Balancing Utility and Scalability in Metric Differential Privacy

Abstract

Metric differential privacy (mDP) is a modification of differential privacy that is more suitable when records can be represented in a general metric space, such as text data represented as word embeddings or geographical coordinates on a map. We consider the task of releasing elements of the metric space under metric differential privacy where utility is measured as the distance of the released element to the original element. Linear programming (LP) can be used to construct a mechanism that achieves the optimal utility for a particular mDP constraint. However, these LPs suffer from a polynomial explosion of variables and constraints that render them impractical for solving real-world problems. An important question is how to design rigorous mDP mechanisms that balance the utility-scalability tradeoff.

Our main contribution is a new method for reducing the LP size used to generate mDP mechanisms by constraining the search space such that certain input and output pairs have transition probabilities derived from the exponential mechanism. Our method produces mDP mechanisms whose LPs are smaller that all prior work in this area. We also provide a lower bound on the best possible mechanism utility. Our experiments on real-world metric spaces highlight the superior utility-scalability tradeoff of our mechanism.

1 INTRODUCTION

Privacy has emerged as a topic of strategic consequence across all computational fields. Differential privacy (DP), a mathematical formulation of privacy proposed by Dwork et al. [2006], provides provable protection guarantees against adversaries with arbitrary side information and computational power. See the book by Dwork and Roth [2013] for a primer on differential privacy and a survey of different techniques proposed in the literature.

More recently, researchers have noted that differential privacy does not take the underlying metric space of the data domain into account. Differential privacy provides the same level of protection to all perturbations of a single user's data which makes it inflexible when these perturbations are not all the same. For example, if the data consists of locations on earth, there is a large difference between discerning whether a user is in a 1 or a 100-mile radius. In many scenarios, the former type of privacy beach is more significant because a user's location is more accurately determined. This has led to the development of *metric* DP (mDP) which provides different protections depending on an underlying metric space, and has been adopted in applications involving releasing sensitive geolocation data [Andrés et al., 2013, Bordenabe et al., 2014] and textual data [Fernandes et al., 2019, Feyisetan et al., 2019, 2020, Xu et al., 2020, Feyisetan and Kasiviswanathan, 2021].

Mechanism utility for mDP is less well-understood than that of general DP, as the metric strongly influences the permitted behavior of the mechanism. While it is possible to design an optimal mechanism under mDP, it is also a computationally challenging task that requires solving a linear program (LP) with $O(n^2)$ variables and $O(n^3)$ constraints [Bordenabe et al., 2014], where n is the size of the metric space (i.e., cardinality of the set). In fact, most mDP mechanisms [Feyisetan et al., 2020, 2019, Xu et al., 2020] do not provide any rigorous guarantees on the utility. Our key contributions are as follows:

(a) We present a general framework for designing mDP mechanisms which have a better tradeoff between mechanism utility and the size of the LP used to compute the mechanism (Section 3). This framework is based on adding to the optimal mDP LP new constraints that certain transition probabilities are equal to a weighted version of those arising from the *exponential mechanism.*[1] As a concrete instantiation of this framework, we

[1]Exponential mechanism [McSherry and Talwar, 2007] is a popular approach for differentially private selection.

Submitted to the 38th Conference on Uncertainty in Artificial Intelligence (UAI 2022). **To be used for reviewing only**.

construct a LP based on r nearest neighbors of each point, under which the resulting LP has just $O(nr)$ variables and $O(n^2r)$ constraints, and in practice, r can be set as a small constant. Therefore, our new mechanism substantially increases the size of the metric space on which mDP mechanisms can be practically applied.

(b) We prove a lower bound, depending on the underlying metric space, on the loss that any mechanism must make (Section 4). This provides the first non-trivial loss lower bound on any mDP mechanism, including the optimal one. This lower bound is valuable, especially in situations when the LP for the optimal mechanism is infeasible to solve.

(c) We perform extensive experiments comparing the utility and privacy of our proposed mechanism and existing mechanisms in text and geolocation applications (Section 5). These experiments indicate that our proposed mechanism performs more closely to the optimal mechanism than others tested and can result in a utility improvement of about 25% compared to the non-optimal mechanisms. In terms of scalability, our results indicate that our proposed mechanism can scale to metric spaces four times larger than the optimal mechanism.

Related Work in Metric DP. Metric DP originated in the context of location privacy where given a dataset of geolocation coordinates (longitude and latitude) on a plane, the notion of adjacency could be better captured using the Euclidean distance between the coordinates [Andrés et al., 2013, Chatzikokolakis et al., 2013]. Metric DP mechanisms have been investigated for various choices of metrics, including Euclidean, Manhattan, and Chebyshev metrics, among others [Chatzikokolakis et al., 2013, Andrés et al., 2013, Chatzikokolakis et al., 2015, Fernandes et al., 2019, Feyisetan et al., 2019, 2020]. Unlike our focus here, none of these results compare the loss of their proposed mechanisms to the optimal loss.

The most directly related work to ours is that of Bordenabe et al. [2014]. This paper proposes finding the optimal mDP mechanism using linear programming. They propose a method based on spanner graphs to reduce the size of the LP (outlined in Appendix A). A δ-spanner graph is a set of edges between points in a metric space such that the distance between two points in the graph approximates the metric up to a multiplicative factor δ. Bordenabe et al. [2014] use a 3-spanner, for which a construction using just $O(n^{1.5})$ edges exists, to reduce the number of constraints in the LP from $O(n^3)$ to $O(n^{2.5})$.

Related Work in Privately Releasing Text Embeddings. Vector representations of words, sentences, and documents have all become basic building blocks in NLP pipelines and algorithms. Hence, it is natural to consider privacy mechanisms that target these representations in the underlying metric space [Fernandes et al., 2019, Feyisetan et al., 2019, Xu et al., 2020, Feyisetan et al., 2020]. The most relevant result to our setting is the mechanism of Feyisetan et al. [2020] (referred to as *Madlib*). In Section 5, we compare our mechanism to Madlib as a baseline.[2]

2 TECHNICAL PRELIMINARIES

Throughout this paper, we consider data that comes from a finite metric space $(\mathcal{W}, d_{\mathcal{W}})$ where $\mathcal{W}$ is the set of values the data may take. For example, in the text release usecase, $\mathcal{W}$ consists of a vocabulary set, and in the geo-locations usecase, $\mathcal{W}$ consists of a set of locations. The metric $d_{\mathcal{W}} : \mathcal{W} \times \mathcal{W} \to \mathbb{R}$ captures dissimilarity between elements in the set. In NLP applications, it is very common to represent words via a high-dimensional text embedding $\phi : \mathcal{W} \to \mathcal{W}' \subseteq \mathbb{R}^d$. [3] Then we can define the distance between the words as the distance between the embedded words: i.e., for all $w_1, w_2 \in \mathcal{W}$, we define $d_{\mathcal{W}}(w_1, w_2) = d'_{\mathcal{W}}(\phi(w_1), \phi(w_2))$.

2.1 PRIVACY ON METRIC SPACES

Informally, a mechanism $\mathcal{M}$ satisfies metric DP[4] if its behavior is nearly the same on inputs that are close together in the metric space. This is formalized by the following notion of ϵ-$d_{\mathcal{W}}$ privacy.

Definition 1 (Metric DP (mDP))**.** *Given a finite set $\mathcal{W}$, a metric $d_{\mathcal{W}} : \mathcal{W} \times \mathcal{W} \to \mathbb{R}$, and a privacy parameter $\epsilon > 0$, a mechanism $\mathcal{M} : \mathcal{W} \to \mathcal{W}$ satisfies ϵ-$d_{\mathcal{W}}$ privacy if for all $w_1, w_2, w \in \mathcal{W}$:*

$$\mathbb{P}r\left[\mathcal{M}(w_1) = w\right] \leq \exp\left(\epsilon d_{\mathcal{W}}(w_1, w_2)\right) \mathbb{P}r\left[\mathcal{M}(w_2) = w\right].$$

The above definition is closely related to the definition of *local DP* [Kasiviswanathan et al., 2011] in that we apply $\mathcal{M}$ to each element of some database $D \in \mathcal{W}^m$ independently. The difference of mDP over local DP is that, because of the $d_{\mathcal{W}}$ term (which is absent in the local DP formulation), mDP mechanism guarantees indistinguishability for those $w_1, w_2 \in \mathcal{W}$ based on the distance $d_{\mathcal{W}}(\phi(w_1), \phi(w_2))$ between them. Similar to traditional differential privacy, mDP is preserved under post-processing and composition of mechanisms [Koufogiannis et al., 2016]. In metric spaces, a natural definition of mechanism loss on an element $w \in \mathcal{W}$ is the expected distance between w

[2] While our mDP mechanism is applicable to any metric space, our experiments are over word embeddings and geolocations in the Euclidean space. Therefore, we do not directly compare with [Fernandes et al., 2019, Feyisetan et al., 2019, Xu et al., 2020] which work with embeddings in non-Euclidean spaces.

[3] Our results do not depend on the choice of the embedding.

[4] Metric DP is sometimes referred to as Lipschitz privacy [Koufogiannis et al., 2016], motivated by the fact that the privacy guarantee can be viewed as a Lipschitz condition on the mechanism, $|\ln(\mathbb{P}r\left[\mathcal{M}(w_1) = w\right]) - \ln(\mathbb{P}r\left[\mathcal{M}(w_2) = w\right])| \leq \epsilon d_{\mathcal{W}}(w_1, w_2)$.

and $\mathcal{M}(w)$: $\mathcal{L}(\mathcal{M}, w) = \mathbb{E}_{\mathcal{M}}[d_{\mathcal{W}}(w, \mathcal{M}(w))]$. Here, the expectation is over the random bits in $\mathcal{M}$. We define the loss of $\mathcal{M}$ to be the worst-case loss of $\mathcal{M}$ on any particular element $w \in \mathcal{W}$:

$$\mathcal{L}(\mathcal{M}) = \max_{w \in \mathcal{W}} \mathcal{L}(\mathcal{M}, w) \tag{1}$$

Notice that $\mathcal{L}(\mathcal{M})$ is non-negative due to $d_{\mathcal{W}}$ being a metric. Considering the loss as a worst-case instead of an average has the advantage that there cannot exist "adversarial" elements $w \in \mathcal{W}$ such that $\mathcal{L}(\mathcal{M}, w)$ is much higher than $\mathcal{L}(\mathcal{M})$. Similar loss functions have been studied in other DP settings such as in [Hardt and Talwar, 2009].

