Privacy-Preserving Public Release of Datasets for Support Vector Machine Classification

Farhad Farokhi, *Senior Member, IEEE*

Abstract—We consider the problem of publicly releasing a dataset for support vector machine classification while not infringing on the privacy of data subjects (i.e., individuals whose private information is stored in the dataset). The dataset is systematically obfuscated using an additive noise for privacy protection. Motivated by the Cramér-Rao bound, inverse of the trace of the Fisher information matrix is used as a measure of the privacy. Conditions are established for ensuring that the classifier extracted from the original dataset and the obfuscated one are close to each other (capturing the utility). The optimal noise distribution is determined by maximizing a weighted sum of the measures of privacy and utility. The optimal privacy-preserving noise is proved to achieve local differential privacy. The results are generalized to a broader class of optimization-based supervised machine learning algorithms. Applicability of the methodology is demonstrated on multiple datasets.

Index Terms—Database Privacy, Support Vector Machine Classification, Fisher Information, Coding and Information Theory.

1 INTRODUCTION

ADVANCES in communication and information processing have opened new possibilities for data mining and big data analysis to answers important challenges facing society. Public and private entities have thus scrambled to capitalize on these possibilities for improving the quality of offered services in a data-oriented manner. However, these enhancements come at the expense of the erosion of privacy within society. Therefore, there is a need for developing algorithms balancing utility and privacy.

Most often, governments would like to provide entire datasets to the public (or select private entities) in a *de-identified (anonymized) manner* so that academics, analysts, and researchers can utilize them for gaining valuable insights and developing new technologies. Release of de-identified data can still infringe on the privacy of people [1]. Therefore, methods, such as suppression and generalization, have been previously used for private data release by guaranteeing k -anonymity. These methods also may not provide adequate individual privacy guarantees [2].

This paper proposes a novel method for *privacy-preserving release of an entire dataset while maintaining useful properties, such as statistics required for reconstructing a support vector machine classifier*. This is done by balancing privacy and utility guarantees using an explicit optimization problem. The dataset is systematically obfuscated using an additive noise and the inverse of the trace of the Fisher information matrix is used as a measure of privacy for the entries of the dataset. By the use of the Cramér-Rao bound [3, p.169], it can be seen that the defined measure of privacy provides a lower bound on the ability of an adversary estimating the individual entries of the dataset. The use of the Fisher information matrix makes the privacy metric independent of the sophistication of the adversary, thus making it a universal measure of privacy. Further, the Cramér-Rao bound provides a practical/operational interpretation of the measure of privacy to the data owners,

i.e., how much someone can learn about an individual in the dataset based on the publicly released obfuscated data. Conditions are provided for ensuring that the classifier extracted from the original data and the noisy data are “close” to each other. This serves as a measure of utility. To find the optimal noise distribution, a weighted sum of the measure of privacy and the measure of utility is maximized. In some cases, it is proved that *the classifier extracted from the noisy privacy-preserving dataset is identical to the classifier extracted from the original private data*. Therefore, the utility is fully preserved while preserving the privacy. Finally, we show that the optimal additive noise from the presented framework also ensures (ϵ, δ) -local differential privacy. This is an interesting observation because (i) the new methodology inherits the strong guarantees that come with the local differential privacy and (ii) we can provide an operative interpretation for the parameters of (ϵ, δ) -local differential privacy using the new measure of privacy in this paper and with the aid of the Cramér-Rao bound.

In summary, this paper makes the following contributions in privacy-preserving release of datasets:

- Using the Cramér-Rao bound, an adversary’s estimation error for reconstructing the individual entries of a private dataset is related to the inverse of the trace of the Fisher information matrix. This enables the use of the Fisher information as a measure of privacy.
- The effect of additive privacy-preserving noise on the utility of machine learning models is captured. The utility of the machine learning models is negatively impacted by the variance of the additive noise. This is first proved for support vector machines and then generalized to optimization-based machine learning.
- The optimal privacy-preserving noise distribution is computed by maximizing a weighted sum of the measures of privacy and utility.
- The optimal privacy-preserving is shown to be Gaussian and proved to satisfy (ϵ, δ) -local differential privacy. For additive privacy-preserving noises that can be correlated among the entries of the dataset,

• F. Farokhi is with the CSIRO’s Data61 and the University of Melbourne. e-mails: farhad.farokhi@unimelb.edu.au; farhad.farokhi@data61.csiro.au

a special multi-dimensional Gaussian noise can be realized that does not degrade the utility of linear support vector machines while providing privacy.

- An operative interpretation for the parameters of (ϵ, δ) -local differential privacy is provided by proving that the ability of an adversary estimating the individual entries of the dataset is bounded by $\mathcal{O}(\ln^2(\delta^{-1})\epsilon^{-2})$.
- The effect of the optimal privacy-preserving noise is demonstrated on the Wisconsin Breast Cancer Diagnostic dataset and the Adult dataset using linear support vector machines and the Lending Club using linear regression. The utility of the machine learning models trained on the data obfuscated by the optimal noise is always better than the utility of machine learning models trained on the data obfuscated by the Laplace mechanism for providing local differential privacy with the same privacy guarantee.

1.1 Related Work

Most often k -anonymity and its extensions, such as ℓ -diversity and m -invariance, have been utilized for releasing datasets in a privacy-preserving manners [4]–[6]. These methods are however vulnerable to attacks; see, e.g., [2], [7], [8]. This demonstrate the need for the use of privacy-preserving noise when releasing datasets publicly.

Differential privacy is an important methodology for privacy-preserving data release and handling because of possessing strong guarantees and post-processing properties [9]. Differential privacy has been successfully utilized in machine learning problems with applications to support vector machines [10], [11], logistic regression [12], [13], and deep learning [14], [15]. These studies, however, concentrate on responding to aggregate queries, such as releasing trained machine learning models in a privacy-preserving manner or reporting statistics of a dataset without jeopardising privacy of individuals. Another relevant notion of privacy is local differential privacy in which the data of individuals in perturbed *locally* to ensure privacy, perhaps due to the lack of trust in the aggregator. Although powerful in preserving privacy, differential privacy and local differential privacy often lack systematic approaches for selecting their parameters, also known as the privacy budget, in order to balance utility and privacy. This has caused concerns in effectiveness of implementations of differentially-private mechanisms in practice [16]. In this paper, with the aid of the Fisher information and the Cramér-Rao bound, we also provide a systematic approach for setting the privacy budgets in local differential privacy when using the Gaussian mechanism.

An alternative to differential privacy is information-theoretic privacy, dating back to wiretap channels [17] and their extension [18]. Information-theoretic measures of privacy use mean square estimation error [19], mutual information [20], and the Fisher information [21] for measuring private information leakage. The use of entropy as a measure of privacy forces the private dataset to be statistically distributed with known distributions. For instance, it is assumed that the density of the private dataset is known [22] or that the data is Gaussian in [23]. This weakens the results

as the density of the data might not be available in advance or that data might follow the underlying assumptions. Therefore, in this paper, we use the Fisher information as a measure of private information leakage because it does not impose statistical assumptions on the private dataset. In fact, when using the Fisher information, the private dataset is treated as an unknown arbitrary deterministic object with no assumptions about its origin. The Fisher information has been previously used for privacy preservation in [21], [24]–[26]. In those studies, however, the utility was not tailored to preserving quality of machine learning models as in this paper. For instance, in this paper, we show that a special multi-dimensional Gaussian noise can be used for privacy-preservation that does not degrade the utility of linear support vector machines. This would not have been possible following those studies.

Other approaches for privacy-preserving machine learning have been proposed previously [27]. These approaches also rely on privacy-preserving release of the machine learning model rather than the dataset.

In image datasets, partial obfuscation of images to hide important private features, such as faces or license plates, have been considered [28], [29]. However, these studies do not easily generalize to other dataset.

The results of this paper are close, in spirit, to construction of synthetic datasets in which a dataset is generated to match statistics of the original dataset perturbed in a differentially-private manner; see, e.g., [30], [31]. Synthetic datasets are known to contain implausible entries (e.g., smoking infant) due to only considering low order statistics (e.g., the first and second order statistics). This is because of the underlying computational complexity associated with realising synthetic datasets. Furthermore, generation of synthetic datasets requires stochasticity assumptions.

1.2 Organization

The rest of the paper is organized as follows. The methodology is developed in Sections 2. Numerical results are demonstrated in Section 3. Finally, Section 4 concludes the paper and provides some directions for future research.

2 METHODOLOGY

We start with presenting the results for the simple case of linear support vector machines. At the end of this section, in Subsection 2.4, we show that the results hold for more general optimization-based machine learning algorithms.

