

# Differentially Private Partial Set Cover with Applications to Facility Location

George Z. Li<sup>1</sup>, Dung Nguyen<sup>2</sup>, and Anil Vullikanti<sup>2</sup>

<sup>1</sup>University of Maryland

<sup>2</sup>Biocomplexity Institute and Initiative, and Department of Computer Science, University of Virginia  
gzli929@gmail.com, dungn@virginia.edu, vsakumar@virginia.edu

## Abstract

Set Cover is a fundamental problem in combinatorial optimization which has been studied for many decades due to its various applications across multiple domains. In many of these domains, the input data consists of locations, relationships, and other sensitive information of individuals which may leaked due to the set cover output. Attempts have been made to design privacy-preserving algorithms to solve the Set Cover under privacy constraints. Under differential privacy, it has been proved that the Set Cover problem has strong impossibility results and no explicit forms of the output can be released to the public.

In this work, we observe that these hardness results dissolve when we turn to the Partial Set Cover problem, where we only need to cover a  $\rho \in (0, 1)$  fraction of the elements. We show that this relaxation enables us to avoid the impossibility results, and give the first algorithm which outputs an explicit form of set cover with non-trivial utility guarantees under differential privacy. Using our algorithm as a subroutine, we design a differentially private bicriteria algorithm to solve a recently proposed facility location problem for vaccine distribution which generalizes the  $k$ -supplier with outliers. Our analysis shows that relaxing the covering requirement to serve only a  $\rho \in (0, 1)$  fraction of the population/universe also allows us to circumvent the inherent hardness of  $k$ -supplier and give the first non-trivial guarantees.

## 1 Introduction

Data privacy is a fundamental challenge in many real world applications of data-driven decision making where there is a risk of inadvertently revealing private information. Differential privacy, introduced in [Dwork *et al.*, 2006], has emerged as a widely accepted formalization of privacy, which gives rigorous parameterized guarantees on the privacy loss while simultaneously enabling non-trivial utility in algorithmic and statistical analysis. Differential privacy is defined in terms of datasets which differ by one individual, called neighboring

datasets, and requires that the output of a mechanism is (approximately) indistinguishable when run on any two neighboring datasets. Formally, it is defined as:

**Definition 1.1.**  $M : \mathcal{X}^n \rightarrow \mathcal{Y}$  is  $(\epsilon, \delta)$ -differentially private if for any neighboring datasets  $X, X' \in \mathcal{X}^n$  and  $S \subseteq \mathcal{Y}$ ,

$$\Pr[M(X) \in S] \leq \exp(\epsilon) \Pr[M(X') \in S] + \delta.$$

If  $\delta = 0$ , we say  $M$  is  $\epsilon$ -differentially private.

Differentially private algorithms have now been developed for a number of problems ranging from statistics [Canonne *et al.*, 2020; Brown *et al.*, 2021], machine learning and deep learning [Lee and Kifer, 2018; Ghazi *et al.*, 2021], social network analysis [Nissim *et al.*, 2007; Hay *et al.*, 2009; Karwa *et al.*, 2011], and combinatorial optimization [Mitrovic *et al.*, 2017; Esencayi *et al.*, 2019; Nguyen and Vullikanti, 2021]. See [Dwork and Roth, 2014; Vadhan, 2017] for a survey on the techniques used.

In this work, we consider a fundamental problem in combinatorial optimization: the Set Cover problem [Williamson and Shmoys, 2011], which involves choosing the smallest subset of a set system  $\mathcal{S} = \{S_1, \dots, S_m\} \subset 2^{\mathcal{U}}$  that covers a universe  $\mathcal{U} = \{u_1, \dots, u_n\}$ . In many settings, the elements of the universe  $\mathcal{U}$  are private (e.g., clients wish to be private in facility location problems). [Gupta *et al.*, 2010] first studied the problem of Set Cover with privacy, and showed that outputting an explicit solution to Set Cover has strong hardness results, even for the special case of Vertex Cover: any differentially private algorithm must output a set cover of size  $m - 1$  with probability 1 on any input, a useless result.

As a result, the authors designed a mechanism which outputs an implicit set cover via a privacy-preserving set of instructions for the elements to reconstruct the set cover. While the implicit solutions are useful in some limited settings, it cannot replace the *explicit* solutions needed in many important applications, such as public health [Eubank *et al.*, 2004; Li *et al.*, 2022]. In particular, explicit solutions are necessary when using a Set Cover algorithm as a subroutine when solving a more complicated problem. As a result, we turn to the Partial Set Cover problem, where we only need to cover a  $\rho$ -fraction of the elements in  $\mathcal{U}$ , for  $\rho \in (0, 1)$ . Due to the space limit, we present our core results in the main paper and defer some analyses and experimental results to the Appendix. We maintain a full, updated version of this work here<sup>1</sup>.

