

# Towards Provably Unlearnable Examples via Bayes Error Optimisation

Ruihan Zhang, Jun Sun, Ee-Peng Lim, Peixin Zhang\*

Singapore Management University  
80 Stamford Rd, Singapore, Singapore 178902  
pxzhang@smu.edu.sg

## Abstract

The recent success of machine learning models, especially large-scale classifiers and language models, relies heavily on training with massive data. These data are often collected from online sources. This raises serious concerns about the protection of user data, as individuals may not have given consent for their data to be used in training. To address this concern, recent studies introduce the concept of unlearnable examples, i.e., data instances that appear natural but are intentionally altered to prevent models from effectively learning from them. While existing methods demonstrate empirical effectiveness, they typically rely on heuristic trials and lack formal guarantees. Besides, when unlearnable examples are mixed with clean data, as is often the case in practice, their unlearnability disappears. In this work, we propose a novel approach to constructing unlearnable examples by systematically maximising the Bayes error, a measurement of irreducible classification error. We develop an optimisation-based approach and provide an efficient solution using projected gradient ascent. Our method provably increases the Bayes error and remains effective when the unlearning examples are mixed with clean samples. Experimental results across multiple datasets and model architectures are consistent with our theoretical analysis and show that our approach can restrict data learnability, effectively in practice.

## 1 Introduction

In recent years, machine learning models have demonstrated remarkable performance across a wide range of tasks, driven by large-scale datasets and powerful computational resources (Kaplan et al. 2020; He et al. 2016). However, this data-centric mechanism also raises serious concerns regarding intellectual property and control over data usage (Yeom et al. 2018; Huang et al. 2021). Once data is shared or collected, often without users’ full awareness, it becomes extremely difficult to restrict its subsequent usage or to prevent it from being exploited by unintended parties. For example, social media users are happy to share personal photos on blogs, yet still wish to prevent unauthorised third parties from leveraging these images for model training.

Motivated by the need for data protection at the source, recent studies have introduced the concept of unlearnable

examples, a proactive defence mechanism that applies imperceptible perturbations to clean data, rendering it unlearnable by machine learning models (Huang et al. 2021). Returning to the earlier example, a user can apply subtle, human-imperceptible perturbations to their photos before sharing them, thereby ensuring that any model trained on these perturbed images performs poorly when attempting to recognise or analyse the individual’s face in unaltered images (Wen et al. 2023). This technique not only enables data owners to regain control over their data but also introduces a novel perspective in machine learning by designing data that actively resists learning.

Prior works for creating unlearnable examples take two main directions. One line of research perturbs data to increase training loss, aiming to obstruct the model from extracting useful patterns (Fowl et al. 2021). Another line encourages overfitting by minimising training error, thereby causing models to overfit non-generalisable features (Fu et al. 2022). While existing methods show empirical effectiveness, they typically require altering the entire training set. In practice, however, unlearnable examples are mixed with clean data from diverse sources (Liu, Peng, and Tang 2023; Nakashima et al. 2022; Seddik and Tamaazousti 2022). In such cases, for example, training a model on unlearnable examples (CIFAR-10 (Krizhevsky, Hinton et al. 2009)) may yield 20.50% test accuracy. But when these unlearnable examples are mixed with the same number of clean examples for training, the test accuracy can be boosted to 92.51%. This greatly limits the effectiveness of existing methods (Huang et al. 2021).

In addition, a formal guarantee is underrated in existing methods, i.e., these strategies are heuristic instead of from a systematic derivation of unlearnability. This absence of theoretical grounding poses risks to interpretability and reliability, especially in safety-critical settings (Wu et al. 2018). Therefore, there is a pressing need for a clear framework that not only provides theoretical guarantees for unlearnability but also remains effective under realistic conditions where clean and unlearnable data coexist.

In this work, we set out from a statistical learning concept of Bayes error, which is the classification error rate of the Maximum A Posterior rule (Fukunaga 1990). Essentially, Bayes error quantifies the inherent difficulty of classification with a given data distribution. Here, we leverage Bayes error as an effective measurement of data unlearnability, i.e.,

\*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

higher Bayes error means more unlearnable. Then, by systematically increasing the Bayes error through constrained perturbations, we can consequently ensure that the perturbed examples become provably harder to learn from, regardless of the training algorithm. Furthermore, this method remains effective even when unlearnable examples are mixed with clean samples.

To increase the Bayes error, we set up an optimisation problem to maximise the estimated Bayes error under norm-bounded constraints (to maintain the data quality). When solving this optimisation problem, we present an efficient solution with projected gradient ascent (Bertsekas 2003). We show through a theoretical analysis that the Bayes error is guaranteed to increase. We then evaluate our method and implementation through extensive empirical studies. Experimental results demonstrate that our approach consistently leads to a test accuracy drop. For example, on a CIFAR-10 training set with 50% clean and 50% unlearnable examples, training on just the clean half gives 91.16% test accuracy. Adding unlearnable examples from existing methods increases accuracy to 92.51%, which defeats the purpose of being unlearnable. In contrast, using ours drops it to 69.68%. This shows our method effectively produces unlearnable examples. Our code is available at <https://github.com/cat-claws/unlearnable>.

## 2 Preliminaries and Problem Definition

We first describe the machine learning background, including distributions, learning algorithms, and the concept of unlearnable examples. Then, we define our research problem.

**Notations in machine learning** Let  $D$  denote a joint probability distribution for random variables  $\mathbf{x} \in \mathbb{X}$  and  $\mathbf{y} \in \mathbb{Y}$ , such that for some positive integer  $n$ , a sample from  $D$  is captured by  $\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$  or  $\{(x_i, y_i)\}_{i=1}^n \sim D$ , where  $n$  is the sample size.

When it comes to machine learning regarding  $D$ , we consider  $\mathbb{X}$  an input feature space, and consider  $\mathbb{Y}$  a label space. Further, we would denote a model (learner) function  $h : \mathbb{X} \rightarrow \mathbb{Y}$  that predicts outputs  $h(x) \in \mathbb{Y}$  based on a (possibly high-dimensional) input point  $x \in \mathbb{X}$ . The quality of a model can be measured through  $\mathbb{E}_{\mathbf{x}, \mathbf{y} \sim D} [\ell(h, \mathbf{x}, \mathbf{y})]$  with a problem-dependent loss function  $\ell(h, x, y)$  (Zhang 2004).

