

Robust Optimization for Local Differential Privacy

Jasper Goseling
University of Twente,
the Netherlands

Milan Lopuhaä-Zwakenberg
University of Twente,
the Netherlands

Abstract—We consider the setting of publishing data without leaking sensitive information. We do so in the framework of Robust Local Differential Privacy (RLDP). This ensures privacy for all distributions of the data in an uncertainty set. We formulate the problem of finding the optimal data release protocol as a robust optimization problem. By deriving closed-form expressions for the duals of the constraints involved we obtain a convex optimization problem. We compare the performance of four possible optimization problems depending on whether or not we require robustness in i) utility and ii) privacy.

I. INTRODUCTION

We consider a setting in which users have data (S, U) with $U \in \mathbb{R}$. A user wants to publish U , but does not want to disclose information about their sensitive data S , which may be correlated with U . Therefore, the users release an obfuscated version Y of U , such that Y is as close to U as possible, measured by mean squared distortion, without leaking too much information about S . This scenario and closely related ones have been studied in, for instance, [1]–[10].

To measure the information leakage about S in Y we use a form of Local Differential Privacy [11] that was introduced in [12], which states that for each possible s_1, s_2 and y the following should hold:

$$\mathbb{P}(Y = y|S = s_1) \leq e^\epsilon \mathbb{P}(Y = y|S = s_2). \quad (1)$$

This condition is less strict than regular LDP, reflecting the fact that only S , and not U itself, needs to be protected. Note that (1) depends on the joint probability distribution P_{SU} . A privacy protocol is given by the conditional probability distribution $P_{Y|SU}$, and given such a protocol, the distortion is measured by $\mathbb{E}(U - Y)^2$, which also depends on P_{SU} .

Thus both privacy and utility depend on the distribution of the data. However, a user may not know this distribution exactly and needs to estimate it. The odds ratio interpretation of differential privacy [13] tells us that an attacker with a better estimate may obtain more information about S than the privacy protocol that is developed based on this estimate would indicate. Furthermore, a good utility under the estimated distribution may not imply a good utility under the actual distribution.

In this paper, therefore, we demand stronger, *robust* privacy and utility guarantees. More concretely, we assume that there exists a publicly available dataset from which the users produce an estimate distribution \hat{P}_{SU} , and a set \mathcal{F} of probability distributions that do not differ significantly from




Fig. 1: Overview. For clarity we write $P^* = P_{SU}^*$ and $\hat{P} = \hat{P}_{SU}$. Since \mathcal{F} is a confidence interval P^* may or may not be an element of \mathcal{F} .

$\hat{P}_{S,U}$ (for a chosen significance level). Then robustness can be incorporated in the following two ways:

- 1) We demand robust privacy by requiring (1) to hold for all $P_{SU} \in \mathcal{F}$;
- 2) We obtain robust utility by minimising $\max_{P_{SU} \in \mathcal{F}} \mathbb{E}(U - Y)^2$.

Robustness allows us to guarantee privacy and utility for all probability distributions one can reasonably expect, without sacrificing too much utility to account for unlikely distributions. It is important to note that: i) robustness in utility is w.r.t. our own uncertainty about $P_{S,U}$ and ii) robustness in privacy is w.r.t. our uncertainty about current knowledge by an attacker. The robustness of privacy and utility can be incorporated independently, leaving us with four possible optimization problems depending on for which of the two we want robust guarantees.

In recent work [9], [10] we introduced the robust privacy framework that we also use here. An important difference is that in [9], [10] the utility measure is $I(S, U; Y)$, leading to an optimization problem is not convex. Therefore, similar to [2], the resulting techniques for analysis in [9], [10] are combinatorial in nature and heuristics are developed. In the current work the utility measure is the mean-squared error, leading to a non-robust optimization problem that is convex.

One of the main contributions of this paper is to provide convex formulations of the corresponding robust optimization problems. These formulations can be handled by standard convex optimization solvers in order to obtain privacy protocols that are provably optimal. Our numerical results demonstrate that: i) without including robustness on the privacy constraints

the effective privacy guarantees that are obtained are very weak, ii) including robustness on privacy leads to a significant penalty on utility, and iii) including a robustness constraint on utility does not have a large impact if privacy robustness is already imposed.