<sup>1</sup><https://arxiv.org/abs/2207.10240>

Our primary contributions are:

- We observe that the impossibility results for outputting an explicit set cover under differential privacy are alleviated when considering the Partial Set Cover problem. When the number sets isn't too large (i.e.,  $m = O(n)$ ), we give a  $O(\log^2(m) \log(1/\delta)/\epsilon(1-\rho))$ -approximation algorithm (see Theorem 2.4). Alternatively, when the optimal partial set cover isn't too large (i.e.,  $\text{OPT} \lesssim \frac{n\epsilon}{\log^3 n \log(1/\delta)}$ ), we give a  $O(\log(\frac{1}{1-\rho}))$ -pseudo approximation algorithm (see Theorem 2.6). Note that both of our guarantees break down as  $\rho \rightarrow 1$ .
- To illustrate the importance of explicit solutions, we use our differentially private Partial Set Cover algorithm as a subroutine to give a differentially private approximation algorithm for a vaccine distribution problem which in particular, generalizes  $k$ -supplier with outliers. We emphasize that this is the first differentially private algorithm for  $k$ -supplier type problems with non-trivial approximation guarantees, which was thought to be impossible due to the high sensitivity of min-max objectives.
- Finally, we evaluate the our private algorithms on real life datasets for set cover and vaccine distribution, and observe a good privacy-utility tradeoff: the empirical approximation factors are much better than the worst case bounds we give in our theorem statements. We view the vaccine distribution problem as a concrete practical contribution of our work (see Figure 1 for an example).

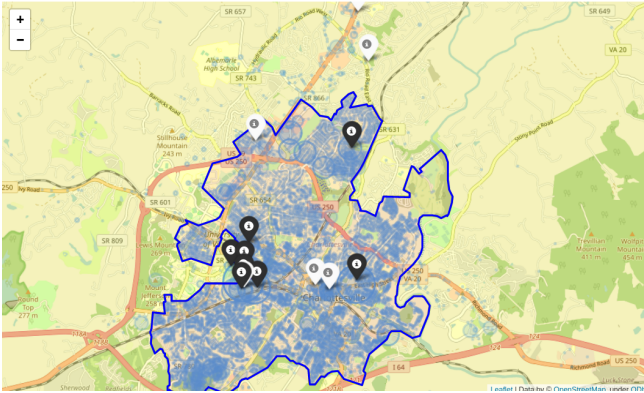


Figure 1: Visualization of locations of chosen facilities. The black and white spots indicate locations chosen by the non-private and private algorithms, respectively. As we see, the private algorithm yields qualitatively different distributions when compared to the non-private algorithm while not losing too much in objective value (see Section 4).

## 1.1 Related Work

The Set Cover problem and its various generalizations have been studied by combinatorial optimization community for several decades [Wolsey, 1982; Alon *et al.*, 2003]. For the simplest version, there exists a greedy algorithm which achieves a  $(\log n + 1)$ -approximation which is best possible unless  $P=NP$  [Moshkovitz, 2015]. [Gupta *et al.*, 2010]

first considered the Set Cover problem under differential privacy, showing impossibility results of outputting explicit set covers. They then gave approximation algorithms via outputting implicit set covers, which we argue is insufficient for many applications. The only other work which outputs explicit set covers under differential privacy is [Hsu *et al.*, 2014], which approach the set cover problem via private linear programming. They give approximation guarantees but ignore  $O(\text{OPT}^2 \cdot \text{polylog}(n, 1/\delta))$  elements. One of our algorithms also has guarantees of this form, but ignores only  $\rho n$  elements, which is better when the size of the optimal set cover is larger. We also give a true approximation for Partial Set Cover, which is the first in the literature.