Optimal Mechanism with LP. It is easy to see that the constraints of mDP are linear. For a mechanism $\mathcal{M}$, we can consider its *stochastic matrix*[5] M given by $M = \{M_{uv} : u, v \in \mathcal{W}\}$ with $M_{uv} = \Pr[\mathcal{M}(u) = v]$. Then, $\mathcal{M}$ satisfies mDP if and only if M is stochastic and satisfies the following constraints

$$M_{uw} \leq M_{vw} \cdot \exp\left(d_{\mathcal{W}}(u, v)\epsilon\right) \qquad \forall u, v, w \in \mathcal{W} \tag{2}$$

Since the constraints are linear, mDP constrains M to be in a polytope. We will overload notation and write $\mathcal{L}(M, w)$ and $\mathcal{L}(M)$ as the losses of the mechanism given by transition matrix M. These losses are given by:

$$\begin{aligned}\mathcal{L}(M) &= \max_{u \in \mathcal{W}} \mathcal{L}(M, w) \\ &= \max_{u \in \mathcal{W}} \sum_{v \in \mathcal{W}} d_{\mathcal{W}}(u, v) M_{uv}. \end{aligned} \tag{3}$$

Over the variables M_{uv}, $\mathcal{L}(\mathcal{M})$ is a maximum of linear functions. The optimal mechanism is given by the stochastic matrix M that minimizes $\mathcal{L}(\mathcal{M})$ subject to the privacy constraints (2). Using standard techniques in linear programming, we can compute the best mechanism with the following LP over the variables M, k:

$$\begin{aligned}\mathcal{P}_{\text{OPTMECH}}(\epsilon) \ &= \textbf{minimize } k \textbf{ subject to} \\ &\quad \mathcal{L}(M, w) \leq k, \ \ \forall w \in \mathcal{W} \\ &\quad M \text{ stochastic} \\ &\quad M \text{ satisfies (2)}\end{aligned}$$

This LP problem has $O(n^2)$ variables and $O(n^3)$ constraints, where $n = |\mathcal{W}|$. Therefore, even with the state-of-the-art LP approaches, which all require $\Omega(N^2)$ time, where N is the number of variables [Jiang et al., 2021], scalability is problematic (here $N = n^2$). This is the central motivation for our work.

3 BALANCING UTILITY-SCALABILITY

Given the scalability issues in solving $\mathcal{P}_{\text{OPTMECH}}(\epsilon)$, a natural idea is to reduce the LP size. In this section, we present a new method to reduce the size of the LP in $\mathcal{P}_{\text{OPTMECH}}$ while still maintaining the mDP guarantee (Definition 1). Our method is based on adding *exponential mechanism* (EXPMECH) [McSherry and Talwar, 2007] equality constraints to the LP. Before we do this, we make use of an observation that the EXPMECH is provably not optimal in mDP which may be of independent interest. Thus, the constraints we add come from a "weighted" version of the EXPMECH. All missing proofs are collected in Appendix B.

3.1 IMPROVING THE EXPMECH IN MDP

Informally, the exponential mechanism [McSherry and Talwar, 2007] is a method for deferentially private selection from a discrete set of candidate outputs. Due to its flexibility the EXPMECH has become a popular tool for designing DP mechanisms. Furthermore, the EXPMECH is known to be optimal in DP for many choices of utility function [Hardt and Talwar, 2009, Aldà and Simon, 2017].

However, in the mDP setting, the exponential mechanism can be fooled by outlier elements. Informally, the dense areas of the metric space can act as a "black hole" where the EXPMECH will output elements in the dense area with high probability, even for the outlier elements. This drives up the loss for the outlier elements.

For a metric space $\mathcal{W} = \{w_1, \ldots, w_n\}$ and metric $d_{\mathcal{W}}$, the EXPMECH has the following transition probability

$$\mathbb{P}r\left[\text{EXPMECH}(w_i) = w_j\right] = \frac{e^{-\epsilon d_{\mathcal{W}}(w_i, w_j)/2}}{\sum_{k=1}^{n} e^{-\epsilon d_{\mathcal{W}}(w_i, w_k)/2}}.$$

To illustrate on a concrete example, consider the metric space where $\mathcal{W} = \{w_1, \ldots, w_n\}$ and where $d_{\mathcal{W}}$ satisfies (1) $d_{\mathcal{W}}(w_1, w_i) = 1$ for $i \geq 2$, and (2) $d_{\mathcal{W}}(w_i, w_j) < \delta$ when $i, j \geq 2$, where δ is a small constant. For each $j \geq 2$, the EXPMECH satisfies $\Pr[\text{EXPMECH}(w_1) = w_j] = \frac{e^{-\epsilon}}{1+ne^{-\epsilon}}$, and thus $\Pr[\text{EXPMECH}(w_1) \neq w_1] = \frac{ne^{-\epsilon}}{1+ne^{-\epsilon}}$. As n grows, the probability of this occurring approaches 1, and the loss does as well. The elements w_i for $i \geq 2$ are acting as a "black hole".

This can be fixed by considering a more general mechanism that assigns weights to the output probabilities. For positive weights $\mathbf{Y} = (Y_1, \ldots, Y_n) \in (\mathbb{R}^+)^n$, consider the more general mechanism given by

$$\mathbb{P}r\left[\text{EXPMECH}_{\mathbf{Y}}(w_i) = w_j\right] = \frac{Y_j e^{-\epsilon d_{\mathcal{W}}(w_i, w_j)/2}}{\sum_{k=1}^{n} Y_k e^{-\epsilon d_{\mathcal{W}}(w_i, w_k)/2}}.$$

This mechanism can be shown to satisfy ϵ-$d_{\mathcal{W}}$ privacy.

Proposition 1. *For any metric space and any* $\mathbf{Y} \in (\mathbb{R}^+)^n$, *the mechanism* $\text{EXPMECH}_{\mathbf{Y}}$ *satisfies* ϵ-$d_{\mathcal{W}}$ *privacy.*

[5] A stochastic matrix is a square matrix whose rows are probability vectors.

In our example, to avoid the problem encountered by the regular EXPMECH, we can weight w_1 higher than $w_2, \ldots, w_n$. For example, there exists a weighting $\mathbf{Y}$ such that the following loss is possible:

Lemma 1. *With $\mathcal{W} = \{w_1, \ldots, w_n\}$ and $d_{\mathcal{W}}$ defined as above, when $\mathbf{Y} = (1, 1/(n-1), 1/(n-1), \ldots)$, we have that*

$$\mathcal{L}(\text{EXPMECH}_{\mathbf{Y}}) \leq \frac{\frac{1}{n-1} + e^{-\epsilon/2}}{1 + e^{-\epsilon/2}} \mathcal{L}(\text{EXPMECH}).$$

When $\epsilon \geq 2 \log n$, we have $\mathcal{L}(\text{EXPMECH}_{\mathbf{Y}}) \leq 2/(n-1)\mathcal{L}(\text{EXPMECH})$.

This establishes that the EXPMECH is provably not optimal on our example metric space. However, one problem with the more general $\text{EXPMECH}_{\mathbf{Y}}$ is that it is not clear how to set $\mathbf{Y}$ to optimize the loss other than the rule of thumb that dense elements should be weighted less. In the next section, we leave it to the LP solver to optimize these weights.

3.2 BALANCING LP LOSS AND SCALABILITY

To reduce the number of LP constraints required to find the optimal mechanism, our key idea is to add equality constraints in such a way that many of the original constraints in $\mathcal{P}_{\text{OPTMECH}}$ are trivially satisfied. This results in potentially a much smaller LP; however, optimality is no longer guaranteed. The balance between optimality and LP size is decided by the number of equality constraints. In fact, we will develop a general framework for balancing this tradeoff, which we call ConstOPTMech (Algorithm 1).

Specifically, to obtain the LP describing ConstOPTMech, we start with $\mathcal{P}_{\text{OPTMECH}}(\epsilon)$ and add non-negative variables $\{Y_w : w \in \mathcal{W}\}$. Then, for certain variables M_{uv}, we add additional "weighted exponential mechanism-like" constraints: $M_{uv} = Y_v e^{-\epsilon d_{\mathcal{W}}(u,v)}$. We leave the weights Y_v to be optimized by the LP solver.

We allow deviations from the additional constraints in the form of a replacement function $I(v) : \mathcal{W} \to 2^{\mathcal{W}}$ that returns the elements $u \in \mathcal{W}$ for which the weighted exponential mechanism should not be used to set M_{uv}. To encode this, we add the following constraints:

$$M_{uv} = Y_v e^{-\epsilon d_{\mathcal{W}}(u,v)} \ \ \forall u, v \in \mathcal{W}; u \notin I(v) \tag{4}$$

The replacement function $I(v)$ indicates where we do not want to use the exponential mechanism, and there are many candidates. We will later consider the following instantiation of this replacement function

$$I_{\text{NN},r}(v) = \{u \in \mathcal{W} | v \text{ is an } r\text{-nearest neighbor of } u.\}, \tag{5}$$

which returns the elements u such that v is one of the r nearest neighbors of u. We employ this function $I_{\text{NN},r}$ because the exponential mechanism already assigns exponentially-low probabilities of returning the farthest elements from a given element, and we conjecture there is not much improvement to be made for such scenarios.