Our techniques are rooted in robust optimization [14], [15], an important aspect of which is to use Fenchel duality on constraints like (1) that need to hold for all distributions in \mathcal{F} . In [16] and [17] duals are derived for constraints involving probabilities and uncertainty sets based on the χ^2 -test. A major difference in the current work, and a contribution on the technical side, is that (1) involves two conditional distributions from \mathcal{F} .

The structure of this paper is as follows. In Section II we present the details of our model. We provide background on robust optimization in Section III and present our analytical results in Section IV. Insights obtained by numerical experiments are given Section V. Finally, in Section VI we provide a discussion and outlook on future work. Due to space constraints proofs are omitted from this version, and are presented in the ArXiv version [18].

II. MODEL

An overview of our model, the details of which are given in this section, is given in Figure 1. There is a publicly accessible dataset $\overrightarrow{(s, u)}$ of size n , in which each entry (s_i, u_i) is drawn independently from a probability distribution P_{SU}^* on a set $\mathcal{S} \times \mathcal{U}$, where \mathcal{S} and $\mathcal{U} \subset \mathbb{R}$ are finite alphabets. New data items (S, U) are also drawn from P_{SU}^* . The user's aim is to create a release protocol $P_{Y|SU}$ such that $Y \in \mathcal{U}$ is as close as possible to U , while not leaking too much information about S . More precisely, the goal is to minimize the distortion, measured as $\mathbb{E}[d(U, Y)]$, where d is any distance function on \mathcal{U} .

The distribution P_{SU}^* is not known exactly. The uncertainty set $\mathcal{F} \subset \mathcal{P}_{\mathcal{S} \times \mathcal{U}}$, where $\mathcal{P}_{\mathcal{S} \times \mathcal{U}}$ denotes the probability simplex over $\mathcal{S} \times \mathcal{U}$, captures the user's uncertainty about P_{SU}^* . It is constructed from the dataset $\overrightarrow{(s, u)}$. More specifically, we let \mathcal{F}_B be the $(1 - \alpha)$ -confidence set for P_{SU} in a χ^2 -test, i.e.,

$$\mathcal{F}_B = \left\{ P_{SU} \in \mathcal{P}_{\mathcal{S} \times \mathcal{U}} \left| \sum_{\substack{s \in \mathcal{S}, \\ u \in \mathcal{U}}} \frac{(\hat{P}_{s,u} - P_{s,u})^2}{P_{s,u}} \leq B \right. \right\}, \quad (2)$$

where

$$B = \frac{F_{|\mathcal{S} \times \mathcal{U}|-1}^{-1}(1 - \alpha)}{n}, \quad (3)$$

\hat{P}_{SU} is the empirical probability distribution of (S, U) and F_d is the CDF of the χ^2 -distribution with d degrees of freedom. For the remainder of this paper we fix B and write $\mathcal{F} := \mathcal{F}_B$.

The user creates $P_{Y|SU}$ in such a way that a Local Differential Privacy-like privacy standard is guaranteed when $P_{SU} \in \mathcal{F}$. We will denote this as robust local differential privacy (RLDP) [10]; it is defined as follows.

Definition 1. Let $\varepsilon \geq 0$ and $\mathcal{F} \subset \mathcal{P}_{\mathcal{S} \times \mathcal{U}}$. We say that $P_{Y|SU}$ satisfies $(\varepsilon, \mathcal{F})$ -RLDP if for all $s_1, s_2 \in \mathcal{S}$, all $y \in \mathcal{Y}$, and all $P_{SU} \in \mathcal{F}$ we have

$$\mathbb{P}(Y = y | S = s_1) \leq e^\varepsilon \mathbb{P}(Y = y | S = s_2). \quad (4)$$

This is less strict than regular LDP [11], which requires that not just S but also U be protected. Unlike regular LDP it depends on P_{SU} , which is why we demand robustness. In terms of coefficients we can write (4) as

$$\forall y, s_1, s_2 \forall P_{SU} \in \mathcal{F}: L_{y, s_1, s_2}(P_{SU}, P_{Y|SU}) \leq 0, \quad (5)$$

$$\begin{aligned} & L_{y, s_1, s_2}(P_{SU}, P_{Y|SU}) \\ &= \sum_{u \in \mathcal{U}} (P_{u|s_1} P_{y|s_1, u} - e^\varepsilon P_{u|s_2} P_{y|s_2, u}). \end{aligned} \quad (6)$$

