

Farzad Zafarani* and Chris Clifton

Differentially Private Naïve Bayes Classifier Using Smooth Sensitivity

Abstract: There is increasing awareness of the need to protect individual privacy in the training data used to develop machine learning models. Differential Privacy is a strong concept of protecting individuals. Naïve Bayes is a popular machine learning algorithm, used as a baseline for many tasks. In this work, we have provided a differentially private Naïve Bayes classifier that adds noise proportional to the *smooth sensitivity* of its parameters. We compare our results to Vaidya, Shafiq, Basu, and Hong [1] which scales noise to the global sensitivity of the parameters. Our experimental results on real-world datasets show that smooth sensitivity significantly improves accuracy while still guaranteeing ϵ -differential privacy.

Keywords: differential privacy, Naïve Bayes classifier, privacy, data mining, smooth sensitivity

DOI 10.2478/popets-2021-0077

Received 2021-02-28; revised 2021-06-15; accepted 2021-06-16.

1 Introduction

With the growth of user data across the internet, it has become more important to protect users' sensitive information. One solution to this problem is privacy-preserving data analysis, providing ability to share information while protecting users' data. Dwork, McSherry, Nissim, and Smith [2] introduced differential privacy, providing a strong privacy guarantee for statistical data release. At a high level, Differential Privacy guarantees that the outcome of a differentially private algorithm would be similar no matter if a particular individual contributes personal data to the database or not. There are several common approaches to differential privacy, including the Laplace Mechanism, which perturbs the parameters of the model with noise that

is drawn from the Laplace distribution, scaled to the impact of a single individual on the result. The exponential Mechanism is another important mechanism to guarantee (ϵ, δ) -differential privacy [3].

A model generated by a machine learning algorithm, when trained on a dataset, can reveal information about the training dataset. There are a series of recent works that guarantee that the output of a machine learning model satisfies differential privacy. These include differentially private Decision Trees [4], SVM [5], Deep Neural Networks [6], and Logistic Regression [7]. Differential privacy is particularly relevant for ensuring that machine learning models do not disclose individual information and even has the promise of improving generalization [8]. Naïve Bayes is a baseline for many classification tasks. Vaidya, Shafiq, Basu, and Hong [1] provided a differentially private algorithm for the Naïve Bayes classifier. They use the Laplace Mechanism to provide this guarantee based on computing the global sensitivity of the parameters. One of the main drawbacks of using global sensitivity is that the amount of the noise added to the output can be high if their could be a dataset where an individual would have a large impact on an outcome. Nissim, Raskhodnikova, and Smith [9] provided a general solution for this problem. In their paper they compute the *Smooth Sensitivity* for a given function f building on the definition of *local sensitivity*. The local sensitivity of f is the maximum amount of change in f if we change a single element in a *particular* dataset x . It is obvious that the local sensitivity of a given function is not greater than its global sensitivity. Ideally, we would like to add noise proportional to the local sensitivity of f , but this does not satisfy the definition of differential privacy (the amount of noise needed reveals too much about the data), hence, in [9] they compute a β -smooth function which is the smallest upper bound for the local sensitivity that provides ϵ -differential privacy. In this paper, we show how this approach can be used to provide a ϵ -differentially private algorithm for the Naïve Bayes classifier based on the smooth sensitivity of the parameters of the model.

Bun and Steinke [10] also provide an algorithm for estimating the mean of a distribution using i.i.d. sample x using smooth sensitivity. They first assume a crude

*Corresponding Author: Farzad Zafarani: Department of Computer Science, Purdue University, USA., E-mail: farzad@purdue.edu

Chris Clifton: Department of Computer Science, Purdue University, USA., E-mail: clifton@cs.purdue.edu

bound on the $\mu \in [a, b]$, then they truncate the samples, i.e. they remove the largest m samples and smallest m samples from x , and compute the mean of the $n - 2m$ samples. Finally they project the estimated mean to the range $[a, b]$. We compare our algorithm to Vaidya, Shafiq, Basu, and Hong [1] and Naïve Bayes using Bun and Steinke [10]’s mean estimation algorithm.

2 Preliminaries

We first give a brief overview of the Naïve Bayes classifier, then provide an overview of differential privacy. More specifically, we state the definitions for ϵ -differential privacy and smooth sensitivity.

2.1 Naïve Bayes Classifier

The Naïve Bayes classifier is a family of probabilistic classifiers that uses Bayes’ theorem and assumes independence between features. That is, it assumes that the value of a particular feature is unrelated to any other features. The Naïve Bayes classifier can handle an arbitrary number of independent variables, whether continuous or categorical, and classifies an instance to one of a finite number of classes.

To train the Naïve Bayes model, a set of training examples with a corresponding target label is provided. The task is to assign a new class c_{MAP} to an unseen instance $X = \langle X_1, X_2, \dots, X_m \rangle$. Thus, the learning task would be that for each instance $X = \langle X_1, X_2, \dots, X_m \rangle$ consists of m features, it assigns probability

$$\Pr(c_j | X_1, X_2, \dots, X_m)$$

for each of K possible label classes $C = \{c_1, c_2, \dots, c_K\}$.

$$c_{MAP} = \arg \max_{j \in \{1, \dots, K\}} \Pr(c_j | X_1, X_2, \dots, X_m)$$

By using Bayes’ theorem, we can further decompose the conditional probability to:

$$\begin{aligned} c_{MAP} &= \arg \max_{j \in \{1, \dots, K\}} \left(\frac{\Pr(X_1, X_2, \dots, X_m | c_j) \Pr(c_j)}{\Pr(X_1, X_2, \dots, X_m)} \right) \\ &= \arg \max_{j \in \{1, \dots, K\}} (\Pr(X_1, X_2, \dots, X_m | c_j) \Pr(c_j)) \end{aligned}$$

The Naïve Bayes classifier makes the further simplifying assumption that the attribute values are conditionally independent, given the target value. Therefore:

$$c_{NB} = \arg \max_{j \in \{1, \dots, K\}} \Pr(c_j) \prod_{i=1}^m \Pr(X_i | c_j)$$

where c_{NB} denotes the final class label for the instance $X = \langle X_1, X_2, \dots, X_m \rangle$.

