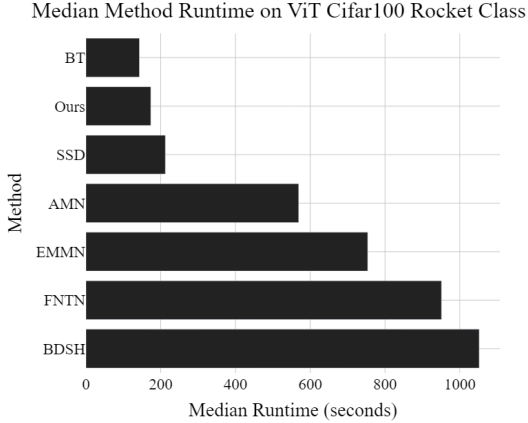


Figure 5: Median method runtime for ViT full-class forgetting on class rocket in seconds. For visual clarity we exclude GKT (~ 3000 seconds).

Table 1: VGG Full-class unlearning performance on PinsFaceRecognition class 1

METHOD	\mathcal{D}_r ACC.	\mathcal{D}_f ACC.	MIA	ZS
BSLN	94.0 \pm 0.0	93.9 \pm 0.0	13.82 \pm 0.0	×
RTRN	100.0 \pm 0.0	0.0 \pm 0.0	2.6 \pm 0.8	×
FNTN	97.6 \pm 0.7	36.9 \pm 9.9	4.3 \pm 2.7	×
AMN	99.7 \pm 0.1	0.0 \pm 0.0	1.4 \pm 1.33	×
SCRUB	98.8 \pm 0.0	97.1 \pm 0.0	8.8 \pm 0.76	×
SSD	55.8 \pm 0.0	0.0 \pm 0.0	4.0 \pm 0.0	×
BT	93.7 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	×
UNSIR	99.5 \pm 0.1	74.4 \pm 9.2	13.6 \pm 8.9	×
GKT	2.0 \pm 0.6	0.0 \pm 0.0	23.9 \pm 30.3	✓
EMMN	51.0 \pm 13.5	69.3 \pm 25.7	26.9 \pm 17.8	✓
BDSH	93.6 \pm 0.4	79.4 \pm 0.0	42.4 \pm 0.4	✓
OURS	91.4 \pm 0.1	1.9 \pm 0.2	4.7 \pm 0.5	✓

class from the dataset must be unlearned, ii) Sub-class forgetting, where a related subset from a class (e.g. all rockets from class vehicle) is forgotten, and iii) Random forgetting, where a subset is sampled uniformly from the entire training distribution. We evaluate our method in all three scenarios.

Comparison methods: We compare JiT against the following methods: i) *Baseline* (BSLN): that has not been unlearned, ii) *Retrain* (RTRN): trained on only the retain data, iii) *Finetune* (FNTN): , where the model is fine-tuned on \mathcal{D}_r for 5 epochs, iv) *Selective Synaptic Dampening* SSD (Foster et al., 2023), v) GKT (Chundawat et al., 2023b), vi) EMMN (Chundawat et al., 2023b), vii) SCRUB (Kurmanji et al., 2023), viii) *Bad Teacher* (BT) (Chundawat et al., 2023a) ix) *Amnesiac* (AMN) (Graves et al., 2021), x) UNSIR (Tarun et al., 2023), xi) *Boundary Shrinking* (BDSH) (Chen et al., 2023). Since GKT, EMMN, and UNSIR are theoretically limited to forgetting just a full-class, these cannot be evaluated in sub-class or random scenarios. Finally, we note that due to VRAM constraints, we could not benchmark SCRUB on ViT.

Datasets: As with previous work, we benchmark JiT on a range of image classification benchmarks. We make use of the CIFAR suite (Krizhevsky and Hinton, 2010), and the Pins Facial Recognition dataset (Burak, 2020), which consists of 17,534 images of 105 celebrity faces.

Models: We evaluate methods on Vision Transformer (ViT) (Dosovitskiy et al., 2021) and VGG11 (Simonyan and Zisserman, 2014), trained on an NVidia RTX 4090 using Stochastic Gradient Descent with an initial learning rate of 0.1, and the OneCycle learning rate scheduler (Smith and Topin, 2019). Additionally, we compare the performance of JiT to BDSH on a ViT-L ($\sim 300m$ parameters) trained on the ILSVRC Imagenet dataset to demonstrate our method can scale to larger problem spaces.

Evaluation metrics: We evaluate model performance according to four key metrics: i) \mathcal{D}_r accuracy, ii) \mathcal{D}_f accuracy, iii) MIA score, and iv) method runtime. For all metrics bar runtime, the objective is not to minimise/maximise them, but rather to be as close to the retrained model as possible. This is important, as performing worse than the retrained model implies insufficient performance, but as noted in Chundawat et al. (2023a), Foster et al. (2023) and Kurmanji et al. (2023), significant deviation from the retrained model (e.g. *over-forgetting*) may leak information about the fact a sample has been forgotten. To remain consistent with existing unlearning literature we use the same logistic regression MIA evaluation as Chundawat et al. (2023a) and Foster et al. (2023).

JiT hyper-parameters: We conduct a hyper-parameter search for η and σ using 250 runs of the TPE search from Optuna (Akiba et al., 2019), for each unlearning scenario. For VGG11, we use the following parameters: full-class unlearning uses $\eta = 0.0003, \sigma = 0.5$, sub-class and random both use $\eta = 0.0003, \sigma = 0.01$. For ViT, the selected parameters are: full-class $\eta = 1.5, \sigma = 0.8$, sub-class $\eta = 0.5, \sigma = 1.5$, and random $\eta = 0.01, \sigma = 0.5$. ViT and VGG use considerably different learning rates, since only a single epoch is used during the unlearning step. If minimising the runtime is a looser constraint, a smaller learning rate can be used for ViT with extra epochs of training. We stress that when selecting hyper-parameters, we selected values that yielded promising results, without rigorously fitting our results to the retrained model.

6 Results

6.1 Benchmark Evaluation

Compute Comparison. Figure 5 shows the runtime of JiT compared to other methods. JiT is very fast, especially in comparison to the other ZS methods, performing more than 5 times faster than Boundary Shrinking. Sufficiently short runtimes are an important desideratum of unlearning, one which JiT empirically satisfies. JiT has a computational complexity of $O(N|\mathcal{D}_f|)$, where N is the number of perturbed samples and $|\mathcal{D}_f|$ is the cardinality of the forget set. Requiring only \mathcal{D}_f and processing each sample just once makes JiT efficient.

Full-class Unlearning. We begin by comparing full class performance to that of the existing ZS methods. As seen in Tables 2 (a), 2 (b), and 1, JiT demonstrates significantly superior performance over GKT and EMMN, and is competitive with Boundary Shrinking. The authors of Chundawat et al. (2023b) note the poor scalability of both EMMN and GKT, which is evident in our results. Failing to scale to large problems or models is a significant barrier, since the value of unlearning is found mostly in large models that are expensive to train, or large datasets that are expensive to store. JiT performance is competitive with Boundary Shrinking despite having a fraction of the compute cost and, even compared to non-ZS SOTA, JiT performs reasonably; dropping only 0.6% retain set performance compared to the baseline on ViT and outperforming both UNSIR and SSD on the MIA. Our performance also holds for VGG and, when using the same hyper-parameters for the face dataset, JiT generalizes well, outperforming SCRUB, SSD, and UNSIR.

Sub-class Unlearning. Tables 3 (a) and 3 (b) show the performance of JiT on sub-class unlearning for ViT and VGG11. For both, JiT is able to comfortably unlearn. For ViT, it actually *over-forgets*, however this is typically an easier problem to correct, since more conservative values can always be selected. \mathcal{D}_r performance also drops slightly more than usual, which a more conservative parameter set could also correct. For VGG11, however, the method is comfortably amongst the SOTA, outperforming methods that are granted access to the retained data. For ViT, JiT better minimises \mathcal{D}_f compared to BDSH.

Random Unlearning. Tables 4 (a) and 4 (b) show method performance when forgetting 100 samples uniformly distributed across the training set. JiT is able to comfortably rival existing ZS and non-ZS SOTA methods; despite slight over-forgetting, \mathcal{D}_r accuracy is almost unchanged ($\sim 1\%$ for both models).

ImageNet Evaluation. Finally, table 5 validates our method on a larger scale problem, with JiT achieving SOTA performance for ZS methods. As larger pretrained models can be robust to noise, we found larger

Table 2: (a) ViT Full-class unlearning performance on CIFAR-100 class Rocket. (b) VGG11 Full-class unlearning performance on CIFAR-100 class Rocket.