We distinguish four possible optimization problems depending on whether or not we require robustness in i) utility and ii) privacy. This gives:

$$\text{NUNP: minimize } \mathbb{E}_{(S,U) \sim \hat{P}_{SU}} [d(U, Y)] \quad (7)$$

subject to

$$L_{y, s_1, s_2}(\hat{P}_{SU}, P_{Y|SU}) \leq 0, \forall y, s_1, s_2,$$

$$\text{NURP: minimize } \mathbb{E}_{(S,U) \sim \hat{P}_{SU}} [d(U, Y)] \quad (8)$$

subject to

$$L_{y, s_1, s_2}(P_{SU}, P_{Y|SU}) \leq 0,$$

$$\forall P_{SU} \in \mathcal{F}, \forall y, s_1, s_2,$$

$$\text{RUNP: minimize } D \quad (9)$$

subject to

$$\mathbb{E}_{(S,U) \sim P_{SU}} [d(U, Y)] \leq D, \quad \forall P_{SU} \in \mathcal{F},$$

$$L_{y, s_1, s_2}(\hat{P}_{SU}, P_{Y|SU}) \leq 0, \forall y, s_1, s_2,$$

$$\text{RURP: minimize } D \quad (10)$$

subject to

$$\mathbb{E}_{(S,U) \sim P_{SU}} [d(U, Y)] \leq D, \quad \forall P_{SU} \in \mathcal{F},$$

$$L_{y, s_1, s_2}(P_{SU}, P_{Y|SU}) \leq 0, \forall P_{SU} \in \mathcal{F}, \forall y, s_1, s_2.$$

In all these problems the optimization variable is $P_{Y|SU}$. Note that NUNP corresponds to the 'naive' approach of assuming $P_{SU}^* = \hat{P}_{SU}$ and doing nonrobust optimization.

We note that one can make different choices for the utility [9], privacy [3], [12], and uncertainty set [17]. As we will see below, robust optimization is a general framework that works for many choices, but one needs problem-specific analytic results (Theorems 2 & 4 and Lemmas 1 & 2 below) in order to reformulate (7)–(10) as convex optimization problems, which can be fed to a solver.

III. ROBUST OPTIMIZATION

Problems (8)–(10) are robust optimization problems [14], [15]. In this section we provide some background on robust optimization.

Let $f : \mathbb{R}^m \times \mathbb{R}^n \rightarrow \bar{\mathbb{R}}$ be defined on the extended real line $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$ and for a given $x \in \mathbb{R}^n$

define $\text{dom}(f(\cdot, x)) = \{a \mid f(a, x) > -\infty\}$. Consider an optimization problem with constraint

$$f(a, x) \leq 0, \forall a \in \mathcal{A}, \quad (11)$$

where $x \in \mathbb{R}^n$ is the optimization variable, $a \in \mathbb{R}^m$ is an uncertain parameter and \mathcal{A} is the uncertainty set for which we need the constraint to hold. The robustness constraints in (8)–(10) are of the form (11) with P_{SU} and $P_{Y|SU}$ in the roles of a and x , respectively.

The main tool that is used in this paper is based on Fenchel duality. In order to present this, let

$$f_*(v, x) = \inf_{a \in \mathbb{R}^m} v^T a - f(a, x) \quad (12)$$

and

$$\delta^*(v \mid \mathcal{A}) = \sup_{a \in \mathcal{A}} v^T a, \quad (13)$$

which are known as the partial concave conjugate of f and the support set of \mathcal{A} , respectively. Also, let $\text{ri}(\cdot)$ denote the relative interior of a set. We will extensively use the following result.

Theorem 1 ([15]). *Let $x \in \mathbb{R}^n$. If $f(a, x)$ is closed concave in a and $\text{ri}(\mathcal{A}) \cap \text{ri}(\text{dom}(f(\cdot, x))) \neq \emptyset$, then x satisfies (11) if and only if*

$$\exists v \in \mathbb{R}^m \text{ s.t. } \delta^*(v \mid \mathcal{A}) \leq f_*(v, x). \quad (14)$$

Since $\delta^*(v \mid \mathcal{A})$ and $f_*(v, x)$ are convex and concave functions, respectively, Theorem 1 provides a means to efficiently solve robust optimization problems once these functions are found in closed form. Another result we use is that Theorem 1 can be applied independently to each uncertain constraint [14], and so we can handle robust utility and robust privacy separately.