Adding the constraints (4) to $\mathcal{P}_{\text{OPTMECH}}(\epsilon)$ will satisfy the privacy constraints (2), but it may be impossible for M to be stochastic. One can see this in the extreme example by setting $I(v) = \emptyset$; then (4) holds for all $u, v \in \mathcal{W}$, and non-negative assignment to Y_v that makes M stochastic need not exist. To fix this, we relax the constraint that M be stochastic, and only insist that its rows sum to values more than 1. We add a penalization term involving a constant $\lambda > 0$ to the loss of each element to penalize the extent to which the rows sum to more than 1. As such, our loss function now takes the following form:

$$\tilde{\mathcal{L}}(M, w) = \sum_{u \in \mathcal{W}} M_{wu} d_{\mathcal{W}}(u, w) + \lambda \sum_{u \in \mathcal{W}} M_{wu} \tag{6}$$

With the modified loss function and relaxed stochasticity requirement, we obtain the LP giving ConstOPTMech.

$$\begin{aligned}
\mathcal{P}_{\text{CONSTOPTMECH}}(\epsilon) = \textbf{minimize } k \textbf{ subject to} \\
& \tilde{\mathcal{L}}(M, w) \leq k \ \ \forall w \in \mathcal{W} \\
& M_{uv} \geq 0, \ \sum_{v \in \mathcal{W}} M_{uv} \geq 1 \ \ \forall u, v \in \mathcal{W} \\
& M \text{ satisfies (2) and (4)} \\
& Y_v \geq 0 \ \ \forall v \in \mathcal{W}
\end{aligned}$$

Notice that this LP is always feasible because one valid solution is $M_{uv} = Y_v e^{-\epsilon d_{\mathcal{W}}(u,v)}$ for all u, v—since there are no restrictions on the Y_v variables, we can set them high enough so that $\sum_{v \in \mathcal{W}} M_{uv} \geq 1$.

The benefit of the equality constraints (4) is that we can drop a large number of the constraints in (2), as they are trivially satisfied. This allows us to find a solution to $\mathcal{P}_{\text{CONSTOPTMECH}}(\epsilon)$ much faster than $\mathcal{P}_{\text{OPTMECH}}(\epsilon)$.

Theorem 1. *$\mathcal{P}_{\text{CONSTOPTMECH}}$ is feasible, and it is possible to solve it using a linear program with $n + 1 + \sum_{v \in \mathcal{W}} |I(v)|$ variables and $2n + \sum_{v \in \mathcal{W}} 2|I(v)|^2 + 3|I(v)|$ constraints. The number of non-zero coefficients in the LP is at most $2n^2 + \sum_{v \in \mathcal{W}} 2|I(v)|^2 + 5|I(v)|$.*

Our choice to drop the stochasticity requirement of M gave us feasibility of $\mathcal{P}_{\text{CONSTOPTMECH}}$, but the solution M is no longer a mechanism because it is not stochastic. We obtain ConstOPTMech by normalizing the solution to $\mathcal{P}_{\text{CONSTOPTMECH}}(\epsilon)$. Furthermore, any choice of λ gives rise to a valid mechanism.

Mechanism CONSTOPTMECH uses half of the privacy budget for solving $\mathcal{P}_{\text{CONSTOPTMECH}}(\frac{\epsilon}{2})$ because normalization may increase the privacy guarantee by a factor of 2. We

Algorithm 1: Mechanism ConstOPTMech

Data: Universe $\mathcal{W}$, metric $d_{\mathcal{W}}$, budget ϵ, replacement function $I(v)$, $\lambda \in \mathbb{R}^+$.
Result: Transition matrix H.
$M, \text{loss} \leftarrow \texttt{Solve}(\mathcal{P}_{ConstOPTMech}(\frac{\epsilon}{2}))$ with λ;
for $u, v \in \mathcal{W}$ **do**
  $H_{uv} \leftarrow \frac{M_{uv}}{\sum_{w \in \mathcal{W}} M_{uw}}$;
return H

are able to show the following privacy guarantee on CONSTOPTMECH. The proof is a generalization of Proposition 1.

Theorem 2. *For any set $\mathcal{W}$, metric $d_{\mathcal{W}}$, and budget ϵ and replacement function $I(v) : \mathcal{W} \to 2^{\mathcal{W}}$, Mechanism* CONSTOPTMECH *satisfies ϵ-$d_{\mathcal{W}}$ privacy.*

Setting the Replacement Function. In Theorem 1 when the replacement function $I(v) = I_{\text{NN},r}(v)$, then the number of variables in $\mathcal{P}_{\text{ConstOPTMech}}$ is at most $nr + n + 1$, the number of constraints is at most $n^2r + 3nr + 2n$, and the number of non-zero elements is at most $2n^2 + 5nr + 2n^2r$ (see Corollary 1).

When we apply with $r = n - 1$, we add no equality constraints, and (4) becomes equivalent to (2). The LP size of $\mathcal{P}_{\text{ConstOPTMech}}$ is $O(n^3)$, the same as that of $\mathcal{P}_{\text{OPTMech}}$. When r is a constant much less than n, then the number of variables is $O(nr)$ and the number of constraints is $O(n^2r)$, each term saving a factor of n. We note that $O(n^2r)$ is a worst-case bound that is not always tight. If one assumes that there is no $v \in \mathcal{W}$ such that at least $10r$ other elements of $\mathcal{W}$ count v as an r-nearest neighbor, then $\sum_{v \in \mathcal{W}} |I(v)|^2 + |I(v)| \leq \sum_{v \in \mathcal{W}} 110r^2 \leq O(nr^2)$.

In Table 1, we compare the number of variables, constraints, and non-zero coefficients arising in $\mathcal{P}_{\text{ConstOPTMech}}$ as compared to those in the optimal mechanism ($\mathcal{P}_{\text{OPTMech}}$) and the LP based of spanner graphs [Bordenabe et al., 2014] (referred to as $\mathcal{P}_{\text{SpannerMech}}$, see Appendix A). We see that $\mathcal{P}_{\text{ConstOPTMech}}$ improves on all three of these quantities compared to the other two LPs when $r \ll n$. However, we note that these are worst-case upper bounds, and in practice the LP complexity measures may smaller. Based on setting of r, we have a tradeoff between scalability and increase to the loss compared to the optimal mechanism. We perform an empirical analysis of this tradeoff in Section 5.

4 LOWER BOUNDS

In this section, we propose an easy to compute lower bound for mechanism loss. Our lower bound builds on the intuition that for an element $w \in \mathcal{W}$, if there are many elements that are far, but not too far, from w, then mDP forces the distribution $\mathcal{M}(w)$ to place significant mass on the elements which are farther away. This gives a lower bound on the loss of $\mathcal{M}$. To make this intuition formal, we define a packing of $\mathcal{W}$ to be a set of elements which are at least a certain distance from each other. In the following, let $B(x, r)$ denote the elements $y \in \mathcal{W}$ such that $d_{\mathcal{W}}(x, y) \leq r$.

Definition 2. *Let $\mathcal{W}$ be a set. A finite set $S \subseteq \mathcal{W}$ is called a (c, r, Q)-packing w.r.t. metric $d_{\mathcal{W}}$ if the following hold: $|S| = c$; for all $x, x' \in S$, $B(x, r) \cap B(x', r) = \emptyset$, i.e. the balls around the elements in S of radius r are disjoint; and for all $x, x' \in S$, $d_{\mathcal{W}}(x, x') \leq Q$, i.e. the maximum distance between any two elements in S is at most Q.*

The lower bound we derive holds for any (c, r, Q)-packing of the metric space $\mathcal{W}$. The catch is that if a packing with a small r or c is used, the bound will not be strong. Our lower bound involves the quantity $N(w, S)$ that depends on a $w \in \mathcal{W}$ and a (c, r, Q)-packing S. $N(w, S)$ is given by:

$$N(w, S) = \sum_{s \in S} \exp(-d_{\mathcal{W}}(w, s)\epsilon),$$

We also have the lower bound $N(w, S) \geq 1 + (c - 1)\exp(-Q\epsilon)$ which follows because S is a (c, r, Q)-packing. Notice that $N(w, S) \geq 1$ (because $w \in S$ and $d_{\mathcal{W}}(w, w) = 0$), and $N(w, S)$ grows linearly with the number of elements in S and grows exponentially when the elements in S are closer to w or when ϵ decreases. This represents the increasing amount of mass that must be placed on these elements according to mDP. Our lower bound will grow stronger with increasing $N(w, S)$. The lower bound is as follows (proof in Appendix B).