Also directly related to our work are the differentially private facility location problems, which [Gupta *et al.*, 2010] also first considered. For the uniform facility location problem, they showed that a  $\Omega(\sqrt{n})$ -approximation is needed under differential privacy, an essentially useless result, and devised a way to implicitly output the facilities. [Esencayi *et al.*, 2019] built upon their work, improving the approximation guarantees to  $O(\frac{\log n}{\epsilon})$  for general metrics. [Cohen-Addad *et al.*, 2022b] considered the facility location problem under the local differential privacy model, gave tight approximation algorithms up to polynomial factors of  $\epsilon$ . A particular interesting quality of their algorithms is that it extends to non-uniform facility location. [Gupta *et al.*, 2010] also considered the  $k$ -median problem, and developed approximation algorithms which guaranteed that the service cost is at most  $6 \cdot \text{OPT} + O(\frac{k^2 \log^2 n}{\epsilon})$ . Since then, there has been an abundance of work on this problem improving the approximation guarantees, practical performance, and efficiency of the differentially private algorithms [Balcan *et al.*, 2017; Ghazi *et al.*, 2020; Blocki *et al.*, 2021; Jones *et al.*, 2021; Cohen-Addad *et al.*, 2022a]. Despite the abundance of work on facility location-type problems,  $k$ -supplier remains untouched; our work is the first to overcome impossibility results and give approximation guarantees for the problem.

## 1.2 Differential Privacy Background

In our algorithms, we will make extensive use of the following basic mechanisms and properties, the proofs of which can be found in [Dwork and Roth, 2014]. The post-processing and composition properties enable us to easily combine smaller differentially private mechanisms to create a more complicated one:

**Theorem 1.2.** *Let  $\mathcal{M}_1 : \mathcal{X}^n \rightarrow \mathcal{Y}_1$  and  $\mathcal{M}_2 : \mathcal{X}^n \rightarrow \mathcal{Y}_2$  be  $(\epsilon_1, \delta_1)$  and  $(\epsilon_2, \delta_2)$ -differentially private algorithms, respectively. The following properties hold:*

- **post-processing:** *Let  $f : \mathcal{Y}_1 \rightarrow \mathcal{Z}$  be an arbitrary (potentially randomized) mapping. Then  $f \circ \mathcal{M}_1 : \mathcal{X}^n \rightarrow \mathcal{Z}$  is  $(\epsilon_1, \delta_1)$ -differentially private.*
- **composition:** *Let  $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}_1 \times \mathcal{Y}_2$  be defined as  $\mathcal{M} = (\mathcal{M}_1, \mathcal{M}_2)$ . Then  $\mathcal{M}$  is  $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -differentially private.*

We next state the Laplace Mechanism which, by adding Laplace noise, provides a simple way to privately output a statistic that depends on a private database:

**Theorem 1.3.** Given a function  $f : \mathcal{X}^n \rightarrow \mathbb{R}^k$ , the  $\ell_1$ -sensitivity is defined as  $\Delta_f = \max_{x \sim x'} \|f(x) - f(x')\|_1$ . The Laplace Mechanism, given the function  $f : \mathcal{X}^n \rightarrow \mathbb{R}^k$ , outputs  $f(x) + (Y_1, \dots, Y_k)$ , where  $Y_i$  are i.i.d. random variables drawn from  $\text{Lap}(\Delta_f/\epsilon)$ . We claim the Laplace Mechanism is  $\epsilon$ -differentially private.

Finally, we state the Exponential Mechanism which approximately optimizes the utility function over some set of candidates choices  $\mathcal{R}$  while preserving privacy:

**Definition 1.4.** Given a utility function  $u : \mathcal{X}^n \times \mathcal{R} \rightarrow \mathbb{R}$ , let  $\Delta_u = \max_{r \in \mathcal{R}} \max_{x \sim x'} |u(x, r) - u(x', r)|$  be the sensitivity of  $u$ , where  $x, x'$  are neighboring datasets. The exponential mechanism  $M(x, u, \mathcal{R})$  outputs an element  $r \in \mathcal{R}$  with probability  $\propto \exp(\frac{\epsilon u(x, r)}{2\Delta_u})$ .

**Theorem 1.5.** The exponential mechanism is  $\epsilon$ -differentially private. Furthermore, if we fix a dataset  $x$  and let  $\text{OPT} = \max_{r \in \mathcal{R}} u(x, r)$ , we have  $\Pr[u(x, M(x, u, \mathcal{R})) \leq \text{OPT} - \frac{2\Delta_u}{\epsilon} (\ln |\mathcal{R}| + t)] \leq e^{-t}$ .

## 2 Differentially Private Partial Set Cover

Formally, we wish to solve the following problem:

**Definition 2.1.** Let  $\mathcal{U} = \{u_1, \dots, u_n\}$  be the universe of elements and let  $\mathcal{S} = \{S_1, \dots, S_m\}$  be a set system where each  $S_i \subseteq \mathcal{U}$  and  $\bigcup_{i=1}^m S_i = \mathcal{U}$ . Finally, let  $\rho < 1$  be the covering requirement. The Partial Set Cover problem asks for the minimal size subset  $\{\pi_1, \dots, \pi_k\}$  of  $\mathcal{S}$  such that  $|\bigcup_{i=1}^k \pi_i| \geq \rho|\mathcal{U}|$  (i.e., the subset covers a  $\rho$  fraction of  $\mathcal{U}$ ).