IV. CLOSED-FORM EXPRESSIONS FOR ROBUSTNESS

While Theorem 1 is a powerful tool to deal with the universal quantifiers in the robust constraints, the downside is that one has to find closed forms of the functions $\delta^*(\bullet \mid \mathcal{A})$ and $f_*(\bullet, x)$. In this section, we show how to do this for the robust utility and privacy constraints of (7)–(10). As will become apparent from the discussion below, the solutions we find are specific to our choice of privacy, utility, and uncertainty set.

A. Robust utility

The robust utility constraint in RUNP and RURP is of the form

$$\mathbb{E}_{(S,U) \sim P_{SU}}[d(U, Y)] \leq D, \quad \forall P_{SU} \in \mathcal{F}. \quad (15)$$

We can write this as

$$E_D(P_{SU}, P_{Y|SU}) \leq 0, \quad \forall P_{SU} \in \mathcal{F}, \quad (16)$$

$$E_D(P_{SU}, P_{Y|SU}) = \mathbb{E}[d(U, Y)] - D \quad (17)$$

$$= \sum_{y,s,u} P_{s,u} P_{y|s,u} d(u, y) - D. \quad (18)$$

Using Theorem 1 we can write this as

$$\exists v \in \mathbb{R}^{|\mathcal{S} \times \mathcal{U}|}: \delta^*(v \mid \mathcal{F}) \leq E_{D*}(v, P_{Y|SU}). \quad (19)$$

Thus we need to find expressions for $\delta^*(v \mid \mathcal{F})$ and $E_{D*}(v, P_{Y|SU})$. Since E_D is linear in P_{SU} , one easily derives the following [14]:

Lemma 1. *Let $E_D(P_{SU}, P_{Y|SU}) = \mathbb{E}[d(U, Y)] - D$, with uncertain parameter P_{SU} and optimization variable $P_{Y|SU}$. Then*

$$E_{D*}(v, P_{Y|SU}) = \begin{cases} D, & \text{if } \forall s, u: v_{s,u} = \sum_y P_{y|s,u} d(y, u), \\ -\infty, & \text{otherwise.} \end{cases} \quad (20)$$

The expression for $\delta^*(v \mid \mathcal{F})$ is a little bit more work, but can be found using established results in convex analysis [15]. A related result is found in [17, Thm. 5].

Theorem 2.

$$\delta^*(v \mid \mathcal{F}) = \min_{w \geq v, c \geq 0} \left\{ -2\sqrt{c} \sum_{s,u} \hat{P}_{s,u} \sqrt{w_{s,u} - v_{s,u}} + \max_{s,u} w_{s,u} + c(B+1) \right\}. \quad (21)$$

Combining these results, we can write the robust utility constraint (15) as

$$\delta^* \left(\left(\sum_y P_{y|s,u} d(y, u) \right)_{s \in \mathcal{S}, u \in \mathcal{U}} \middle| \mathcal{F} \right) \leq D, \quad (22)$$

where $\delta^*(\bullet \mid \mathcal{F})$ is as in Theorem 2.

B. Robust privacy

Fix y, s_1, s_2 . The robust privacy constraint in NURP and RURP is of the form

$$L_{y,s_1,s_2}(P_{SU}, P_{Y|SU}) \leq 0, \forall P_{SU} \in \mathcal{F}. \quad (23)$$

One would like to apply Theorem 1; however, L_{y,s_1,s_2} is not convex in P_{SU} . Therefore we rewrite it as follows. Let $\mathcal{F}_{U|s_1,s_2}$ be the projection of \mathcal{F} onto $\mathcal{P}_{\mathcal{U}} \times \mathcal{P}_{\mathcal{U}}$ by the map $P \mapsto (P_{U|s_1}, P_{U|s_2})$. We can write (23) as

$$\forall (P_{U|s_1}, P_{U|s_2}) \in \mathcal{F}_{U|s_1,s_2}: \tilde{L}_{y,s_1,s_2}(P_{U|s_1}, P_{U|s_2}, P_{Y|SU}) \leq 0, \quad (24)$$

where

$$\begin{aligned} \tilde{L}_{y,s_1,s_2}(P_{U|s_1}, P_{U|s_2}, P_{Y|SU}) \\ = \sum_{u \in \mathcal{U}} (P_{u|s_1} P_{y|s_1,u} - e^\epsilon P_{u|s_2} P_{y|s_2,u}). \end{aligned} \quad (25)$$