(a)					(b)				
METHOD	\mathcal{D}_r ACC.	\mathcal{D}_f ACC.	MIA	ZS	METHOD	\mathcal{D}_r ACC.	\mathcal{D}_f ACC.	MIA	ZS
BSLN	88.9 ± 0.0	94.7 ± 0.0	94.4 ± 0.0	×	BSLN	66.3 ± 0.0	77.0 ± 0.0	97.4 ± 0.0	×
RTRN	90.1 ± 0.0	0.0 ± 0.0	3.2 ± 0.5	×	RTRN	63.2 ± 0.5	0.0 ± 0.0	10.4 ± 1.1	×
FNTN	80.8 ± 1.4	0.6 ± 0.7	19.0 ± 8.7	×	FNTN	59.7 ± 0.4	3.9 ± 3.0	13.2 ± 4.2	×
AMN	87.9 ± 0.9	0.0 ± 0.0	1.4 ± 0.9	×	AMN	64.3 ± 0.4	0.0 ± 0.0	1.8 ± 0.8	×
SSD	88.90 ± 0.0	0.0 ± 0.0	1.8 ± 0.0	×	SCRUB	66.2 ± 0.1	0.0 ± 0.0	8.2 ± 1.7	×
BT	87.5 ± 0.5	4.2 ± 5.2	0.0 ± 0.1	×	SSD	63.79 ± 0.0	0.0 ± 0.0	8.6 ± 0.0	×
UNSIR	88.5 ± 0.4	65.3 ± 9.1	29.1 ± 6.1	×	BT	65.5 ± 0.2	0.1 ± 0.3	0.0 ± 0.1	×
GKT	1.0 ± 0.6	0.0 ± 0.0	60.0 ± 51.6	✓	UNSIR	64.6 ± 0.4	42.9 ± 14.3	40.7 ± 12.1	×
EMMN	84.6 ± 0.4	94.3 ± 1.5	93.7 ± 2.2	✓	GKT	2.3 ± 0.2	0.0 ± 0.0	56.2 ± 20.0	✓
BDSH	87.6 ± 0.0	0.0 ± 0.0	5.0 ± 0.0	✓	EMMN	26.9 ± 7.7	24.3 ± 23.7	58.2 ± 14.5	✓
OURS	87.5 ± 0.0	51.9 ± 2.13	4.3 ± 0.38	✓	BDSH	66.2 ± 0.1	13.0 ± 0.0	2.9 ± 0.1	✓
					OURS	66.2 ± 0.3	14.2 ± 0.6	2.9 ± 0.3	✓

Table 3: (a) VGG-16 Sub-class unlearning performance on CIFAR-20 sub-class Rocket. (b) ViT Sub-class unlearning performance on CIFAR-20 sub-class Rocket

(a)					(b)				
METHOD	\mathcal{D}_r ACC.	\mathcal{D}_f ACC.	MIA	ZS	METHOD	\mathcal{D}_r ACC.	\mathcal{D}_f ACC.	MIA	ZS
BSLN	75.3 \pm 0.0	79.0 \pm 0.0	83.1 \pm 0.0	×	BSLN	95.7 \pm 0.0	94.5 \pm 0.0	80.4 \pm 0.0	×
RTRN	72.9 \pm 0.2	11.5 \pm 2.8	14.1 \pm 1.3	×	RTRN	94.6 \pm 0.1	22.3 \pm 8.3	3.4 \pm 1.1	×
FNTN	65.5 \pm 0.7	6.2 \pm 3.7	22.3 \pm 5.5	×	FNTN	85.7 \pm 3.1	6.2 \pm 6.0	16.0 \pm 2.7	×
AMN	73.8 \pm 0.2	2.4 \pm 2.4	3.0 \pm 0.9	×	AMN	93.5 \pm 0.2	0.8 \pm 1.7	0.8 \pm 0.3	×
SCRUB	62.4 \pm 28.4	10.1 \pm 22.48	16.7 \pm 21.7	×	SSD	95.1 \pm 0.0	5.12 \pm 0.0	5.4 \pm 0.0	×
SSD	75.0 \pm 0.0	4.2 \pm 0.0	11.0 \pm 0.0	×	BT	93.6 \pm 0.3	3.3 \pm 2.9	0.0 \pm 0.1	×
BT	74.9 \pm 0.2	48.4 \pm 16.9	0.1 \pm 0.1	×	UNSIR	93.3 \pm 0.4	74.9 \pm 10.1	27.3 \pm 13.8	×
UNSIR	74.1 \pm 0.2	57.5 \pm 10.3	57.4 \pm 8.6	×	BDSH	95.7 \pm 0.0	48.4 \pm 0.0	0.1 \pm 0.0	✓
BDSH	74.4 \pm 0.0	17.535 \pm 0.0	12.9 \pm 0.1	✓	OURS	92.2 \pm 0.0	0.0 \pm 0.0	14.66 \pm 8.8	✓
OURS	73.7 \pm 0.8	19.3 \pm 18.3	11.2 \pm 7.8	✓					

perturbations were required to induce forgetting. To keep the input in-domain, we apply normalization to the noised image via: $\frac{(x+\xi)}{\sqrt{(1+\sigma^2)}}$.

7 Discussion

JiT is by far the fastest ZS unlearning method benchmarked, a critical characteristic for satisfying the unlearning task. JiT is competitive with state-of-the-art performance in the ZS unlearning domain, as well as competing with non-ZS methods in the sub-class and random unlearning tasks despite their easier task. The entropy experiments highlight that JiT is able to replicate the output entropy of a retrained model over a forget set, while preserving retain set performance. When compared to existing ZS methods, JiT can be considered a strong baseline. It is fast and performant, and performed acceptably across all benchmarks implemented. If time constraints are ignored, BDSH is more stable and less sensitive to hyper-parameter selection, on account of taking the true gradients of the model with respect to the input, rather than the approximation we employ with JiT. However, in practice the poor time complexity of BDSH will likely prove prohibitive when trying to unlearn from internet-scale models, whereas JiT is amongst the fastest methods we benchmarked. Future work could explore the efficacy of using JiT with exact gradients, or a more specialised gradient approximation. JiT has the potential for positive societal impacts, aiding the preservation of individual privacy. However, due to a lack of certification, poor use of JiT could result in organizations believing they have removed the influence of an individual’s data when they haven’t.

Table 4: (a) VGG11 Random unlearning performance for 100 samples from CIFAR-10. (b) ViT Random unlearning performance for 100 samples from CIFAR-10.

(a)					(b)				
METHOD	\mathcal{D}_r ACC.	\mathcal{D}_f ACC.	MIA	ZS	METHOD	\mathcal{D}_r ACC.	\mathcal{D}_f ACC.	MIA	ZS
BSLN	87.0 \pm 0.0	92.0 \pm 3.6	70.1 \pm 5.4	×	BSLN	98.9 \pm 0.0	100.0 \pm 0.0	90.8 \pm 3.5	×
RTRN	87.7 \pm 0.2	91.0 \pm 2.5	78.9 \pm 3.5	×	RTRN	98.6 \pm 0.1	98.8 \pm 0.8	91.8 \pm 1.8	×
FNTN	84.4 \pm 0.8	86.4 \pm 4.4	70.8 \pm 4.7	×	FNTN	97.3 \pm 0.3	97.2 \pm 1.0	86.1 \pm 2.1	×
AMN	86.8 \pm 0.3	51.3 \pm 4.4	13.1 \pm 2.9	×	AMN	97.6 \pm 0.3	73.5 \pm 5.1	10.4 \pm 4.9	×
SCRUB	87.7 \pm 0.1	92.7 \pm 2.9	71.8 \pm 5.2	×	SSD	98.0 \pm 1.6	98.1 \pm 2.4	85.5 \pm 0.1	×
SSD	85.6 \pm 2.7	90.8 \pm 3.7	66.7 \pm 5.9	×	BT	97.6 \pm 0.4	86.7 \pm 3.6	33.5 \pm 5.6	×
BT	86.9 \pm 0.2	82.5 \pm 4.9	40.8 \pm 6.3	×	BDSH	98.0 \pm 0.29	97.9 \pm 1.6	78.8 \pm 0.0	✓
BDSH	86.9 \pm 0.1	92.2 \pm 3.4	69.8 \pm 5.1	✓	OURS	98.0 \pm 0.3	98.0 \pm 1.5	78.8 \pm 4.0	✓
OURS	86.3 \pm 0.3	88.7 \pm 3.9	64.2 \pm 5.2	✓					

Table 5: Zero-shot methods performance on a ViT-L trained on ILSVRC Imagenet.