From the training dataset, we can pre-compute the conditional probabilities $\Pr(X_i | c_j)$. Also, $\Pr(c_j)$ can be computed by counting the number of items that are labeled c_j in the training dataset. As with Vaidya, Shafiq, Basu, and Hong’s work [1], we deal with both categorical and numerical attributes. The way that we estimate the probability is different for each class:

- *Categorical Value*: For a categorical attribute X_i with J possible attribute values a_1, a_2, \dots, a_J , the probability $\Pr(X_i = a_k | c_j) = \frac{\tau\{X_i=a_k \wedge C=c_j\}}{\tau\{C=c_j\}}$, where the $\tau\{x\}$ operator returns the number of elements in the training set D that satisfy property x . To prevent division by zero, we use *Laplace smoothing* which adds 1 to all counts.
- *Numerical Value*: For a numerical attribute X_i , one standard approach is to assume that for each possible discrete value c_k of C , the distribution of each continuous X_i is Gaussian, and is defined by a mean and standard deviation specific to X_i and c_k [11]. To train such a Naïve Bayes classifier we must therefore estimate the mean and standard deviation of these Gaussians,

$$\mu_{ik} = E[X_i | C = c_k]$$

$$\sigma_{ik}^2 = E[(X_i - \mu_{ik})^2 | C = c_k]$$

, for each numerical attribute X_i and each possible value c_k of C .

If the numerical values are bounded, one can use the Truncated normal distribution and estimate its parameters. The probability density function of the Truncated normal distribution for $a \leq x \leq b$ is:

$$f(x; \mu, \sigma, a, b) = \frac{1}{\sigma} \frac{\phi(\frac{x-\mu}{\sigma})}{\phi(\frac{b-\mu}{\sigma}) - \phi(\frac{a-\mu}{\sigma})}$$

Where

$$\phi(\tau) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}\tau^2)$$

After estimating the values for mean and variance, the probability that an instance is of class c_j can be directly computed from the density function.

2.2 Differential Privacy

Dwork, McSherry, Nissim, and Smith [2] defined the notion of differential privacy. At a high level, differential privacy guarantees that if your data is a part of a database from which we release information, then the released information will be similar if your data is a

part of the database or not. That is, your data will have a negligible impact on the released information. Hence, no meaningful information can be inferred about individuals. The definitions below come from their work.

Definition 2.1. (Laplace Distribution) The probability density function (p.d.f) of the Laplace distribution $Lap(\mu, \lambda)$ is $f(x|\mu, \lambda) = \frac{1}{2\lambda} e^{-|x-\mu|/\lambda}$ with mean μ and standard deviation $\sqrt{2}\lambda$.

Definition 2.2. (Cauchy Distribution) The probability density function (p.d.f) of the Cauchy distribution $Cauchy(x_0, \lambda)$ is $f(x|x_0, \lambda) = \frac{1}{\lambda\pi(1+((x-x_0)/\lambda)^2)}$ with location parameter x_0 and scale λ .

Let $\mathcal{D}_1, \dots, \mathcal{D}_m$ denote domains, each of which could be categorical or numerical. A database D consists of n rows, $\{X^{(1)}, X^{(2)}, \dots, X^{(n)}\}$, where each $X^{(i)} \in \mathcal{D}_1 \times \dots \times \mathcal{D}_m$.

We say two databases D_1 and D_2 are at distance k of each other and we write it as $d(D_1, D_2) = k$ if they differ by k rows. Two database D_1 and D_2 are called *neighbors* if $d(D_1, D_2) = 1$.

Definition 2.3. (Global Sensitivity). For $f : \mathcal{D} \rightarrow \mathcal{R}$, the global sensitivity of f with respect to ℓ_1 metric is:

$$GS_f = \max_{x, y: d(x, y)=1} \|f(x) - f(y)\|_1$$

Definition 2.4. (Differential Privacy). A randomized Mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ with domain \mathcal{D} and range \mathcal{R} is ϵ -differentially private if for all $D_1, D_2 \in \mathcal{D}$ satisfying $d(D_1, D_2) = 1$, and for all sets $\mathcal{S} \subseteq \mathcal{R}$ of possible outputs:

$$\Pr[\mathcal{M}(D_1) \in \mathcal{S}] \leq e^\epsilon \Pr[\mathcal{M}(D_2) \in \mathcal{S}]$$

We also introduce *concentrated differential privacy*; while we do not require this, our comparison with [10] involves situations where [10] satisfies concentrated differential privacy rather than Definition 2.4.

Definition 2.5. (Concentrated Differential Privacy). A randomized Mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ with domain \mathcal{D} and range \mathcal{R} is $\frac{1}{2}\epsilon^2$ -CDP if for all $D_1, D_2 \in \mathcal{D}$ satisfying $d(D_1, D_2) = 1$:

$$\forall \alpha > 1 \quad D_\alpha(\mathcal{M} || \mathcal{M}(x')) \leq \frac{1}{2}\epsilon^2 \alpha$$

where D_α is the Rényi divergence and α is a divergence parameter.

We also make use of a couple of properties of the way differentially private mechanisms combine. *Sequential com-*

position states that privacy loss is additive: If we take “multiple looks” at the data, the privacy budget ϵ expended is the sum of the privacy budgets ϵ_i of each “look”.

Theorem 1. (Sequential Composition [12][13]) Let $\mathcal{M}_1 : \mathcal{D} \rightarrow \mathcal{R}_1$ be an ϵ_1 -differentially private algorithm, and let $\mathcal{M}_2 : \mathcal{D} \rightarrow \mathcal{R}_2$ be an ϵ_2 -differentially private algorithm. Then their combination, defined to be $\mathcal{M}_{1,2} : \mathcal{D} \rightarrow \mathcal{R}_1 \times \mathcal{R}_2$ by the mapping: $\mathcal{M}_{1,2}(x) = (\mathcal{M}_1(x), \mathcal{M}_2(x))$ is $\epsilon_1 + \epsilon_2$ -differentially private.

Dwork et al. [2] showed how to calibrate the noise to the global sensitivity of the function f such that it satisfies ϵ -differential privacy. In their work, they have shown that the magnitude of the noise is proportional to $Lap(0, GS_f/\epsilon)$. Intuitively, whenever we add noise proportional to the global sensitivity of f , we are adding noise proportional to the maximum magnitude of changes in f .

3 Smooth Sensitivity

One drawback of computing the global sensitivity of f is that for many functions, there may be some possible datasets where changing one individual can make a dramatic change in the outcome. For example, suppose we are computing median on a value ranging from 0..1. The dataset consisting of individuals with values 0, 0, and 1 has median 0, but by changing one individual the median goes to 1 - so the added noise must essentially obscure the entire result.

In practice, most databases will not have this property. Nissim, Raskhodnikova, and Smith [9] showed that we can add noise based on the actual dataset we have rather than a worst-case dataset, and still satisfy ϵ -differential privacy. We now outline their result.