2.1 Soft Margin Support Vector Machine

Consider a set of training data $\{(x_i, y_i)\}_{i=1}^q \subseteq \mathbb{R}^p \times \{-1, +1\}$. In binary linear support vector machine classification, it is desired to obtain a separating hyper plane of the form $\{x \in \mathbb{R}^p : f(x) = \alpha^\top x + \beta = 0\}$, with its corresponding classification rule $\text{sign}(f(x))$ to group the training data into two sets (of $y = +1$ and $y = -1$), i.e., ensure that $\text{sign}(f(x_i)) = y_i$ for all i . Up to re-scaling of α and β , this problem can be cast as

$$\begin{aligned} \underset{\alpha \in \mathbb{R}^p, \beta \in \mathbb{R}, \xi \in \mathbb{R}^q}{\text{argmin}} \quad & \frac{1}{2} \alpha^\top \alpha + \theta \mathbf{1}^\top \xi, & (1a) \\ \text{s.t.} \quad & y_i (\alpha^\top x_i + \beta) \geq 1 - \xi_i, \quad \forall i \in \mathcal{Q}, & (1b) \\ & \xi_i \geq 0, \quad \forall i \in \mathcal{Q}. & (1c) \end{aligned}$$

where $\mathcal{Q} := \{1, \dots, q\}$ and $\theta > 0$. Note that $y_i(\alpha^\top x_i + \beta) > 1$ implies that y_i and $f(x) = \alpha^\top x_i + \beta$ have the same sign, i.e., the classifier correctly groups the training data (x_i, y_i) . Since it might not be possible to perfectly classify the data into two sets, there might be cases in which $\xi_i > 1$ for some i . In those cases, the classifier miss-categorizes some data points. It is desired to keep the number of such entries as low as possible by adding the term $\theta \mathbf{1}^\top \xi$ to the cost function. The optimization problem in (1) may not admit a unique solution. To alleviate this issue, an alternative problem can be posed:

$$(\alpha^*, \beta^*, \xi^*) := \underset{\substack{\alpha \in \mathbb{R}^p, \beta \in \mathbb{R}, \xi \in \mathbb{R}^q}}{\operatorname{argmin}} \frac{1}{2} \alpha^\top \alpha + \frac{\rho}{2} (\beta^2 + \xi^\top \xi) + \theta \mathbf{1}^\top \xi, \quad (2a)$$

$$\text{s.t.} \quad y_i(\alpha^\top x_i + \beta) \geq 1 - \xi_i, \forall i \in \mathcal{Q}, \quad (2b)$$

$$\xi_i \geq 0, \forall i \in \mathcal{Q}, \quad (2c)$$

where $\rho > 0$. The uniqueness of the solution of (2) is guaranteed by the strict convexity of the cost and the convexity of the constraint set. Note that, by reducing ρ , the solution of the optimization problem in (2) can be made to arbitrarily closely approximate a solution of (1). Therefore, the proposed alteration is without loss of generality; however, it simplifies the derivation of subsequent results.

Proposition 2.1. Let $\aleph$ denote the set of solutions of (1). Then, $\lim_{\rho \searrow 0} (\alpha^*, \beta^*, \xi^*) \in \aleph$.

Proof: See Appendix A. $\square$

2.2 Independently and Identically Distributed Noise

Solving (2), or as a matter of fact (1), requires access to the training data, which can infringe on the privacy of individuals whose data is gathered for machine learning training. Therefore, the data owners are inclined to provide a noisy version of the data $\{(\bar{x}_i, y_i)\}_{i \in \mathcal{Q}} \subseteq \mathbb{R}^p \times \{-1, +1\}$ in which

$$\bar{x}_i = x_i + n_i, \quad (3)$$

where $n_i \in \mathbb{R}^p$ is a zero-mean random additive noise with probability density function $\gamma : \mathbb{R}^p \rightarrow \mathbb{R}_{\geq 0}$. Note that $n_i, \forall i \in \mathcal{Q}$, are assumed to be drawn from the same distribution. Therefore, they are independently and identically distributed (i.i.d.) random variables. This makes process of generating the noise computationally friendly for very large databases (i.e., when q is very large); however, as explained later, it also restricts the set of applicable noises. In the next subsection, the noise is generalized by removing the i.i.d. assumption.

Assumption 2.1. γ is twice continuously differentiable.

Under this assumption, a lower-bound for an adversary's estimation error of the private database (based on the noisy data) can be determined. The lower bound is independent of the actions of the adversary and is thus immune to unrealistic assumptions (e.g., that the adversary is rational in case of game-theoretic approaches to privacy analysis). In what follows, $\operatorname{supp}(\gamma)$ is defined to be $\{n \mid \gamma(n) > 0\}$. Further, for any continuously differentiable function $g(x)$, the notation $\partial g(x)/\partial x$ is reserved to denote a column vector containing the partial derivatives of the function¹.

1. Note that there are other conventions in which the gradient is a row vector; however, in this paper, the gradient is a column vector.

Further, for any twice continuously differentiable function $g(x)$, $\mathfrak{D}^2 g(x)$ is the Hessian matrix (i.e., a symmetric matrix containing all the second order partial derivative).

Proposition 2.2. For any unbiased estimate of x_i denoted by $\hat{x}_i$, it holds that

$$\mathbb{E}\{\|\Pi(x_i - \hat{x}_i)\|_2^2\} \geq \operatorname{Tr}(\Pi \mathcal{I}^{-1}) \geq 1/\operatorname{Tr}(\Pi^{-1} \mathcal{I}), \quad (4)$$

where $\Pi \in \mathbb{R}^{p \times p}$ is a positive definite matrix and $\mathcal{I}$ is the Fisher information matrix defined as

$$\mathcal{I} = \int_{n \in \operatorname{supp}(\gamma)} \gamma(n) \left[\frac{\partial \log(\gamma(n))}{\partial n} \right] \left[\frac{\partial \log(\gamma(n))}{\partial n} \right]^\top dn. \quad (5)$$

Proof: The first inequality in (4) follows from the Cramér-Rao bound [3, p.169]. The rest follows from the fact that $\operatorname{Tr}(\Pi \mathcal{I}^{-1}) = \operatorname{Tr}(\Pi^{1/2} \mathcal{I}^{-1} \Pi^{1/2}) \geq 1/\operatorname{Tr}(\Pi^{-1/2} \mathcal{I} \Pi^{-1/2}) = 1/\operatorname{Tr}(\Pi^{-1} \mathcal{I})$ with the inequality stemming from Item (11) in [32, p.44]. $\square$

Therefore, by minimizing $\operatorname{Tr}(\Pi^{-1} \mathcal{I})$, the privacy of the participants can be guaranteed. This is because, irrespective of the adversaries behaviour (computational resources, intelligence, etc.), it cannot estimate the original entries of the database with an estimation error smaller than $1/\operatorname{Tr}(\Pi^{-1} \mathcal{I})$. Note that the positive definite scaling matrix Π is used here because not all entries of x_i have the same scaling (e.g., some might be fractional numbers between zero and one while others are large integers).

The addition of the noise potentially changes the classifier. Therefore, a measure of quality must be established (as otherwise the most private decision is to use an additive noise with co-variance approaching infinity). The optimal linear support vector machine can be extracted from

$$(\bar{\alpha}^*, \bar{\beta}^*, \bar{\xi}^*) := \underset{\substack{\alpha \in \mathbb{R}^p, \beta \in \mathbb{R}, \xi \in \mathbb{R}^q}}{\operatorname{argmin}} \frac{1}{2} \alpha^\top \alpha + \frac{\rho}{2} (\beta^2 + \xi^\top \xi) + \theta \mathbf{1}^\top \xi, \quad (6a)$$

$$\text{s.t.} \quad y_i(\alpha^\top \bar{x}_i + \beta) \geq 1 - \xi_i, \forall i \in \mathcal{Q}, \quad (6b)$$

$$\xi_i \geq 0, \forall i \in \mathcal{Q}. \quad (6c)$$

Note that the difference between (2) and (6) is access to the noisy or original data. The solution of (6) is in fact a random variable. The quality of the classifier extracted from the noisy data can be guaranteed by ensuring that the following measure of quality is large:

$$\mathcal{D} := \mathbb{P} \left\{ \left\| \begin{bmatrix} \bar{\alpha}^* - \alpha^* \\ \bar{\beta}^* - \beta^* \\ \bar{\xi}^* - \xi^* \end{bmatrix} \right\|_2 \leq \varepsilon \right\}. \quad (7)$$

The following result relates $\mathcal{D}$ to the second moment of the distribution of the additive noise.

Proposition 2.3. There exist $c \geq 1$ and $\varepsilon_0 > 0$ such that $\mathcal{D} \geq 1 - (qc^2/\varepsilon^2) \operatorname{Tr}(V_{nn}), \forall \varepsilon \in (0, \varepsilon_0)$ with $V_{nn} := \mathbb{E}\{nn^\top\}$.

Proof: See Appendix B. $\square$

Remark 2.1. The proof of Proposition 2.3, stemming from the application of [33, Theorem 4.4], requires that the cost function of the classification optimization problem is positive definite. This is clearly the case for all choices of $\rho > 0$. We can still select ρ to be very small to reduce its impact on the classifier. This follows from Proposition 2.1 proving that the classifier with $\rho > 0$ is close to the classifier with $\rho = 0$ if ρ is small enough. Therefore, restricting the classification problem to $\rho > 0$ will not

Fig. 1. The gray region illustrates the set of δ and ϵ that meet the condition of Theorem 2.5 with $\sqrt{\lambda}\|\Pi^{1/2}\|_{\infty,2} = 0.1$.