Theorem 3. *Consider an arbitrary set $\mathcal{W}$ and metric $d_{\mathcal{W}} : \mathcal{W} \times \mathcal{W} \to \mathbb{R}$. Then, for any $\epsilon > 0$, any mechanism $\mathcal{M}$ satisfying ϵ-$d_{\mathcal{W}}$ privacy and any (c, r, Q)-packing S of $\mathcal{W}$, it holds that*

$$\mathcal{L}(\mathcal{M}) \geq \max_{w \in \mathcal{W}} r \left(1 - \frac{1}{N(w, S)}\right). \tag{7}$$

It follows that

$$\mathcal{L}(\mathcal{M}) \geq r \left(1 - \frac{1}{1 + (c - 1)\exp(-Q\epsilon)}\right). \tag{8}$$

Both (7) and (8) have a simple interpretation. The r term represents the minimum loss that must be incurred when a mechanism returns an element that is in a different ball than the starting ball in the packing. The r term is multiplied by $P = 1 - \frac{1}{N(w,S)}$, which can be interpreted as a probability since it is between 0 and 1 (since $N(w, S) \geq 1$). As we show in the theorem proof, P is a lower bound on the probability that the mechanism returns an element in a different ball from w and thus incurs the error r. P increases with $N(w, S)$, which depends on the packing in the ways we identified above. P is small only when $N(w, S)$ approaches 1, and using the bound $N(w, S) \geq 1 + (c - 1)\exp(-Q\epsilon)$, we see the central term controlling its closeness to 1 is

	# Variables	# Constraints	# Non-Zeroes
Optimal LP: $\mathcal{P}_{\text{OPTMECH}}$	$O(n^2)$	$O(n^3)$	$O(n^3)$
Spanner-based LP: $\mathcal{P}_{\text{SPANNERMECH}}$ [Bordenabe et al., 2014]	$O(n^2)$	$O(n^{2.5})$	$O(n^{2.5})$
Our Method: $\mathcal{P}_{\text{CONSTOPTMECH}}$ with function $I_{\text{NN},r}$	$O(nr)$	$O(n^2 r)$	$O(n^2 r)$

Table 1: Comparison of the number of variables and constraints for the various LP-based methods achieving mDP. Note that $\mathcal{P}_{\text{CONSTOPTMECH}}$ improves on existing methods when $r \ll n$.

$(c-1)\exp(-Q\epsilon)$. Here the parameter Q crucially comes into play because if the elements in the packing are too far apart, then mDP is a weak privacy guarantee, $N(w, S)$ will approach 1, and the lower bound will weaken.

An important special case of our theorem occurs when we take S to be the two farthest elements $w^1_{max}, w^2_{max} \in \mathcal{W}$. In this case, S is a $(2, \frac{r^*}{2}, r^*)$-packing where $r^* = d_{\mathcal{W}}(w^1_{max}, w^2_{max})$. Our lower bound then reads $\mathcal{L}(\mathcal{M}) \geq \frac{r^*}{2}\left(\frac{\exp(-r^*\epsilon)}{1+\exp(-r^*\epsilon)}\right)$.

5 EXPERIMENTAL RESULTS

We investigate through experiments how the loss of our proposed mechanism, ConstOPTMech, compares to other state-of-the-art mDP mechanisms.[6] We also include comparisons to our loss lower bound (derived in Section 4). Furthermore, we perform studies to compare ConstOPTMech to SpannerMech, as it is the most directly related work. To do this, we experimentally evaluate the complexity of solving the LPs used to compute both mechanisms. We focus on text embeddings and geolocation metric spaces because, as noted in Section 1, mDP mechanisms have primarily been used for privately releasing text and location data.

5.1 EXPERIMENTAL SETUP

Our experiments consist of generating metric spaces in both application domains and then running two types of experiments. The first evaluates privacy vs. loss on a fixed metric space. The second evaluates scalability as the size of the metric space grows.

We measure utility (loss) of a mechanism based on (1). Since this loss is agnostic to any downstream modeling task performed on these private releases, we do not focus on any specific downstream task.

Metric Space Generation. To produce a metric space of a specific size, we sample metric spaces differently depending on the application.

For text embeddings, we sample from a base metric space consisting of the set of English words and the metric $d_{\mathcal{W}}$ induced by a text embedding $\phi : \mathcal{W} \to \mathbb{R}^d$. Precisely, we used $d_{\mathcal{W}}(u, v) = d(\phi(u), \phi(v))$, where d is the Euclidean distance. We used both the FastText [Bojanowski et al., 2017] and the GloVe embedding [Pennington et al., 2014] for our embedding ϕ. To sample a metric space from a base metric space, we selected a subset $\mathcal{W}'$ of English words. Instead of selecting $\mathcal{W}'$ at random, which would likely produce a set of completely unrelated words with roughly the same distance between each pair, we used a clustered approach. First, we let $\mathcal{W}'$ consist of one random word. To sample another word, we add a random word to $\mathcal{W}'$ with 50% chance. Otherwise we select a random word w in $\mathcal{W}'$ and add one of w's 50 closest English words according to $d_{\mathcal{W}}$. This allows us to produce samples of the larger metric space that have representative clusters of words. We repeat this process until the metric space has the desired size.

For geolocation application, we used the method of Bordenabe et al. [2014], which uses the Geolife [Zheng et al., 2010] dataset. The Geolife dataset consists of 17621 location traces clustered around Beijing. We divide Beijing into rectangular regions of $0.005°$ (about 0.6 km) in width and height. For each trace, we consider its top 30 regions, and we form a histogram of top regions across all traces. To form a metric space of size n, we take the n most popular regions in the histogram. Our metric is the Euclidean distance between the centers of the regions.

Performance Benchmarks. We use the following benchmarks to measure mechanism privacy, loss, and scalability.

Privacy: In practice it is usually acceptable to use (ϵ, δ)-DP for some small δ. We adopt (ϵ, δ)-mDP for our experiments, as we do not want to penalize an algorithm for having some small probability of two elements u, v being distinguished. We say M satisfies (ϵ, δ)-mDP if for any $u, v \in \mathcal{W}$ and $S \subseteq \mathcal{W}$, we have

$$\Pr[M(u) \in S] \leq e^{\epsilon d_{\mathcal{W}}(u,v)} \Pr[M(v) \in S] + \delta. \quad (9)$$

For a fixed δ, we let ϵ_{tight} be the smallest ϵ such that M satisfies (ϵ, δ)-MDP:

$$\epsilon_{\text{tight}}(M) = \inf_{\epsilon \geq 0} M \text{ satisfies } (\epsilon, \delta)\text{-mDP} \quad (10)$$

In our experiments, we set $\delta = 0.001$.

Loss: For practical considerations, we use a more robust measurement of loss in our experiments, where it may not be problematic if the mechanism performs poorly on a small fraction of elements. Instead of using the maximum loss over all elements $\mathcal{L}(M)$ (1), we use the qth-quantile over

[6] In this section, we use ConstOPTMech to denote Algorithm 1 invoked with replacement function $I_{\text{NN},r}$.

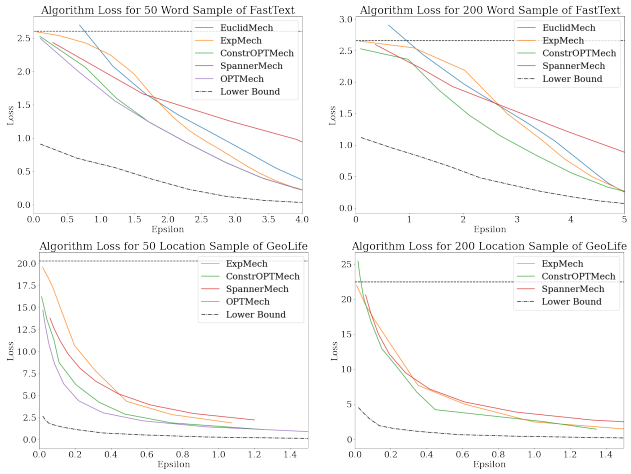

Figure 1: Loss of Madlib (EuclidMech), ExpMech, ConstOPTMech, SpannerMech, and OPTMech versus ϵ_{tight} on 50 and 200-size metric spaces generated from FastText and Geolife, along with the lower bound. The horizontal line indicates the loss of returning a uniform random element.

the set $\{\mathcal{L}(M, w) : w \in \mathcal{W}\}$:

$$\mathcal{L}_q(M) = \text{quantile}_q(\{\mathcal{L}(M, w) : w \in \mathcal{W}\}) \tag{11}$$

This loss estimate allows mechanisms to perform poorly on a small subset of the metric space, which in practice may be outlier or noisy data. In all experiments, we use $q = 95\%$ so that mechanisms are evaluated based on their losses on the best 95% of elements.

LP Scalability: To measure scalability of our mechanisms, we measured the time and number of nonzero coefficients (NNZ) used in the LPs. We use the number of nonzero coefficients over the number of variables or constraints since LP solvers tend to be optimized toward solving sparse LPs. We also consider computation time to be an important measure, as it captures complexity beyond the NNZ. For mechanisms that do not require linear programs, the computational requirements are trivial and we do not test them.

Specific Details for Each Mechanism. We tested five mechanisms: a) Our proposed ConstOPTMech, b) OPTMech (based on solving $\mathcal{P}_{\text{OPTMech}}$), c) SpannerMech [Bordenabe et al., 2014], d) Madlib mechanism [Feyisetan et al., 2020], and e) ExpMech [McSherry and Talwar, 2007]. Madlib, ExpMech, and OPTMech have no further parameters other than ϵ. For SpannerMech, we implement the algorithm as it is described in Bordenabe et al. [2014].

Mechanism ConstOPTMech takes λ and $I(v)$ as parameters (see Algorithm 1). We optimize over λ with possible values in $\{0.001, 0.1, 1.0\}$. For the replacement function $I(v)$, we use $I_{\text{NN},r}(v)$ (5). We try both $r = 5$ and $r = 10$, and we will designate these values in our results.

Evaluating Lower Bound. For each metric space, we computed our lower bound according to Theorem 3. This theorem produces a lower bound for any (c, r, Q) packing of $\mathcal{W}$ and any ϵ. However, it is infeasible to try out every possible (c, r, Q), packing. Instead, we generated candidate (c, r, Q)-packings using a k-center algorithm, using values of k that varied from 1 to the size of the metric space. For each value of ϵ that we tested, we used the strongest lower bound given by one of our generated (c, r, Q)-packings.

Experimental Outline. The first experiments we conducted are *utility experiments*. We test which mechanisms are better at minimizing loss, subject to privacy constraints. To do this, we plot $\mathcal{L}_q$ versus ϵ_{tight} for a metric space consisting of 50 and 200 elements generated from FastText, GloVe, or Geolife. We also plot the lower bound. For ConstOPTMech, we use $r = 10$. We do not run OPTMech for metric spaces of size 200, as the number of constraints would be 200^3 which is too large. We do not include comparisons to Madlib for the Geolife dataset, as it is designed for text embeddings.