**Unlearnable examples** In this work, we explore a scenario in which some of the training data needs to be protected from being learned by the model. A defender can read the original training dataset  $\{(x_i, y_i)\}_{i=1}^n \sim D$ , and can apply small perturbations  $\Delta x_i$  to these samples. Typically, we have  $\|\Delta x_i\|_p \leq \epsilon$  ( $\epsilon > 0$  caps the perturbation range within an  $L^p$ -norm) such that the data quality is not affected for human perception. This results in dataset  $\{(x'_i, y_i)\}_{i=1}^n$ , where  $\forall i \in 1, \dots, n, x'_i = x_i + \Delta x_i$ . The goal is that when an arbitrary model trains on  $\{(x'_i, y_i)\}_{i=1}^n$  (and possibly with additional data other than  $\{(x_i, y_i)\}_{i=1}^n$ ), the resultant model performs badly in testing (or at least worse than a model trained without the perturbed examples). Formally, let  $h$  be a model selected by arbitrary learning algorithm  $\Gamma(\cdot)$  such that  $h = \Gamma(x_1 + \Delta x_1, y_1, \dots, x_n + \Delta x_n, y_n)$ , and this goal

could be expressed as

$$\begin{aligned} \max \min_{h \in \mathbb{H}} \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim D} [\ell(h, \mathbf{x}, \mathbf{y})], \\ \text{s.t. } \forall i \in 1, 2, \dots, n, \quad \|\Delta x_i\|_p \leq \epsilon \end{aligned} \quad (1)$$

where  $\mathbb{H}$  is the space of possible models that can be selected. Optimisation problem (1) represents a bilevel optimisation formulation, i.e., the minimisation here trains  $h$ , and maximisation manipulates  $\Delta x_i$  to prevent trained  $h$  from performing well in testing.

In practice, instead of solving optimisation problem (1), two alternative optimisation problems are tackled instead, as shown in optimisation problem (2) and optimisation problem (3).

$$\max_{\Delta x_i} \min_{h \in \{\mathbb{X} \rightarrow \mathbb{Y}\}} \frac{1}{n} \sum_{i=1}^n \ell_{\text{train}}(h, x_i + \Delta x_i, y_i) \quad (2)$$

$$\min_{\Delta x_i} \min_{h \in \{\mathbb{X} \rightarrow \mathbb{Y}\}} \frac{1}{n} \sum_{i=1}^n \ell_{\text{train}}(h, x_i + \Delta x_i, y_i) \quad (3)$$

The intuition behind optimisation problem (2) is that if the training loss remains high and the model fails to fit the training data, then naturally not much information has been transferred to the model. Conversely, for optimisation problem (3), the idea is that if the training loss is very low and the model overfits the training data, then what has been learned may not generalise well.

### Problem definition

Next, we define our problem. Although existing methods are shown to be relatively effective in some settings, they suffer from two main shortcomings. First, and most importantly, they work only if the entire training dataset is an unlearnable example, which is hardly practical. Second, these strategies do not come with a formal guarantee for unlearnability growth. To address the above-mentioned shortcoming, we thus aim to develop a theoretically grounded method such that it works even if only a small portion of the data is perturbed to be unlearnable.

## 3 Unlearnable Example Construction

We propose a novel approach for constructing unlearnable examples. We begin by formalising data unlearnability based on Bayes error, then introduce a differentiable estimate from samples. We further design a method to provably increase the Bayes error to hinder learning, and discuss whether such an approach works in the practical scenario where clean and unlearnable data coexist.

### 3.1 Bayes Error as Unlearnability Measurement

We focus on the classification problem in this work. In such a setting, the Bayes error represents the minimal inevitable error of any classifier. It is defined as the expected error probability of the maximum a posteriori rule. Formally, given a distribution  $D$  over  $\mathbb{X} \times \mathbb{Y}$ , the Bayes error (rate) of  $D$  can be expressed (Fukunaga 1990; Garber and Djouadi 1988) as:

$$\mathbb{E}_{\mathbf{x} \sim D_{\mathbf{x}}} \left[ 1 - \max_c p(\mathbf{y} = c | \mathbf{x}) \right] \quad (4)$$

As shown in the optimisation problem (1), a model attempts to learn from data by minimising the prediction error, which is the objective that the defender aims to counteract by increasing this error. Since the Bayes error represents the minimal inevitable error of any classifier, it provides a fundamental upper bound on the wild model’s performance. Therefore, by increasing the Bayes error, the defender effectively lower the upper bound of the model’s performance, which equivalently increases the unlearnability of the data, thereby enhancing the unlearnability of the data.

### 3.2 Increasing Bayes Error Given Finite Samples

While Bayes error computation is straightforward with a known distribution (Eq. (4)), it is not clear how to compute Bayes error from samples. To this end, we first present how to approximate Bayes error from samples, and then propose a method to systematically increase the Bayes error.