Fig. 2. The lower bound on $\mathbb{E}\{\|\Pi(x_i - \hat{x}_i)\|_2^2\}$ for the case where $\text{Tr}(\Pi^{1/2})\|\Pi^{1/4}\|_{\infty,2}^2 = 1$.

significantly or adversely influence the classifier if ρ is small. Regarding privacy, we should note that the addition of Gaussian noise will always protect the privacy of the data; see Theorem 2.5 for guarantees in the sense of differential privacy. However, by selecting $\rho = 0$, we cannot analyze the utility of privacy-preserving dataset. Therefore, we recommend selecting an infinitesimally small ρ to also provide utility guarantees.

Proposition 2.3 shows that, if the co-variance of the additive noise is kept small, $\mathcal{D}$ remains large. This motivates the following formulation for finding the optimal privacy-preserving noise:

$$\mathbf{P}(\lambda) : \min_{\gamma} \text{Tr}(\Pi^{-1}\mathcal{I}) + \lambda \text{Tr}(V_{nn}). \quad (8)$$

We can also cast the problem of finding the optimal privacy-preserving noise as $\max_{\gamma} [1/\text{Tr}(\Pi^{-1}\mathcal{I}) - \bar{\lambda} \text{Tr}(V_{nn})]$. However, these two problems are equivalent in the sense that, for any $\bar{\lambda} > 0$, there exists $\lambda > 0$ such that they admit the same solution. Now, one of the main results of the paper regarding the optimal privacy-preserving policy can be presented.

Theorem 2.4. The solution of $\mathbf{P}(\lambda)$ is given by

$$\gamma^*(n) = \frac{1}{\sqrt{\det(2\pi\Pi^{-1/2}/\sqrt{\lambda})}} \exp\left(-\frac{\sqrt{\lambda}}{2} n^\top \Pi^{1/2} n\right). \quad (9)$$

Proof: See Appendix C. $\square$

The optimal additive noise in Theorem 2.4, which is a Gaussian noise, meets the condition of (ϵ, δ) -local differential privacy, presented below.

Definition 2.1. The reporting policy in (3) is (ϵ, δ) -locally differential private if $\mathbb{P}\{x_i + n_i \in \mathcal{X}\} \leq \exp(\epsilon) \mathbb{P}\{x'_i + n_i \in \mathcal{X}\} + \delta$, where $\mathcal{X} \subseteq \mathbb{R}^p$ is a Borel-measurable set and x_i, x'_i are possible entries in the dataset.

Theorem 2.5. The optimal additive noise in Theorem 2.4 is (ϵ, δ) -locally differential private for all ϵ, δ satisfying $\lambda^{1/4}\|\Pi^{1/4}\|_{\infty,2}(1 + \sqrt{2\ln(1/\delta)}) \leq \epsilon$, where $\|A\|_{\infty,2} = \max_{x \neq 0} \|Ax\|_2 / \|x\|_\infty$ is the induced norm of A .

Proof: The proof follows from combining the result of Theorem 2.4 and the differential privacy property of Gaussian additive noises in [34]. $\square$

Theorem 2.5 proves that the optimal noise in Theorem 2.4 is (ϵ, δ) -locally differential private. This is an important result as it shows that the new measure of privacy inherits the strong guarantees of differential privacy. The gray region in Fig. 1 illustrates the set of δ and ϵ that meet the condition of Theorem 2.5 with $\lambda^{1/4}\|\Pi^{1/4}\|_{\infty,2} = 0.1$. For (ϵ, δ) -locally differential private Gaussian mechanism, with the aid of Proposition 2.8, we can see that $\mathbb{E}\{\|\Pi(x_i - \hat{x}_i)\|_2^2\} \geq \text{Tr}(\Pi^{1/2})\|\Pi^{1/4}\|_{\infty,2}^2(1 + \sqrt{2\ln(1/\delta)})^2/\epsilon^2 = \mathcal{O}(\ln^2(\delta^{-1})\epsilon^{-2})$. This inequality provides an operative iteration to the parameters of the differential privacy, i.e., we can bound on the error of estimating entries of the private database $\mathbb{E}\{\|\Pi(x_i - \hat{x}_i)\|_2^2\}$ for a specific selections of (ϵ, δ) . Fig. 2 illustrates the lower bound on $\mathbb{E}\{\|\Pi(x_i - \hat{x}_i)\|_2^2\}$ for the case where $\text{Tr}(\Pi^{1/2})\|\Pi^{1/4}\|_{\infty,2}^2 = 1$. As expected, the adversary's estimation error grows by decreasing ϵ and δ .

In general, it might be desired to use noises whose support set is constrained to be a subset of $\mathcal{N} \subseteq \mathbb{R}^p$. This is because entries of database must be within some range to make sense (e.g., age or height cannot be negative). Note that, in this case, one cannot necessarily ensure that the additive noise has a zero mean (due to the arbitrary nature of the set $\mathcal{N}$). In this case, the problem can be cast as

$$\bar{\mathbf{P}}(\lambda) : \min_{\gamma} \text{Tr}(\Pi^{-1}\mathcal{I}) + \lambda \text{Tr}(V_{nn}), \quad (10a)$$

$$\text{s.t. } \text{supp}(\gamma) \subseteq \mathcal{N}. \quad (10b)$$

Theorem 2.6. The solution of $\bar{\mathbf{P}}(\lambda)$ is $\gamma^*(n) = \sqrt{\det(\Pi)}u(n)^2$, where

$$\nabla^2 u(n) + (\mu - (\lambda/4)n^\top \Pi^{-1}n)u(n) = 0, \quad (11a)$$

$$u(n) = 0, \forall n \in \partial\bar{\mathcal{N}}, \quad u(n) > 0, \forall n \in \text{int}(\bar{\mathcal{N}}), \quad (11b)$$

$$\int_{n \in \mathcal{N}} u(n)^2 dn = 1, \quad (11c)$$

for some $\mu > 0$ with $\bar{\mathcal{N}} := \{\Pi^{1/2}n, \forall n \in \mathcal{N}\}$.

Proof: The proof follows from the same line of reasoning as in the proof of Theorem 2.4 with the change of variable $\gamma(n) = \sqrt{\det(\Pi)}u(\Pi^{1/2}n)^2$. $\square$

Note that the partial differential equation (11a) is a stationary (time-independent) Schrödinger equation. Therefore, it admits a unique solution for each μ [35]. Here, μ is the dual variable corresponding to the equality constraint in (11c) and can be computed using the dual ascent.

Corollary 2.7. Let $\Pi = \text{diag}(\vartheta_1, \dots, \vartheta_p)$, $\mathcal{N} = \prod_i [\underline{n}_i, \bar{n}_i]$. The solution of $\bar{\mathbf{P}}(\lambda)$ is $\gamma^*(n) = \sqrt{\prod_i \vartheta_i} (\prod_i u_i(n_i))^2$, where, for all $1 \leq i \leq p$,

$$d^2 u_i(n_i)/dn_i^2 + (\mu_i - (\lambda/(4\vartheta_i))n_i^2)u_i(n_i) = 0, \quad (12a)$$

$$u_i(n_i) = 0, n_i \in \{\underline{n}_i, \bar{n}_i\}, u_i(n_i) > 0, \forall n_i \in (\underline{n}_i, \bar{n}_i), \quad (12b)$$

$$\int_{\underline{n}_i}^{\bar{n}_i} u_i(n_i)^2 dn_i = 1, \quad (12c)$$

for some $\mu_i > 0$.

Proof: The proof follows from the method of separation of variables [35]. $\square$

The differential equations in Corollary 2.7 are known as the Airy differential equations and their solution can be characterized using the Airy functions [36]. So far, the additive noise are somewhat conservative as they are i.i.d. This assumption is removed in the next subsection.

2.3 Correlated Noise

In this section, it is assumed that the additive noise for various data points is correlated. Therefore, $[\bar{x}_1^\top \cdots \bar{x}_q^\top]^\top = [x_1^\top \cdots x_q^\top]^\top + w$, where $w \in \mathbb{R}^{qp}$ is a zero-mean random additive noise with probability density function $\tilde{\gamma} : \mathbb{R}^{qp} \rightarrow \mathbb{R}_{\geq 0}$. Similarly, the following standing assumption holds.

Assumption 2.2. $\tilde{\gamma}(w)$ is twice continuously differentiable.

Again, the Cramér-Rao bound can be used to establish a lower bound on the ability of the adversary for inferring the private database. In this paper, $\otimes$ and $\circ$, respectively, denote the Kronecker and the Hadamard products.