For differential privacy, we consider two Partial Set Cover instances neighboring if the set systems  $(\mathcal{U}_1, \mathcal{S}_1)$  and  $(\mathcal{U}_2, \mathcal{S}_2)$  are two neighboring datasets if they differ by exactly one element  $u$  in the universe and the sets  $S_{i_1} \in \mathcal{S}_1$  and  $S_{i_2} \in \mathcal{S}_2$  differ only by  $u$  (or are the same). Note that our definition of privacy can recover the one considered in [Gupta *et al.*, 2010] and [Hsu *et al.*, 2014] as a special case, but is in general stronger. We elaborate on this subtlety further in the Appendix.

At first glance, our (and [Gupta *et al.*, 2010]’s) definition of neighboring datasets for (Partial) Set Cover may not be entirely intuitive. As an example of an application where such a privacy definition makes sense, let’s consider the problem studied by [Eubank *et al.*, 2004] of placing sensors in people-location graphs to detect the spread of a disease. Formally, we have a bipartite graph where nodes in one part of the graph represent locations and nodes in the other part represent people. An edge exists between a person and a location if the person visits that location. We wish to place sensors at the locations so that infected individuals can be detected. Since placing sensors is an expensive process, we wish to place the fewest sensors as possible to cover a  $\rho$ -fraction of the population. This can be formulated as a Partial Set Cover problem. For this example, our notion of privacy corresponds to node privacy for people in the graph. In addition to this example, we show in the following section that our privacy definition here also lines up nicely with client-privacy in facility location problems when using Partial Set Cover to solve  $k$ -supplier with outliers.

In the following subsections, we’ll present our algorithms for differentially private Partial Set Cover. Since Partial Set Cover is NP-Hard, we will focus on giving approximation algorithms while preserving differential privacy. We say an algorithm gives an  $\alpha$ -approximation to the optimal solution if it outputs a partial set cover of size at most  $\alpha \cdot \text{OPT}$ , where  $\text{OPT}$  is the size of an optimal partial set cover.

### 2.1 A Private Variant of the Greedy Algorithm

The general outline of our first algorithm is as follows: we use a private version of the classical greedy algorithm for Partial Set Cover to output a permutation of the sets  $\pi_1, \dots, \pi_m$ . Then, we use an offline implementation of the AboveThreshold mechanism to choose a threshold  $k$  such that  $\pi_1, \dots, \pi_k$  covers a  $\rho$  fraction of the elements.

---

#### Algorithm 1 $\mathcal{M}_{\text{PartialSetCover}}(\mathcal{U}, \mathcal{S}, \rho, \epsilon, \delta)$

---

**Input:** Set system  $(\mathcal{U}, \mathcal{S})$ , covering requirement  $\rho$ , and privacy parameters  $(\epsilon, \delta)$

let  $\mathcal{U}_1 \leftarrow \mathcal{U}, \mathcal{S}_1 \leftarrow \mathcal{S}, \epsilon' \leftarrow \frac{\epsilon}{2 \ln(\epsilon/\delta)}$ .

for  $i = 1, \dots, m$  do

    pick set  $S \in \mathcal{S}_i$  with probability  $\propto \exp(\epsilon' |S \cap \mathcal{U}_i|)$ .

    let  $\pi_i \leftarrow S, \mathcal{U}_{i+1} \leftarrow \mathcal{U}_i - S, \mathcal{S}_{i+1} \leftarrow \mathcal{S}_i - \{S\}$ .

end for

let  $T \leftarrow \rho n + \frac{12 \log m}{\epsilon}, \hat{T} \leftarrow T + \text{Lap}(2/\epsilon)$ .

for  $i = 1, \dots, m$  do

    let  $f_i \leftarrow |\pi_1 \cup \dots \cup \pi_i|, \gamma_i \leftarrow f_i + \text{Lap}(4/\epsilon)$ .

end for

let  $k$  be first index such that  $\gamma_k \geq \hat{T}$ .