**Estimating Bayes error through local posterior averaging**  
Given only sampled data  $\{(x_i, y_i)\}_{i=1}^n \sim D$ , rather than the distribution  $D$ , we do not have access to the true posterior  $p_D(y | \mathbf{x})$ . The posterior may appear quantifiable through the following expression.

$$\hat{p}_D(y | \mathbf{x} = x) = \frac{\sum_j \mathbb{I}[(x_i = x) \wedge (y = y_j)]}{\sum_i \mathbb{I}[x_i = x]} \quad (5)$$

Eq. (5) measures the frequency of each outcome. Yet, it is likely that some outcomes are never observed, e.g., when the  $\mathbb{X}$  space is continuous, making the Eq. (5) undefined. To address this issue, we present Algorithm 1 to estimate Bayes error, where the similarity function  $s$  can be defined according to the joint distribution pattern, e.g., a Gaussian kernel function

$$s(x_1, x_2) = \exp(-\|x_1 - x_2\|^2 / 2\sigma^2). \quad (6)$$

---

#### Algorithm 1: Local posterior averaging

---

**Require:** Sampled data  $\{(x_i, y_i)\}_{i=1}^n$ ,  $\mathbb{Y}$  space, similarity function  $s : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$

**Ensure:** Posterior values  $\{\hat{p}_D(y | \mathbf{x} = x_i)\}_{i=1}^n$

- 1: **for**  $i \leftarrow 1, 2, \dots, n$  **do**
  - 2:     **for**  $c \in \mathbb{Y}$  **do**
  - 3:          $\hat{p}_D(y = c | \mathbf{x} = x_i) \leftarrow \frac{\sum_{j=1, j \neq i}^n \mathbb{I}[y_j = c] \cdot s(x_j, x_i)}{\sum_{k=1, k \neq i}^n s(x_k, x_i)}$
  - 4:     **end for**
  - 5: **end for**
- 

Essentially, Algorithm 1 takes sampled data as input and for each sampled  $x$ , it provides the posterior at that  $x$ . This then provides the sampled posterior values, and therefore the sampled maximum posteriors. Bayes error can thus be estimated as

$$\hat{\beta}_D = 1 - \frac{1}{n} \sum_i \max_c \hat{p}_D(y = c | \mathbf{x} = x_i), \quad (7)$$

which is the fixed-sample form of Eq. (4).




Figure 1: A joint distribution of two truncated normal distributions. (Left) The density function of this distribution. The grey-shaded area represents the Bayes error. (Right) A sample is drawn from this distribution. Note that the samples are randomly placed in the vertical direction.

*Example 1.* Consider a binary classification problem in  $\mathbb{R}$  with two truncated normal distributions, as illustrated in the Figure 1. The Bayes error of such a distribution is 0.1427, which can be analytically calculated from the probability density function (PDF) of truncated normal distributions. Figure 1 (Right) illustrates sampled data drawn from this distribution. When estimating Bayes error from samples, a naive estimation like Eq. (5) yields an undefined Bayes error, because there are no overlapping sampled outcomes. In contrast, when applying Algorithm 1 and adopting the radial basis function (Chang et al. 2010) for similarity, we get 0.1426 as an estimate of Bayes error, which is within a  $\pm 0.1\%$  range of the analytic result.

In the following, we establish that the presented Bayes error estimate through local posterior averaging well approximates the analytic result.

**Lemma 3.1.** *Let  $s$  be a symmetric, non-negative similarity function satisfying the following conditions,*

$$\begin{aligned} \forall x'. \quad & \int s(x, x') dx = 1, \\ \forall \alpha > 0. \quad & \int_{\|x-x'\| > \alpha} s(x, x') dx \rightarrow 0 \quad \text{as } n \rightarrow \infty, \quad (8) \\ \exists \kappa > 0. \quad & \forall x, x', n. \quad 0 \leq s(x, x') \leq \kappa. \end{aligned}$$

*The posterior estimate from Algorithm 1 converges to the true posterior, i.e.,*

$$\mathbb{E}[|\hat{p}(y = c | \mathbf{x}) - p(y = c | \mathbf{x})|] \leq C_2 \cdot \sigma^2 + \frac{C_1}{\sqrt{n\sigma^d/2}}. \quad (9)$$

where  $C_1$  and  $C_2$  are constants.

Proof for Theorem 3.1, which leverages the bias and variance derivation for the Nadaraya-Watson estimator (Nadaraya 1964; Watson 1964), is provided in the extended version (Zhang et al. 2025).

**Theorem 3.2.** *With the posterior estimate from Algorithm 1, the Bayes error estimate converges to the true Bayes error at the rate  $2 \cdot C_2 \cdot |\mathbb{Y}| \sigma^2 + \frac{2 \cdot C_1 \cdot |\mathbb{Y}|}{\sqrt{n\sigma^d}}$ .*

*Proof.* From Eq. (4), the excess classification error

$|\hat{\beta}_D - \beta_D|$  can be expressed as

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}} \left[ 1 - \max_c \hat{p}(\mathbf{y} = c | \mathbf{x}) \right] - \mathbb{E}_{\mathbf{x}} \left[ 1 - \max_c p(\mathbf{y} = c | \mathbf{x}) \right] \\ & \leq 2 \cdot \mathbb{E}_{\mathbf{x}} \left[ \sum_c |p(\mathbf{y} = c | \mathbf{x}) - \hat{p}(\mathbf{y} = c | \mathbf{x})| \right] \end{aligned} \quad (10)$$

Now apply the posterior estimation error bound  $C_2\sigma^2 + \frac{C_1}{\sqrt{n\sigma^d}}$  in Theorem 3.1, where  $C_1$  and  $C_2$  are constants. Then, we get  $|\hat{\beta}_D - \beta_D| \leq 2 \cdot C_2 \cdot |\mathbb{Y}| \sigma^2 + \frac{2 \cdot C_1 \cdot |\mathbb{Y}|}{\sqrt{n\sigma^d}}$ , still within the same asymptotic order.  $\square$

**Bayes error maximisation under  $L^p$ -norm perturbation constraint** Given sampled data  $\{(x_i, y_i)\}_{i=1}^n \sim D$ , we would like to find a distribution  $D'$  close to  $D$  such that we have  $\hat{\beta}_{D'} > \hat{\beta}_D$ . Specifically, let  $\{(x'_i, y'_i)\}_{i=1}^n \sim D'$  be sampled data from  $D'$ , and an  $L^p$  closeness constraint between  $D'$  and  $D$  can be expressed as

$$\forall i \in [1, n] \cap \mathbb{N}. \quad (y'_i = y_i) \wedge (\|x'_i - x_i\|_p \leq \epsilon) \quad (11)$$

where  $\epsilon > 0$  is the  $x$  perturbation range. Since  $y'_i$  always equals  $y_i$  for all  $i$ , we use  $y_i$  whenever  $y'_i$  is needed. Applying Eq. (7), we can estimate Bayes error of  $D'$  as

$$\hat{\beta}_{D'} = \frac{1}{n} \sum_{i=1}^n \left( 1 - \max_c \hat{p}_{D'}(\mathbf{y} = c | \mathbf{x} = x'_i) \right). \quad (12)$$