Definition 3.1. (Local Sensitivity). For $f : \mathcal{D} \rightarrow \mathcal{R}$ and $x \in \mathcal{D}$ the local sensitivity of f at x (with respect to the ℓ_1 metric) is:

$$LS_f(x) = \max_{y: d(x, y)=1} \|f(x) - f(y)\|_1$$

Note that $GS_f = \max_x LS_f(x)$. We would like to be able to add noise proportional to local sensitivity. However, the local sensitivity may itself be high sensitivity, i.e., noise magnitude may compromise privacy. Nissim, Raskhodnikova, and Smith define a *Smooth bound* that

addresses this issue by looking not just at neighbors of the current dataset, but also their neighbors, etc.

Definition 3.2. (A Smooth bound). For $\beta > 0$, a function $S : \mathcal{D} \rightarrow \mathbb{R}^+$ is a β -smooth upper bound on the local sensitivity of f if it satisfies the following requirements:

- $\forall x \in \mathcal{D} : \quad S(x) \geq LS_f(x)$
- $\forall x, y \in \mathcal{D}, d(x, y) = 1 : \quad S(x) \leq e^\beta S(y)$

Definition 3.3. (Smooth Sensitivity). For $\beta > 0$, the β -smooth sensitivity of f is:

$$S_{f,\beta}^*(x) = \max_{y \in \mathcal{D}} (LS_f(y) \cdot e^{-\beta d(x,y)})$$

The smooth sensitivity $S_{f,\beta}^*(x)$ is the smallest function that satisfies Definition 3.2.

Nissim et al. [9] showed that one could do much better than scaling the noise to the global sensitivity of f by adding noise proportional to the *Smooth sensitivity* of f where it will give much higher output accuracy.

3.1 Computing Smooth Sensitivity

We now describe how to compute the smooth sensitivity of a function.

Definition 3.4. The sensitivity of f at distance k is:

$$A^k(x) = \max_{y \in \mathcal{D}: d(x,y) \leq k} LS_f(y)$$

We can express the smooth sensitivity of f in terms of A^k as follows:

$$\begin{aligned} S_{f,\varepsilon}^*(x) &= \max_{k=0,1,\dots,n} e^{-k\varepsilon} \left(\max_{y: d(x,y)=k} LS_f(y) \right) \\ &= \max_{k=0,1,\dots,n} e^{-k\varepsilon} A^{(k)}(x) \end{aligned}$$

3.2 Calibrating Noise to the Smooth Sensitivity

To release a function f of the database D , the curator computes f and publishes $\mathcal{M}(D) = f(D) + \lambda Z$ where Z is a random variable drawn from a noise distribution, and λ is the scaling parameter.

Theorem 2. (Nissim et al. [9]) Let $f : \mathcal{D} \rightarrow \mathbb{R}$ be any real-valued function and let $S : \mathcal{D} \rightarrow \mathbb{R}$ be a β -smooth upper bound on the local sensitivity of f . Then we have:

if $\beta \leq \frac{\varepsilon}{2(\gamma+1)}$ and $\gamma > 1$, the algorithm $x \rightarrow f(x) + \frac{2(\gamma+1)S(x)}{1+|z|^\gamma} \eta$ where η is sampled with distribution $h(z) \approx \frac{1}{1+|z|^\gamma}$, is ε -differentially private.

Nissim et al. [9] showed that to scale the noise to the smooth sensitivity of f , it is sufficient to sample from an *admissible noise distribution*, defined as follows.

For a subset \mathcal{S} of \mathbb{R}^d , we write $\mathcal{S} + \xi$ for the set $\{z + \xi | z \in \mathcal{S}\}$, and $\varepsilon^\lambda \cdot \mathcal{S}$ for the set $\{\varepsilon^\lambda \cdot z | z \in \mathcal{S}\}$. We also write $a \pm b$ for the interval $[a - b, a + b]$.

Definition 3.5. (Nissim et al. [9]) A probability distribution on \mathbb{R}^n , given by a density function h , is (α, β) -admissible (with respect to ℓ_1), if for $\alpha = \alpha(\varepsilon, \delta)$, $\beta = \beta(\varepsilon, \delta)$, the following two conditions hold for all $\delta \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}$ satisfying $\|\xi\|_1 \leq \alpha$ and $|\lambda| \leq \beta$, and for all measurable subsets $\mathcal{S} \subset \mathbb{R}^n$:

- Sliding Property:

$$\Pr_{Z \sim h} [Z \in \mathcal{S}] \leq e^{\frac{\varepsilon}{2}} \cdot \Pr_{Z \sim h} [Z \in \mathcal{S} + \xi] + \delta/2$$

- Dilation Property:

$$\Pr_{Z \sim h} [Z \in \mathcal{S}] \leq e^{\varepsilon/2} \cdot \Pr_{Z \sim h} [Z \in e^\lambda \cdot \mathcal{S}] + \delta/2$$

Theorem 3. (Nissim et al. [9]) For any $\gamma > 1$, the distribution with density $h(z) \approx \frac{1}{1+|z|^\gamma}$ is $(\frac{\varepsilon}{2(\gamma+1)}, \frac{\varepsilon}{2(\gamma+1)})$ -admissible. Moreover, the n -dimensional product of independent copies of h is $(\frac{\varepsilon}{2(\gamma+1)}, \frac{\varepsilon}{2(\gamma+1)})$ admissible.

Nissim et al. show that a Cauchy distribution satisfies Theorem 3. They also show that approximate differential privacy can be satisfied under smooth sensitivity using noise from a gaussian or laplace distribution. While we show only pure $(\varepsilon, 0)$ -differential privacy below, it is easily extended to approximate differential privacy; we show how this compares empirically in Section 7.

4 Differentially Private Naïve Bayes

While smooth sensitivity has been known for some time, it is often challenging to apply to practical problems. Unlike global sensitivity, which requires only a worst-case analysis, to use Theorem 2 with a naive application of Definition 3.4 is exponential in dataset size. We now show how smooth sensitivity can be applied to create a differentially private Naïve Bayes classifier. We first compute the Naïve Bayes parameters (Section 2.1). We

then perturb the parameters with noise that preserves ε -differential privacy. As stated in Section 2.1, a standard approach for fitting a machine learning model to a numerical attribute is to assume that the underlying distribution is Gaussian. We also assume that the numerical feature values are bounded. Hence, we start by computing the smooth sensitivity for estimating the parameters of the Truncated normal distribution.

Note that we only train using a subset of the data. The way this subset is defined enables a smooth bound, as neighboring databases result in (at worst) training on a different subset rather than completely new data.

Definition 4.1. (Trimmed Sample) Let $X = x_1, x_2, \dots, x_n$ denote the sample in sorted order and m be a trimming parameter. The trimmed sequence of sample X is:

$$x_{m+1}, x_{m+2}, \dots, x_{n-m}$$

In other words, we draw a window of size $n - 2m$ on the data.

4.1 Smooth Sensitivity of the Mean

We start by computing the Smooth sensitivity of the mean of a dataset.