Next, we conduct *scalability experiments* on the mechanisms that involve solving LPs. We do this by fixing a privacy constraint ϵ_{tight} and, for each mechanism, testing the NNZ and time taken as the size of the metric space grows. We increased the number of samples in the metric space starting at 50 and increasing in increments of 50 until we reached 400 elements or the time spent solving the LP exceeded 1800 seconds. We fix $\epsilon_{\text{tight}} = 2.0$ for metric spaces sampled from FastText, at 1.0 for those sampled from GloVe, and at 0.3 for those sampled from Geolife.

5.2 RESULTS

We discuss the experimental results for metric spaces generated from FastText and Geolife. The results for GloVe are similar, and they appear in Appendix C.

Utility Experiments: Plots appear in Figure 1. In all tests, ConstOPTMech has lower loss than all other non-optimal mechanisms at all values of ϵ. This includes high privacy regimes, where the loss is near the loss of returning a uniformly random element, and low privacy regimes, where the loss approaches zero. The most pronounced improvement in loss occurs in middle ranges of ϵ (about $[1.0, 3.5]$ for FastText and about $[0.3, 0.8]$ for Geolife) where ConstOPTMech offers an improvement of about 15-30% over all other non-optimal mechanisms. For example, when $\epsilon = 3$ on the 200-size FastText sample, ConstOPTMech offers a loss of 1.0, while the next-best mechanism, the ExpMech, offers a loss of 1.3. This represents a 23% reduction. On the metric space sampled from Geolife of size 50, the loss reduction is as high as 50% at $\epsilon = 0.3$. The middle ranges of ϵ where ConstOPTMech is superior are the values with the most practical importance, since at these ranges, the losses are far from the random baseline yet are still nonzero—the mechanisms are offering both utility and privacy.

On the sampled metric spaces of size 50, ConstOPTMech attains only slightly worse loss than OPTMech, the optimal mechanism. As ϵ grows past 1.5 for FastText and 0.7 for

FastText	50 elements	150 elements	300 elements	400 elements
ConstOPTMech (10)	1.89 sec, 3.23e4 nnz	24.91 sec, 1.47e5 nnz	163.75 sec, 3.79e5 nnz	234.63 sec, 5.71e5 nnz
ConstOPTMech (5)	0.47 sec, 1.21e4 nnz	12.85 sec, 6.60e4 nnz	65.56 sec, 2.19e5 nnz	141.90 sec, 3.70e5 nnz
SpannerMech	1.25 sec, 2.12e4 nnz	62.16 sec, 2.56e5 nnz	1001.82 sec, 1.19e6 nnz	$>$ 1800 sec, — nnz
OPTMech	8.85 sec, 2.50e5 nnz	$>$ 1800 sec, — nnz		
Geolife	**50 elements**	**150 elements**	**300 elements**	**400 elements**
ConstOPTMech (10)	1.41 sec, 1.74e4 nnz	13.85 sec, 8.05e4 nnz	80.24 sec, 2.51e5 nnz	169.52 sec, 4.13e5 nnz
ConstOPTMech (5)	0.40 sec, 8.30e3 nnz	12.53 sec, 5.47e4 nnz	66.28 sec, 1.99e5 nnz	145.74 sec, 3.46e5 nnz
SpannerMech	0.95 sec, 1.56e4 nnz	40.38 sec, 1.49e5 nnz	928.72 sec, 6.05e5 nnz	$>$ 1800 sec, — nnz
OPTMech	25.39 sec, 2.50e5 nnz	$>$ 1800 sec, — nnz		

Table 2: Computation times and memory requirements for computing ConstOPTMech when $r = 10, 5$; SpannerMech; and OPTMech, for varying metric space sizes generated from FastText and Geolife datasets. Each mechanism satisfies $\epsilon_{\text{tight}} = 2.0$ (FastText) and 0.3 (Geolife).

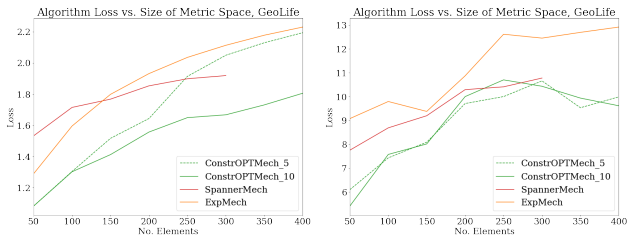

Figure 2: Loss of ExpMech, ConstOPTMech (with $r = 5, 10$), and SpannerMech versus size of metric space for metric spaces generated from FastText and Geolife. Here, ϵ_{tight} is fixed at 2.0 (FastText) and 0.3 (Geolife).

Geolife, their losses become virtually the same. Because we are using $r = 10$, this means just $50 \times 10 = 500$ entries out of the 2500 entries in the transition matrix are not fixed. This suggests that the 10 nearest neighbors to an element play the largest role in minimizing the element's loss.

In all scenarios tested, there is a large gap between the lower bound and the losses of the mechanisms, even the optimal mechanism. Hence, it is uncertain how close ConstOPTMech is to OPTMech on the metric spaces of size 200.

Scalability Experiments: We were able to run ConstOPTMech until 400 elements, whereas OPTMech timed out at 100 elements and SpannerMech timed out at 350 elements. Table 2 shows some of the time and NNZ data for the mechanisms. These results indicate that computing ConstOPTMech is faster than computing SpannerMech. This is particularly evident for the metric spaces with size 150 (resp. 300), where ConstOPTMech with $r = 10$ uses at most 40% (resp. 16%) as much time as SpannerMech, and ConstOPTMech with $r = 5$ uses at most 20% (resp. 6.5%) as much time. These faster times come despite our optimization of λ in ConstOPTMech, which requires solving three LPs. In other words, the actual time to solve one LP used in ConstOPTMech is one third as high as the reported times.

In terms of NNZ, ConstOPTMech with $r = 10$ uses at most 57% (resp. 32%) as much time as SpannerMech, and ConstOPTMech with $r = 5$ uses at most 37% (resp 18%) as many non-zero coefficients on the metric spaces with size $150, 300$. Note that one of the reasons why these savings are less than that observed for the time improvements is because ConstOPTMech uses LPs which are simpler than the LPs used by SpannerMech, which for example have more variables ($O(n^2)$ variables versus $O(nr)$).

All the previous mechanisms have improved performance over OPTMech, which uses moderate time and NNZ on metric spaces with 50 elements and does not scale to 100 elements and beyond.

In Figure 2, we show our plots of loss versus the size of the metric space when ϵ_{tight} is fixed. The plots confirms that ConstOPTMech also has lower loss than SpannerMech. At metric space sizes below 200, ConstOPTMech with $r = 10$ has approximately a 22% (resp. 12%) reduction in loss compared to SpannerMech on FastText and Geolife, though this reduces to about 10% on FastText (resp. $< 5\%$ on Geolife) for metric space sizes greater than 200. ConstOPTMech with $r = 5$ similarly outperforms SpannerMech on smaller metric spaces, though on FastText, ConstOPTMech with $r = 5$ performs worse than SpannerMech, a symptom of it using too few nearest neighbors. On large metric spaces, more nearest neighbors must be used to maintain low loss.

6 CONCLUSION

We tackle the problem of designing scalable metric differential privacy mechanisms that achieve near optimal utility. Our new mechanism combines the optimal LP-based mechanism and the exponential mechanism to achieve a better utility-scalability tradeoff than existing mechanisms. We also provide a simple to compute lower bound that improves our understanding of the optimal utility. Our experiments show that our mechanism is computationally tractable on larger metric spaces while also almost matching the utility of the optimal LP-based mechanism. While our mechanism operates on any metric space, an interesting question is whether the geometry of the metric space can be leveraged to improve either utility or scalability.

References

Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen Kobourov, and Richard Spence. Graph Spanners: A Tutorial Review. *arXiv:1909.03152 [cs, math]*, March 2020. URL http://arxiv.org/abs/1909.03152. arXiv: 1909.03152.

Francesco Aldà and Hans Ulrich Simon. On the Optimality of the Exponential Mechanism. In Shlomi Dolev and Sachin Lodha, editors, *Cyber Security Cryptography and Machine Learning*, volume 10332, pages 68–85. Springer International Publishing, Cham, 2017. ISBN 978-3-319-60079-6 978-3-319-60080-2. doi: 10.1007/978-3-319-60080-2_5. URL http://link.springer.com/10.1007/978-3-319-60080-2_5. Series Title: Lecture Notes in Computer Science.

Miguel E Andrés, Nicolás E Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 901–914, 2013.

Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *TACL*, 5, 2017.

Nicolás E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Optimal Geo-Indistinguishable Mechanisms for Location Privacy. *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 251–262, November 2014. doi: 10.1145/2660267.2660345. URL http://arxiv.org/abs/1402.5029. arXiv: 1402.5029.

Konstantinos Chatzikokolakis, Miguel E Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi. Broadening the scope of differential privacy using metrics. In *PETS*, 2013.

Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Marco Stronati. Constructing elastic distinguishability metrics for location privacy. *PETS*, 2015.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4), 2013.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, pages 265–284. Springer, 2006.

Natasha Fernandes, Mark Dras, and Annabelle McIver. Generalised differential privacy for text document processing. *Principles of Security and Trust*, 2019.

Oluwaseyi Feyisetan and Shiva Kasiviswanathan. Private release of text embedding vectors. In *Proceedings of the First Workshop on Trustworthy Natural Language Processing*, pages 15–27, 2021.