**Output:**  $(\pi, k)$  //  $\pi_1, \dots, \pi_k$  is a set cover for  $(\mathcal{U}, \mathcal{S})$

---

**Lemma 2.2.** Let  $k^*$  be the first index such that  $\pi_1, \dots, \pi_{k^*}$  is a  $(\rho n + \frac{24 \log m}{\epsilon})$ -partial covering of  $\mathcal{U}$ . Then with probability  $1 - O(\frac{1}{m})$ , for  $m = O(n)$ , we have  $k^* \leq O\left(\frac{\ln(m)^2}{\epsilon'(1-\rho)}\right) \cdot \text{OPT} = O\left(\frac{\ln(m)^2 \ln(1/\delta)}{\epsilon(1-\rho)}\right) \cdot \text{OPT}$ .

*Proof.* For iteration  $i \in [m]$ , let  $L_i$  be the size of the set which covers the most additional elements (i.e.,  $L_i = \max_{S \in \mathcal{S}_i} |S \cap \mathcal{U}_i|$ ). For an iteration where  $L_i \geq 6 \ln m/\epsilon'$ , the probability of selecting a set which covers less than  $L_i - 3 \ln m/\epsilon'$  is at most  $\frac{1}{m^2}$ . Hence, over all iterations where  $L_i \geq 6 \ln m/\epsilon'$ , we will choose a set which covers at least  $L_i/2$  elements with probability at least  $1 - \frac{1}{m}$ . By a standard argument [Williamson and Shmoys, 2011], this will only use at most  $\text{OPT} \cdot \ln n$  sets. Consider the last iteration  $t$  such that  $L_t \geq 6 \ln m/\epsilon'$ . If the number of elements covered through iteration  $t$  is at least  $\rho'n$ , then we are done. The rest of the proof deals with the case where less than  $\rho'n$  are covered.

Next, we analyze what happens when  $L_j < 6 \ln m/\epsilon'$  for  $j = t + 1, \dots, m$ . The utility guarantees of the exponential mechanism are essentially useless from this iteration onwards. Notice that the number of remaining elements  $|U_j|$  is at most  $\text{OPT} \cdot |L_j|$ . Unfortunately, we cannot claim that

each set chosen covers at least one element; this is simply not true. Instead, we analyze the probability that a set covering at least one element is chosen. Let  $\rho' = \frac{\rho+1}{2}$  and note that  $\rho' < 1$  and  $\rho'n \geq \rho n + \frac{24 \log m}{\epsilon}$  for sufficiently large  $n$ . Since there are at least  $(1 - \rho')n$  uncovered elements remaining and  $m = O(n)$ , there necessarily exists some constant  $\rho''$  such that the probability of not covering anything is at most  $[1 - (1 - \rho'')]^{2 \ln m}$ . Thus, the probability of not covering anything over the course of  $\frac{2 \ln m}{1 - \rho''}$  iterations is at most

$$[1 - (1 - \rho'')]^{\frac{2 \ln m}{1 - \rho''}} \leq \exp(-2 \ln m) = \frac{1}{m^2},$$

where we used  $1 - x \leq \exp(-x)$ . Thus, each of the  $|U_j|$  remaining elements is covered using at most  $\frac{2 \ln m}{1 - \rho''}$  sets, with probability at least  $1 - \frac{1}{m}$ . Since there are at most  $\text{OPT} \cdot |L_j|$  elements remaining which need to be covered, at most  $\text{OPT} \cdot \frac{2 \ln(m)^2}{\epsilon'(1 - \rho'')}$  sets are used.  $\square$

**Lemma 2.3.** (Lemma A.1) *With probability  $1 - O(\frac{1}{m})$ , the threshold  $k$  satisfies  $|\pi_1 \cup \dots \cup \pi_k| \in [\rho n, \rho n + 24 \log m / \epsilon]$ .*

**Theorem 2.4.** *For  $\epsilon \in (0, 1)$ ,  $\delta < \frac{1}{e}$ , and  $m = O(n)$ , the following are true for Partial Set Cover*

- Algorithm 1 preserves  $(2\epsilon, \delta)$ -differential privacy.
- With probability  $1 - O(\frac{1}{m})$ , Algorithm 1 is an  $O\left(\frac{\ln(m)^2 \ln(1/\delta)}{\epsilon(1-\rho)}\right)$ -approximation algorithm.

*Proof.* Let's first consider the privacy. Outputting the permutation of sets was shown to be  $(\epsilon, \delta)$ -differentially private in [Gupta et al., 2010]. Our mechanism for outputting the threshold  $k$  can be viewed as an offline implementation of the AboveThreshold mechanism from [Dwork and Roth, 2014]. Since switching to a neighboring set system changes the number of elements covered by a family of sets by at most 1, the analysis of [Dwork and Roth, 2014] applies and outputting the threshold is  $(\epsilon, 0)$ -differentially private. By basic composition of adaptive mechanisms [Dwork and Roth, 2014], we have  $(2\epsilon, \delta)$ -differential privacy.