Theorem 4. Given a list of n bounded real numbers $V = \{x_1, x_2, \dots, x_n\}$ in the range of the interval $[L, U]$, the Smooth sensitivity of the mean of V can be computed in $\mathcal{O}(n^2)$.

Proof. Without loss of generality we assume that V is in non-decreasing order. Now consider the set $V = \{x_1, x_2, \dots, x_n\}$ with mean $\mu_V = \frac{x_1 + x_2 + \dots + x_n}{n}$ and a set $V' = \{x'_1, x'_2, \dots, x'_n\}$ with mean $\mu_{V'} = \frac{x'_1 + x'_2 + \dots + x'_n}{n}$ where $d(V, V') = k$, i.e., it differs from V by k elements such that $|\mu_V - \mu_{V'}|$ is maximized. In the case that $\mu_{V'} > \mu_V$, it is easy to see that we have to replace x_1, x_2, \dots, x_k with U (See Fig. 1). Similarly, in the case that $\mu_V > \mu_{V'}$, we should replace $x_{n-k+1}, x_{n-k+2}, \dots, x_n$ with L . Iterating through all possible choices of $1 \leq k \leq n$ and changing one last element to its extreme case (L or U) would give use the smooth sensitivity. \square

4.2 Smooth Sensitivity of Variance

In this section, we will describe how to compute the smooth sensitivity of the variance of a dataset.




Fig. 1. (top): blue dots represents points on the x axis where the red cross mark represents their corresponding mean. (middle): shifting the mean towards the right-side by picking the k -smallest numbers and replacing them with U . (bottom): shifting the mean towards the left-side by picking the k -largest numbers and replacing them with L .

Definition 4.2. (k -maximal variance subset) Given a set $V = \{x_1, x_2, \dots, x_n\}$ of n real numbers in ascending order and an integer k ($k < n$), a subset $Q \subset V$ is called a k -maximal variance subset of V if $|Q| = k$ and $\forall Q' \subset V, Q' \neq Q$ where $|Q'| = k$, $\text{Var}[Q'] \leq \text{Var}[Q]$.

Theorem 5. Given a list of n bounded real numbers $V = \{x_1, x_2, \dots, x_n\}$ in the range of interval $[L, U]$ and an integer k ($k < n$), the k -maximal variance subset can be computed in $\mathcal{O}(n^2)$.

Proof. Without loss of generality assume that V is in non-decreasing order. Let $Q = \{\Gamma_1, \Gamma_2, \dots, \Gamma_k\}$ be the k -maximal variance of V . We define the mean of Q to be $\mu_Q = \frac{1}{k}(\Gamma_1 + \Gamma_2 + \dots + \Gamma_k)$. Let $Q' = \{\Gamma_1, \Gamma_2, \dots, \Gamma_k + \Delta\}$, i.e., it only differs from Q by adding Δ to Γ_k . Let $\mu_{Q'}$ be the mean of Q' , we have:

$$\mu_{Q'} = \mu_Q + \frac{1}{k}\Delta$$

Let σ_Q^2 and $\sigma_{Q'}^2$ be the variance of Q and Q' , respectively. The variance of $Q = \{\Gamma_1, \Gamma_2, \dots, \Gamma_k\}$ (multiplied by k) is:

$$k\sigma_Q^2 = (\Gamma_1^2 + \dots + \Gamma_k^2) - k\mu_Q^2$$

Now we will see the impact of adding Δ to Γ_k on the variance.

$$\begin{aligned} k\sigma_{Q'}^2 - k\sigma_Q^2 &= [(\Gamma_k + \Delta)^2 - \Gamma_k^2] - k[\mu_{Q'}^2 - \mu_Q^2] \\ &= [\Gamma_k^2 + \Delta^2 + 2\Gamma_k\Delta - \Gamma_k^2] - k(\mu_Q + \frac{\Delta}{k})^2 \\ &\quad + k\mu_Q^2 \\ &= [\Gamma_k^2 + \Delta^2 + 2\Gamma_k\Delta - \Gamma_k^2] - k\mu_Q^2 \\ &\quad - \frac{\Delta^2}{k} - 2\mu_Q\Delta + k\mu_Q^2 \\ &= 2\Delta(\Gamma_k - \mu_Q) + \frac{k-1}{k}\Delta^2 \\ &\geq 2\Delta(\Gamma_k - \mu_Q) \end{aligned}$$

This value will be non-negative whenever the sign of Δ and $(\Gamma_k - \mu_Q)$ are the same. That is, the difference

in the variance will increase if we move Γ_k further from the mean μ_Q .

Now assume that we are given a variance-maximizing sequence S of k values chosen from $V = \{x_1, x_2, \dots, x_n\}$. Assume for the contradiction that S contains an element x_m where x_m is not at the tail of V . That means there exists x_r ($r \neq m$), where if we replace x_m with x_r by the given inequalities we will increase the variance and $\Delta = x_m - x_r$ which contradicts that S is a variance maximizing sequence. So given $V = \{x_1, x_2, \dots, x_n\}$, we know that the k -maximal variance subset selects elements from the tail of V . Iterating through all possible cases of $0 \leq i \leq k$, $j = k - i$, where the first i elements are selected from the beginning of V , i.e., (x_1, x_2, \dots, x_i) , and j elements are selected from the end of the sequence, i.e., $(x_{n-j+1}, \dots, x_{n-1}, x_n)$, selecting the sequence that gives the maximum variance will be the solution to the k -maximal variance subset. \square

Corollary 5.1. Given a list of n bounded real numbers $V = \{x_1, x_2, \dots, x_n\}$ in the range $[L, U]$ and an integer k ($k < n$), the k -maximal variance subset can be achieved by removing $n - k$ consecutive elements in V .

Definition 4.3. (k -minimal variance subset) Given a set $V = \{x_1, x_2, \dots, x_n\}$ of n real numbers in ascending order and an integer k ($k < n$), a subset $Q \subset V$ is called k -minimal variance subset of V if $|Q| = k$ and $\forall Q' \subset V, Q' \neq Q$ where $|Q'| = k$, $\text{Var}[Q'] \geq \text{Var}[Q]$.

Theorem 6. Given a list of n real numbers $V = \{x_1, x_2, \dots, x_n\}$ in the range of the interval $[L, U]$ and an integer k ($k < n$), the k -minimal variance subset can be computed in $\mathcal{O}(n^2)$.