Oluwaseyi Feyisetan, Tom Diethe, and Thomas Drake. Leveraging hierarchical representations for preserving privacy and utility in text. In *IEEE ICDM*, 2019.

Oluwaseyi Feyisetan, Borja Balle, Thomas Drake, and Tom Diethe. Privacy-and utility-preserving textual analysis via calibrated multivariate perturbations. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, pages 178–186, 2020.

Moritz Hardt and Kunal Talwar. On the Geometry of Differential Privacy. *arXiv:0907.3754 [cs]*, November 2009. URL http://arxiv.org/abs/0907.3754. arXiv: 0907.3754.

Shunhua Jiang, Zhao Song, Omri Weinstein, and Hengjie Zhang. A faster algorithm for solving general lps. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 823–832, 2021.

Shiva Kasiviswanathan, Homin Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3), 2011.

Fragkiskos Koufogiannis, Shuo Han, and George J Pappas. Gradual release of sensitive data under differential privacy. *Journal of Privacy and Confidentiality*, 7(2), 2016.

Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *FOCS*, volume 7, pages 94–103, 2007.

Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543, 2014.

Zekun Xu, Abhinav Aggarwal, Oluwaseyi Feyisetan, and Nathanael Teissier. A differentially private text perturbation method using regularized mahalanobis metric. In *Proceedings of the Second Workshop on Privacy in NLP at EMNLP 2020*, pages 7–17, 2020.

Yu Zheng, Xing Xie, and Wei-Ying Ma. Geolife: A collaborative social networking service among user, location and trajectory. *IEEE Data(base) Engineering Bulletin*, June 2010. URL https://www.microsoft.com/en-us/research/publication/geolife-a-collaborative-social-networking-serv

A REDUCING THE SIZE OF THE LP WITH SPANNER GRAPHS

Bordenabe et al. [2014] observed that the number of privacy constraints in (2) can be reduced using a spanner graph with far fewer than $O(n^2)$ edges to approximate the metric $d_{\mathcal{W}}$. Thus, the privacy constraints (2) only need to be applied for elements u, v such that (u, v) is an edge in the spanner graph, and the spanner property ensures that all privacy constraints will approximately hold. In a spanner graph with $|E|$ edges, this results in $O(n|E|)$ privacy constraints, rather than $O(n^3)$. We now provide the formal details.

Definition 3. *(Spanner graph) Given a metric space $(\mathcal{W}, d_{\mathcal{W}})$, a graph $G = (\mathcal{W}, E)$ induces a metric $d_G(u, v)$ on $\mathcal{W}$ defined by the shortest path from u to v in G. The graph G is called a δ-spanner if for all $u, v \in \mathcal{W}$, $d_G(u, v) \leq d_{\mathcal{W}}(u, v) \leq \delta d_G(u, v)$.*

To construct a spanner, one can use the folklore greedy algorithm (e.g. [Ahmed et al., 2020]). For a positive integer k, the greedy algorithm uses constructs a $(2k-1)$-spanner with $O(n^{1+1/k})$ edges. Instead of requiring that mDP holds for all pairs $u, v \in \mathcal{W}$, the spanner method only requires it to hold for edges in G:

$$M_{uw} \leq M_{vw} \exp(d_{\mathcal{W}}(u, v)\epsilon) \qquad \forall (u, v) \in G, w \in \mathcal{W} \tag{12}$$

The modified linear program for finding the optimal mechanism under just the spanner property is:

$$\begin{aligned} \mathcal{P}_{\textsc{SpannerMech}}(\epsilon, G) = \textbf{minimize} \;\; & k \textbf{ subject to} \\ & \mathcal{L}(M, w) \leq k \;\; \forall w \in \mathcal{W} \\ & M \text{ stochastic} \\ & M \text{ satisfies (12)} \end{aligned}$$

Bordenabe et al. [2014] prove that the mechanism produced by $\mathcal{P}_{\textsc{SpannerMech}}(\epsilon, G)$ satisfies $\delta\epsilon$ $d_{\mathcal{W}}$ privacy. This proof follows simply from the definition of a δ-spanner.

Theorem 4. *[Bordenabe et al., 2014]: Let $G = (V, E)$ be a δ-spanner. The mechanism M produced by $\mathcal{P}_{\textsc{SpannerMech}}(\epsilon, G)$ satisfies $\delta\epsilon$ $d_{\mathcal{W}}$ privacy.*

The number of variables in $\mathcal{P}_{\textsc{SpannerMech}}$ is about the same as that of $\mathcal{P}_{\textsc{OptMech}}$, but the number of constraints is reduced from n^3 to $n|E|$. Using the greedy spanner algorithm, the number of constraints and non-zero coefficients is $O(n^{2+1/k})$, while the mDP guarantee is $(2k-1)\epsilon$. Following [Bordenabe et al., 2014], we instantiate the algorithm with a greedy spanner with $k = 2$. Thus, the number of constraints and non-zero coefficients is bounded by $O(n^{2.5})$.

B OMITTED PROOFS

The following proof involved proving that a weighted form of the exponential mechanism satisfies mDP.

Proposition 1. *For any metric space and any $\mathbf{Y} \in (\mathbb{R}^+)^n$, the mechanism $\textsc{ExpMech}_{\mathbf{Y}}$ satisfies ϵ-$d_{\mathcal{W}}$ privacy.*

Proof. Let $\mathbf{Y} \in (\mathbb{R}^+)^n$ be aribtrary. Our goal is to show, for all $u, u', v \in \mathcal{W}$,

$$\Pr[\textsc{ExpMech}_{\mathbf{Y}}(u) = v] \leq \exp(\epsilon d_{\mathcal{W}}(u, u')) \Pr[\textsc{ExpMech}_{\mathbf{Y}}(u') = v]$$

For positive real numbers A, B, δ, let $A \approx_\delta B$ denote that $A \leq \exp(\delta)B$ and that $B \leq \exp(\delta)A$. We will show that

$$\Pr[\textsc{ExpMech}_{\mathbf{Y}}(u) = v] \approx_{\epsilon d_{\mathcal{W}}(u,u')} \Pr[\textsc{ExpMech}_{\mathbf{Y}}(u') = v], \tag{13}$$

and this will imply our goal.

It is easy to see that if $A_1 \approx_\delta B_1$ and $A_2 \approx_\delta B_2$, then (i) $A_1 + A_2 \approx_\delta B_1 + B_2$, (ii) $A_1 A_2 \approx_{2\delta} B_1 B_2$, and (iii) $\frac{1}{A_1} \approx_\delta \frac{1}{B_1}$.

By the triangle inequality, we have that

$$\begin{aligned}
d_{\mathcal{W}}(u', v) &\leq d_{\mathcal{W}}(u, u') + d_{\mathcal{W}}(u, v) \\
\frac{-d_{\mathcal{W}}(u, v)}{2} &\leq \frac{d_{\mathcal{W}}(u, u') - d_{\mathcal{W}}(u', v)}{2} \\
e^{-\epsilon d_{\mathcal{W}}(u,v)/2} &\leq e^{\epsilon d_{\mathcal{W}}(u,u')/2} e^{-\epsilon d_{\mathcal{W}}(u',v)/2} \\
d_{\mathcal{W}}(u, v) &\leq d_{\mathcal{W}}(u', v) + d_{\mathcal{W}}(u, u') \\
\frac{-d_{\mathcal{W}}(u', v)}{2} &\leq \frac{d_{\mathcal{W}}(u, u') - d_{\mathcal{W}}(u, v)}{2} \\
e^{-\epsilon d_{\mathcal{W}}(u',v)/2} &\leq e^{\epsilon d_{\mathcal{W}}(u,u')/2} e^{-\epsilon d_{\mathcal{W}}(u,v)/2}
\end{aligned}$$

This implies that $e^{-\epsilon d_{\mathcal{W}}(u,v)/2} \approx_{\epsilon d_{\mathcal{W}}(u,u')/2} e^{-\epsilon d_{\mathcal{W}}(u',v)/2}$ for any choice of $u, u', v \in \mathcal{W}$. Thus,

$$Y_v e^{-\epsilon d_{\mathcal{W}}(u,v)/2} \approx_{\epsilon d_{\mathcal{W}}(u,u')/2} Y_v e^{-\epsilon d_{\mathcal{W}}(u',v)/2} \tag{14}$$

Applying property (i), we also have

$$\sum_{w \in \mathcal{W}} Y_w e^{-\epsilon d_{\mathcal{W}}(u,w)/2} \approx_{\epsilon d_{\mathcal{W}}(u,u')/2} \sum_{w \in \mathcal{W}} Y_w e^{-\epsilon d_{\mathcal{W}}(u',w)/2} \tag{15}$$

Dividing (14) and (15) and applying property (ii) and (iii), we obtain (13). □

Lemma 1. *With $\mathcal{W} = \{w_1, \ldots, w_n\}$ and $d_{\mathcal{W}}$ defined as above, when $\mathbf{Y} = (1, 1/(n-1), 1/(n-1), \ldots)$, we have that*

$$\mathcal{L}(\textsc{ExpMech}_{\mathbf{Y}}) \leq \frac{\frac{1}{n-1} + e^{-\epsilon/2}}{1 + e^{-\epsilon/2}} \mathcal{L}(\textsc{ExpMech}).$$

When $\epsilon \geq 2 \log n$, we have $\mathcal{L}(\textsc{ExpMech}_{\mathbf{Y}}) \leq 2/(n-1)\mathcal{L}(\textsc{ExpMech})$.

Proof. Recall that $d_{\mathcal{W}}(w_1, w_j) = 1$ for all $j \geq 2$ and $d_{\mathcal{W}}(w_i, w_j) = \delta$ for $i, j \geq 2$.