Method	\mathcal{D}_r Acc.	\mathcal{D}_f Acc	MIA
BSLN	86.0	100.0	94.0
BDSH	85.9	60.0	0.0
OURS	83.6	10.0	0.0

8 Limitations

The appendix provides a sensitivity analysis (10.2). It demonstrates that JiT can be sensitive to hyperparameter selection. This is a by-product of having no access to \mathcal{D}_r to finetune, and also due to our gradient approximation method. Exact gradients are slower, but may prove more stable. Importantly, we note that better stability can be achieved through gradient clipping, though this came at the cost of method performance.

Since we minimise the gradient over each forget sample independently, we advise caution when using a model with batch normalization; since this changes the model’s mapping of single input/output to batch input/output (Gulrajani et al., 2017). This limitation can be mitigated by selecting an alternative choice of model or normalization strategy (e.g. layer norms).

JiT, like most SOTA methods, is not certified. The impact of this will vary by application domain, but may preclude its use in especially sensitive areas. Finally, we note that our method is specifically tailored to classification tasks. While it is possible a variation of this approach could work for large generative models, we restrict our focus and our claims to larger classifiers and leave generative applications to future work.

9 Conclusion

Unlearning is an important, challenging problem. The ZS setting is amongst the hardest, requiring delicate treatment of the unlearning process to ensure model performance is protected. In this work, we approached this challenge from an information theoretic perspective, deriving an unlearning algorithm directly from the notion of minimising information gained from a sample. We demonstrate empirically the geometric insights behind why JiT can effectively tackle the ZS unlearning problem, alongside showing experimentally that JiT can reconstruct behaviour analogous to that of a model retrained from scratch. JiT achieves performance competitive with state-of-the-art ZS and non-ZS methods. We evaluate JiT on a range of benchmarks, demonstrating its efficacy in full-class, sub-class, and random unlearning, across multiple models. Future work is needed to establish a stronger theoretical relationship between forgetting and information theory, as well as exploring whether this can be formalized to provide guarantees on forgetting using information theoretic approaches.

References

- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631, 2019.
- Devansh Arpit, Stanisław Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International conference on machine learning*, pages 233–242. PMLR, 2017.
- Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 141–159. IEEE, 2021.
- Burak. Pinterest face recognition dataset. kaggle.com/datasets/hereisburak/pins-face-recognition, 2020. Accessed: 2023-08-09.
- Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015 IEEE symposium on security and privacy*, pages 463–480. IEEE, 2015.
- Min Chen, Weizhuo Gao, Gaoyang Liu, Kai Peng, and Chen Wang. Boundary unlearning: Rapid forgetting of deep networks via shifting the decision boundary. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7766–7775, 2023.
- Vikram S Chundawat, Ayush K Tarun, Murari Mandal, and Mohan Kankanhalli. Can bad teaching induce forgetting? unlearning in deep networks using an incompetent teacher. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 7210–7217, 2023a.
- Vikram S Chundawat, Ayush K Tarun, Murari Mandal, and Mohan Kankanhalli. Zero-shot machine unlearning. *IEEE Transactions on Information Forensics and Security*, 2023b.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- Vitaly Feldman. Does learning require memorization? a short tale about a long tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 954–959, 2020.
- Jack Foster, Stefan Schoepf, and Alexandra Brintrup. Fast machine unlearning without retraining through selective synaptic dampening. *arXiv preprint arXiv:2308.07707*, 2023.
- Antonio Ginart, Melody Guan, Gregory Valiant, and James Y Zou. Making ai forget you: Data deletion in machine learning. *Advances in neural information processing systems*, 32, 2019.
- Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Eternal sunshine of the spotless net: Selective forgetting in deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9304–9312, 2020.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Laura Graves, Vineel Nagesetty, and Vijay Ganesh. Amnesiac machine learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11516–11524, 2021.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. *Advances in neural information processing systems*, 30, 2017.
- Varun Gupta, Christopher Jung, Seth Neel, Aaron Roth, Saeed Sharifi-Malvajerdi, and Chris Waites. Adaptive machine unlearning. *Advances in Neural Information Processing Systems*, 34:16319–16330, 2021.

- Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.
- Jenny E Jeong and Peng Qiu. Quantifying the relative importance of experimental data points in parameter estimation. *BMC systems biology*, 12:1–14, 2018.
- Alex Krizhevsky and Geoff Hinton. Convolutional deep belief networks on cifar-10. *Unpublished manuscript*, 40(7):1–9, 2010.
- Meghdad Kurmanji, Peter Triantafillou, and Eleni Triantafillou. Towards unbounded machine unlearning. *arXiv preprint arXiv:2302.09880*, 2023.
- Dennis V Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27(4):986–1005, 1956.
- Ronak Mehta, Sourav Pal, Vikas Singh, and Sathya N Ravi. Deep unlearning via randomized conditionally independent Hessians. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10422–10431, 2022.
- Seth Neel, Aaron Roth, and Saeed Sharifi-Malvajerdi. Descent-to-delete: Gradient-based methods for machine unlearning. In *Algorithmic Learning Theory*, pages 931–962. PMLR, 2021.
- J. Ross Quinlan. Induction of decision trees. *Machine learning*, 1:81–106, 1986.
- Ayush Sekhari, Jayadev Acharya, Gautam Kamath, and Ananda Theertha Suresh. Remember what you want to forget: Algorithms for machine unlearning. *Advances in Neural Information Processing Systems*, 34:18075–18086, 2021.
- Vedant Shah, Frederik Träuble, Ashish Malik, Hugo Larochelle, Michael Mozer, Sanjeev Arora, Yoshua Bengio, and Anirudh Goyal. Unlearning via sparse representations. *arXiv preprint arXiv:2311.15268*, 2023.
- Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pages 3–18. IEEE, 2017.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, volume 11006, pages 369–386. SPIE, 2019.
- Ayush K Tarun, Vikram S Chundawat, Murari Mandal, and Mohan Kankanhalli. Fast yet effective machine unlearning. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- Simon Tong and Daphne Koller. Support vector machine active learning with applications to text classification. *Journal of machine learning research*, 2(Nov):45–66, 2001.
- Eleni Triantafillou, Fabian Pedregosa, Jamie Hayes, Peter Kairouz, Isabelle Guyon, Meghdad Kurmanji, Gintare Karolina Dziugaite, Peter Triantafillou, Kairan Zhao, Lisheng Sun Hosoya, Julio C. S. Jacques Junior, Vincent Dumoulin, Ioannis Mitliagkas, Sergio Escalera, Jun Wan, Sohier Dane, Maggie Demkin, and Walter Reade. Neurips 2023 - machine unlearning, 2023. URL <https://kaggle.com/competitions/neurips-2023-machine-unlearning>.
- Alexander Warnecke, Lukas Pirch, Christian Wressnegger, and Konrad Rieck. Machine unlearning of features and labels. *arXiv preprint arXiv:2108.11577*, 2021.
- Robert F Woolson. Wilcoxon signed-rank test. *Wiley encyclopedia of clinical trials*, pages 1–3, 2007.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.

10 Appendix

10.1 Method Algorithm

Algorithm 1 JiT UNLEARNING

INPUT: The trained model $f_\theta(\cdot)$ and the forget set \mathcal{S} .

PARAMETER: η, σ, N

OUTPUT: $f_{\hat{\theta}}(\cdot) = \mathcal{U}_{\mathcal{S}}(f_\theta(\cdot))$

```

1: Initialise  $\text{optim}(\theta, lr = \eta)$ 
2: for  $x$  in  $\mathcal{S}$  do
3:    $\ell = 0$ 
4:   for  $i$  in  $\text{range}(N)$  do
5:      $\mathbf{x}' = \mathbf{x} + \boldsymbol{\xi}$  for  $\boldsymbol{\xi} \sim \mathcal{N}(0, \sigma^2)$ 
6:      $k = \frac{\|f_\theta(\mathbf{x}) - f_\theta(\mathbf{x}')\|_2}{\|\boldsymbol{\xi}\|_2}$ 
7:      $\ell = \ell + k$ 
8:   end for
9: end for
10:  $\ell = \ell / N$ 
11:  $\hat{\theta} \leftarrow \text{optim}\{\nabla_\theta \ell\}$ 
12: return  $f_{\hat{\theta}}(\cdot)$ 
```

Figure 6: Pseudocode algorithm for JiT Unlearning

10.2 Sensitivity Analysis




Figure 7: Plot of the \mathcal{D}_r sensitivity to change in hyper-parameters for VGG11 full-class forgetting.