Now, we turn to the utility guarantee. Consider the threshold  $k$  selected; by Lemma 2.3, the threshold is such that  $|\pi_1 \cup \dots \cup \pi_k|$  in the interval  $\left[\rho n, \rho n + \frac{24 \log m}{\epsilon}\right]$  with probability at least  $1 - \frac{1}{m}$ . Hence, it is a  $\rho n$ -partial cover and by Lemma 2.2, uses at most  $O\left(\frac{\ln(m)^2 \ln(1/\delta)}{\epsilon(1-\rho)}\right) \cdot \text{OPT}$  sets.  $\square$

**Remark 1.** *In Theorem 2.4 (and in future Theorems 2.6 and 3.2)), we can reduce the probability of failure to an arbitrary polynomial in  $m$  by losing constant factors in the approximation guarantee.*

## 2.2 Algorithm via Maximum Coverage

In this section, we give another algorithm for Partial Set Cover under differential privacy. Our algorithm here will give a pseudo-approximation for the problem (i.e., our approximation factor will be with respect to the optimal Set Cover solution instead of the optimal Partial Set Cover solution). Such a

guarantee is similar to the one given in [Hsu et al., 2014]. To give our algorithm, we first need to define the Differentially Private Maximum Coverage problem:

**Definition 2.5.** *Let  $\mathcal{U} = \{u_1, \dots, u_n\}$  be the universe of elements and let  $\mathcal{S} = \{S_1, \dots, S_m\}$  be a set system where each  $S_i \subseteq \mathcal{U}$  and  $\bigcup_{i=1}^m S_i = \mathcal{U}$ . Finally, let  $k$  be the budget. The Maximum Coverage problem asks us to find a size  $k$  subset  $\{\pi_1, \dots, \pi_k\}$  of  $\mathcal{S}$  such that  $|\bigcup_{i=1}^k \pi_i|$  is maximized.*

As in the Partial Set Cover problem, we view the elements of the universe as the private information and we view two set systems as neighbors if they differ by exactly one element  $u$  in the universe. Since the objective here is submodular and monotone [Williamson and Shmoys, 2011], we can apply the following result from [Mitrovic et al., 2017] for submodular maximization. We remark that the result we state is stronger than the one given in [Mitrovic et al., 2017], since we can use a specialized privacy analysis for maximum coverage (like in [Gupta et al., 2010]) which is not possible for general submodular functions. This has shown up in other works such as [Jones et al., 2021].

**Lemma 2.6.** *There exists an  $(\epsilon, \delta)$ -differentially private algorithm for the maximum coverage problem which such that the expected number of elements covered is  $(1 - \frac{1}{e}) \text{OPT} - \frac{2k \ln n}{\epsilon_0}$ , where  $\epsilon_0 = \frac{\epsilon}{2 \ln(\epsilon/\delta)}$ .*

The main idea for our algorithm for Partial Set Cover is that under some restrictions on the set system and budget, the Maximum Coverage problem can be approximated within a constant factor under differential privacy via the algorithm in Lemma 2.6. Then, iteratively applying the algorithm for the maximum coverage problem with budget  $k$  set to the size of the optimal Partial Set Cover suffices to obtain a good approximation algorithm for Partial Set Cover.

**Lemma 2.7.** (Lemma A.2) *There exists an  $(2\epsilon, \delta)$ -differentially private algorithm for the maximum coverage problem such that for some constant  $C$ , if we have*

$$k \leq \frac{C \epsilon_0}{\ln^2(n)} \cdot \text{OPT},$$

*then the algorithm is a 0.15-approximation with probability  $1 - O(\frac{1}{n})$ , where  $\epsilon_0 = \frac{\epsilon}{2 \ln(e \ln(n)/\delta \ln(1+\alpha))}$ .*

Given this result, we can state our algorithm for Partial Set Cover. For simplicity of notation, let's denote the Algorithm referenced in Lemma 2.7 by  $\text{MAXCOVER}(\mathcal{U}, \mathcal{S}, k, \epsilon, \delta)$ . As mentioned before, the idea is to guess the size of the optimal Partial Set Cover via binary search. Then, assuming we have  $\text{OPT}$ , we can run  $\text{MAXCOVER}$  approximately  $O(\log(1 - \rho))$  times in order to cover  $\rho n$  elements. Though the idea is very simple, there are many subtleties in the algorithm due to privacy. First, our binary search must guarantee that our guess  $\text{OPT}'$  is at most upper; this is because the guarantee in Lemma 2.7 only applies when the budget isn't too large. Additionally, when binary searching for  $\text{OPT}$ , we need to decide if the guess is too large or too small, based on the number of elements covered by the output. However, we cannot do this directly since the elements are considered private; as a result, we need to add Laplace noise before making the comparison. This makes the analysis slightly more complicated.