We are now ready to define our constrained optimisation problem of finding  $D'$  as follows. Essentially, we aim to find  $x_1, \dots, x_n$  that maximise the Bayes error estimate, subject to an  $L^p$  norm constraint.

$$\begin{aligned} & \max_{\{x'_i\}_{i=1}^n} \frac{1}{n} \sum_{i=1}^n \left( 1 - \max_c \frac{\sum_{j=1, j \neq i}^n \mathbb{I}[y_i = c] \cdot s(x'_j, x'_i)}{\sum_{k=1, k \neq i}^n s(x'_k, x'_i)} \right) \\ & \text{s.t.} \quad \|x'_i - x_i\|_p \leq \epsilon, \quad \forall i = 1, 2, \dots, n \end{aligned} \quad (13)$$

Note that the objective function in Eq. (13) is an expanded form of Eq. (12) using Algorithm 1, and thus fully explicit given the original data  $x_i, y_i$  and the similarity function  $s$ . This objective function is also computationally tractable and differentiable almost everywhere.

Next, we present the solution to the optimisation problem above. Direct gradient ascent is not applicable in constrained settings, as it does not guarantee that the variable in each iteration remains within the constrained set. To address this, we employ projected gradient ascent (PGA), a first-order iterative method that enforces feasibility via projection. PGA operates by performing a standard gradient ascent step followed by a projection onto the constrained set. The details of our approach are shown in Algorithm 2

Essentially, each iteration of the gradient ascent algorithm requires computing gradients of the objective function (Bayes error estimate) with respect to all variable inputs  $x'_i$  (line 4). Then, we use this gradient to shift the current  $x'_i$  to the direction where the Bayes error grows (line 5). After that, we project the shift to the  $L^p$  constraint to obtain a feasible next-step  $x'_i$  (line 6). Line 6 in Algorithm 2 gives a general form

---

**Algorithm 2:** Projected gradient ascent for Bayes error

---

**Require:** Sampled data  $\{(x_i, y_i)\}_{i=1}^n$ ,  $\mathbb{Y}$  space, similarity function  $s : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ , step size  $\eta > 0$ ,  $L^p$  constraint parameter  $\epsilon$ , maximum iterations  $T$

**Ensure:**  $\{(x'_i, y_i)\}_{i=1}^n$

- 1: Initialize  $x'_i \leftarrow \mathbf{0}, \forall i \in 1, 2, \dots, n$
- 2: **for**  $t \leftarrow 0, 1, \dots, T - 1$  **do**
- 3:     **for**  $i = 1 \dots n$  **do**
- 4:          $\mathbf{g}_i^{(t)} \leftarrow -\nabla_{x'_i} \left( \sum_{j=1}^n \frac{\max_c \hat{p}_{D'}(c | x'_i)}{n} \right) \Big|_{\forall j. x'_j = x'_j^{(t)}}$
- 5:          $\delta_i^{(t+1)} \leftarrow x'_i^{(t)} + \eta \cdot \mathbf{g}_i^{(t)} - x_i$
- 6:          $\delta_i^{(t+1)} \leftarrow \arg \min_{\delta', \|\delta'\|_p \leq \epsilon} \left\| \delta' - \delta_i^{(t+1)} \right\|_2^2$
- 7:          $x'_i \leftarrow \delta_i^{(t+1)} + x_i$
- 8:     **end for**
- 9: **end for**

---




Figure 2: Perturbation of two-dimensional points.

expression for projection, and for  $L^\infty$ -norm, a closed-form projection can be expressed as  $\max(-\epsilon, \min(\delta_i^{(t+1)}, \epsilon))$ .

Note that gradient computation (Line 4 in Algorithm 2) is a fundamental step for PGA to work. Eq. (16) below shows the partial derivative of the objective function in Eq. (13) with respect to each  $x_i$ .

*Example 2.* To illustrate the effect of perturbations, we present an example in Figure 2. Given sampled data from Moons (a simple Scikit-learn toy dataset to visualise classification (Pedregosa et al. 2011; Chen et al. 2023)), our goal is to perturb each point within an  $L^\infty$  norm, e.g.,  $\|\Delta x_i\|_\infty < 0.15$  in this example. Next, we apply the PGA procedure described earlier to obtain perturbations  $\delta_i^*$  for each sample, constrained by  $\|\delta_i\|_\infty \leq \epsilon$ , with  $\epsilon = 0.15$ . Then, according to Algorithm 2, we obtain the perturbed points. Figure 2 illustrates the perturbation, as well as the perturbed points. We can then use Algorithm 1 to find that the new Bayes error value is 0.1832, 28% higher than the original one. Apart from the Bayes error growth, we can observe that while generally the movement of each point is towards the other class (increasing classification uncertainty), the travelled distance does not necessarily equal the perturbation constraint, i.e., 0.15 here. This suggests that in some cases, it is not that a larger shift is more optimal.

In the following, we establish that applying Algorithm 2 can lead to a guaranteed increase in the Bayes error.

**Lemma 3.3.** *Let  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  be a differentiable function with a  $\kappa$ -Lipschitz continuous gradient. Let  $\mathbb{C} \subseteq \mathbb{R}^d$  be a*

closed, convex set. Let the PGA update be defined as

$$x^{(k+1)} = \Pi_{\mathbb{C}} \left( x^{(k)} + \eta \nabla f(x^{(k)}) \right) \quad (14)$$

with  $x^{(0)} \in \mathbb{C}$ , and step size  $\eta \in (0, \frac{2}{\kappa})$ , where  $\Pi_{\mathbb{C}}$  finds a point on  $\mathbb{C}$  closest to a vector. Then for all  $k \geq 0$ ,

$$f(x^{(k)}) \geq f(x^{(0)}). \quad (15)$$

Proof for Theorem 3.3 is in the extended version (Zhang et al. 2025). Next, we establish that our objective function satisfies the  $\kappa$ -Lipschitz condition.