Figures 7 and 8 show the sensitivity of key metrics to changes in the hyper-parameters η and σ . In general, the approach is robust to small perturbations of the learning rate, but naturally as it increases by orders

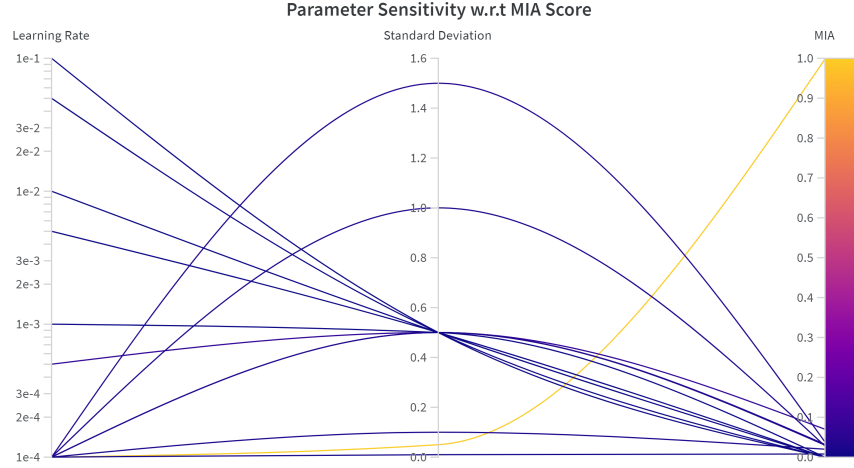


Figure 8: Plot of the MIA sensitivity to change in hyper-parameters for VGG11 full-class forgetting.

of magnitude, performance varies significantly. For σ , in general increasing the noise actually reduced the forgetting/increased \mathcal{D}_r accuracy. This is because VGG was so sensitivity to noise, and so small values of σ did little to reduce the divergence of model output, and thus for $\sigma < 1$ the loss increases due to the division in the loss term. For models that are robust to additive noise, the relationship between σ and performance is often parabolic.

Choice of how many perturbed variants is simple. The more samples the more stable the forgetting, and the only trade-off is compute time.

10.3 Class Breakdown of Method Performance

Below are the forget-class breakdowns of each method’s performance for full-class and sub-class unlearning. For each unlearning scenario, the same parameters are used across all classes; this can lead to significant variance in method performance.

Table 6: ViT CIFAR-100 full-class unlearning breakdown

METHOD	FORGET CLASS	\mathcal{D}_r	\mathcal{D}_f	MIA	METHOD RUNTIME	ZS
BASILINE	BABY	88.9 ± 0.0	90.2 ± 0.0	75.6 ± 0.0	459.7 ± 83.0	×
BASILINE	LAMP	88.8 ± 0.0	97.2 ± 0.0	95.6 ± 0.0	504.1 ± 131.7	×
BASILINE	MUSHROOM	88.9 ± 0.0	94.9 ± 0.0	92.8 ± 0.0	474.6 ± 60.2	×
BASILINE	ROCKET	88.9 ± 0.0	94.7 ± 0.0	94.4 ± 0.0	470.3 ± 85.7	×
BASILINE	SEA	88.9 ± 0.0	90.5 ± 0.0	80.4 ± 0.0	549.3 ± 81.8	×
RETRAIN	BABY	90.3 ± 0.1	0.0 ± 0.0	21.5 ± 2.8	2029.1 ± 169.7	×
RETRAIN	LAMP	90.1 ± 0.2	0.0 ± 0.0	2.3 ± 0.7	2115.0 ± 196.3	×
RETRAIN	MUSHROOM	90.0 ± 0.2	0.0 ± 0.0	0.7 ± 0.4	1948.8 ± 123.1	×
RETRAIN	ROCKET	90.1 ± 0.1	0.0 ± 0.0	3.2 ± 0.5	1939.7 ± 109.7	×
RETRAIN	SEA	90.3 ± 0.2	0.0 ± 0.0	8.4 ± 2.2	2176.1 ± 92.8	×
FINETUNE	BABY	80.7 ± 1.4	0.0 ± 0.0	26.8 ± 12.7	1461.8 ± 96.2	×
FINETUNE	LAMP	80.2 ± 1.5	0.4 ± 0.9	11.8 ± 4.3	1537.4 ± 133.2	×
FINETUNE	MUSHROOM	81.1 ± 0.8	2.3 ± 2.4	7.1 ± 1.9	1421.7 ± 80.5	×
FINETUNE	ROCKET	80.8 ± 1.4	0.5 ± 0.7	19.0 ± 8.7	1404.9 ± 82.8	×
FINETUNE	SEA	80.8 ± 1.4	0.0 ± 0.0	22.0 ± 7.1	1525.0 ± 147.3	×
AMNESIAC	BABY	88.4 ± 0.7	0.0 ± 0.0	1.8 ± 0.3	1132.0 ± 147.1	×
AMNESIAC	LAMP	88.4 ± 0.6	0.0 ± 0.0	2.7 ± 0.4	1146.5 ± 97.5	×
AMNESIAC	MUSHROOM	88.3 ± 0.7	0.0 ± 0.0	0.5 ± 0.2	1071.9 ± 98.1	×
AMNESIAC	ROCKET	87.9 ± 0.9	0.0 ± 0.0	1.0 ± 0.6	1029.7 ± 70.8	×
AMNESIAC	SEA	88.3 ± 0.3	0.0 ± 0.0	0.8 ± 0.2	1176.2 ± 136.7	×
SSD	BABY	88.59 ± 0.0	0.0 ± 0.0	0.60 ± 0.0	685.40 ± 97.35	×
SSD	LAMP	89.06 ± 0.0	36.89 ± 0.0	0.40 ± 0.0	741.23 ± 170.57	×
SSD	MUSHROOM	88.82 ± 0.0	0.0 ± 0.0	3.80 ± 0.0	650.93 ± 76.43	×
SSD	ROCKET	88.90 ± 0.0	0.0 ± 0.0	1.80 ± 0.0	655.64 ± 65.36	×
SSD	SEA	87.95 ± 0.0	0.0 ± 0.0	3.20 ± 0.0	767.35 ± 177.77	×
TEACHER	BABY	87.5 ± 0.4	23.8 ± 22.5	0.0 ± 0.0	610.0 ± 98.1	×
TEACHER	LAMP	87.5 ± 0.4	25.2 ± 12.5	0.1 ± 0.2	668.6 ± 135.1	×
TEACHER	MUSHROOM	87.4 ± 0.4	12.8 ± 5.9	0.0 ± 0.1	602.9 ± 33.2	×
TEACHER	ROCKET	87.5 ± 0.5	4.2 ± 5.2	0.0 ± 0.1	602.6 ± 63.2	×
TEACHER	SEA	87.7 ± 0.2	51.1 ± 17.4	0.0 ± 0.0	661.8 ± 117.1	×
UNSIR	BABY	88.8 ± 0.4	2.0 ± 1.2	14.3 ± 6.1	954.6 ± 110.3	×
UNSIR	LAMP	88.5 ± 0.4	70.9 ± 4.4	29.4 ± 4.8	1002.2 ± 135.1	×
UNSIR	MUSHROOM	88.4 ± 0.6	83.9 ± 2.9	21.3 ± 2.7	891.8 ± 69.9	×
UNSIR	ROCKET	88.5 ± 0.4	65.3 ± 9.1	29.1 ± 6.1	868.8 ± 47.9	×
UNSIR	SEA	88.8 ± 0.2	13.9 ± 6.2	9.1 ± 4.7	986.1 ± 149.6	×
GKT	BABY	1.00 ± 0.07	0.00 ± 0.00	70.00 ± 48.30	1074.86 ± 791.54	✓
GKT	LAMP	1.01 ± 0.06	0.00 ± 0.00	60.00 ± 51.64	1924.50 ± 1269.59	✓
GKT	MUSHROOM	1.01 ± 0.06	0.00 ± 0.00	60.00 ± 51.64	1793.25 ± 1100.70	✓
GKT	ROCKET	1.00 ± 0.06	0.00 ± 0.00	60.00 ± 51.64	1943.99 ± 1227.35	✓
GKT	SEA	1.01 ± 0.06	0.00 ± 0.00	70.00 ± 48.30	1788.32 ± 1150.91	✓
EMMN	BABY	84.63 ± 0.30	87.86 ± 2.98	74.82 ± 6.54	1064.73 ± 202.59	✓
EMMN	LAMP	84.99 ± 0.37	93.03 ± 2.04	87.58 ± 2.95	1100.77 ± 212.08	✓
EMMN	MUSHROOM	84.65 ± 0.33	93.11 ± 1.84	91.14 ± 2.23	1068.89 ± 132.31	✓
EMMN	ROCKET	84.62 ± 0.40	94.33 ± 1.53	93.68 ± 2.16	1115.30 ± 179.26	✓
EMMN	SEA	84.68 ± 0.33	89.47 ± 4.26	74.64 ± 7.44	1075.82 ± 216.62	✓
BDSH	BABY	83.87 ± 0.0	0.0 ± 0.0	8.6 ± 0.0	1388.14 ± 310.5	✓
BDSH	LAMP	86.94 ± 0.0	2.0 ± 0.0	28.44 ± 0.0	1386.44 ± 304.9	✓
BDSH	MUSHROOM	87.8 ± 0.0	0.0 ± 0.0	0.8 ± 0.0	1398.19 ± 305.3	✓
BDSH	ROCKET	87.62 ± 0.0	0.0 ± 0.0	5.0 ± 0.0	1396.85 ± 306.4	✓
BDSAH	SEA	87.34 ± 0.0	3.0 ± 0.0	0.8 ± 0.0	1414.41 ± 146.2	✓
OURS	BABY	87.16 ± 0.03	38.70 ± 0.95	0.40 ± 0.00	607.94 ± 130.32	✓
OURS	LAMP	88.39 ± 0.01	93.90 ± 0.32	49.86 ± 2.09	596.99 ± 143.57	✓
OURS	MUSHROOM	87.70 ± 0.03	77.90 ± 0.57	4.78 ± 0.37	656.20 ± 141.64	✓
OURS	ROCKET	87.46 ± 0.02	51.90 ± 2.13	4.26 ± 0.38	629.38 ± 116.30	✓
OURS	SEA	83.78 ± 0.08	24.00 ± 1.15	16.32 ± 0.36	596.28 ± 151.89	✓