Again we use Theorem 1 to write this as

$$\exists v_1, v_2 \in \mathbb{R}^{|\mathcal{U}|}: \delta^*(v_1, v_2 \mid \mathcal{F}_{U|s_1,s_2}) \leq \tilde{L}_{y,s_1,s_2*}(v_1, v_2, P_{Y|SU}). \quad (26)$$

Thus we need to find expressions for these functions. As $\tilde{L}_{y,s_1,s_2}(P_{U|s_1}, P_{U|s_2}, P_{Y|SU})$ is linear in its first two arguments, we find the following:

Lemma 2. *One has*

$$\begin{aligned} & \tilde{L}_{y,s_1,s_2^*}(v_1, v_2, P_{Y|SU}) \\ &= \begin{cases} 0, & \text{if } \forall u: v_{1,u} = P_{y|s_1,u}, \\ & \text{and } v_{2,u} = -e^\epsilon P_{y|s_2,u}, \\ -\infty, & \text{otherwise.} \end{cases} \end{aligned} \quad (27)$$

The expression for $\delta^*(v_1, v_2 | \mathcal{F}_{U|s_1,s_2})$ takes considerably more effort, in part because we first have to find a closed expression for $\mathcal{F}_{U|s_1,s_2}$. This is given in the following result:

Theorem 3.

$$\mathcal{F}_{U|s_1,s_2} = \left\{ \left(R_{U|s_1}, R_{U|s_2} \right) \in \mathcal{P}_{\mathcal{U}}^2 \mid \sum_{i=1}^2 \sqrt{\sum_u \frac{\hat{P}_{s_i,u}^2}{R_{u|s_i}}} \leq \sqrt{B+1} - 1 + \sum_{i=1}^2 \hat{P}_{s_i} \right\}. \quad (28)$$

Similar to, but more in a more complicated way than, Theorem 2, one can prove:

Theorem 4. *The support function of $\mathcal{F}_{U|s_1,s_2}$ is given by*

$$\begin{aligned} \delta^*(v_1, v_2 | \mathcal{F}_{U|s_1,s_2}) &= \min_{c \geq 0, w_1 \geq v_1, w_2 \geq v_2} \left\{ \right. \\ & - (2^{-2/3} + 2^{1/3}) c^{2/3} \sum_{i=1}^2 \left(\sum_u \hat{P}_{s_i,u} \sqrt{w_i(u) - v_i(u)} \right)^{2/3} \\ & \left. + \sum_{i=1}^2 \max_u w_i(u) + c \left(\sqrt{B+1} - 1 + \sum_{i=1}^2 \hat{P}_{s_i} \right) \right\}. \end{aligned} \quad (29)$$

Thus we can write (23) as

$$\delta^* \left((P_{y|s_1,u})_{u \in \mathcal{U}}, (P_{y|s_2,u})_{u \in \mathcal{U}} \mid \mathcal{F}_{U|s_1,s_2} \right) \leq 0, \quad (30)$$

where $\delta^*(\bullet | \mathcal{F}_{U|s_1,s_2})$ is as in Theorem 4. Replacing (15) by (22), and (23) by (30) allows us to provide convex formulations of (7)–(10). In the next section we provide some insights that are obtained through numerical experiments.

V. NUMERICAL EXPERIMENTS

The performance measures that we consider are

$$D^* = \mathbb{E}_{(S,U) \sim P_{SU}^*} [d(U, Y)] \quad (31)$$

and

$$\epsilon^* = \log \max_{y, s_1, s_2} \frac{P_{y|s_1}}{P_{y|s_2}}, \quad (32)$$

interpreting $0/0$ as 1 and with $\epsilon^* = \infty$ if there exist y, s_1 and s_2 for which $P_{y|s_2} = 0$ and $P_{y|s_1} > 0$. These measures give the performance of the $P_{Y|SU}$ that is obtained from optimization under the actual distribution P_{SU}^* .

In all experiments we use $\mathcal{S} = \{0, 1, 2\}$, $\mathcal{U} = \{0, 1, 2, 3, 4\}$, $d(x, y) = (x - y)^2$, $\alpha = 0.05$ and $\epsilon = 0.5$. We draw P_{SU}^* according to the Jeffreys prior on $\mathcal{P}_{\mathcal{S} \times \mathcal{U}}$, i.e., the symmetric Dirichlet distribution with parameter $\frac{1}{2}$ [19].