**Theorem 3.4.** *There exists a constant  $\kappa$  such that for all  $\eta \in (0, \frac{2}{\kappa})$ , we have that for any  $\{(x'_i, y'_i)\}_{i=1}^n$  from Algorithm 2, it is guaranteed that  $\hat{\beta}_{D'} \geq \hat{\beta}_D$ .*

*Proof.* We consider the objective function in Eq. (13) as a function  $f_1 : \mathbb{X}^n \times \mathbb{Y}^n \rightarrow \mathbb{R}$ , and the perturbation range as a set  $\mathbb{C} \subseteq \mathbb{X}^n$ . First, we show that this objective function is a differentiable function with a  $\kappa$ -Lipschitz continuous gradient. Specifically, the partial derivative (of  $f_1$ ) with respect to each variable can be expressed as Eq. (16).

$$\begin{aligned} \nabla_{x'_i} & \left( \frac{1}{n} \sum_{i=1}^n \left( 1 - \max_c \frac{\sum_{j=1, j \neq i}^n \mathbb{[}y_i = c\mathbb{]} \cdot s(x'_j, x'_i)}{\sum_{k=1, k \neq i}^n s(x'_k, x'_i)} \right) \right) \\ & = -\frac{1}{n} \sum_{j=1, j \neq i}^n \left( \frac{1}{(\sum_k s(x'_i, x'_k))^2} \cdot \max_c \sum_k (\mathbb{[}y_j = c\mathbb{]} - \right. \\ & \quad \left. \mathbb{[}y_k = c\mathbb{]}) s(x_i, x_k) + \max_c \sum_k (\mathbb{[}y_i = c\mathbb{]} - \right. \\ & \quad \left. \mathbb{[}y_k = c\mathbb{]}) \frac{1}{(\sum_k s(x'_j, x'_k))^2} \cdot s(x_j, x_k) \right) \nabla_{x'_i} s(x'_i, x'_j) \end{aligned} \quad (16)$$

Observe that the function  $s$  satisfies  $\exists \kappa_1 > 0. \forall x, x'. 0 \leq s(x, x') \leq \kappa_1$ . Thus, as long as it is not the case that  $\forall i, j, s(x_i, x_j) = 0$ , then we get that

$$(\exists \kappa_2. \|\nabla_{x'_i} s(x'_i, x'_j)\| \leq \kappa_2) \rightarrow (\exists \kappa_3. \|\nabla_{x'_i} f_1\| \leq \kappa_3). \quad (17)$$

That is, if the similarity function is  $\kappa_2$ -Lipschitz, then the objective function of Eq. (16) is  $\kappa_3$ -Lipschitz. Next, we show that the perturbation range  $\mathbb{C}$  is a closed, convex set (Ver-shynin 2010). Namely, for all orders greater than 0,  $L^p$  norm is closed, and for all orders greater than or equal to 1,  $L^p$  norm is convex. Therefore, applying Theorem 3.3, we show that for all  $k \geq 0$ :

$$f_1(x_1^{(k)}, \dots, x_n^{(k)}, y_1, \dots) \geq f_1(x_1^{(0)}, \dots, x_n^{(0)}, y_1, \dots). \quad (18)$$

□

**Dealing with embedded inputs** The previous discussion is based on a discrete representation of samples whose distance is easy to measure. As for complicated inputs, e.g., coloured images, we could first compute their embedding vector and then solve the Bayes error maximisation problem. Specifically, let  $m : \mathbb{X}_{\text{ori}} \rightarrow \mathbb{X}_{\text{emb}}$  denote the embedding function from original input space  $\mathbb{X}_{\text{ori}}$  to the embedding space  $\mathbb{X}_{\text{emb}}$ . In this case, their similarity is measured in the embedding

space, and the constraint is still applied to the original input space, i.e., Eq. (13) becomes Eq. (19) and  $m$  needs to be known.

$$\begin{aligned} \min_{\{x_i\}_{i=1}^n} & \sum_{i=1}^n \left( \max_c \frac{\sum_{j=1, j \neq i}^n \mathbb{[}y_i = c\mathbb{]} \cdot s(m(x'_j), m(x'_i))}{\sum_{k=1, k \neq i}^n s(m(x'_k), m(x'_i))} \right) \\ \text{s.t.} & \|x'_i - x_i\|_p \leq \epsilon, \quad \forall i = 1, 2, \dots, n \end{aligned} \quad (19)$$

### 3.3 Mixing Clean and Unlearnable Examples

In practice, models are unlikely to be trained solely on the perturbed unlearnable data. To model this situation, we adjust Eq. (13) by letting some perturbation be 0. That is, for some (but not all)  $i$  in  $1, 2, \dots, n$ , it is restricted that  $x'_i = x_i$ . Intuitively, this could be a case where the adversary could leverage publicly available data to augment the training, whereas the defender cannot modify all the public data. In the following, we show in Theorem 3.5 that Algorithm 2 can still effectively increase Bayes error.

**Corollary 3.5.** *Given  $\{(x_i, y_i)\}_{i=1}^n$ , let  $\mathbb{I}_n \subset \mathbb{N} \cap [1, n]$ ,  $\mathbb{I}_n \neq$  be a subset of all indices. Suppose for all indices  $j$  in  $\mathbb{I}_n$ , it is restricted that  $\Delta x_j = 0$ . It follows that the remaining perturbations can increase Bayes error; i.e., the following condition is satisfied.*

$$\begin{aligned} & (\forall j \in \mathbb{I}_n. \Delta x_j = 0) \wedge \\ & \left( 1 - \frac{1}{n} \sum_i \max_c \hat{p}_D(y = c \mid \mathbf{x} = x'_i) \geq \right. \\ & \quad \left. 1 - \frac{1}{n} \sum_i \max_c \hat{p}_D(y = c \mid \mathbf{x} = x_i) \right) \end{aligned} \quad (20)$$

*Proof for Theorem 3.5 can be found in the extended version (Zhang et al. 2025).*

## 4 Experiments

We conduct experiments to evaluate the empirical effectiveness of the proposed method, aiming to answer the following research questions (RQs).