Table 7: ViT CIFAR-20 sub-class unlearning breakdown

METHOD	FORGET CLASS	\mathcal{D}_r	\mathcal{D}_f	MIA	METHOD RUNTIME	ZS
BASILINE	BABY	95.7 \pm 0.0	96.4 \pm 0.0	91.6 \pm 0.0	443.0 \pm 48.7	\times
BASILINE	LAMP	95.8 \pm 0.0	89.6 \pm 0.0	81.0 \pm 0.0	419.7 \pm 25.2	\times
BASILINE	MUSHROOM	95.7 \pm 0.0	97.0 \pm 0.0	77.8 \pm 0.0	525.7 \pm 124.3	\times
BASILINE	ROCKET	95.7 \pm 0.0	94.5 \pm 0.0	80.4 \pm 0.0	535.0 \pm 75.3	\times
BASILINE	SEA	95.7 \pm 0.0	99.2 \pm 0.0	88.4 \pm 0.0	475.5 \pm 120.3	\times
RETRAIN	BABY	94.5 \pm 0.2	93.2 \pm 1.1	77.4 \pm 3.4	2148.7 \pm 112.5	\times
RETRAIN	LAMP	94.7 \pm 0.1	34.5 \pm 8.6	5.6 \pm 1.6	2149.3 \pm 124.4	\times
RETRAIN	MUSHROOM	94.6 \pm 0.1	26.6 \pm 6.4	2.3 \pm 0.5	2161.5 \pm 81.6	\times
RETRAIN	ROCKET	94.6 \pm 0.1	22.3 \pm 8.3	3.4 \pm 1.1	2168.3 \pm 118.4	\times
RETRAIN	SEA	94.6 \pm 0.2	95.1 \pm 0.8	66.0 \pm 3.8	2142.0 \pm 85.9	\times
FINETUNE	BABY	87.6 \pm 0.8	85.4 \pm 4.5	66.6 \pm 7.1	1426.9 \pm 72.7	\times
FINETUNE	LAMP	87.7 \pm 0.5	16.9 \pm 10.4	14.7 \pm 3.9	1424.9 \pm 71.4	\times
FINETUNE	MUSHROOM	87.4 \pm 0.9	15.7 \pm 12.1	9.2 \pm 4.1	1442.4 \pm 113.2	\times
FINETUNE	ROCKET	85.7 \pm 3.1	6.2 \pm 6.0	16.0 \pm 2.7	1436.6 \pm 117.9	\times
FINETUNE	SEA	87.6 \pm 1.6	89.2 \pm 4.2	65.0 \pm 12.9	1512.9 \pm 129.8	\times
AMNESIAC	BABY	93.3 \pm 0.3	38.8 \pm 7.4	0.9 \pm 0.7	1025.0 \pm 35.2	\times
AMNESIAC	LAMP	93.7 \pm 0.5	0.6 \pm 1.5	2.0 \pm 1.0	1009.0 \pm 37.7	\times
AMNESIAC	MUSHROOM	93.4 \pm 0.5	0.2 \pm 0.4	1.5 \pm 0.5	1131.8 \pm 137.6	\times
AMNESIAC	ROCKET	93.5 \pm 0.2	0.8 \pm 1.7	0.8 \pm 0.3	1186.6 \pm 107.5	\times
AMNESIAC	SEA	93.3 \pm 0.2	21.4 \pm 8.5	0.4 \pm 0.2	1070.4 \pm 138.1	\times
SSD	BABY	95.54 \pm 0.0	94.10 \pm 0.0	77.20 \pm 0.0	736.00 \pm 12.01	\times
SSD	LAMP	95.54 \pm 0.0	14.58 \pm 0.0	3.2 \pm 0.0	728.98 \pm 73.07	\times
SSD	MUSHROOM	95.51 \pm 0.0	6.68 \pm 0.0	0.40 \pm 0.0	718.83 \pm 73.41	\times
SSD	ROCKET	95.13 \pm 0.0	5.12 \pm 0.0	5.40 \pm 0.0	699.33 \pm 72.47	\times
SSD	SEA	95.57 \pm 0.0	97.05 \pm 0.0	82.20 \pm 0.0	645.74 \pm 53.38	\times
TEACHER	BABY	93.0 \pm 0.5	46.7 \pm 17.9	0.0 \pm 0.1	553.9 \pm 65.9	\times
TEACHER	LAMP	93.6 \pm 0.7	8.2 \pm 7.1	0.1 \pm 0.2	558.1 \pm 63.9	\times
TEACHER	MUSHROOM	93.6 \pm 0.4	13.0 \pm 9.1	0.0 \pm 0.0	620.1 \pm 111.7	\times
TEACHER	ROCKET	93.6 \pm 0.3	3.3 \pm 2.9	0.0 \pm 0.1	631.9 \pm 115.2	\times
TEACHER	SEA	93.6 \pm 0.3	26.0 \pm 14.0	0.2 \pm 0.1	586.7 \pm 89.7	\times
UNSIR	BABY	93.2 \pm 0.3	94.5 \pm 0.8	88.0 \pm 3.1	871.1 \pm 60.7	\times
UNSIR	LAMP	93.4 \pm 0.5	76.5 \pm 5.2	36.5 \pm 11.7	899.8 \pm 72.0	\times
UNSIR	MUSHROOM	93.1 \pm 0.6	79.8 \pm 7.6	19.0 \pm 7.4	925.7 \pm 117.5	\times
UNSIR	ROCKET	93.3 \pm 0.4	74.9 \pm 10.1	27.3 \pm 13.8	983.1 \pm 143.2	\times
UNSIR	SEA	93.3 \pm 0.3	94.3 \pm 2.3	77.0 \pm 7.2	1024.5 \pm 144.8	\times
BDSH	BABY	95.36 \pm 0.0	93.32 \pm 0.0	18.8 \pm 0.0	1163 \pm 49.76	\checkmark
BDSH	LAMP	95.76 \pm 0.0	89.58 \pm 0.0	80.8 \pm 0.0	1152.48 \pm 53.9	\checkmark
BDSH	MUSHROOM	95.72 \pm 0.0	88.37 \pm 0.0	2.6 \pm 0.0	1087.87 \pm 200.2	\checkmark
BDSH	ROCKET	95.66 \pm 0.0	48.44 \pm 0.0	1.4 \pm 0.0	1087.15 \pm 212.6	\checkmark
BDSH	SEA	95.09 \pm 0.0	78.91 \pm 0.0	4.6 \pm 0.0	1101.91 \pm 215.9	\checkmark
OURS	BABY	87.40 \pm 0.02	0.00 \pm 0.00	0.80 \pm 0.00	532.20 \pm 112.52	\checkmark
OURS	LAMP	90.20 \pm 0.02	0.00 \pm 0.00	22.86 \pm 0.19	560.02 \pm 58.54	\checkmark
OURS	MUSHROOM	93.73 \pm 0.01	0.00 \pm 0.00	1.40 \pm 0.00	563.76 \pm 69.07	\checkmark
OURS	ROCKET	92.15 \pm 0.01	0.00 \pm 0.00	14.66 \pm 0.10	535.14 \pm 105.65	\checkmark
OURS	SEA	87.39 \pm 0.02	0.00 \pm 0.00	3.40 \pm 0.00	542.19 \pm 87.80	\checkmark