In our first experiment we draw $K = 30$ instances of P_{SU}^* . For each instance, we draw n samples (s, u) from which we

estimate \hat{P}_{SU} . For each of the $K = 30$ instances we report ϵ^* and D^* in a scatter plot. We do so for each combination of nonrobust/robust utility and privacy (nomenclature in (7)–(10)). In Figures 2a and 2b the results are depicted for $n = 5|\mathcal{S}||\mathcal{U}| = 75$ and for $n = 10^3|\mathcal{S}||\mathcal{U}| = 15,000$, respectively. Instances that have $\epsilon^* = \infty$ are omitted from the figures.

We observe that without robustness in privacy, i.e., NUNP and RUNP, almost all instances have significantly lower privacy performance than what they are designed for, i.e., $\epsilon^* \gg \epsilon$. Also, by introducing robustness in privacy as a hard constraint, we significantly suffer in terms of utility. Note though that even with robust privacy we do not always get $\epsilon^* \leq \epsilon$; this is because $P^* \in \mathcal{F}$ only with probability $1 - \alpha$. In Figure 2a NURP has more than the expected number of outliers, which can be explained from the fact that the χ^2 test only asymptotically gives a confidence interval as $n \rightarrow \infty$. Another interesting observation (which will be confirmed by the next experiments) is that once we have imposed robust privacy, imposing robust utility does not make much difference, i.e., NURP and RURP are very similar. This shows that the utility cost of demanding robust privacy is considerably greater than the utility benefit of demanding robust utility. Note, that by comparing to Figure 2b, some of the large ϵ^* values for NURP in Figure 2a seem to be small sample artifacts. A final observation is that in both figures there are outliers, especially in ϵ^* .

Our second experiment illustrates the influence of n . We give the mean D^* and ϵ^* over $K = 10^3$ instances as a function of n in Figures 2c and 2d, respectively. In addition to the mean we report the standard deviation. Thick lines depict the mean over $K = 10^3$ samples, the width of the bands correspond to ± 1 standard deviation. Both are computed by first removing outliers based on the 1.5IQR rule.

We observe that for small n , due to the large uncertainty on P_{SU} , RURP is very conservative with $\epsilon^* \ll \epsilon$. The figures also confirm that imposing robustness in utility has relatively little impact.

Our implementations are in CVX [20] and are solved using Mosek [21].

VI. DISCUSSION

We have given convex formulations of optimization problems for finding robust data release protocols $P_{Y|SU}$. In addition we also studied the non-robust form. Numerical experiments revealed that the non-robust forms achieve privacy levels that are much worse than anticipated. In particular, the naive approach of assuming $P_{SU}^* = \hat{P}_{SU}$ leads to undesirable privacy leakage.

Our convex formulations and corresponding implementations have a number of variables and constraints that grows exponentially in $|\mathcal{S}|$ and $|\mathcal{U}|$. It would be of great interest to develop bounding methods that have reduced complexity, but that still provide strong guarantees.

The current model imposes that the output alphabet of $P_{Y|SU}$ is equal to the input alphabet \mathcal{U} . It would be of interest to leverage this assumption.




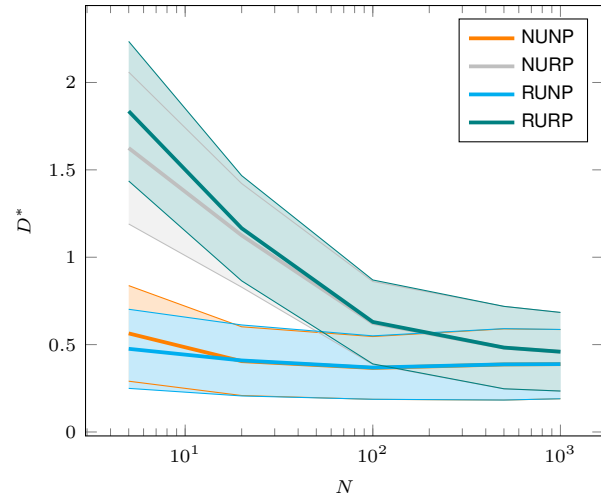
(a) Impact of robustness on utility ($K = 30$, $n = 75$).(b) Impact of robustness on utility ($K = 30$, $n = 15,000$).(c) This figure depicts ϵ^* as a function of $n = N \cdot |\mathcal{S}| |\mathcal{U}|$.(d) This figure depicts D^* as a function of $n = N \cdot |\mathcal{S}| |\mathcal{U}|$.