1. How does the proposed perturbation affect Bayes error?
2. Does training on our unlearnable examples lead to reduced test accuracy? What happens when only a subset of the training data is made unlearnable?
3. How does our method perform against adaptive attacks such as adversarial training?

**Experimental settings** We follow standard experimental settings as established in previous work (Huang et al. 2021; Wen et al. 2023), and evaluate our method on three image classification benchmarks: CIFAR-10, CIFAR-100 (Krizhevsky, Hinton et al. 2009), and Tiny ImageNet (Le and Yang 2015). Perturbations are constrained within an  $L^\infty$ -norm bound of  $\epsilon = 8/255$  per image. For feature extraction, we adopt the commonly used ResNet-18 model (He et al. 2016; Robey et al. 2022; Zhang et al. 2019; Wang et al. 2020). All models are trained for 100 epochs using stochastic gradient descent (Robbins and Monro 1951) with a learning rate of 0.1,

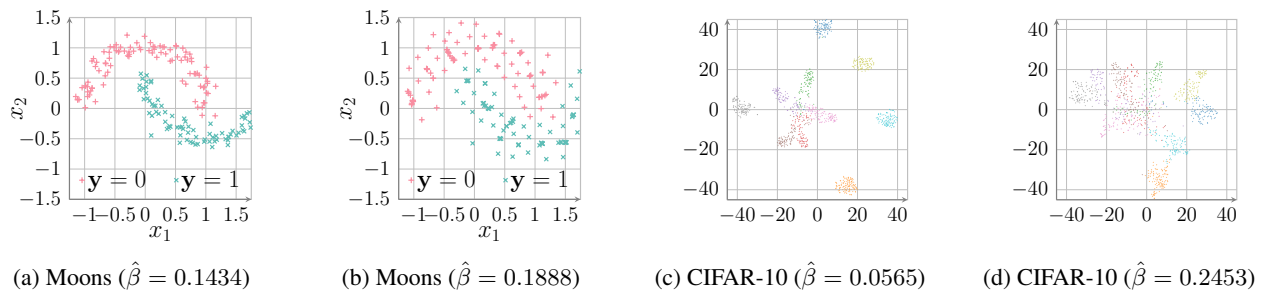


Figure 3: Examples before and after perturbation




Figure 4: Test accuracy on each setting. “No add-on” means only training on part of the original training set.

momentum of 0.9, weight decay of  $5 \times 10^{-4}$ , and cosine annealing learning rate scheduling (Loshchilov and Hutter 2017). For our optimisation implementation, we formulate an auto-differentiable Bayes error objective in PyTorch (Paszke et al. 2019), enabling gradient computation in the input space via back propagation. Additional experimental details are provided in the corresponding subsections.

#### 4.1 RQ1: The change in Bayes error

We empirically study how the Bayes error changes after perturbing data inputs using our method. As an intuitive illustration, we first conduct the experiment on Moons (a simple toy dataset to visualise classification (Pedregosa et al. 2011; Chen et al. 2023)), which consists of two interleaving half circles. A sample of 200 points is shown in Figure 3a. We then apply Algorithm 2 with  $\epsilon = 0.25$  to construct unlearnable examples, resulting in the distribution shown in Figure 3b. The perturbed examples appear more scattered within each class, and examples from different classes become closer. This reduces inter-class separation, leading to a 24.3% increase in Bayes error, from 0.1434 (original) to 0.1888 (unlearnable).

Next, we extend our experiment to high-dimensional image data. As shown in Figure 3c, the t-SNE visualization of the original CIFAR-10 training set exhibits well-separated class features. After applying Algorithm 2, the resulting unlearnable dataset (visualized in Figure 3d) shows significant feature overlap across classes, making the data intuitively less learnable. Correspondingly, the Bayes error increases significantly from 0.0565 to 0.2453, representing more than a 3.3x increase. These results confirm that our method effectively increases the Bayes error.

#### 4.2 RQ2: Effectiveness of Unlearnable Examples

To evaluate the effectiveness of our unlearnable examples, we use them for training and then measure the test accuracy. Test accuracy is an empirical measurement of unlearnability. Intuitively, if the test accuracy remains high, the training data may not be sufficiently unlearnable. As a reference, training on the full original CIFAR-10 (training) set yields 95.12% test accuracy.

We first train on our entire unlearnable set and observe a test accuracy of 28.12% on CIFAR-10, showcasing an over 70% accuracy drop from 95.12%. Then, we explore mixed-data settings, e.g., replacing half of the original training samples with unlearnable ones results in 69.68% test accuracy, with a 27% drop from 95.12%. To further isolate the effect of unlearnable examples, we train only on the original half set without any unlearnable data add-on, which achieves 91.16% accuracy, over 30% higher than 69.68%. This suggests including unlearnable examples degrades model performance.

We iteratively adjust the proportion of original/unlearnable data in the mixture to reflect realistic scenarios where protected data forms a minority. As shown in Figure 4, we report test accuracy under varying ratios of kept original samples (# original / # original and unlearnable) ratios. We include two representative baseline methods, one by minimising training loss and the other by maximising training loss. Specifically, Huang et al. (2021) introduce label-correlated noise to create shortcut features that hinder meaningful learning. In contrast, Wen et al. (2023) identify the embedding centre of each class and push individual examples away from it, reducing intra-class coherence.

Compared with baseline methods, our method consistently induces greater accuracy drops, on average, 8% more than Huang et al. (2021) and 9% more than Wen et al. (2023). More importantly, our examples cause a consistent accuracy decline across all settings. In contrast, previous methods may occasionally act as extra training data and even improve accuracy when mixed in (e.g., 50% (Huang et al. 2021)). These results demonstrate that our unlearnable examples are not only effective in full training scenarios but also degrade more when only used as a part.