Proof. The proof is similar to the proof of Theorem 5. Without loss of generality assume that V is in non-decreasing order. Hence, given $V = \{x_1, x_2, \dots, x_n\}$ the optimal solution would remove $\{x_1, x_2, \dots, x_i\}$ and $\{x_{n-k+i}, \dots, x_{n-1}, x_n\}$. Iterating through all possible $0 \leq i \leq k$ would give the optimal solution. \square

Theorem 7. Given a list of n bounded real numbers $V = \{x_1, x_2, \dots, x_n\}$ in the range of interval $[L, U]$, the Smooth sensitivity of the variance of V can be computed in $\mathcal{O}(n^2)$.




Fig. 2. Blue dots represents points on the x axis where the red cross mark represents their corresponding mean. For a fixed $k = 6$, we use a sliding window (shaded box) and move it through x axis. We assign t of the numbers to L and $k - t$ to U and save the maximum variance of the resulted sequence.

Proof. By Corollary 5.1, the k -maximal variance subset can be achieved by removing $n - k$ consecutive elements. Without loss of generality assume that V is in non-decreasing order. We can iterate through all k consecutive elements in V and assign t of them to be L , and $k - t$ of them to be U . The maximum over all possible cases would be the maximal variance. Similarly, for minimizing the variance by using Theorem 6, we replace $\{x_1, x_2, \dots, x_t\}$, i.e., the first t elements and $\{x_{n-k+t}, \dots, x_n\}$, i.e., the last $k - t$ elements with $\mu = \frac{x_{t+1} + x_{t+2} + \dots + x_{n-k+t-1}}{(n-k-2)}$. Changing one element to its extreme case (U or L), or to the mean of the sequence would give us the smooth sensitivity. \square

Fig. 3 shows an example for the extreme change of the global sensitivity.




Fig. 3. Maximum change of the global sensitivity. (top): all points are on L , by changing $k = 3$ points to the other extreme point (U), the variance and mean have the maximum change.

Note that another way to compute a differentially private variance of bounded dataset is to use $\text{Var}(X) = E[X^2] - E[X]^2$. In this case, we would need to compute a differentially private $E[X]$ and $E[X^2]$. While we already have the mean $E[X]$, we would need to use the privacy budget used for $\text{Var}(X)$ to instead calculate $E[X^2]$. However, this results in $\text{Var}(X)$ being computed from two noisy values, resulting in a less accurate result than our method of computing it directly.

5 Bun and Steinke's Mean Estimation Using Smooth Sensitivity

Bun and Steinke [10] provided an algorithm for privately estimating the mean of a distribution. Their algorithm estimates the mean from an i.i.d. sample x of the dataset. They first assume a crude bound on the mean $\mu \in [a, b]$, then they truncate the samples, i.e. they remove the largest m samples and smallest m samples from x , and compute the mean of the $n - 2m$ samples. Finally they project the estimated mean to the range $[a, b]$. They have defined several admissible distributions such that adding noise proportional to them would guarantee $\frac{1}{2}\epsilon^2$ -CDP (Concentrated Differential Privacy). These include Student's T, Laplace log-normal, uniform log-normal, and arsinh-normal. Theorem 8 is the main result of their paper:

Theorem 8. (Bun and Steinke [10]) Let $n \geq O(\log((b-a)\sigma)/\epsilon)$, then there exist a ϵ -DP (or $\frac{1}{2}\epsilon^2$ -CDP) algorithm $\mathcal{M} : R^n \rightarrow R$ such that, for all $\mu \in [a, b]$, we have:

$$\mathbb{E}[(\mathcal{M}(x) - \mu)^2] \leq \frac{\sigma^2}{n} + \frac{\sigma^2}{n^2} \cdot O\left(\frac{\log \frac{b-a}{\sigma}}{\epsilon} + \frac{\log n}{\epsilon^2}\right)$$

In Theorem 8, the first part is the non-private optimal mean-squared error and the additional term is the cost of the privacy. The key difference between our approach and that of [10] is that we assume structural or data-independent bounds on values (e.g., age between 0 and 125). While less general than [10], in practical cases this allows for better differentially private estimates. This also enables an independent private estimate of variance (Section 4.2) that provides better results than basing variance on estimates of the mean. The approach of [10] also requires choosing a smoothing parameter; it is not clear how to do this in a data independent or differentially private manner (although in our experiments and those of [10], the results were not that sensitive to the choice of smoothing parameter.)

Lemma 9 in [10] shows that concentrated-differential privacy holds for the Laplace log-normal distribution. From [10] and the smooth sensitivity paper [9], it follows that pure differential privacy holds (the above Theorem) when using the Cauchy distribution. To give a fair comparison with the pure differential privacy of our approach, we use the Cauchy distribution with Bun and Steinke's approach in our comparisons.

6 Algorithm

We now give pseudocode to describe the algorithm for computing the smooth sensitivity of the Naïve Bayes classifier. At a high level, we first compute the parameters of the Naïve Bayes model, then compute the smooth sensitivity of each parameter and perturb the parameters with noise drawn from a Cauchy distribution. From Theorem 4 and Theorem 7, one can compute the smooth sensitivity of the parameters for fitting a Gaussian distribution to the continuous data. For discrete variables the sensitivity is 1 and can be perturbed by adding the small amount of noise $1/\epsilon$.

We use an equal division of privacy budget between all accesses to data in keeping with [1]. Our goal for this paper is to show the value of smooth sensitivity, so we have kept with their division; we briefly discuss other allocations of privacy budget in Section 10.1.

Similar to Vaidya et al.'s [1] approach, when Cauchy noise is added, it is possible to make mean, standard deviation, counts, and class prior negative. To prevent this, we truncate the negative values to zero (a postprocessing step that does not impact the differential privacy guarantee.)

6.1 Complete Privacy Guarantee

Theorem 9. Algorithm 1 provides ϵ -differential privacy.

Proof. Each step i of the Algorithm 1 is ϵ'_i -differentially private. By Theorem 1, the composition of finite number of ϵ'_i -differentially private algorithm is itself $(\sum_i \epsilon'_i)$ -differentially private. \square

6.2 Runtime Analysis of Algorithm 1

Bounding the smooth sensitivity for a given dataset does come at a cost:

Theorem 10. Algorithm 1 computes differentially private Naïve Bayes in $\mathcal{O}(nk + n^2)$.