---

**Algorithm 2**  $\mathcal{M}_{MaxCoverage}(\mathcal{U}, \mathcal{S}, \rho, \epsilon, \delta)$ 


---

**Input:** Set system  $(\mathcal{U}, \mathcal{S})$ , covering requirement  $\rho$ , privacy parameters  $(\epsilon, \delta)$

**let** upper =  $\lfloor \frac{C(1-\frac{\rho}{2})n\epsilon_0}{\ln^3(n)} \rfloor$ ,  $t = \lceil \log_{0.85}(1 - \rho') \rceil$ .

**Binary Search** on  $\{1, \dots, \text{upper}\}$ , and let the current guess be  $\text{OPT}'$

**let** SOL =  $\emptyset$ ,  $\mathcal{U}_1 \leftarrow \mathcal{U}$ ,  $\mathcal{S}_1 \leftarrow \mathcal{S}$

**let**  $\epsilon' \leftarrow \frac{\epsilon}{t \log_2(n)}$ ,  $\delta' \leftarrow \frac{\delta}{t \log_2(n)}$ .

**for**  $i = 1, \dots, t$  **do**

**run** MAXCOVER( $\mathcal{U}_i, \mathcal{S}_i, \text{OPT}', \epsilon', \delta'$ ) to obtain sets

$\pi_i = \{\pi_{i,1}, \dots, \pi_{i,\text{OPT}'}\}$ .

**let** SOL  $\leftarrow$  SOL  $\cup \pi_i$

**let**  $\mathcal{U}_{i+1} \leftarrow \mathcal{U}_i - \bigcup_{j=1}^{\text{OPT}'}$   $\pi_{i,j}$ ,  $\mathcal{S}_{i+1} \leftarrow \mathcal{S}_i - \pi_i$ .

**endfor**

**let**  $\gamma$  be the number of elements covered by SOL

**let**  $\hat{\gamma} = \gamma + \frac{\log n}{\epsilon'} + \text{Lap}(1/\epsilon')$

**if**  $\hat{\gamma} \geq \rho n$  **increase**  $\text{OPT}'$ ; otherwise, **decrease**  $\text{OPT}'$

**Output:** SOL for minimum  $\text{OPT}'$  satisfying  $\hat{\gamma} \geq \rho n$

---

**Theorem 2.8.** Algorithm 2 is  $(\epsilon, \delta)$ -DP. Furthermore, assuming the optimal set cover has size  $\text{OPT} \leq \frac{C(1-\frac{\rho}{2})n\epsilon_0}{\ln^3(n)}$ , where  $C$  is from Lemma 2.7 and  $\epsilon_0 = O\left(\frac{\epsilon/t}{\ln \ln(n) + \ln(t/\delta)}\right)$ , Algorithm 2 outputs a  $(1 - \rho)$ -Partial Set Cover using at most  $O(\log(\frac{1}{1-\rho})) \cdot \text{OPT}$  sets with probability  $1 - \tilde{O}(\frac{1}{n})$ .

*Proof.* First, let's consider the privacy guarantee. For each iteration of the binary search, we run MAXCOVER  $t$  times. By basic composition, we this is  $(t\epsilon', t\delta')$ -differentially private. Additionally,  $\hat{\gamma}$  is the output of the Laplace Mechanism, so it is  $(t\epsilon', 0)$ -differentially private. The binary search takes at most  $\log_2(\text{upper}) \leq \log_2(n)$  iterations to converge, so  $(2\epsilon, \delta)$ -differential privacy follows by basic composition once again. Finally, the outputted solution SOL is  $(2\epsilon, \delta)$ -differentially private by post-processing.

Now, we will turn to the utility guarantee. Let  $\rho' = \frac{\rho+1}{2}$ ,  $\beta = 0.15$ , and let us first consider the algorithm when our guess  $\text{OPT}'$  is at least  $\text{OPT}$ . We claim that running MAXCOVER  $t$  times, as in Algorithm 2, covers at least  $\rho'n \geq \rho n + \frac{2 \log n}{\epsilon'}$  elements. Since the output of the Laplace mechanism  $\hat{\gamma}$  will be at least  $\rho n$  with probability at least  $1 - \frac{1}{n}$ , the binary search will converge to some  $\text{OPT}' \leq \text{OPT}$ . Now, consider the partial set cover output by the algorithm; we know that the number of sets used is  $t \cdot \text{OPT}' \leq t \cdot \text{OPT}$ , hence a  $t$ -approximation. Next, we will prove our claim.