Observe that as $\delta \to 0$, we have

$$\mathcal{L}(\textsc{ExpMech}, w_1) \to \frac{(n-1)e^{-\epsilon/2}}{1 + (n-1)e^{-\epsilon/2}}, \quad \text{and} \quad \forall i \geq 2.\ \mathcal{L}(\textsc{ExpMech}, w_i) \to \frac{e^{-\epsilon/2}}{(n-1) + e^{-\epsilon/2}}. \tag{16}$$

On the other hand, for the weighted EM, we have as $\delta \to 0$,

$$\mathcal{L}(\textsc{ExpMech}_{\mathbf{Y}}, w_1) \to \frac{e^{-\epsilon/2}}{1 + e^{-\epsilon/2}}, \quad \text{and} \quad \forall i \geq 2.\ \mathcal{L}(\textsc{ExpMech}_{\mathbf{Y}}, w_i) \to \frac{e^{-\epsilon/2}}{1 + e^{-\epsilon/2}}. \tag{17}$$

Dividing the two losses, we obtain

$$\frac{\mathcal{L}(\textsc{ExpMech}_{\mathbf{Y}})}{\mathcal{L}(\textsc{ExpMech})} = \frac{1 + (n-1)e^{-\epsilon/2}}{(n-1)(1 + e^{-\epsilon/2})} = \frac{\frac{1}{n-1} + e^{-\epsilon/2}}{1 + e^{-\epsilon/2}}.$$

When $e^{-\epsilon/2} \leq \frac{1}{n-1}$, we have

$$\frac{\frac{1}{n-1} + e^{-\epsilon/2}}{1 + e^{-\epsilon/2}} \leq \frac{2}{n-1}.$$

This establishes the results of the lemma. □

Theorem 1. *$\mathcal{P}_{\textsc{ConstOptMech}}$ is feasible, and it is possible to solve it using a linear program with $n + 1 + \sum_{v \in \mathcal{W}} |I(v)|$ variables and $2n + \sum_{v \in \mathcal{W}} 2|I(v)|^2 + 3|I(v)|$ constraints. The number of non-zero coefficients in the LP is at most $2n^2 + \sum_{v \in \mathcal{W}} 2|I(v)|^2 + 5|I(v)|$.*

Proof. We will exhibit a valuation of M_{uv} and $\mathbf{Y} = \{Y_w : w \in \mathcal{W}\}$ variables such that $\mathcal{P}_{\text{ConstOptMech}}$ is feasible. Our valuation will set $M_{uv} = Y_v e^{-\epsilon d_{\mathcal{W}}(u,v)}$, and we will let $\mathbf{Y}$ be non-negative with values to be decided later. Notice that each M_{uv} is at least 0, and (4) is satisfied for all choices of $I(v)$. Furthermore, we can show that (2) is satisfied for any choice of $\mathbf{Y}$:

$$\begin{aligned} d_{\mathcal{W}}(u,w) &\geq d_{\mathcal{W}}(v,w) - d_{\mathcal{W}}(u,v) \\ e^{-\epsilon d_{\mathcal{W}}(u,w)} &\leq e^{-\epsilon(d_{\mathcal{W}}(v,w) - d_{\mathcal{W}}(u,v))} \\ Y_w e^{-\epsilon d_{\mathcal{W}}(u,w)} &\leq Y_w e^{-\epsilon d_{\mathcal{W}}(v,w)} e^{\epsilon d_{\mathcal{W}}(u,v)} \\ M_{uw} &\leq M_{vw} e^{\epsilon d_{\mathcal{W}}(u,v)} \end{aligned}$$

Finally, we can set the entries of $\mathbf{Y}$ high enough so that $\sum_{v \in \mathcal{W}} M_{uv} \geq 1$ for all $u \in \mathcal{W}$. Since all the constraints are satisfied, this produces a feasible solution to $\mathcal{P}_{\text{ConstOptMech}}$.

Now, we will show that by doing variable substitution into $\mathcal{P}_{\text{ConstOptMech}}(\epsilon)$, we can eliminate many of the variables and constraints and obtain an LP of much smaller size. We will do this by directly substituting the equalities of (4) into the remaining constraints. The constraints that $\tilde{\mathcal{L}}(M,w) \leq k$ for all $w \in \mathcal{W}$ become

$$\sum_{v \in \mathcal{W}, w \in I(v)} M_{wv} d_{\mathcal{W}}(u,v) + \sum_{v \in \mathcal{W}, w \notin I(v)} Y_v e^{-\epsilon d_{\mathcal{W}}(w,v)} d_{\mathcal{W}}(w,v) \leq k \ \ \forall w \in \mathcal{W} \tag{18}$$

The constraints that $M_{uv} \geq 0$ are implied by $Y_v \geq 0$ if $u \notin I(v)$. Thus, the only remaining constraints are:

$$M_{uv} \geq 0 \ \ \forall u \in \mathcal{W}, u \in I(v) \tag{19}$$

Substituting into $\sum_{v \in \mathcal{W}} M_{uv} \geq 1$, we obtain

$$\sum_{v \in \mathcal{W}, u \in I(v)} M_{uv} + \sum_{v \in \mathcal{W}, u \notin I(v)} Y_v e^{-\epsilon d_{\mathcal{W}}(u,v)} \geq 1 \ \ \forall u \in \mathcal{W} \tag{20}$$

The difficult case is substituting into the privacy constraints (2). We split them into four sets:

$$Y_w e^{-\epsilon d_{\mathcal{W}}(u,w)} \leq Y_w e^{-\epsilon d_{\mathcal{W}}(v,w)} e^{\epsilon d_{\mathcal{W}}(u,v)} \ \ \forall w \in \mathcal{W}, u \notin I(w), v \notin I(w) \tag{21}$$

$$M_{uw} \leq M_{vw} e^{\epsilon d_{\mathcal{W}}(u,v)} \ \ \forall w \in \mathcal{W}, u \in I(w), v \in I(w) \tag{22}$$

$$M_{uw} \leq Y_w e^{-\epsilon d_{\mathcal{W}}(v,w)} e^{\epsilon d_{\mathcal{W}}(u,v)} \ \ \forall w \in \mathcal{W}, u \in I(w), v \notin I(w) \tag{23}$$

$$Y_w e^{-\epsilon d_{\mathcal{W}}(u,w)} \leq M_{vw} e^{\epsilon d_{\mathcal{W}}(u,v)} \ \ \forall w \in \mathcal{W}, u \notin I(w), v \in I(w) \tag{24}$$

The constraints in (21) are always satisfied by following an argument identical to what we used to show feasibility. Thus, we may drop these constraints. Focusing on (23), if we let $w \in \mathcal{W}$ be arbitrary and $u \in I(w)$ be arbitrary, then the set of inequalities $\{M_{uw} \leq Y_w e^{-\epsilon d_{\mathcal{W}}(v,w)} e^{\epsilon d_{\mathcal{W}}(u,v)} : v \notin I(w)\}$ is equivalent to $M_{uw} \leq \min_{v \notin I(w)} Y_w e^{-\epsilon d_{\mathcal{W}}(v,w)} e^{\epsilon d_{\mathcal{W}}(u,v)}$. Thus, (23) is equivalent to:

$$M_{uw} \leq Y_w \min_{v \notin I(w)} e^{-\epsilon d_{\mathcal{W}}(v,w)} e^{\epsilon d_{\mathcal{W}}(u,v)} \ \ \forall w \in \mathcal{W}, u \in I(w) \tag{25}$$

Similarly, (24) is equivalent to:

$$Y_w \max_{u \notin I(w)} e^{-\epsilon d_{\mathcal{W}}(u,w)} \leq M_{vw} e^{\epsilon d_{\mathcal{W}}(u,v)} \ \ \forall w \in \mathcal{W}, v \in I(w) \tag{26}$$

After substituting, we have the following LP (which no longer involves the variables M_{uv} such that $u \notin I(v)$):

$$\begin{aligned} \mathcal{P}_{\text{ConstOptMech},red}(\epsilon) \ &= \textbf{minimize } k \textbf{ subject to} \\ & M \text{ satisfies (18)} \\ & M \text{ satisfies (19), (20)} \\ & M \text{ satisfies (22), (25), (26)} \\ & Y_v \geq 0 \ \ \forall v \in \mathcal{W} \end{aligned}$$

This LP has the variables Y_v and M_{uv} when $v \in \mathcal{W}$ and $u \in I(v)$, for a total of $n + \sum_{w\in\mathcal{W}} |I(w)| + 1$ variables. The number of constraints in (18), and (20) is $2n$; the number of constraints in (19) is $\sum_{u\in\mathcal{W}} |I(u)|$, the number of constraints in (25) and (26) is $2\sum_{w\in\mathcal{W}} |I(w)|$, and the number of constraints in (22) is $\sum_{w\in\mathcal{W}} |I(w)|^2$.