Proposition 2.8. For any unbiased estimate of x_i denoted by $\hat{x}_i$, it holds that

$$\min_{1 \leq i \leq q} \mathbb{E}\{\|\Pi(x_i - \hat{x}_i)\|_2^2\} \geq \min_{1 \leq i \leq q} \text{Tr}(\Pi[(e_i^\top \otimes I_p) \tilde{\mathcal{I}}(e_i \otimes I_p)]^{-1}) \geq 1/\text{Tr}((I_q \otimes \Pi^{-1})^{-1} \tilde{\mathcal{I}}), \quad (13)$$

where $\tilde{\mathcal{I}}$ is the Fisher information matrix defined as

$$\tilde{\mathcal{I}} = \int_{w \in \mathbb{R}^{qp}} \tilde{\gamma}(w) \left[\frac{\partial \log(\tilde{\gamma}(w))}{\partial w} \right] \left[\frac{\partial \log(\tilde{\gamma}(w))}{\partial w} \right]^\top dw. \quad (14)$$

Proof: See Appendix D. $\square$

The Karush–Kuhn–Tucker (KKT) conditions for (2) are

$$\alpha^* = \sum_{i=1}^q (\omega^*)_i y_i x_i, \quad \beta^* = \sum_{i=1}^q (\omega^*)_i y_i / \rho, \quad (15a)$$

$$(\xi^*)_i = \max((\omega^*)_i + (\varsigma^*)_i - 1) / \rho, \quad \forall i \in \mathcal{Q}, \quad (15b)$$

$$0 = (\omega^*)_i (y_i (\alpha^{*\top} x_i + \beta^*) - 1 + (\xi^*)_i), \quad \forall i \in \mathcal{Q}, \quad (15c)$$

$$0 = (\xi^*)_i (\varsigma^*)_i, \quad \forall i \in \mathcal{Q}, \quad (15d)$$

where $\omega^*, \varsigma^* \in \mathbb{R}_{\geq 0}^q$ are Lagrange multipliers. The set of equations in (15) have at least one solution due to the strict convexity of the cost and the convexity of the constraint set.

Proposition 2.9. $\mathbb{P}\{\bar{\alpha}^* = \alpha^*, \bar{\beta}^* = \beta^*, \bar{\xi}^* = \xi^*\} = \mathbb{P}\{\Omega w = 0\}$ with $\Omega := [((\lambda^* \circ y)^\top \otimes I_p)^\top (I_q \otimes \alpha^{*\top})^\top]^\top$.

Proof: See Appendix E. $\square$

This motivates the following problem formulation to find most privacy-preserving noise for which the classifiers extracted from the original data and the noisy data are identical:

$$\tilde{\mathbf{P}} : \min_{\tilde{\gamma}} \quad \text{Tr}((I_q \otimes \Pi^{-1}) \tilde{\mathcal{I}}), \quad (16a)$$

$$\text{s.t.} \quad V_{ww} \leq mI, \quad (16b)$$

$$\mathbb{P}\{\Omega w = 0\} = 1, \quad (16c)$$

where V_{ww} is the co-variance of the additive noise, $m > 0$, and I is the identity matrix. The constraint $V_{ww} \leq mI$ is added to ensure that a noise with finite moment (thus realizable) is obtained. This is without any practical consequences as m can be selected arbitrarily large.

Theorem 2.10. Let Ψ be a matrix whose columns form an orthonormal basis for the null space of Ω . The solution of $\tilde{\mathbf{P}}$ is $\tilde{\gamma}(w) = \mathbb{1}_{w \in \text{im}(\Psi)} \hat{\gamma}(\Psi^\top w)$, where $\hat{\gamma}(\bar{w}) = ((2\pi m)^{\dim(\text{im}(\Psi))})^{-1/2} \exp(-\bar{w}^\top \bar{w} / (2m))$.

Proof: See Appendix F. $\square$

Theorem 2.10 provides the optimal noise for corrupting the dataset without any adverse effect on the utility of the linear support vector machine (i.e., no utility degradation).

2.4 Optimization-Based Machine Learning

Here, we generalize the problem formulation of the previous subsections to more general optimization-based machine learning algorithms. Consider a set of training data $\{(x_i, y_i)\}_{i=1}^q \subseteq \mathbb{R}^{p_x} \times \mathbb{R}^{p_y}$. A general optimization-based machine learning problem can be cast as

$$\argmin_{\varphi \in \mathbb{R}^{p_\varphi}} \ell(\varphi; \{(x_i, y_i)\}_{i=1}^q), \quad (17)$$

where ℓ is a fitness function. For instance, in nonlinear soft margin support vector machine $\ell(\varphi; \{(x_i, y_i)\}_{i=1}^q) := 0.5\varphi^\top \text{diag}(I, 0)\varphi + \theta \sum_{i=1}^q \max(1 - y_i [\kappa(x_i)^\top y_i] \varphi, 0)$, where $\kappa(x_i)$ denotes the transformed data points using nonlinear mapping $\kappa(\cdot)$. Another example is the artificial neural networks with $\ell(\varphi; \{(x_i, y_i)\}_{i=1}^q) := \frac{1}{q} \sum_{i=1}^q \|y_i - \text{ANN}(x_i, \varphi)\|_2^2$, where $\text{ANN}(x, \varphi)$ denotes the output of the artificial neural network.

Proposition 2.11. Assume ℓ is twice continuously differentiable. There exist $c \geq 1$ and $\varepsilon_0 > 0$ such that $\mathcal{D} \geq 1 - (q^2 c^2 / \varepsilon^2) \text{Tr}(V_{nn})$, $\forall \varepsilon \in (0, \varepsilon_0)$ with $V_{nn} := \mathbb{E}\{nn^\top\}$.

Proof: See Appendix G. $\square$

The result of Proposition 2.11 shows that the problem formulation (8) is still relevant for optimization-based machine learning. Therefore, Theorem 2.4 still provides the optimal additive noise when releasing datasets for more general machine learning algorithms.

Remark 2.2 (Differentiability of Fitness Function). Proposition 2.11 requires the fitness function to be twice differentiable. This is not the case generally. For instance, in deep learning, a popular non-differentiable activation function is rectified linear unit (ReLU). However, other activation functions, such as Sigmoid, softplus [37], GELU [38], SoftExponential [39], and SGNL [40], can be used to ensure twice-differentiability of fitness function. Again, note that the addition of a Gaussian noise will always protect the privacy of the data; see Theorem 2.5. However, selecting a non-twice-differentiable loss function restricts our ability in analyzing the utility.

3 NUMERICAL EXAMPLE

3.1 Illustrative Example

First, consider a simple example for the sake of the illustration of the results. A random dataset is generated with 50 entries drawn from a two-dimensional zero-mean Gaussian distribution with unit co-variance (corresponding to $y = +1$) and another 50 entries drawn from a two-dimensional Gaussian distribution with unit co-variance and mean $[0 \ 5]^\top$ (corresponding to $y = -1$). This dataset can be easily classified (visually), which is beneficial for

Fig. 3. The blue dots illustrate the original data $(x_i)_{i \in \mathcal{Q}}$ and the blue line is the classifier extracted from (2) using the original data. The red, yellow, and purple dots show the noisy data $(\tilde{x}_i)_{i \in \mathcal{Q}}$ with the optimal noise in Theorem 2.4 for $\lambda = 10^{-4}$, $\lambda = 10^{-2}$, and $\lambda = 10^0$, respectively. The solid lines illustrate the classifiers extracted from (6) using the noisy data of the corresponding color.

Fig. 4. The blue dots illustrate the original data $(x_i)_{i \in \mathcal{Q}}$ and the blue line is the classifier extracted from (2) using the original data. The red dots illustrate the noisy data $(\tilde{x}_i)_{i \in \mathcal{Q}}$ with the optimal noise in Theorem 2.10. The classifier extracted from (6) using the noisy data remains the same.

the illustration of the effect of the noise. In what follows $\rho = 10^{-2}$ so that the solutions of (1) and (2) are sufficiently close. Further, $\theta = 1$ and $\Pi = I$.

First, the case with i.i.d. noise is considered. This is illustrated in Fig. 3. The blue dots show the original data $(x_i)_{i \in \mathcal{Q}}$ and the blue line is the classifier extracted from (2) using the original data. The red, yellow, and purple dots show the noisy data $(\tilde{x}_i)_{i \in \mathcal{Q}}$ with the optimal noise in Theorem 2.4 for $\lambda = 10^{-4}$, $\lambda = 10^{-2}$, and $\lambda = 10^0$, respectively. The solid lines illustrate the classifiers extracted from (6) using the noisy data of the corresponding color. As expected, the classifiers constructed using the noisy data approaches the classifier extracted from the original data as λ increases. Note that $\mathbb{E}\{\|x_i - \hat{x}_i\|_2^2\}$ is lower bounded by 200, 20, and 2 for $\lambda = 10^{-4}$, $\lambda = 10^{-2}$, and $\lambda = 10^0$, respectively.

Now, the case with correlated additive noise is considered. Fig. 4 shows the results in this case. The blue dots illustrate the original data $(x_i)_{i \in \mathcal{Q}}$ and the blue line is the classifier extracted from (2) using the original data. The red dots illustrate the noisy data $(\tilde{x}_i)_{i \in \mathcal{Q}}$ with the optimal noise in Theorem 2.10 for $m = 100$. Note that using this noise, the extracted classifier from (6) is the same as the classifier extracted from the original data. In this case, $\mathbb{E}\{\|x_i - \hat{x}_i\|\} \geq 80.5$. This error can be worsened by increasing m . The specific structure of the noise spreads the data along the classifier (and not towards it).