Fig. 2: Experimental results on K instances for optimization under Nonrobust/Robust Utility and Nonrobust/Robust Privacy. ($\alpha = 0.05$, $\varepsilon = 0.5$, $|\mathcal{S}| = 3$, $|\mathcal{U}| = 5$)

REFERENCES

- [1] D. Rebollo-Monedero, J. Forne, and J. Domingo-Ferrer, "From t -closeness-like privacy to postrandomization via information theory," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 11, pp. 1623–1636, 2010.
- [2] P. Kairouz, S. Oh, and P. Viswanath, "Extremal mechanisms for local differential privacy," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 492–542, 2016.
- [3] A. Makhdoumi, S. Salamatian, N. Fawaz, and M. Médard, "From the information bottleneck to the privacy funnel," in *2014 IEEE Information Theory Workshop (ITW 2014)*. IEEE, 2014, pp. 501–505.
- [4] S. Salamatian, A. Zhang, F. du Pin Calmon, S. Bhamidipati, N. Fawaz, B. Kveton, P. Oliveira, and N. Taft, "Managing your private and public data: Bringing down inference attacks against your privacy," *J. Sel. Topics Signal Processing*, vol. 9, no. 7, pp. 1240–1255, 2015.
- [5] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Information extraction under privacy constraints," *Information*, vol. 7, no. 1, p. 15, 2016.
- [6] S. Kung, "A compressive privacy approach to generalized information bottleneck and privacy funnel problems," *Journal of the Franklin Institute*, vol. 355, no. 4, pp. 1846–1872, 2018.
- [7] N. Ding and P. Sadeghi, "A submodularity-based clustering algorithm for the information bottleneck and privacy funnel," in *2019 IEEE Information Theory Workshop (ITW)*. IEEE, 2019, pp. 1–5.
- [8] S. Salamatian, F. P. Calmon, N. Fawaz, A. Makhdoumi, and M. Médard, "Privacy-utility tradeoff and privacy funnel," 2020, preprint, http://www.mit.edu/~salmansa/files/privacy_TIFS.pdf.
- [9] M. Lopuhaä-Zwakenberg and J. Goseling, "The privacy-utility tradeoff of robust local differential privacy," *arXiv:2101.09139*, 2021, preprint.
- [10] —, "Robust local differential privacy," in *2021 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2021, pp. 557–562.
- [11] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith, "What can we learn privately?" *SIAM Journal on Computing*, vol. 40, no. 3, pp. 793–826, 2011.
- [12] M. Lopuhaä-Zwakenberg, H. Tong, and B. Škorić, "Data sanitisation for the privacy funnel with differential privacy guarantees," *arXiv:2008.13151*, 2020.
- [13] D. Kifer and A. Machanavajjhala, "Pufferfish: A framework for mathematical privacy definitions," *ACM Transactions on Database Systems (TODS)*, vol. 39, no. 1, pp. 1–36, 2014.
- [14] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust optimization*. Princeton university press, 2009.
- [15] A. Ben-Tal, D. Den Hertog, and J.-P. Vial, "Deriving robust counterparts of nonlinear uncertain inequalities," *Mathematical programming*, vol. 149, no. 1-2, pp. 265–299, 2015.

- [16] A. Ben-Tal, D. Den Hertog, A. De Waegenaere, B. Melenberg, and G. Rennen, "Robust solutions of optimization problems affected by uncertain probabilities," *Management Science*, vol. 59, no. 2, pp. 341–357, 2013.
- [17] D. Bertsimas, V. Gupta, and N. Kallus, "Data-driven robust optimization," *Mathematical Programming*, vol. 167, no. 2, pp. 235–292, 2018.
- [18] J. Goseling and M. Lopushaa-Zwakenberg, "Robust optimization for local differential privacy," *arXiv:2205.05015*, 2022, preprint.
- [19] C. P. Robert, N. Chopin, and J. Rousseau, "Harold jeffreys's theory of probability revisited," *Statistical Science*, vol. 24, no. 2, pp. 141–172, 2009.
- [20] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," <http://cvxr.com/cvx>, Mar. 2014.
- [21] MOSEK ApS, *MOSEK Optimization Suite version 9.2*, <https://docs.mosek.com/9.2/intro/index.html>, 2019.