Proof. The most time consuming part of computing differentially private Naïve Bayes classifier is Theorem 5 that computes the smooth sensitivity of the variance in $\mathcal{O}(n^2)$ where n is the number of rows in dataset. The rest of the computation can be done in linear time, hence, the pre-processing time is $\mathcal{O}(n^2)$. The running time of

Algorithm 1: DIFFERENTIALLY PRIVATE
NAÏVE BAYES CLASSIFIER

Input:

- Labeled training data $D = \{(X^{(i)}, y_i)\}_{i=1}^n$
- ε : the privacy parameter
- $Lap(\alpha, \beta)$: samples the Laplace distribution with mean α and scale β
- $Cauchy(\alpha, \beta)$: samples the Cauchy distribution with location parameter α and scale β
- Bound $[a_i, b_i]$ for each numerical attribute X_i

```

1  $num\_count \leftarrow$  number of numerical attributes
  of  $D$ 
2  $cat\_count \leftarrow$  number of categorical attributes
  of  $D$ 
3  $\varepsilon' \leftarrow \varepsilon / (2 * num\_count + cat\_count + 1)$ 
4 for each attribute  $X_j$  do
5   if  $X_j$  is categorical then
6     sensitivity,  $s \leftarrow 1$ 
7     scale factor,  $sf \leftarrow s / \varepsilon'$ 
8      $\forall$  counts  $\tau\{X_j = a_k \wedge C = c_l\}' =$ 
        $\tau\{X_j = a_k \wedge C = c_l\} + Laplace(0, sf)$ 
9     Use  $\tau\{X_j = a_k \wedge C = c_l\}'$  to compute
        $Pr(a_k | c_l)$ 
10  else if  $X_j$  is numeric then
11    Trim the given sample based on
      Definition 4.1.
12    Compute the Smooth Sensitivity,  $s$  for
      mean  $\mu_{jk}$  as per Theorem 4 with
      bound  $[a_j, b_j]$ 
13    scale factor,  $sf \leftarrow \sqrt{2}s / \varepsilon'$ 
14     $\mu'_{jk} \leftarrow \mu_{jk} + Cauchy(0, sf)$ 
15    Compute the Smooth Sensitivity,  $s$  for
      the standard deviation  $\sigma_{jk}$  as per
      Theorem 7 with bound  $[a_j, b_j]$ 
16    scale factor,  $sf \leftarrow \sqrt{2}s / \varepsilon'$ 
17     $\sigma'_{jk} \leftarrow \sigma_{jk} + Cauchy(0, sf)$ 
18    Use  $\mu'_{jk}$  and  $\sigma'_{jk}$  to compute  $Pr(x_i | c_k)$ 
19  for each class  $c_j$  do
20    count  $\tau\{C = c_j\}' \leftarrow \tau\{C = c_j\} +$ 
       $Cauchy(0, \sqrt{2}S / \varepsilon')$ 
21    Use  $\tau\{C = c_j\}'$  to compute the prior
       $Pr(c_j)$ 

```

the standard Naïve Bayes is $\mathcal{O}(nk)$, therefore, the total running time of differentially private Naïve Bayes with the added pre-processing is $\mathcal{O}(nk + n^2)$. \square

7 Empirical Analysis

We have implemented the Naïve Bayes classifier of Algorithm 1 in Python. We show a comparison with the results presented in [1], as well as a more realistic example that was not used in their work. We have also compared our algorithm to Bun and Steinke's [10] private mean estimation. Since they did not specify an algorithm for computing the variance of a distribution, and our variance computation requires known bounds on values, we have naively used their mean estimation algorithm to estimate the variance using $Var(x) = E[X^2] - E[X]^2$. This demonstrates the practical improvements realized with smooth sensitivity. We also show practical computational costs required to achieve this benefit.

7.1 Datasets

The datasets used in our experiment include those from the UCI repository [14] used in [1]: Adult, Mushroom, Skin, Seed and Glass. We also give results from a much more realistic dataset: the IPUMS USA: Version 8.0 Extract of 1940 Census for U.S. Census Bureau Disclosure Avoidance Research.

The UCI **Adult dataset** is drawn from 1994 census data of the United States. It consists of a 48K record subset drawn from a stratified sample of the U.S. population, the binary classification task is to predict if the income of an individual is less than or equal to 50K or not.

The UCI **Mushroom dataset** includes descriptions of hypothetical samples corresponding to 23 species of gilled mushrooms in the Agaricus and Lepiota Family. Each species is identified as definitely edible, definitely poisonous, or of unknown edibility and not recommended.

The UCI **Skin dataset** [15] is collected by randomly sampling B,G,R values from face images of various age groups (young, middle, and old), race groups (white, black, and asian), and genders obtained from the FERET [16] and PAL databases.

The UCI **Seed dataset** includes group comprised kernels belonging to three different varieties of wheat:

Kama, Rosa and Canadian, 70 elements each, randomly selected for the experiment.

The UCI **Glass dataset** contains the description of 214 fragments of glass originally collected for a study in the context of criminal investigation. Each fragment has a measured reflectivity index and chemical composition (weight percent of Na, Mg, Al, Si, K, Ca, Ba and Fe).

The IPUMS **1940 Census** dataset is a sample drawn uniformly at random from the U.S. population, taken from the 1940 Census. The Adult dataset, and most other IPUMS microdata sets, are stratified datasets intended to be used with weighted values, and as such are not representative of real populations when used as unweighted values. Using with weighted values poses additional difficulties for computing sensitivity that are beyond the scope of this paper. When used as training data for machine learning, they are not representative of performance on real populations. The 1940 Census data is a uniform sample of the population, and as such is an appropriate representation for a machine learning task.¹ We used the 13 attributes that are included in the adult dataset, and construct a binary classification task to predict whether the income of an individual is less than or equal to the mean income of the population, a similar prediction task to that used with the Adult dataset (although with a different threshold, due to inflation between 1940 and 1994.) We discarded individuals with unknown values. To give an idea of the variance across different subpopulations, we report values for different U.S. States. Table 1 shows the detailed description of the datasets that we have used for this experiment.

Dataset	No. of Records	Attributes	Classes
Adult	48K	14	2
Mushroom	8K	22	2
Seed	210	7	3
Skin	245K	3	2
Glass	214	9	7
Wyoming	250K	13	2
Nevada	110K	13	2
Washington	1.7M	13	2
Oregon	1M	13	2

Table 1. Description of the datasets used for our experiment.

¹ The IPUMS data is available at <https://usa.ipums.org/usa/1940CensusDASTestData.shtml>.

7.2 Experimental Results

Since there is randomness in our algorithm for adding noise, we have run all algorithms five iterations with 10-fold cross-validation; we show mean and error bars across the iterations. We use two baselines: a standard Naïve Bayes classifier and the constant “predict the majority class” classifier. Probably the most interesting comparison is with the Differentially Private classifier by Vaidya, Shafiq, Basu, and Hong [1] and Bun and Steinke’s mean estimation algorithm [10] as this shows the specific gains achieved through using smooth sensitivity rather than global sensitivity.