3.2 Experimental Results

Here, the applicability of the presented methodology on practical datasets is demonstrated with two machine learn-

Breast Cancer Wisconsin (Diagnostic) dataset

Fig. 5. (top) The blue line illustrates the classification success rate of the classifier generated from the original data. The red curve depicts the classification success rate using the classifier extracted from (6) using the noisy data corrupted with the optimal noise in Theorem 2.4 versus λ . The yellow curve shows the classification success rate based on the dataset obfuscated using Laplace noise to achieve local differential privacy (with the same privacy guarantee). (bottom) The privacy guarantee $\min_i \mathbb{E}\{\|x_i - \hat{x}_i\|_2^2\}$ versus λ .

ing algorithms. Particularly, we use the Breast Cancer Wisconsin (Diagnostic) dataset and the Adult dataset with linear support vector machines and the Lending Club dataset with linear regression.

3.2.1 Medical Dataset

The Breast Cancer Wisconsin (Diagnostic) Data Set is openly available in [41] containing features computed from a digitized image of a fine needle aspirate of breast mass. The features include parameters, such as radius, texture, smoothness, and symmetry, of cell nucleus. For this dataset, we are interested in training a support vector machine for cancer diagnosis. We systematically corrupt the entries of the dataset to avoid the unintentional private data leakages.

First, the case of i.i.d. additive noise is demonstrated. The red curve in Fig. 5 (top) shows $\mathbb{E}\{\sum_{i=1}^q \mathbb{1}_{y_i(\tilde{\alpha}^* \top x_i + \tilde{\beta}^*) > 0}\}/n$,

Fig. 6. The blue line illustrates the classification success rate of the classifier generated from the original data. The red curve depicts the classification success rate using the classifier extracted from (6) using the noisy data corrupted with the optimal noise in Theorem 2.4 versus λ . The yellow curve shows the classification success rate based on the dataset obfuscated using Laplace noise to achieve local differential privacy (with the same privacy guarantee).

empirically computed by averaging 100 simulations, versus λ . In this case, the classifier is extracted from the noisy data with the optimal noise in Theorem 2.4. Note that $\mathbb{E}\{\sum_{i=1}^q \mathbb{1}_{y_i(\bar{\alpha}^* \top x_i + \bar{\beta}^*) > 0}\}/n$ is the classification success rate of the original data using the classifier generated from the noisy data. The blue line illustrates the success rate of the classifier extracted from the original data. As λ increases, the success rate of the classifier from the noisy data approaches the classification rate of the classifier based on the original data. The yellow curve shows the classification rate for the case where the dataset is corrupted using Laplace noise to achieve local differential privacy. The level of the differential privacy is set so that the privacy guarantee $\min_i \mathbb{E}\{\|x_i - \hat{x}_i\|\}$ remains the same for both noises. This is to ensure that methods are compared in a fair manner. When using the Laplace mechanism, the success rate of the classifier is always below the optimal noise in Theorem 2.4. The red curve in Fig. 5 (bottom) illustrates the level of privacy $\min_i \mathbb{E}\{\|x_i - \hat{x}_i(\bar{x})\|\}$ versus λ . Fig. 5 captures the trade-off between privacy and utility.

The correlated additive noise can also be utilized in this example. Upon using the noisy data $(\bar{x}_i)_{i \in Q}$ with the optimal noise in Theorem 2.10, $\min_{i \in Q} \mathbb{E}\{\|x_i - \hat{x}_i\|\} \geq 28.4m$. Therefore, for $m = 100$, $\min_{i \in Q} \mathbb{E}\{\|x_i - \hat{x}_i\|\} \geq 2840$. Note that, in this case, the classifier remains exactly the same even with such a high privacy guarantee.

3.2.2 Adult Dataset

This dataset contains nearly 49,000 records from the 1994 Census database [41]. The records contain features, such as age, education, and work. The aim is to train a classifier that identifies individuals earning more than \$50,000. Here, we obfuscate the dataset in order to eliminate private information leakages about the individuals in the dataset.

We illustrate the effect of the optimal privacy-preserving noise in Theorem 2.4 in Fig. 6. The red curve in Fig. 5 shows $\mathbb{E}\{\sum_{i=1}^q \mathbb{1}_{y_i(\bar{\alpha}^* \top x_i + \bar{\beta}^*) > 0}\}/n$, empirically computed by averaging 100 simulations, versus λ . The blue line illustrates the success rate of the classifier extracted from the original data. The success rate of the classifier from the noisy data approaches the classification rate of the classifier based on the original data as λ increases. The yellow curve shows the

Fig. 7. The blue line illustrates the fitness of regression model generated from the original data. The red curve depicts the fitness of the regression model extracted from (17) using the noisy data corrupted with the optimal noise in Theorem 2.4 versus λ . The yellow curve shows the fitness of regression model based on the dataset obfuscated using Laplace noise to achieve local differential privacy (with the same privacy guarantee).

classification success rate for the case where the dataset is corrupted using Laplace noise to achieve local differential privacy. Again, the level of the differential privacy is set so that the privacy guarantee, $\min_i \mathbb{E}\{\|x_i - \hat{x}_i\|\}$, remains the same for both noises. Evidently, the optimal privacy-preserving noise in Theorem 2.4 outperforms the Laplace noise in terms of the utility of the trained support vector machine model.

3.2.3 Finance Dataset

This dataset contains information about approximately 890,000 loans made on the Lending Club, which is a peer-to-peer lending platform [42]. The dataset contains information about the loans, such as total amount, information about the borrower, such as age and credit rating, and the interest rates of the loans. We encode categorical data, such as state of residence and loan grade, with integer numbers. We also remove unique identifier attributes, such as member id, and unrelated attributes, such as the URL for the listing data. We are interested in training a linear regression $\varphi^\top x_i$ to estimate the loan interest rate y_i . Training the regression model can be cast as (17) with loss function $\ell(\varphi; \{(x_i, y_i)\}_{i=1}^q) := \sigma \varphi^\top \varphi + \sum_{i=1}^q (y_i - \varphi^\top x_i)^2 / q$ for $\sigma > 0$. In what follows, we set $\sigma = 10^{-5}$.

Fig. 7 shows the expected loss $\mathbb{E}\{\ell(\varphi^*; \{(x_i, y_i)\}_{i=1}^q)\}$, empirically computed by averaging 100 simulations, as a function of λ . The red curve captures the fitness of the regression model generated from the noisy data with the optimal noise in Theorem 2.4. The blue line illustrates the fitness of the regression model extracted from the original data without privacy preserving noise. As λ , or the emphasis on utility, increases the fitness of the regression model from the noisy data approaches the fitness of the regression model computed based on the original data. The yellow curve shows the fitness rate for the case where the dataset is corrupted using Laplace noise to achieve local differential privacy. The fitness of the regression model with the optimal noise in Theorem 2.4 is always better than the one with Laplace noise.

4 CONCLUSIONS AND FUTURE WORK

A private dataset was corrupted by an additive noise for privacy protection in such a way that the classifier extracted

from it remains close to the classifier computed based on the original dataset. The applicability of the methodology was demonstrated on practical datasets in medicine, demography, and finance. Future work can focus on implementation of constrained additive noises for practical datasets.