Similarly, the number of nonzero coefficients in (18) is n^2, in (19) is $\sum_{w\in\mathcal{W}} |I(w)|$, in (20) is n^2, in (25) and (26) is $4\sum_{w\in\mathcal{W}} |I(u)|$, and in (22) is $2\sum_{w\in\mathcal{W}} |I(w)|^2$. □

Corollary 1. *When $I(v) = I_{\mathrm{NN},r}(v)$, then the number of variables in $\mathcal{P}_{\textsc{ConstOptMech}}$ is at most $nr + n + 1$, the number of constraints is at most $n^2r + 3nr + 2n$, and the number of non-zero elements is at most $2n^2 + 5nr + 2n^2r$.*

Proof. When $I(w) = I_{\mathrm{NN},r}(w)$, we have $\sum_{w\in\mathcal{W}} |I(w)| \leq rn$. Subject to this inequality, the maximum value of $\sum_{w\in\mathcal{W}} |I(w)|^2$ is n^2r, achieved when r of the $I(w)$'s have size n. Thus, the total number of constraints is at most $2n + 3nr + n^2r$. The total number of nonzero coefficients is $2n^2 + 5nr + 2n^2r$. □

Theorem 2. *For any set $\mathcal{W}$, metric $d_{\mathcal{W}}$, and budget ϵ and replacement function $I(v) : \mathcal{W} \to 2^{\mathcal{W}}$, Mechanism* ConstOptMech *satisfies ϵ-$d_{\mathcal{W}}$ privacy.*

Proof. Notice that the variables M_{uv} outputted by $\mathcal{P}_{\textsc{ConstOptMech}}(\epsilon)$ satisfy $M_{uw} \leq e^{\epsilon d_{\mathcal{W}}(u,v)/2} M_{vw}$ for all $u, v, w \in \mathcal{W}$, which are exactly the constraints to ensure $\frac{\epsilon}{2}$ mDP. Using a method identical to the proof of Proposition 1, we can show that the quantities $H_{uv} = \frac{M_{uv}}{\sum_{w\in\mathcal{W}} M_{uw}}$ satisfy $H_{uw} \leq e^{\epsilon d_{\mathcal{W}}(u,v)} H_{vw}$ for all $u, v, w \in \mathcal{W}$. Thus, the mechanism outputted by Algorithm 1 satisfies ϵ mDP. □

Theorem 3. *Consider an arbitrary set $\mathcal{W}$ and metric $d_{\mathcal{W}} : \mathcal{W} \times \mathcal{W} \to \mathbb{R}$. Then, for any $\epsilon > 0$, any mechanism $\mathcal{M}$ satisfying ϵ-$d_{\mathcal{W}}$ privacy and any (c, r, Q)-packing S of $\mathcal{W}$, it holds that*

$$\mathcal{L}(\mathcal{M}) \geq \max_{w\in\mathcal{W}} r\left(1 - \frac{1}{N(w,S)}\right). \tag{7}$$

It follows that

$$\mathcal{L}(\mathcal{M}) \geq r\left(1 - \frac{1}{1 + (c-1)\exp(-Q\epsilon)}\right). \tag{8}$$

Proof. Let S be any (c, r, Q)-packing of $\mathcal{W}$, and let $\kappa = \max_{w\in\mathcal{W}} 1 - \frac{1}{N(w,S)}$. Suppose for the sake of contradiction that there exists a mechanism $\mathcal{M}$ satisfying ϵ-mDP such that $\mathcal{L}(\mathcal{M}) < r\kappa$. By definition, this implies $\mathcal{L}(\mathcal{M}, w) < r\kappa$ for all $w \in \mathcal{W}$. By Markov's inequality, this means $\Pr[\mathcal{M}(w) \in B(w, r)] > 1 - \kappa$. Using the mDP property, the following holds for any $w, w' \in \mathcal{W}$:

$$\begin{aligned}\Pr[\mathcal{M}(w) \in B(s,r)] &\geq \exp(-d_{\mathcal{W}}(w,s)\epsilon)\Pr[\mathcal{M}(s) \in B(s,r)] \\ &> \exp(-d_{\mathcal{W}}(w,s)\epsilon)(1-\kappa).\end{aligned}$$

For all $s, s' \in S$, the events $\mathcal{M}(w) \in B(s, r)$ and $\mathcal{M}(w) \in B(s', r)$ are disjoint by the properties of a (c, r, Q)-packing. This implies, for all $w \in \mathcal{W}$,

$$1 \geq \sum_{s\in S} \Pr[\mathcal{M}(w) \in B(s,r)].$$

We know from above that $\Pr[\mathcal{M}(w) \in B(s, r)] > \exp(-d_{\mathcal{W}}(w, s)\epsilon)(1 - \kappa)$. Thus, for all w,

$$\begin{aligned}1 &> (1-\kappa)\sum_{s\in S}\exp(-d_{\mathcal{W}}(w,s)\epsilon) = (1-\kappa)N(w,S) \\ \kappa &> 1 - \frac{1}{N(w,S)}\end{aligned} \tag{27}$$

However, since $\kappa = \max_{w\in\mathcal{W}} 1 - \frac{1}{N(w,S)}$, there is some w such that $\kappa = 1 - \frac{1}{N(w,S)}$. This contradicts (27).

□

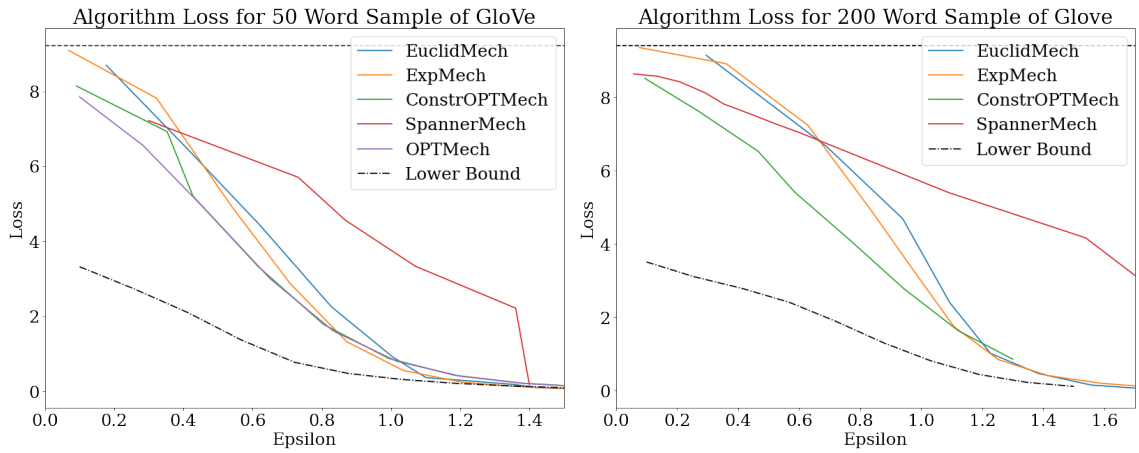

Figure 3: Loss of Madlib (EuclidMech), EM, ConstOPTMech, SpannerMech, and OPTMech versus ϵ_{actual} on 50 and 200-size metric spaces generated from GloVe, along with the lower bound. The horizontal dotted line indicates the loss of returning a uniform random element.

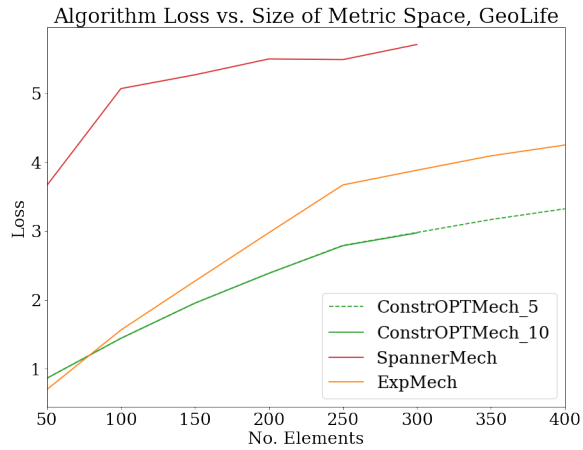

Figure 4: Loss of EM, ConstOPTMech (with $r = 5, 10$), and SpannerMech versus size of metric space for metric spaces generated from GloVe. Here, ϵ_{actual} is fixed at 1.0.

C ADDITIONAL EXPERIMENTAL RESULTS

Here, we include omitted plots for metric spaces generated from the GloVe word embedding. In Figure 3, we plot the losses of the tested mechanisms for generated metric spaces of size 50 and 200. Our results indicate that ConstOPTMech achieves the best error out of all non-optimal mechanisms by about 5 to 10% (size 50) and 20% (size 200) for the middle ranges of epsilon (approximately $[0.5, 1.0]$). This is consistent with the experiments on the other metric spaces. Furthermore, SpannerMech has much higher loss than the other mechanisms, particularly for higher values of ϵ. This is likely due to distortions in the underlying metric space caused by the spanner.

Table 3 shows the time and NNZ required to compute the LPs in the LP-based mechanisms with $\epsilon_{tight} = 1.0$ and increasing metric space sample sizes. SpannerMech was unable to scale past 300 elements and timed out, but unlike previous experiments, ConstOPTMech with $r = 10$ was also unable to scale past 300 elements due to timing out. However, SpannerMech with $r = 5$ was able to scale to 400 elements without timing out. While it executed much more quickly than SpannerMech, there was not a large difference in the NNZ between the two methods. These results still indicate that $ConstOPTMech$ is more scalable than SpannerMech, though with the caveat that one must choose r carefully and the answer depends on the metric space.

Figure 4 indicates that ConstOPTMech with $r = 5, 10$ result in much lower loss of about $50-75\%$ compared to SpannerMech with ϵ_{tight} fixed, with the gap larger for smaller metric spaces. This is consistent with our observations made in Figure 3.

GloVe	50	150	300	400
ConstOPTMech (10)	1.68 sec, 3.33e4 nnz	31.28 sec, 2.71e5 nnz	611.05 sec, 1.10e6 nnz	> 1800 sec, — nnz
ConstOPTMech (5)	0.80 sec, 1.48e4 nnz	15.76 sec, 1.56e5 nnz	127.07 sec, 4.59e5 nnz	447.10 sec, 7.76e5 nnz
SpannerMech	0.85 sec, 1.62e4 nnz	74.77 sec, 1.52e5 nnz	495.96 sec, 6.59e5 nnz	> 1800 sec, — nnz
OPTMech	8.56 sec, 2.50e5 nnz	> 1800 sec, — nnz		

Table 3: Running times and memory requirement of computing ConstOPTMech when $r = 10, 5$; SpannerMech; and OPTMech for varying metric space sizes generated from FastText and Geolife. Each mechanism satisfies $\epsilon_{tight} = 2.0$ (FastText) and 0.3 (Geolife).