Table 8: VGG11 CIFAR-100 class unlearning breakdown

METHOD	FORGET CLASS	\mathcal{D}_r	\mathcal{D}_f	MIA	METHOD RUNTIME	ZS
BASLINE	BABY	66.88 \pm 0.00	52.00 \pm 0.00	88.20 \pm 0.00	74.90 \pm 0.60	×
BASLINE	LAMP	66.79 \pm 0.00	61.00 \pm 0.00	96.80 \pm 0.00	75.10 \pm 0.90	×
BASLINE	MUSHROOM	66.67 \pm 0.00	73.00 \pm 0.00	97.40 \pm 0.00	74.10 \pm 1.10	×
BASLINE	ROCKET	66.63 \pm 0.00	77.00 \pm 0.00	97.40 \pm 0.00	75.00 \pm 1.00	×
BASLINE	SEA	66.65 \pm 0.00	75.00 \pm 0.00	82.60 \pm 0.00	75.50 \pm 1.10	×
RETRAIN	BABY	63.53 \pm 0.47	0.00 \pm 0.00	12.20 \pm 1.90	1100.50 \pm 2.60	×
RETRAIN	LAMP	63.42 \pm 0.45	0.00 \pm 0.00	9.40 \pm 1.20	1099.40 \pm 2.10	×
RETRAIN	MUSHROOM	63.47 \pm 0.25	0.00 \pm 0.00	6.60 \pm 1.00	1100.00 \pm 2.10	×
RETRAIN	ROCKET	63.21 \pm 0.54	0.00 \pm 0.00	10.40 \pm 1.10	1100.00 \pm 3.00	×
RETRAIN	SEA	63.28 \pm 0.34	0.00 \pm 0.00	18.40 \pm 1.00	1100.30 \pm 2.40	×
FINETUNE	BABY	60.02 \pm 0.64	0.00 \pm 0.00	26.50 \pm 5.10	102.80 \pm 0.70	×
FINETUNE	LAMP	59.97 \pm 0.60	1.80 \pm 1.20	18.40 \pm 3.20	102.80 \pm 1.50	×
FINETUNE	MUSHROOM	60.36 \pm 0.58	1.40 \pm 1.30	12.90 \pm 3.00	102.40 \pm 1.20	×
FINETUNE	ROCKET	59.70 \pm 0.43	3.90 \pm 3.00	13.20 \pm 4.20	103.20 \pm 1.40	×
FINETUNE	SEA	59.73 \pm 0.61	0.00 \pm 0.00	28.00 \pm 7.10	103.40 \pm 0.60	×
AMNESIAC	BABY	64.70 \pm 0.36	0.00 \pm 0.00	4.30 \pm 1.10	99.50 \pm 3.40	×
AMNESIAC	LAMP	64.58 \pm 0.37	0.00 \pm 0.00	4.80 \pm 1.20	99.90 \pm 3.00	×
AMNESIAC	MUSHROOM	64.49 \pm 0.52	0.00 \pm 0.00	3.10 \pm 1.10	100.80 \pm 3.30	×
AMNESIAC	ROCKET	64.34 \pm 0.38	0.00 \pm 0.00	1.80 \pm 0.80	99.70 \pm 2.80	×
AMNESIAC	SEA	64.39 \pm 0.40	0.00 \pm 0.00	1.20 \pm 0.30	100.20 \pm 3.50	×
SCRUB	BABY	67.05 \pm 0.11	0.00 \pm 0.00	9.30 \pm 1.50	109.10 \pm 2.60	×
SCRUB	LAMP	66.81 \pm 0.10	0.00 \pm 0.00	8.70 \pm 0.90	108.50 \pm 3.50	×
SCRUB	MUSHROOM	67.05 \pm 0.08	0.00 \pm 0.00	9.40 \pm 0.50	107.60 \pm 2.80	×
SCRUB	ROCKET	66.19 \pm 0.14	0.00 \pm 0.00	8.20 \pm 1.70	108.70 \pm 2.50	×
SCRUB	SEA	66.71 \pm 0.12	0.00 \pm 0.00	6.30 \pm 2.40	107.70 \pm 2.70	×
SSD	BABY	52.68 \pm 0.01	0.00 \pm 0.00	7.40 \pm 0.00	81.30 \pm 1.30	×
SSD	LAMP	65.44 \pm 0.00	0.00 \pm 0.00	5.80 \pm 0.00	81.60 \pm 0.80	×
SSD	MUSHROOM	62.19 \pm 0.00	0.00 \pm 0.00	14.60 \pm 0.00	81.20 \pm 1.00	×
SSD	ROCKET	63.79 \pm 0.01	0.00 \pm 0.00	8.60 \pm 0.00	81.50 \pm 0.90	×
SSD	SEA	32.75 \pm 0.01	0.00 \pm 0.00	7.00 \pm 0.00	81.30 \pm 1.20	×
TEACHER	BABY	66.05 \pm 0.33	1.00 \pm 1.10	0.10 \pm 0.10	79.60 \pm 0.80	×
TEACHER	LAMP	65.96 \pm 0.19	0.40 \pm 1.00	0.00 \pm 0.00	79.10 \pm 1.30	×
TEACHER	MUSHROOM	65.67 \pm 0.27	0.90 \pm 1.50	0.10 \pm 0.10	79.60 \pm 1.40	×
TEACHER	ROCKET	65.51 \pm 0.24	0.10 \pm 0.30	0.00 \pm 0.10	79.10 \pm 1.00	×
TEACHER	SEA	65.57 \pm 0.24	5.20 \pm 4.60	0.00 \pm 0.00	86.70 \pm 24.00	×
UNSIR	BABY	64.88 \pm 0.43	3.80 \pm 2.30	29.40 \pm 11.00	111.10 \pm 2.00	×
UNSIR	LAMP	64.72 \pm 0.35	25.90 \pm 6.10	17.00 \pm 5.90	111.60 \pm 2.50	×
UNSIR	MUSHROOM	64.79 \pm 0.27	21.10 \pm 10.70	11.30 \pm 5.20	111.30 \pm 2.90	×
UNSIR	ROCKET	64.58 \pm 0.39	42.90 \pm 14.30	40.70 \pm 12.10	112.10 \pm 3.50	×
UNSIR	SEA	64.51 \pm 0.33	13.90 \pm 7.50	22.90 \pm 4.50	111.30 \pm 2.90	×
GKT	BABY	2.31 \pm 0.27	0.00 \pm 0.00	47.90 \pm 26.50	634.30 \pm 9.20	✓
GKT	LAMP	2.42 \pm 0.37	0.00 \pm 0.00	45.30 \pm 23.70	630.80 \pm 8.60	✓
GKT	MUSHROOM	2.23 \pm 0.29	0.00 \pm 0.00	47.40 \pm 11.80	628.90 \pm 6.00	✓
GKT	ROCKET	2.31 \pm 0.24	0.00 \pm 0.00	56.20 \pm 20.00	629.00 \pm 6.90	✓
GKT	SEA	2.42 \pm 0.39	0.00 \pm 0.00	60.50 \pm 28.60	632.00 \pm 5.60	✓
EMMN	BABY	30.11 \pm 8.56	7.80 \pm 9.20	54.00 \pm 11.20	274.70 \pm 4.40	✓
EMMN	LAMP	33.02 \pm 9.02	21.50 \pm 17.00	54.70 \pm 14.60	276.00 \pm 4.60	✓
EMMN	MUSHROOM	31.87 \pm 11.81	13.50 \pm 12.40	53.40 \pm 12.80	275.40 \pm 4.60	✓
EMMN	ROCKET	26.91 \pm 7.74	24.30 \pm 23.70	58.20 \pm 14.50	274.80 \pm 3.90	✓
EMMN	SEA	30.28 \pm 8.81	33.30 \pm 21.40	69.10 \pm 16.00	275.40 \pm 4.70	✓
BDSH	BABY	66.88 \pm 0.0	16.9 \pm 0.3	4.3 \pm 0.0	85.16 \pm 1.1	✓
BDSH	LAMP	66.31 \pm 0.0	16.7 \pm 0.5	14.68 \pm 0.0	85.4 \pm 2.2	✓
BDSH	MUSHROOM	66.83 \pm 0.0	21.67 \pm 0.0	12.64 \pm 0.0	85.27 \pm 1.8	✓
BDSH	ROCKET	66.17 \pm 0.0	13.0 \pm 0.0	2.9 \pm 0.0	85.01 \pm 1.6	✓
OURS	BABY	67.06 \pm 0.04	11.10 \pm 1.40	9.90 \pm 0.30	78.70 \pm 2.40	✓