REFERENCES

- [1] L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.
- [2] A. Machanavajhala, J. Gehrke, D. Kifer, and M. Venkatasubramanian, "l-diversity: privacy beyond k-anonymity," in *22nd International Conference on Data Engineering (ICDE'06)*, pp. 24–24, 2006.
- [3] J. Shao, *Mathematical Statistics*. Springer Texts in Statistics, Springer-Verlag New York, 2003.
- [4] X. Xiao and Y. Tao, "M-invariance: towards privacy preserving re-publication of dynamic datasets," in *Proceedings of the 2007 ACM SIGMOD international conference on Management of data*, pp. 689–700, ACM, 2007.
- [5] L. Zou, L. Chen, and M. T. Özsu, "K-automorphism: A general framework for privacy preserving network publication," *Proceedings of the VLDB Endowment*, vol. 2, no. 1, pp. 946–957, 2009.
- [6] A. Machanavajhala, J. Gehrke, D. Kifer, and M. Venkatasubramanian, "l-diversity: Privacy beyond k-anonymity," in *22nd International Conference on Data Engineering (ICDE'06)*, pp. 24–24, 2006.
- [7] F. Farokhi, "Development and analysis of deterministic privacy-preserving policies using non-stochastic information theory," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 10, pp. 2567–2576, 2019.
- [8] N. Li, T. Li, and S. Venkatasubramanian, "t-closeness: Privacy beyond k-anonymity and l-diversity," in *2007 IEEE 23rd International Conference on Data Engineering*, pp. 106–115, IEEE, 2007.
- [9] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [10] B. I. Rubinstein, P. L. Bartlett, L. Huang, and N. Taft, "Learning in a large function space: Privacy-preserving mechanisms for SVM learning," *Journal of Privacy and Confidentiality*, vol. 4, no. 1, 2012.
- [11] H. Li, L. Xiong, L. Ohno-Machado, and X. Jiang, "Privacy preserving rbf kernel support vector machine," *BioMed research international*, vol. 2014, 2014.
- [12] K. Chaudhuri and C. Monteleoni, "Privacy-preserving logistic regression," in *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, pp. 289–296, 2009.
- [13] J. Zhang, Z. Zhang, X. Xiao, Y. Yang, and M. Winslett, "Functional mechanism: regression analysis under differential privacy," *Proceedings of the VLDB Endowment*, vol. 5, no. 11, pp. 1364–1375, 2012.
- [14] 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, 2016.
- [15] R. Shokri and V. Shmatikov, "Privacy-preserving deep learning," in *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pp. 1310–1321, ACM, 2015.
- [16] J. Tang, A. Korolova, X. Bai, X. Wang, and X. Wang, "Privacy loss in Apple's implementation of differential privacy on macOS 10.12," in *3rd Workshop on the Theory and Practice of Differential Privacy at ACM Conference on Computer and Communications Security (CCS)*, 2017.
- [17] A. D. Wyner, "The wire-tap channel," *Bell System Technical Journal*, vol. 54, no. 8, pp. 1355–1387, 1975.
- [18] H. Yamamoto, "A source coding problem for sources with additional outputs to keep secret from the receiver or wiretappers," *IEEE Transactions on Information Theory*, vol. 29, no. 6, pp. 918–923, 1983.
- [19] F. Farokhi, H. Sandberg, I. Shames, and M. Cantoni, "Quadratic Gaussian privacy games," in *2015 54th IEEE Conference on Decision and Control (CDC)*, pp. 4505–4510, IEEE, 2015.
- [20] Y. Liang, H. V. Poor, and S. Shamai, "Information theoretic security," *Foundations and Trends® in Communications and Information Theory*, vol. 5, no. 4–5, pp. 355–580, 2009.
- [21] F. Farokhi and H. Sandberg, "Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4726–4734, 2017.
- [22] S. Asodeh, F. Alajaji, and T. Linder, "Notes on information-theoretic privacy," in *52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1272–1278, IEEE, 2014.
- [23] E. Akyol, C. Langbort, and T. Başar, "Information-theoretic approach to strategic communication as a hierarchical game," *Proceedings of the IEEE*, vol. 105, no. 2, pp. 205–218, 2016.
- [24] H. Anderson, "Efficiency versus protection in a general randomized response model," *Scandinavian Journal of Statistics*, pp. 11–19, 1977.
- [25] F. Farokhi and H. Sandberg, "Ensuring privacy with constrained additive noise by minimizing Fisher information," *Automatica*, vol. 99, pp. 275–288, 2019.
- [26] F. Farokhi and H. Sandberg, "Fisher information privacy with application to smart meter privacy using HVAC units," in *Privacy in Dynamical Systems (F. Farokhi, ed.)*, pp. 3–17, Singapore: Springer, 2020.
- [27] K.-P. Lin and M.-S. Chen, "Releasing the svm classifier with privacy-preservation," in *2008 Eighth IEEE International Conference on Data Mining*, pp. 899–904, IEEE, 2008.
- [28] J. He, B. Liu, D. Kong, X. Bao, N. Wang, H. Jin, and G. Kesidis, "Puppies: Transformation-supported personalized privacy preserving partial image sharing," in *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pp. 359–370, IEEE, 2016.
- [29] A. M. Olteanu, K. Huguenin, I. I. Dacosta Petrocelli, and J.-P. Hubaux, "Consensual and privacy-preserving sharing of multi-subject and interdependent data," in *Network and Distributed System Security Symposium (NDSS)*, 2018.
- [30] A.-S. Charest, "How can we analyze differentially-private synthetic datasets?," *Journal of Privacy and Confidentiality*, vol. 2, no. 2, 2011.
- [31] J. M. Abowd and L. Vilhuber, "How protective are synthetic data?," in *International Conference on Privacy in Statistical Databases*, pp. 239–246, Springer, 2008.
- [32] H. Lütkepohl, *Handbook of Matrices*. John Wiley & Sons, 1997.
- [33] J. W. Daniel, "Stability of the solution of definite quadratic programs," *Mathematical Programming*, vol. 5, no. 1, pp. 41–53, 1973.
- [34] A. Nikolov, K. Talwar, and L. Zhang, "The geometry of differential privacy: the sparse and approximate cases," in *Proceedings of the 45th annual ACM Symposium on Theory of Computing*, pp. 351–360, ACM, 2013.
- [35] A. Mohamed and H. Atia, "Separation of the Schrödinger operator with an operator potential in the Hilbert spaces," *Applicable Analysis*, vol. 84, no. 1, pp. 103–110, 2005.
- [36] F. W. J. Olver, "Airy functions," in *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables (M. Abramowitz and I. A. Stegun, eds.)*, pp. 446–452, New York: Dover, 9 ed., 1972.
- [37] C. Dugas, Y. Bengio, F. Bélisle, C. Nadeau, and R. Garcia, "Incorporating second-order functional knowledge for better option pricing," in *Advances in neural information processing systems*, pp. 472–478, 2001.
- [38] D. Hendrycks and K. Gimpel, "Gaussian error linear units (GELUs)," *arXiv preprint arXiv:1606.08415*, 2016.
- [39] L. B. Godfrey and M. S. Gashler, "A continuum among logarithmic, linear, and exponential functions, and its potential to improve generalization in neural networks," in *2015 7th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K)*, vol. 1, pp. 481–486, IEEE, 2015.
- [40] A. Wuraola and N. Patel, "SQLN: A new computationally efficient activation function," in *2018 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–7, IEEE, 2018.
- [41] D. Dheeru and E. Karra Taniskidou, "UCI machine learning repository," 2017. University of California, Irvine, <http://archive.ics.uci.edu/ml>.
- [42] W. Kan, "Lending club loan data: Analyze lending club's issued loans." <https://www.kaggle.com/wendykan/lending-club-loan-data>.
- [43] A. L. Dontchev and R. T. Rockafellar, *Implicit Functions and Solution Mappings: A View from Variational Analysis*. Springer Series in Operations Research and Financial Engineering, Springer New York, 2014.
- [44] P. Bremaud, *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Texts in Applied Mathematics, Springer NY, 2001.
- [45] C. H. Edwards, *Advanced Calculus of Several Variables*. Academic Press, 1973.

APPENDIX A

PROOF OF PROPOSITION 2.1

Define the set valued mapping $S_{\text{opt}} : \mathbb{R} \rightrightarrows \mathbb{R}^p \times \mathbb{R} \times \mathbb{R}^q$ such that

$$S_{\text{opt}}(p) := \left\{ (\alpha', \beta', \xi') \in S_{\text{feas}} \mid \forall (\alpha, \beta, \xi) \in S_{\text{feas}} : \right. \\ \left. \frac{1}{2} \alpha'^\top \alpha' + \frac{\rho}{2} (\beta'^2 + \xi'^\top \xi') + \theta \mathbf{1}^\top \xi' \right. \\ \left. \leq \frac{1}{2} \alpha^\top \alpha + \frac{\rho}{2} (\beta^2 + \xi^\top \xi) + \theta \mathbf{1}^\top \xi \right\},$$

where

$$S_{\text{feas}} := \left\{ (\alpha, \beta, \xi) \in \mathbb{R}^p \times \mathbb{R} \times \mathbb{R}^q \mid y_i(\alpha^\top x_i + \beta) \geq 1 - \xi_i, \right. \\ \left. \forall i \in \mathcal{Q} \wedge \xi_i \geq 0, \forall i \in \mathcal{Q} \right\}.$$

By definition, $S_{\text{opt}}(p)$ is the set of the solutions of (2) for all p (including $p = 0$). Therefore, for all $p > 0$, $S_{\text{opt}}(p) = \{(\alpha^*, \beta^*, \xi^*)\}$ and $S_{\text{opt}}(0) = \emptyset$. From Theorem 3B.5 in [43], it can be deduced that $\limsup_{p \rightarrow 0} S_{\text{opt}}(p) \subseteq S_{\text{opt}}(0) = \emptyset$. The rest follows from the fact that $\lim_{\rho \searrow 0} (\alpha^*, \beta^*, \xi^*) = \lim_{\rho \searrow 0} S_{\text{opt}}(p) \subseteq \limsup_{\rho \rightarrow 0} S_{\text{opt}}(p) \subseteq \emptyset$ where the first equality follows from the fact that $S_{\text{opt}}(p)$ is a singleton for all $p > 0$.