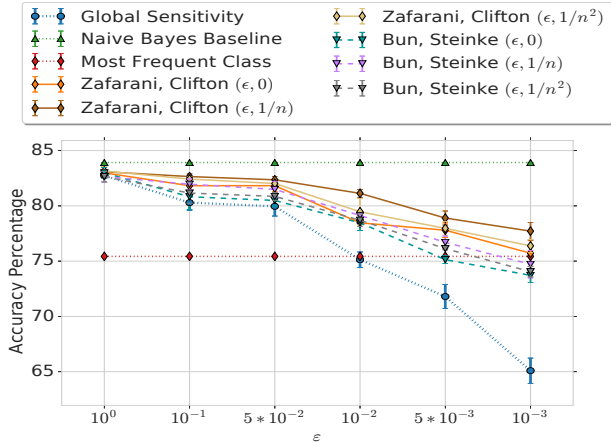





Note that [1] reports results in terms of the privacy budget used for each attribute; we instead report the total privacy budget utilized under sequential composition (Theorem 1). This is simply a scaling of ϵ and does not fundamentally change their reported results.

Results are shown in Figures 4-12. We give the mean value across multiple draws from the differentially privacy mechanisms, as well as standard deviation. Smooth sensitivity gives significant improvements in classifier accuracy, particularly for the datasets using human data. For completeness we also show results for (ϵ, δ) -differential privacy [12], a weaker form of differential privacy, for $\delta = 1/n$ and $1/n^2$. (For clarity, algorithm 1 shows only $(\epsilon, 0)$ -differential privacy; the extension to (ϵ, δ) is straightforward.) As shown by Nissim et al. [9], if we add Laplace or Gaussian noise of magnitude calibrated to the smooth sensitivity, this would give approximate differential privacy. In this work, to achieve approximate differential privacy, instead of using Cauchy distribution in Algorithm 1, we have used Gaussian distribution. As it is shown in Figures 4-12, approximate-DP does give some accuracy improvement, but at a cost of a weaker privacy guarantee.

Since we only use smooth sensitivity with continuous values and the Mushroom dataset contains only categorical values, the results in Figure 5 are the same for all methods; we include for completeness with [1] and to show variance across parameter values.

7.3 Computational Costs

As we showed in Theorem 1, smooth sensitivity does not come for free. The global sensitivity algorithm is $O(n \log n)$, smooth sensitivity takes us to $O(n^2)$. To demonstrate what this means in practical terms, we show how results change as we vary the dataset size. Fig. 13 compares the two approaches, showing runtimes

Fig. 4. Accuracy at various values of ϵ : Adult datasetFig. 7. Accuracy at various values of ϵ : Seed datasetFig. 5. Accuracy at various values of ϵ : Mushroom datasetFig. 8. Accuracy at various values of ϵ : Skin datasetFig. 6. Accuracy at various values of ϵ : Glass datasetFig. 9. Accuracy at various values of ϵ : Wyoming from 1940 Census




Fig. 10. Accuracy at various values of ϵ : Nevada from 1940 Census




Fig. 11. Accuracy at various values of ϵ : Oregon from 1940 Census




Fig. 12. Accuracy at various values of ϵ : Washington from 1940 Census




Fig. 13. Runtime Analysis on datasets provided in Table 1, blue dot represents the Global Sensitivity training time while the red star represent the corresponding Smooth sensitivity training time in seconds (log log scale).

of Python implementations on a computer with an 2.2 GHz Intel Core i7 CPU, 64 GB 1600 MHz DDR3 RAM, running Ubuntu 18.04. While we do see a substantial runtime cost in training the Naïve Bayes classifier, these are reasonable times for many practical applications. It is interesting to note that as the dataset size increases, factors other than sensitivity calculation become increasingly important; the nearly order of magnitude difference in cost with a 5000 instance dataset drops to half an order of magnitude (although still a substantial time difference) with over a million instances.

8 Related Work

We have discussed the work of Vaidya et al. [1], which addresses the same problem we do using global sensitivity. Li et al. [17] proposed a new model for a differentially private Naïve Bayes classifier over multiple data sources. Their proposed method enables a trainer to train a Naïve Bayes classifier over the dataset provided jointly by different data owners, without requiring a trusted aggregator as in our work and [1]. Yilmaz et al. [18] provided a differentially private Naïve Bayes classifier under the local differential privacy setting. With local differential privacy, individuals perturb their data before sending to an *untrusted* aggregator. The stronger adversary model of these two approaches (eliminating the trusted aggregator) results in significantly more noise and reduced accuracy.

There have been studies of privacy-preserving Naïve Bayes under different privacy models. Kantarcioglu et al. [19] proposed a privacy-preserving Naïve Bayes classifier for horizontally partitioned data. Their solution uses secure summation and logarithm to learn a distributed Naïve Bayes classifier securely. Vaidya and Clifton [20] gave a solution to the same problem but for vertically partitioned data under the semi-honest model. These approaches protect the data during training, an orthogonal problem to differential privacy's protection of disclosure via the learned model.

Making machine learning models differentially private has had more general interest, with solutions proposed for several machine learning approaches. Some of the more well known include Jagannathan, Pillaipakkamnatt and Wright [4] that gives a differentially private algorithm for random decision trees, and Abadi et al. [6] that give a differentially private framework for deep learning models. One possible area for future work is to determine if these approaches could benefit from using smooth sensitivity rather than global sensitivity, and if so, how that might be tractably computed.

Beyond machine learning, there are many differentially private algorithms for statistical tests. Campbell et al. [21] gave a differentially private ANOVA test. Task and Clifton [22] provided differentially private significance testing on paired-sample data. In Fig. 14, we show that a non-uniform allocation of privacy budget to different attributes can have an impact; we discuss this further in the Appendix. This is not a trivial problem in the context of differential privacy. Anandan and Clifton [23] provided a differentially private solution for a more basic version of this problem, feature selection for data mining tasks. In their work they analyze the sensitivity of various feature selection techniques used in data mining and show that some of them are not suitable for differentially private analysis due to high sensitivity.

9 Conclusion

In this paper, we have developed a differentially private Naïve Bayes Classifier using Smooth Sensitivity for numerical data, along with global sensitivity for categorical values. For fitting numerical values, we have made the assumption typically used with Naïve Bayes that the data follows a Gaussian distribution. When the features are bounded, we assume that the underlying data follows a truncated normal distribution. We have computed the smooth sensitivity of the parameters

of the Gaussian, μ and σ . To obtain the ϵ -differential private algorithm, we have added noise proportional to the smooth sensitivity of the parameters. Previous work on Naïve Bayes differential private classifier done by Vaidya, Shafiq, Basu, and Hong [1] perturb the parameters of the Naïve Bayes classifier by a noise that is scaled to the global sensitivity of the parameters. We demonstrate on real-world datasets that our method achieves a significant accuracy improvement. While this comes at a computational cost, it is a cost only in model training. The released model is essentially identical to [1], but with higher accuracy, and still satisfying the ϵ -differential privacy definition. Smooth sensitivity provides the same differential privacy guarantee as global sensitivity for given values of ϵ and δ , including $(\epsilon, 0)$. The greater the impact of numerical features on the result, the greater the benefit of smooth sensitivity. We have also compared our result with private mean estimation of Bun and Steinke [10] where they estimate the mean of an unknown distribution based on an i.i.d. sample from the dataset, adding noise proportional to the smooth sensitivity of the truncated mean.