**Effectiveness on complex datasets** Figure 5 illustrates experimental results on more complex datasets, CIFAR-100 and Tiny ImageNet. Across all mixture ratios, including the case when fully training on unlearnable examples, our constructed




Figure 5: Test accuracy for complex datasets

Model architecture	clean	ours
ResNet-18	95.04	28.12
ResNet-34	95.12	29.86
VGG-19	91.70	18.14
DenseNet-121	95.06	25.89
MobileNet v2	93.30	30.54

Table 1: Test accuracy when unlearnable examples are used in training with various architectures

examples consistently cause a noticeable drop in test accuracy. These results further confirm the general effectiveness of our method on larger and more challenging datasets.

**Transferability to unseen model architectures** To transform image data into an embedding space, we use a feature extractor with ResNet-18 (He et al. 2016) architecture. A natural question is whether the constructed unlearnable examples remain effective when a different model architecture is used for training. To investigate this, we follow Wen et al. (2023) and evaluate four alternative architectures, with results summarised in Table 1. Across all architectures, training on our CIFAR-10 unlearnable examples consistently yields lower accuracy compared to training on clean data. The average accuracy drop is 68% ( $\pm 10\%$ ), closely matching the drop observed with ResNet-18 (28.12%,  $\pm 5\%$ ), suggesting that the effectiveness of our examples generalises well across model architectures.

### 4.3 RQ3: Resistance to Adversarial Training

Due to their potential public release, once constructed, unlearnable examples are no longer amendable by the defender. This may place them at a disadvantage, especially when it is known that the data has been intentionally perturbed. In such cases, countermeasures such as adversarial training may be applied to forcibly extract information from these examples. To evaluate the effectiveness of our approach in this scenario, we conduct experiments using PGD-based adversarial training (Madry et al. 2018; Wen et al. 2023), a commonly adopted adaptive strategy, on our constructed unlearnable examples. Figure 6 presents model performance over training iterations. Compared to standard training on clean data, our unlearnable examples consistently lead to reduced accuracy. Notably, under adversarial training, the model achieves only 74.5% test accuracy, which is 21.5% lower than training on original examples. We also include stronger adaptive attack (Yu et al. 2024) and achieves consistently (8.1% on average) higher




Figure 6: Test accuracy at each iteration

performance than baselines. This substantial drop renders the resulting model largely unusable in practice.

## 5 Related Work

Protecting data at the source has gained increasing attention due to the rising concerns over unauthorised data usage. Huang et al. (2021) and Ren et al. (2022) introduced the concept of unlearnable examples, where imperceptible perturbations are applied to make data inherently difficult to learn. Approaches to constructing unlearnable examples have typically followed two heuristics: (1) maximising the training loss to prevent models from fitting useful features (Wen et al. 2023; Fowl et al. 2021), and (2) minimising it to encourage overfitting on non-generalizable patterns (Huang et al. 2021; Fu et al. 2022). There are also other approaches (Liu, Wang, and Gao 2024; Liu et al. 2024; Liu, Peng, and Tang 2023; Mu and Lim 2025), including directly optimising perturbations for specific model parameters (Lu, Kamath, and Yu 2023). Such created unlearnable examples have been found to work ineffectively when mixed with clean data and lack formal guarantees. This work thus aims to tackle the problem of ensuring unlearnability even in the mixed data situation.

Estimating the Bayes error of a given data distribution has long been an interesting topic (Fukunaga and Hostetler 1975). Related methods include estimating its lower or upper bounds, e.g. the Bhattacharyya distance (Fukunaga 1990) or the Henze-Penrose divergence (Berisha et al. 2016; Sekeh, Oselio, and Hero 2020). Alternatively, there are attempts to estimate the Bayes error using generative models (Kingma and Dhariwal 2018; Theisen et al. 2021) or real-world human annotation (Renggli et al. 2021). Indirect Bayes error estimation has also been adopted by quantity substitution, e.g., accuracy to robustness (Zhang and Sun 2024a,b). Compared with these estimates, our Bayes error estimate is easier to compute, such as to model and improve unlearnability.

## 6 Conclusion

We propose a formal objective to increase unlearnability by increasing Bayes error, and a projected gradient ascent strategy to ensure its growth in practice. Experiments confirm our method consistently reduces test accuracy even if the training data is a mixture of clean and unlearnable ones. Overall, it offers an effective approach to user data protection. In the future, an interesting direction is to extend our framework to protect data from generative models, such as generative adversarial networks.

## Acknowledgements

This research is supported by the Ministry of Education, Singapore under its Academic Research Fund Tier 3 (Award ID: MOET32020-0004).