APPENDIX B

PROOF OF PROPOSITION 2.3

First, it is established that there exist $c \geq 1$ and $\epsilon_0 > 0$ such that

$$\left\| \begin{bmatrix} y_1(x_1 - \bar{x}_1)^\top \\ \vdots \\ y_q(x_q - \bar{x}_q)^\top \end{bmatrix} \right\|_2 \leq \epsilon < \epsilon_0 \implies \left\| \begin{bmatrix} \bar{\alpha}^* - \alpha^* \\ \bar{\beta}^* - \beta^* \\ \bar{\xi}^* - \xi^* \end{bmatrix} \right\|_2 \leq c\epsilon.$$

The proof of this claim follows from [33, Theorem 4.4] and reorganizing the inequality constraints in (2) and (6) to be of the form of

$$-\begin{bmatrix} y_1 x_1^\top & y_1 \\ \vdots & \vdots \\ y_q x_q^\top & y_q \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \leq -\mathbf{1}, \quad -\begin{bmatrix} y_1 \bar{x}_1^\top & y_1 \\ \vdots & \vdots \\ y_q \bar{x}_q^\top & y_q \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \leq -\mathbf{1}.$$

Note that

$$\left\| \begin{bmatrix} y_1(x_1 - \bar{x}_1)^\top \\ \vdots \\ y_q(x_q - \bar{x}_q)^\top \end{bmatrix} \right\|_2 = \max_{\|p\|_2=1} \left\| \begin{bmatrix} y_1(x_1 - \bar{x}_1)^\top \\ \vdots \\ y_q(x_q - \bar{x}_q)^\top \end{bmatrix} p \right\|_2 \\ = \max_{\|p\|_2=1} \left\| \begin{bmatrix} y_1 n_1^\top \\ \vdots \\ y_q n_q^\top \end{bmatrix} p \right\|_2 \\ \leq \sqrt{\max_{\|p\|=1} \sum_{i=1}^q (y_i n_i^\top p)^2} \\ = \sqrt{\max_{\|p\|=1} \sum_{i=1}^q (n_i^\top p)^2} \\ \leq \sqrt{\max_{\|p\|=1} \sum_{i=1}^q \|n_i\|_2^2 \|p\|_2^2} = \sqrt{\sum_{i=1}^q \|n_i\|_2^2}.$$

This implies that

$$\mathbb{P} \left\{ \left\| \begin{bmatrix} y_1(x_1 - \bar{x}_1)^\top \\ \vdots \\ y_q(x_q - \bar{x}_q)^\top \end{bmatrix} \right\|_2 \leq \epsilon \right\} \geq \mathbb{P} \left\{ \sqrt{\sum_{i=1}^q \|n_i\|_2^2} \leq \epsilon \right\} \\ = \mathbb{P} \left\{ \sum_{i=1}^q \|n_i\|_2^2 \leq \epsilon^2 \right\}$$

Therefore, it can be proved that

$$\mathbb{P} \left\{ \left\| \begin{bmatrix} \bar{\alpha}^* - \alpha^* \\ \bar{\beta}^* - \beta^* \\ \bar{\xi}^* - \xi^* \end{bmatrix} \right\|_2 \leq c\epsilon \right\} \geq \mathbb{P} \left\{ \left\| \begin{bmatrix} y_1(x_1 - \bar{x}_1)^\top \\ \vdots \\ y_q(x_q - \bar{x}_q)^\top \end{bmatrix} \right\|_2 \leq \epsilon \right\} \\ \geq \mathbb{P} \left\{ \sum_{i=1}^q \|n_i\|_2^2 \leq \epsilon^2 \right\} \\ = 1 - \mathbb{P} \left\{ \sum_{i=1}^q \|n_i\|_2^2 \geq \epsilon^2 \right\} \\ \geq 1 - \frac{1}{\epsilon^2} \mathbb{E} \left\{ \sum_{i=1}^q \|n_i\|_2^2 \right\},$$

where the last inequality follows from the Markov's inequality [44, Theorem 8.3]. Noting that $\mathbb{E} \left\{ \sum_{i=1}^q \|n_i\|_2^2 \right\} = q \text{Tr}(\int_n n n^\top \gamma(n) dn)$ while setting $\varepsilon = c\epsilon$ and $\varepsilon_0 = c\epsilon_0$ completes the proof.

APPENDIX C

PROOF OF THEOREM 2.4

Note that

$$\mathcal{L} := \text{Tr}(\Pi^{-1} \mathcal{I}) + \int_n \lambda n^\top n \gamma(n) dn + \mu \left(-1 + \int_n \gamma(n) dn \right) \\ = \int_n \left(\frac{1}{\gamma(n)} \left[\frac{\partial \gamma(n)}{\partial n} \right]^\top \Pi^{-1} \left[\frac{\partial \gamma(n)}{\partial n} \right] \right. \\ \left. + (\lambda n^\top n + \mu) \gamma(n) \right) dn - \mu.$$

Noting that the cost function and the constraints are convex, the stationary condition of $\mathcal{L}$ (that the variational derivative is equal to zero) is sufficient for optimality. Further, if multiple density functions satisfy the sufficiency conditions, they all exhibit the same cost. Using Theorem 5.3 in [45, p. 440], it can be seen that the extrema must satisfy

$$\lambda n^\top n + \mu + \frac{1}{\gamma(n)^2} \left[\frac{\partial \gamma(n)}{\partial n} \right]^\top \Pi^{-1} \left[\frac{\partial \gamma(n)}{\partial n} \right] \\ - 2 \frac{1}{\gamma(n)} \text{Tr}(\Pi^{-1} \mathfrak{D}^2 \gamma(n)) = 0.$$

With the change of variable $\gamma(n) = u(n)^2$, the condition can be transformed into

$$\lambda n^\top n + \mu - 4 \frac{1}{u(n)} \text{Tr}(\Pi^{-1} \mathfrak{D}^2 u(n)) = 0.$$

It can be seen that the $u(n) = c_0 \exp(-n^\top X n/2)$ with $X = \Pi^{1/2} \lambda^{1/2}/2$ and $c_0 = (1/\det(\pi X^{-1}))^{1/4}$ satisfies this equation with $\mu = -4\text{Tr}(X)$.

APPENDIX D

PROOF OF PROPOSITION 2.8

Define $x = [x_1^\top \dots x_q^\top]^\top$ and $\bar{x} = [\bar{x}_1^\top \dots \bar{x}_q^\top]^\top$. Note that the conditional probability density function of $\bar{x} = x + w$ given x is equal to $\tilde{\gamma}(\bar{x} - x)$. Using the Cramér-Rao bound [3, p. 169], it can be deduced that

$$\begin{aligned} \mathbb{E}\{\|\Pi(x_i - \hat{x}_i(\bar{x}))\|_2^2\} \\ \geq \text{Tr}\left(\Pi\left[\int_{\bar{x}} \tilde{\gamma}(\bar{x} - x) \left[\frac{\partial \log(\tilde{\gamma}(\bar{x} - x))}{\partial x_i}\right] \right. \right. \\ \left. \left. \times \left[\frac{\partial \log(\tilde{\gamma}(\bar{x} - x))}{\partial x_i}\right]^\top d\bar{x}\right]^{-1}\right). \end{aligned}$$

Therefore,

$$\begin{aligned} \text{Tr}\left(\Pi\left[\int_{\bar{x}} \tilde{\gamma}(\bar{x} - x) \left[\frac{\partial \log(\tilde{\gamma}(\bar{x} - x))}{\partial x_i}\right] \right. \right. \\ \left. \left. \times \left[\frac{\partial \log(\tilde{\gamma}(\bar{x} - x))}{\partial x_i}\right]^\top d\bar{x}\right]^{-1}\right) \\ = \text{Tr}(\Pi[(e_i^\top \otimes I_p) \tilde{\mathcal{L}}(e_i \otimes I_p)]^{-1}) \end{aligned}$$

because

$$\frac{\partial \log(\tilde{\gamma}(\bar{x} - x))}{\partial x_i} = (e_i^\top \otimes I_p) \frac{\partial \log(\tilde{\gamma}(w))}{\partial w} \Big|_{w=\bar{x}-x}.$$

Now, by Item (11) in [32, p. 44], it can be seen that

$$\begin{aligned} \text{Tr}(\Pi[(e_i^\top \otimes I_p) \tilde{\mathcal{L}}(e_i \otimes I_p)]^{-1}) \\ \geq 1/\text{Tr}(\Pi^{-1}(e_i^\top \otimes I_p) \tilde{\mathcal{L}}(e_i \otimes I_p)) \\ = 1/\text{Tr}((e_i^\top \otimes I_p)(I_q \otimes \Pi^{-1}) \tilde{\mathcal{L}}(e_i \otimes I_p)) \\ = 1/\text{Tr}((e_i e_i^\top) \otimes I_p)(I_q \otimes \Pi^{-1}) \tilde{\mathcal{L}}). \end{aligned}$$

Thus,

$$\begin{aligned} \max_{1 \leq i \leq q} 1/\mathbb{E}\{\|x_i - \hat{x}_i(\bar{x})\|_2^2\} \\ \leq \max_{1 \leq i \leq q} \text{Tr}(((e_i e_i^\top) \otimes I_p)(I_q \otimes \Pi^{-1}) \tilde{\mathcal{L}}) \\ \leq \sum_{i=1}^q \text{Tr}(((e_i e_i^\top) \otimes I_p)(I_q \otimes \Pi^{-1}) \tilde{\mathcal{L}}) \\ = \text{Tr}((I_q \otimes \Pi^{-1}) \tilde{\mathcal{L}}). \end{aligned}$$

The rest follows from the fact that $\min_{1 \leq i \leq q} \mathbb{E}\{\|x_i - \hat{x}_i(\bar{x})\|_2^2\} = 1/(\max_{1 \leq i \leq q} 1/\mathbb{E}\{\|x_i - \hat{x}_i(\bar{x})\|_2^2\})$.