Acknowledgements

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

- [1] J. Vaidya, B. Shafiq, A. Basu, and Y. Hong, "Differentially private naive bayes classification," in *2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, vol. 1, pp. 571–576, IEEE, 2013.
- [2] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*, pp. 265–284, Springer, 2006.
- [3] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *FOCS*, vol. 7, pp. 94–103, 2007.
- [4] G. Jagannathan, K. Pillaipakkamnatt, and R. N. Wright, "A practical differentially private random decision tree classifier," in *2009 IEEE International Conference on Data Mining Workshops*, pp. 114–121, IEEE, 2009.
- [5] B. I. Rubinstein, P. L. Bartlett, L. Huang, and N. Taft, "Learning in a large function space: Privacy-preserving mechanisms for svm learning," *arXiv preprint arXiv:0911.5708*, 2009.
- [6] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with

- differential privacy,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 308–318, ACM, 2016.
- [7] K. Chaudhuri and C. Monteleoni, “Privacy-preserving logistic regression,” in *Advances in neural information processing systems*, pp. 289–296, 2009.
- [8] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth, “Guilt-free data reuse,” *Communications of the ACM*, vol. 80, pp. 86–93, Apr. 2017.
- [9] K. Nissim, S. Raskhodnikova, and A. Smith, “Smooth sensitivity and sampling in private data analysis,” in *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pp. 75–84, ACM, 2007.
- [10] M. Bun and T. Steinke, “Average-case averages: Private algorithms for smooth sensitivity and mean estimation,” in *Advances in Neural Information Processing Systems*, pp. 181–191, 2019.
- [11] T. M. Mitchell, “Machine learning,” 1997.
- [12] F. McSherry and I. Mironov, “Differentially private recommender systems: Building privacy into the netflix prize contenders,” in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 627–636, ACM, 2009.
- [13] C. Dwork and J. Lei, “Differential privacy and robust statistics,” in *STOC*, vol. 9, pp. 371–380, 2009.
- [14] D. Dua and C. Graff, “UCI machine learning repository,” 2017.
- [15] R. Bhatt and A. Dhall, “Skin segmentation dataset,” *UCI Machine Learning Repository*, 2010.
- [16] P. J. Phillips, H. Wechsler, J. Huang, and P. J. Rauss, “The feret database and evaluation procedure for face-recognition algorithms,” *Image and vision computing*, vol. 16, no. 5, pp. 295–306, 1998.
- [17] T. Li, J. Li, Z. Liu, P. Li, and C. Jia, “Differentially private naive bayes learning over multiple data sources,” *Information Sciences*, vol. 444, pp. 89–104, 2018.
- [18] E. Yilmaz, M. Al-Rubaie, and J. M. Chang, “Locally differentially private naive bayes classification,” *arXiv preprint arXiv:1905.01039*, 2019.
- [19] M. Kantarcioglu, J. Vaidya, and C. Clifton, “Privacy preserving naive bayes classifier for horizontally partitioned data,” in *IEEE ICDM workshop on privacy preserving data mining*, pp. 3–9, 2003.
- [20] J. Vaidya and C. Clifton, “Privacy preserving naive bayes classifier for vertically partitioned data,” in *Proceedings of the 2004 SIAM International Conference on Data Mining*, pp. 522–526, SIAM, 2004.
- [21] Z. Campbell, A. Bray, A. Ritz, and A. Groce, “Differentially private anova testing,” in *2018 1st International Conference on Data Intelligence and Security (ICDIS)*, pp. 281–285, IEEE, 2018.
- [22] C. Task and C. Clifton, “Differentially private significance testing on paired-sample data,” in *Proceedings of the 2016 SIAM International Conference on Data Mining*, pp. 153–161, SIAM, 2016.
- [23] B. Anandan and C. Clifton, “Differentially private feature selection for data mining,” in *Proceedings of the Fourth ACM International Workshop on Security and Privacy Analytics*, pp. 43–53, ACM, 2018.




Fig. 14. Adult dataset: Accuracy at various values of ϵ for varying allocation of privacy budget - numerical:categorical

10 Appendix

10.1 Varying Allocation of the Privacy Budget

We briefly experimented with varying the allocation of the privacy budget between categorical and numeric variables, an example is shown in Fig. 14. With the exception of substantially overweighting categorical variables and giving little privacy budget (high noise) for numeric (line 1:2), the results are inconclusive; the differences are well within the error seen across cross-validation groups. This suggests that numeric attributes are important to this classification problem, and due to the higher sensitivity, giving too little privacy budget to them limits their value.

Algorithm 1 and all other figures correspond to the “2:1” allocation in Fig. 14: $2 \times \epsilon'$ allocated to each numeric attribute, one for the mean estimation, and one for standard deviation.

Further research would be needed to establish reasoning to apply additional weight to numeric or categorical variables. Making the selection empirically before applying differential privacy would constitute a disclosure of information, and those violate the provided ϵ -differential privacy guarantee. (Our choice of a “2:1” allocation outside of Fig. 14 was based on the need to need to gather both mean and standard deviation, rather than based on empirical analysis.) Properly arriving at a weighting in a way that is both private and effective is an interesting problem, and one that comes up generally in differentially private machine learning.




Fig. 15. Accuracy at various values of ϵ for varying allocation of privacy budget - numerical:categorical on synthetic dataset where only the categorical attributes are correlated with the class.




Fig. 16. Accuracy at various values of ϵ for varying allocation of privacy budget - numerical:categorical on synthetic data where only the numerical attributes are correlated with the class.

To measure the effect of privacy budget on numerical and categorical attributes we further experimented on a synthetic dataset. We generated two synthetic datasets of 10000 entries with 5 categorical attributes and 5 numerical. In the first dataset, the numerical attributes are correlated with the label while the categorical values are uncorrelated. In the second dataset, the categorical attributes are correlated with the label while numerical attributes are uncorrelated. The impact of different distributions of privacy budget are presented in Fig. 16 and Fig. 15. Unsurprisingly, the optimal privacy budget distribution is dataset dependent.

The unbalanced weight helps numerical more than it hurts categorical, particularly at higher privacy values. This supports our conjecture that a 2:1 split would be

appropriate since we estimate more values for numerical attributes.