## References

- Berisha, V.; Wisler, A.; Hero, A. O.; and Spanias, A. 2016. Empirically Estimable Classification Bounds Based on a Non-parametric Divergence Measure. *IEEE Transactions on Signal Processing*, 64(3): 580–591.
- Bertsekas, D. P. 2003. On the Goldstein-Levitin-Polyak gradient projection method. *IEEE Transactions on automatic control*, 21(2): 174–184.
- Chang, Y.-W.; Hsieh, C.-J.; Chang, K.-W.; Lin, C.-J.; et al. 2010. Training and testing low-degree polynomial data mappings via linear svm. *Journal of Machine Learning Research*, 11(4).
- Chen, Q.; Cao, F.; Xing, Y.; and Liang, J. 2023. Evaluating Classification Model Against Bayes Error Rate. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8): 9639–9653.
- Fowl, L.; Goldblum, M.; Chiang, P.-y.; Geiping, J.; Czaja, W.; and Goldstein, T. 2021. Adversarial examples make strong poisons. *Advances in Neural Information Processing Systems*, 34: 30339–30351.
- Fu, S.; He, F.; Liu, Y.; Shen, L.; and Tao, D. 2022. Robust unlearnable examples: Protecting data against adversarial learning. *arXiv preprint arXiv:2203.14533*.
- Fukunaga, K. 1990. *Introduction to Statistical Pattern Recognition (2nd Ed.)*. USA: Academic Press Professional, Inc. ISBN 0122698517.
- Fukunaga, K.; and Hostetler, L. 1975. k-nearest-neighbor Bayes-risk estimation. *IEEE Transactions on Information Theory*, 21(3): 285–293.
- Garber, F.; and Djouadi, A. 1988. Bounds on the Bayes classification error based on pairwise risk functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(2): 281–288.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Huang, H.; Ma, X.; Erfani, S. M.; Bailey, J.; and Wang, Y. 2021. Unlearnable Examples: Making Personal Data Unexploitable. In *International Conference on Learning Representations*.
- Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T. B.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; and Amodei, D. 2020. Scaling Laws for Neural Language Models. *CoRR*, abs/2001.08361.
- Kingma, D. P.; and Dhariwal, P. 2018. Glow: Generative Flow with Invertible 1x1 Convolutions. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Le, Y.; and Yang, X. 2015. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7): 3.
- Liu, S.; Peng, F.; and Tang, K. 2023. Reliable robustness evaluation via automatically constructed attack ensembles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 8852–8860.
- Liu, S.; Wang, Y.; and Gao, X.-S. 2024. Game-theoretic unlearnable example generator. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 21349–21358.
- Liu, Y.; Xu, K.; Chen, X.; and Sun, L. 2024. Stable unlearnable example: Enhancing the robustness of unlearnable examples via stable error-minimizing noise. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 3783–3791.
- Loshchilov, I.; and Hutter, F. 2017. SGDR: Stochastic Gradient Descent with Warm Restarts. In *International Conference on Learning Representations*.
- Lu, Y.; Kamath, G.; and Yu, Y. 2023. Exploring the limits of model-targeted indiscriminate data poisoning attacks. In *International Conference on Machine Learning*, 22856–22879. PMLR.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Mu, W.; and Lim, K. H. 2025. Bayesian Privacy Guarantee for User History in Sequential Recommendation Using Randomised Response. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management, CIKM '25*, 5041–5046. New York, NY, USA: Association for Computing Machinery. ISBN 9798400720406.
- Nadaraya, E. A. 1964. On estimating regression. *Theory of Probability & Its Applications*, 9(1): 141–142.
- Nakashima, K.; Kataoka, H.; Matsumoto, A.; Iwata, K.; Inoue, N.; and Satoh, Y. 2022. Can vision transformers learn without natural images? In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 1990–1998.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32*, volume 32, 8024–8035. Curran Associates, Inc.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikitlearn: Machine Learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830.

- Ren, J.; Xu, H.; Wan, Y.; Ma, X.; Sun, L.; and Tang, J. 2022. Transferable unlearnable examples. *arXiv preprint arXiv:2210.10114*.
- Renggli, C.; Rimanic, L.; Hollenstein, N.; and Zhang, C. 2021. Evaluating Bayes Error Estimators on Real-World Datasets with FeeBee. In Vanschoren, J.; and Yeung, S., eds., *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Robbins, H.; and Monro, S. 1951. A stochastic approximation method. *The annals of mathematical statistics*, 400–407.
- Robey, A.; Chamon, L.; Pappas, G. J.; and Hassani, H. 2022. Probabilistically Robust Learning: Balancing Average and Worst-case Performance. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 18667–18686. PMLR.
- Seddik, M. E. A.; and Tamaazousti, M. 2022. Neural networks classify through the class-wise means of their representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 8204–8211.
- Sekeh, S. Y.; Oselio, B.; and Hero, A. O. 2020. Learning to Bound the Multi-Class Bayes Error. *IEEE Transactions on Signal Processing*, 68: 3793–3807.
- Theisen, R.; Wang, H.; Varshney, L. R.; Xiong, C.; and Socher, R. 2021. Evaluating State-of-the-Art Classification Models Against Bayes Optimality. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 9367–9377. Curran Associates, Inc.
- Vershynin, R. 2010. Lectures in functional analysis. *Department of Mathematics, University of Michigan*.
- Wang, Y.; Zou, D.; Yi, J.; Bailey, J.; Ma, X.; and Gu, Q. 2020. Improving Adversarial Robustness Requires Revisiting Misclassified Examples. In *International Conference on Learning Representations*.
- Watson, G. S. 1964. Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A*, 359–372.
- Wen, R.; Zhao, Z.; Liu, Z.; Backes, M.; Wang, T.; and Zhang, Y. 2023. Is Adversarial Training Really a Silver Bullet for Mitigating Data Poisoning? In *The Eleventh International Conference on Learning Representations*.
- Wu, M.; Hughes, M.; Parbhoo, S.; Zazzi, M.; Roth, V.; and Doshi-Velez, F. 2018. Beyond Sparsity: Tree Regularization of Deep Models for Interpretability. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Yeom, S.; Giacomelli, I.; Fredrikson, M.; and Jha, S. 2018. Privacy risk in machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st computer security foundations symposium (CSF)*, 268–282. IEEE.
- Yu, Y.; Wang, Y.; Xia, S.; Yang, W.; Lu, S.; Tan, Y.-P.; and Kot, A. C. 2024. Purify unlearnable examples via rate-constrained variational autoencoders. In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org.
- Zhang, H.; Yu, Y.; Jiao, J.; Xing, E.; Ghaoui, L. E.; and Jordan, M. 2019. Theoretically Principled Trade-off between Robustness and Accuracy. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 7472–7482. PMLR.
- Zhang, R.; and Sun, J. 2024a. Certified Robust Accuracy of Neural Networks Are Bounded Due to Bayes Errors. In Gurfinkel, A.; and Ganesh, V., eds., *Computer Aided Verification*, 352–376. Cham: Springer Nature Switzerland. ISBN 978-3-031-65630-9.
- Zhang, R.; and Sun, J. 2024b. How Does Bayes Error Limit Probabilistic Robust Accuracy. *arXiv preprint arXiv:2405.14923*.
- Zhang, R.; Sun, J.; Lim, E.-P.; and Zhang, P. 2025. Towards Provably Unlearnable Examples via Bayes Error Optimisation. *arXiv:2511.08191*.
- Zhang, T. 2004. Statistical behavior and consistency of classification methods based on convex risk minimization. *The Annals of Statistics*, 32(1): 56 – 85.