APPENDIX E

PROOF OF PROPOSITION 2.9

Note that $\bar{x} = x + w$ with $\bar{x} = [\bar{x}_1^\top \dots \bar{x}_q^\top]^\top$ and $x = [x_1^\top \dots x_q^\top]^\top$, the KKT conditions for (2)–depicted in (15)–and the KKT conditions for (6) become identical to each other if and only if $\sum_{i=1}^q (\omega^*)_{ii} y_i (e_i^\top \otimes I_p) w = 0$ and $\alpha^{*\top} (e_i^\top \otimes I_p) w = 0$. The rest follows from algebraic manipulation of the equations.

APPENDIX F

PROOF OF THEOREM 2.10

If $w = \Psi \bar{w}$ with probability one for some random variable $\bar{w}$, it can be seen that $\Omega w = \Omega \Psi \bar{w} = 0$. Note that $\tilde{\gamma}(w) = 0$ for all $w \in \text{im}(\Psi)^\perp$ and $\tilde{\gamma}(w) = \hat{\gamma}(\Psi^\dagger w)$ for all $w \in \text{im}(\Psi)$, where $\hat{\gamma}$ is the probability density function of $\bar{w}$. Therefore,

$$\begin{aligned} \tilde{\mathcal{L}} &= \int_{w \in \mathbb{R}^p} \tilde{\gamma}(w) \left[\frac{\partial \log(\tilde{\gamma}(w))}{\partial w} \right] \left[\frac{\partial \log(\tilde{\gamma}(w))}{\partial w} \right]^\top dw \\ &= \int_{w \in \text{im}(\Psi)} \hat{\gamma}(\Psi^\dagger w) \left[\frac{\partial \log(\hat{\gamma}(\Psi^\dagger w))}{\partial w} \right] \left[\frac{\partial \log(\hat{\gamma}(\Psi^\dagger w))}{\partial w} \right]^\top dw \\ &= \int_{\bar{w}} \hat{\gamma}(\bar{w}) \Psi^\dagger \left[\frac{\partial \log(\hat{\gamma}(\bar{w}))}{\partial \bar{w}} \right] \left[\frac{\partial \log(\hat{\gamma}(\bar{w}))}{\partial \bar{w}} \right]^\top \Psi^\dagger d\bar{w}. \end{aligned}$$

Further, note that $\Psi^\dagger \Psi^{\dagger\top} = \Psi^\dagger (\Psi^\top)^\dagger = (\Psi^\top \Psi)^\dagger = I$, where the last equality follows from the definition of Ψ . Therefore, the optimization problem in (16) can be translated into

$$\begin{aligned} \min_{\hat{\gamma}} \quad & \text{Tr}\left(\Psi^\dagger (I_q \otimes \Pi^{-1}) \Psi^{\dagger\top} \right. \\ & \left. \times \int_{\bar{w}} \hat{\gamma}(\bar{w}) \left[\frac{\partial \log(\hat{\gamma}(\bar{w}))}{\partial \bar{w}} \right] \left[\frac{\partial \log(\hat{\gamma}(\bar{w}))}{\partial \bar{w}} \right]^\top d\bar{w} \right), \quad (18a) \\ \text{s.t.} \quad & V_{\bar{w}\bar{w}} \leq mI. \quad (18b) \end{aligned}$$

Following the same line of reasoning as in the proof of Theorem 2.4, it can be shown that the solution of (18) is a Gaussian noise with zero mean and co-variance $V_{\bar{w}\bar{w}}$. Therefore, this problem can be further simplified into

$$\begin{aligned} \min_{V_{\bar{w}\bar{w}} \geq 0} \quad & \text{Tr}\left(\Psi^\dagger (I_q \otimes \Pi^{-1}) \Psi^{\dagger\top} V_{\bar{w}\bar{w}}^{-1}\right), \quad (19a) \\ \text{s.t.} \quad & V_{\bar{w}\bar{w}} \leq mI. \quad (19b) \end{aligned}$$

Because of the inequality constraint in the above optimization problem, $V_{\bar{w}\bar{w}}^{-1} \geq (1/m)I$, and, thus, it can be proved that

$$\begin{aligned} \text{Tr}\left(\Psi^\dagger (I_q \otimes \Pi^{-1}) \Psi^{\dagger\top} V_{\bar{w}\bar{w}}^{-1}\right) \\ = \text{Tr}\left((I_q \otimes \Pi^{-1})^{1/2} \Psi^{\dagger\top} V_{\bar{w}\bar{w}}^{-1} \Psi^\dagger (I_q \otimes \Pi^{-1})^{1/2}\right) \\ \geq (1/m) \text{Tr}\left((I_q \otimes \Pi^{-1})^{1/2} \Psi^{\dagger\top} \Psi^\dagger (I_q \otimes \Pi^{-1})^{1/2}\right) \\ = (1/m) \text{Tr}\left(\Psi^\dagger (I_q \otimes \Pi^{-1}) \Psi^{\dagger\top}\right). \end{aligned}$$

The lower bound on the cost function can be clearly achieved by selecting $V_{\bar{w}\bar{w}} = mI$.

APPENDIX G

PROOF OF PROPOSITION 2.11

The optimality condition is $\nabla_\phi \ell(\phi; \{(x_i, y_i)\}_{i=1}^q) = 0$. Small variations of the dataset shows that

$$\mathcal{D}_\phi^2 \ell(\phi; \{(x_i, y_i)\}_{i=1}^q) \nabla_{x_j} \phi + \nabla_{x_j} \nabla_\phi \ell(\phi; \{(x_i, y_i)\}_{i=1}^q) = 0.$$

Therefore, if $\mathcal{D}_\phi^2 \ell(\phi; \{(x_i, y_i)\}_{i=1}^q)$ is invertible, it can be shown that

$$\nabla_{x_j} \phi = -[\mathcal{D}_\phi^2 \ell(\phi; \{(x_i, y_i)\}_{i=1}^q)]^{-1} \nabla_{x_j} \nabla_\phi \ell(\phi; \{(x_i, y_i)\}_{i=1}^q).$$

The Taylor's Theorem shows that

$$\|\bar{\varphi} - \varphi\|_2 \leq \sum_{j=1}^q \left[\left\| [\mathfrak{D}_{\varphi\varphi}^2 \ell(\varphi'; \{(x_i, y_i)\}_{i=1}^q)]^{-1} \right. \right. \\ \left. \left. \times \mathfrak{D}_{x_i\varphi}^2 \ell(\varphi'; \{(x_i, y_i)\}_{i=1}^q) \right\|_2 \|x_i - \bar{x}_i\|_2 \right],$$

for some $\varphi' = s\varphi + (1-s)\bar{\varphi}$ and $s \in [0, 1]$. Therefore, there exist $c \geq 1$ and $\epsilon_0 > 0$ (due to continuity of the second order derivatives) such that if $\sum_{i=1}^q \|x_i - \bar{x}_i\|_2 \leq \epsilon < \epsilon_0$, then $\|\bar{\varphi} - \varphi\|_2 \leq c\epsilon$. The Markov's inequality [44, Theorem 8.3] implies that

$$\begin{aligned} \mathbb{P} \left\{ \sum_{i=1}^q \|x_i - \bar{x}_i\|_2 \leq \epsilon/c \right\} &= \mathbb{P} \left\{ \sum_{i=1}^q \|n_i\|_2 \leq \epsilon/c \right\} \\ &= \mathbb{P} \left\{ \left(\sum_{i=1}^q \|n_i\|_2 \right)^2 \leq \epsilon^2/c \right\} \\ &\geq 1 - \frac{c^2}{\epsilon^2} \mathbb{E} \left\{ \left(\sum_{i=1}^q \|n_i\|_2 \right)^2 \right\}. \end{aligned}$$

Furthermore,

$$\begin{aligned} \mathbb{E} \left\{ \left(\sum_{i=1}^q \|n_i\|_2 \right)^2 \right\} &= \mathbb{E} \left\{ \sum_{i=1}^q n_i^\top n_i + \sum_{i=1}^q \sum_{j \neq i} \sqrt{n_i^\top n_i n_j^\top n_j} \right\} \\ &\leq q \text{Tr}(V_{nn}) + \sum_{i=1}^q \sum_{j \neq i} \sqrt{\mathbb{E} \{ n_i^\top n_i n_j^\top n_j \}} \\ &= q \text{Tr}(V_{nn}) + \sum_{i=1}^q \sum_{j \neq i} \sqrt{\text{Tr}(V_{nn})^2} \\ &= q^2 \text{Tr}(V_{nn}). \end{aligned}$$

where the inequality follows from the Jensen's inequality (and the concavity of the square root function).

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Farokhi, F

Title:

Privacy-Preserving Public Release of Datasets for Support Vector Machine Classification

Date:

2020

Citation:

Farokhi, F. (2020). Privacy-Preserving Public Release of Datasets for Support Vector Machine Classification. IEEE Transactions on Big Data, 7 (5), <https://doi.org/10.1109/tbdata.2019.2963391>.

Persistent Link:

<http://hdl.handle.net/11343/251371>