Table 9: VGG11 face unlearning class breakdown

METHOD	FORGET CLASS	\mathcal{D}_r	\mathcal{D}_f	MIA	METHOD RUNTIME	ZS
BASELINE	1	93.95 \pm 0.00	93.88 \pm 0.00	13.82 \pm 0.00	72.93 \pm 0.52	×
BASELINE	10	93.99 \pm 0.00	95.90 \pm 0.00	11.48 \pm 0.00	72.92 \pm 0.84	×
BASELINE	20	94.06 \pm 0.00	84.48 \pm 0.00	65.19 \pm 0.00	72.82 \pm 0.98	×
BASELINE	30	93.93 \pm 0.00	97.46 \pm 0.00	11.17 \pm 0.00	72.64 \pm 0.85	×
BASELINE	40	94.01 \pm 0.00	92.31 \pm 0.00	15.38 \pm 0.00	72.66 \pm 0.62	×
RETRAIN	1	100.00 \pm 0.00	0.00 \pm 0.00	2.63 \pm 0.76	491.93 \pm 2.02	×
RETRAIN	10	100.00 \pm 0.00	0.00 \pm 0.00	0.74 \pm 0.47	494.31 \pm 1.83	×
RETRAIN	20	100.00 \pm 0.00	0.00 \pm 0.00	2.49 \pm 1.02	492.70 \pm 2.20	×
RETRAIN	30	100.00 \pm 0.00	0.00 \pm 0.00	5.98 \pm 1.15	491.86 \pm 2.70	×
RETRAIN	40	100.00 \pm 0.00	0.00 \pm 0.00	4.79 \pm 1.52	494.67 \pm 2.37	×
FINETUNE	1	97.61 \pm 0.67	36.87 \pm 9.86	4.28 \pm 2.69	84.67 \pm 0.88	×
FINETUNE	10	95.58 \pm 2.93	43.77 \pm 21.26	8.61 \pm 6.74	84.50 \pm 0.51	×
FINETUNE	20	98.08 \pm 0.48	25.10 \pm 7.26	5.41 \pm 1.84	84.76 \pm 0.74	×
FINETUNE	30	97.92 \pm 0.82	23.78 \pm 9.66	6.48 \pm 3.74	84.81 \pm 0.89	×
FINETUNE	40	97.65 \pm 0.57	16.58 \pm 7.11	6.32 \pm 3.23	84.51 \pm 0.77	×
AMNESIAC	1	99.72 \pm 0.13	0.00 \pm 0.00	1.38 \pm 1.33	82.19 \pm 1.10	×
AMNESIAC	10	99.67 \pm 0.14	0.00 \pm 0.00	0.98 \pm 0.75	82.19 \pm 0.61	×
AMNESIAC	20	99.70 \pm 0.12	0.00 \pm 0.00	1.49 \pm 0.64	82.14 \pm 0.87	×
AMNESIAC	30	99.67 \pm 0.10	0.00 \pm 0.00	1.12 \pm 0.65	82.14 \pm 0.74	×
AMNESIAC	40	99.67 \pm 0.17	0.00 \pm 0.00	1.28 \pm 1.16	81.92 \pm 0.70	×
SCRUB	1	98.83 \pm 0.02	97.14 \pm 0.00	8.75 \pm 0.76	85.48 \pm 0.85	×
SCRUB	10	98.85 \pm 0.01	97.54 \pm 0.00	10.49 \pm 0.35	85.50 \pm 0.79	×
SCRUB	20	98.83 \pm 0.03	98.00 \pm 0.32	74.20 \pm 0.74	85.45 \pm 1.13	×
SCRUB	30	98.88 \pm 0.03	95.89 \pm 0.00	81.84 \pm 0.92	85.57 \pm 1.03	×
SCRUB	40	98.86 \pm 0.02	96.50 \pm 0.49	78.03 \pm 1.14	84.88 \pm 0.75	×
SSD	1	55.77 \pm 0.01	0.00 \pm 0.00	3.95 \pm 0.00	92.03 \pm 42.05	×
SSD	10	73.68 \pm 0.01	0.00 \pm 0.00	2.46 \pm 0.00	86.66 \pm 27.23	×
SSD	20	0.82 \pm 0.00	0.00 \pm 0.00	100.00 \pm 0.00	78.18 \pm 1.46	×
SSD	30	85.96 \pm 0.00	0.00 \pm 0.00	10.06 \pm 0.00	82.12 \pm 9.97	×
SSD	40	47.43 \pm 0.01	0.00 \pm 0.00	4.27 \pm 0.00	78.22 \pm 1.07	×
TEACHER	1	93.70 \pm 0.31	0.00 \pm 0.00	0.00 \pm 0.00	75.02 \pm 0.73	×
TEACHER	10	93.72 \pm 0.27	0.00 \pm 0.00	0.00 \pm 0.00	74.64 \pm 0.83	×
TEACHER	20	93.90 \pm 0.26	0.00 \pm 0.00	0.28 \pm 0.39	74.74 \pm 0.84	×
TEACHER	30	93.46 \pm 0.34	0.08 \pm 0.16	0.06 \pm 0.18	74.53 \pm 0.72	×
TEACHER	40	93.80 \pm 0.24	0.43 \pm 0.92	0.09 \pm 0.27	75.04 \pm 0.78	×
UNSIR	1	99.49 \pm 0.12	74.43 \pm 9.17	13.62 \pm 8.87	87.32 \pm 0.45	×
UNSIR	10	99.39 \pm 0.34	87.13 \pm 4.20	44.84 \pm 9.23	87.66 \pm 0.64	×
UNSIR	20	99.44 \pm 0.14	54.92 \pm 15.19	10.33 \pm 4.72	86.78 \pm 0.86	×
UNSIR	30	99.43 \pm 0.14	65.72 \pm 11.97	13.24 \pm 5.21	87.76 \pm 1.04	×
UNSIR	40	99.54 \pm 0.11	48.97 \pm 23.85	6.92 \pm 6.98	87.18 \pm 0.80	×
GKT	1	2.02 \pm 0.59	0.00 \pm 0.00	23.88 \pm 30.25	441.46 \pm 1.81	✓
GKT	10	2.01 \pm 0.57	0.00 \pm 0.00	53.44 \pm 44.14	439.67 \pm 3.07	✓
GKT	20	1.98 \pm 0.59	0.00 \pm 0.00	28.18 \pm 34.79	441.43 \pm 2.98	✓
GKT	30	2.15 \pm 0.66	0.00 \pm 0.00	28.38 \pm 40.11	439.63 \pm 3.38	✓
GKT	40	2.11 \pm 0.63	0.00 \pm 0.00	44.87 \pm 49.71	441.20 \pm 2.28	✓
EMMN	1	50.98 \pm 13.53	69.30 \pm 25.68	26.91 \pm 17.76	283.65 \pm 1.33	✓
EMMN	10	51.78 \pm 14.34	65.00 \pm 21.82	27.79 \pm 26.23	283.64 \pm 1.87	✓
EMMN	20	41.38 \pm 12.94	26.22 \pm 17.24	55.75 \pm 13.54	282.91 \pm 0.70	✓
EMMN	30	45.42 \pm 18.22	63.02 \pm 27.70	32.91 \pm 21.73	283.73 \pm 1.17	✓
EMMN	40	46.60 \pm 16.99	26.50 \pm 12.03	60.94 \pm 7.43	282.83 \pm 1.26	✓
BDSH	1	93.60 \pm 0.0	79.43 \pm 0.0	42.4 \pm 0.0	91.24 \pm 1.7	✓
BDSH	10	94.0 \pm 0.0	93.44 \pm 0.0	79.51 \pm 0.0	91.10 \pm 0.7	✓
BDSH	20	93.68 \pm 0.0	64.18 \pm 0.0	32.93 \pm 0.0	90.85 \pm 1.8	✓
BDSH	30	94.11 \pm 0.0	90.23 \pm 0.0	65.9 \pm 0.0	90.46 \pm 3.5	✓
BDSH	40	94.01 \pm 0.0	90.60 \pm 0.0	73.5 \pm 0.0	90.89 \pm 1.5	✓

Table 10: VGG11 CIFAR-20 sub-class unlearning breakdown

METHOD	FORGET CLASS	\mathcal{D}_r	\mathcal{D}_f	MIA	METHOD RUNTIME	ZS
BASELINE	BABY	75.21 \pm 0.00	82.29 \pm 0.00	74.80 \pm 0.00	76.03 \pm 1.48	\times
BASELINE	LAMP	75.45 \pm 0.00	56.86 \pm 0.00	72.00 \pm 0.00	75.29 \pm 0.80	\times
BASELINE	MUSHROOM	75.28 \pm 0.00	75.61 \pm 0.00	73.40 \pm 0.00	75.73 \pm 0.84	\times
BASELINE	ROCKET	75.26 \pm 0.00	78.99 \pm 0.00	83.00 \pm 0.00	75.39 \pm 1.05	\times
BASELINE	SEA	75.09 \pm 0.00	92.88 \pm 0.00	90.60 \pm 0.00	75.02 \pm 1.19	\times
RETRAIN	BABY	72.60 \pm 0.23	67.69 \pm 3.00	55.36 \pm 1.20	425.76 \pm 1.39	\times
RETRAIN	LAMP	72.99 \pm 0.23	18.01 \pm 2.81	17.90 \pm 2.03	425.73 \pm 1.64	\times
RETRAIN	MUSHROOM	72.98 \pm 0.33	8.78 \pm 2.23	12.14 \pm 1.28	425.62 \pm 1.72	\times
RETRAIN	ROCKET	72.87 \pm 0.21	11.46 \pm 2.80	14.06 \pm 1.30	425.49 \pm 1.86	\times
RETRAIN	SEA	72.61 \pm 0.27	85.53 \pm 2.59	66.76 \pm 2.29	425.21 \pm 1.70	\times
FINETUNE	BABY	65.45 \pm 0.84	60.45 \pm 6.82	53.20 \pm 5.89	120.34 \pm 1.07	\times
FINETUNE	LAMP	65.96 \pm 1.55	17.53 \pm 5.18	22.56 \pm 4.10	121.19 \pm 1.31	\times
FINETUNE	MUSHROOM	65.69 \pm 1.06	6.13 \pm 5.96	18.52 \pm 4.15	121.28 \pm 0.64	\times
FINETUNE	ROCKET	65.51 \pm 0.72	6.23 \pm 3.69	22.26 \pm 5.51	120.56 \pm 0.87	\times
FINETUNE	SEA	64.95 \pm 0.76	82.15 \pm 4.64	70.98 \pm 9.58	120.37 \pm 1.12	\times
AMNESIAC	BABY	73.70 \pm 0.25	53.55 \pm 6.44	4.96 \pm 0.90	97.95 \pm 1.11	\times
AMNESIAC	LAMP	74.18 \pm 0.24	9.20 \pm 2.38	7.08 \pm 1.62	97.62 \pm 0.90	\times
AMNESIAC	MUSHROOM	74.05 \pm 0.24	2.52 \pm 1.60	4.28 \pm 0.94	97.81 \pm 1.00	\times
AMNESIAC	ROCKET	73.80 \pm 0.19	2.43 \pm 2.43	2.98 \pm 0.86	98.39 \pm 0.85	\times
AMNESIAC	SEA	73.65 \pm 0.16	64.00 \pm 17.71	1.22 \pm 0.53	97.64 \pm 1.23	\times
SCRUB	BABY	75.04 \pm 0.26	76.99 \pm 2.26	71.00 \pm 2.92	130.78 \pm 2.81	\times
SCRUB	LAMP	75.44 \pm 0.16	41.22 \pm 2.77	40.22 \pm 6.66	130.45 \pm 2.12	\times
SCRUB	MUSHROOM	75.40 \pm 0.13	15.52 \pm 7.61	11.82 \pm 0.85	130.79 \pm 2.64	\times
SCRUB	ROCKET	62.44 \pm 28.38	10.12 \pm 22.48	16.71 \pm 21.68	131.41 \pm 3.75	\times
SCRUB	SEA	75.05 \pm 0.12	91.18 \pm 1.56	83.13 \pm 3.26	130.69 \pm 2.52	\times
SSD	BABY	71.97 \pm 0.00	0.00 \pm 0.00	8.80 \pm 0.00	83.78 \pm 1.04	\times
SSD	LAMP	73.98 \pm 0.00	3.56 \pm 0.00	9.00 \pm 0.00	84.39 \pm 1.09	\times
SSD	MUSHROOM	71.99 \pm 0.00	0.00 \pm 0.00	6.20 \pm 0.00	84.03 \pm 1.25	\times
SSD	ROCKET	74.96 \pm 0.00	4.17 \pm 0.00	11.00 \pm 0.00	83.91 \pm 0.81	\times
SSD	SEA	74.80 \pm 0.00	69.62 \pm 0.00	56.20 \pm 0.00	83.54 \pm 1.42	\times
TEACHER	BABY	74.87 \pm 0.17	77.90 \pm 1.25	0.38 \pm 0.71	78.74 \pm 0.59	\times
TEACHER	LAMP	75.18 \pm 0.11	35.55 \pm 13.07	0.40 \pm 0.38	79.47 \pm 1.10	\times
TEACHER	MUSHROOM	74.99 \pm 0.13	27.47 \pm 14.49	0.25 \pm 0.09	90.58 \pm 36.91	\times
TEACHER	ROCKET	74.86 \pm 0.20	48.36 \pm 16.87	0.07 \pm 0.10	91.51 \pm 38.05	\times
TEACHER	SEA	74.71 \pm 0.14	85.28 \pm 2.58	0.08 \pm 0.14	78.36 \pm 0.90	\times
UNSIR	BABY	73.85 \pm 0.20	77.66 \pm 2.12	77.60 \pm 2.52	110.77 \pm 0.90	\times
UNSIR	LAMP	74.17 \pm 0.27	44.52 \pm 7.40	55.06 \pm 6.21	110.70 \pm 0.73	\times
UNSIR	MUSHROOM	74.03 \pm 0.32	47.14 \pm 4.00	41.70 \pm 8.15	110.58 \pm 1.31	\times
UNSIR	ROCKET	74.08 \pm 0.24	57.47 \pm 10.25	57.44 \pm 8.61	111.27 \pm 0.66	\times
UNSIR	SEA	73.80 \pm 0.17	90.75 \pm 1.83	85.32 \pm 3.83	109.55 \pm 1.01	\times
BDSH	BABY	74.25 \pm 0.0	59.12 \pm 0.0	31.2 \pm 0.0	86.05 \pm 1.1	\checkmark
BDSH	LAMP	75.3 \pm 0.0	36.46 \pm 0.0	36.8 \pm 0.0	85.24 \pm 1.3	\checkmark
BDSH	MUSHROOM	75.06 \pm 0.0	36.63 \pm 0.0	25.74 \pm 0.0	86.47 \pm 2.5	\checkmark
BDSH	ROCKET	74.41 \pm 0.0	17.54 \pm 0.0	12.88 \pm 0.0	85.91 \pm 1.6	\checkmark
BDSH	SEA	72.88 \pm 0.0	39.15 \pm 0.0	11.84 \pm 0.0	85.254 \pm 1.2	\checkmark
OURS	BABY	73.57 \pm 1.28	49.82 \pm 19.43	33.32 \pm 17.53	79.47 \pm 3.63	\checkmark
OURS	LAMP	74.05 \pm 1.34	32.46 \pm 15.76	32.38 \pm 19.76	79.73 \pm 3.22	\checkmark
OURS	MUSHROOM	74.20 \pm 0.69	33.67 \pm 12.27	23.52 \pm 12.84	79.14 \pm 4.00	\checkmark
OURS	ROCKET	73.68 \pm 0.84	19.31 \pm 18.31	11.20 \pm 7.78	80.14 \pm 3.27	\checkmark
OURS	SEA	73.07 \pm 1.02	37.82 \pm 22.10	17.29 \pm 20.20	79.04 \pm 3.25	\checkmark