If at any iteration  $i < t$ , we have  $\rho'n$  elements are covered by the previously selected sets, we are done. Suppose there remains at least  $(1 - \rho')n$  elements uncovered at all iterations  $i \leq t$ ; we will show that  $\rho'n$  elements are covered after iteration  $t$ . By definition of  $\text{OPT}$  and the fact that  $\text{OPT}' \geq \text{OPT}$ , there exists  $\text{OPT}'$  sets which cover the remaining uncovered elements. Thus, when  $\text{OPT}$  is suitably small as in the theorem statement, running MAXCOVER always covers at least an  $\alpha$ -fraction of the remaining elements, leaving a  $(1 - \alpha)$ -fraction of the remaining elements uncovered. By algebra,

running MAXCOVER  $t$  times suffices to guarantee that at most  $(1 - \rho')n$  elements remain uncovered.  $\square$

**Remark 2.** We emphasize that the guarantees in Theorem 2.8 are pseudo-approximations. That is, the approximation guarantee is with respect to the optimal set cover whereas the algorithm outputs a partial set cover.

### 3 An Application to Facility Location

To show an example where an *explicit* solution for (partial) set cover is a necessary building block for a larger algorithm, we will consider a facility location problem called MOBILEVACCCLINIC, introduced in [Li *et al.*, 2022]. The authors introduced the following generalization of the well known  $k$ -supplier problem for deploying vaccine distribution sites:

**Definition 3.1.** Let  $\mathcal{C}$  be a set of locations in a metric space with distance function  $d : \mathcal{C} \times \mathcal{C} \mapsto \mathbb{R}_{\geq 0}$ . Let  $\mathcal{P}$  be a set of  $n$  people where each person  $p \in \mathcal{P}$  is associated with a set  $S_p \subseteq \mathcal{C}$ , which can be interpreted as the set of locations  $p$  visits throughout the day. Finally, let  $k \in \mathbb{N}$  be a budget on the number of facilities. We want to output a set of locations  $F \subseteq \mathcal{C}$  with  $|F| \leq k$  to place facilities which minimizes  $\max_{p \in \mathcal{P}} d(S_p, F)$ , where  $d(S, F) = \min_{j \in S, j' \in F} d(j, j')$ . For simplicity, we normalize the metric so that the diameter is  $\max_{j, j' \in \mathcal{C}} d(j, j') = 1$ .

We consider the outliers version of the above problem, where we are only required to serve  $\lfloor \rho n \rfloor$  people in the population, for some  $\rho < 1$ . Since the problem was designed for vaccine distribution, the outliers variant is still very interesting: we only need to vaccinate 94% of the population in order to achieve herd immunity. For the differential privacy, the set of locations  $\mathcal{C}$ , covering requirement  $\rho$ , and budget  $k$  is public information and the individuals along with their travel patterns is private information. We call two instances of MOBILEVACCCLINIC neighbors if they differ by exactly one individual  $p$  (along with their travel pattern  $S_p$ ). Thus, we protect the privacy of the individuals: a differentially private algorithm's output will not differ *too much* whether or not any individual and their travel patterns are included in the data.

Unfortunately, we show that this problem is hard in a very strong sense: there cannot exist any (efficient or inefficient) algorithm which provides a finite approximation guarantee. Hence, we turn towards bicriteria algorithms where we violate the budget  $k$ : an algorithm is an  $(\alpha, \beta)$ -bicriteria approximation for MOBILEVACCCLINIC with outliers if it obtains an additive  $\beta$ -approximation to the optimal radius while placing at most  $\alpha k$  facilities. We show that this is possible. The idea of the algorithm is simple: we first guess the optimal radius  $R^*$  via a binary search on the interpoint distances (in practice, we just guess  $R^*$  up to an additive error of  $\gamma$  for faster convergence). Assuming we know  $R^*$ , we consider the *reverse* problem where we wish to place the fewest facilities in order to cover a  $\rho$  fraction of the clients within a radius of  $R^*$ . This is exactly a Partial Set Cover problem, so we can apply Algorithm 1 (it is easy to verify the privacy requirements of MOBILEVACCCLINIC matches that of Partial Set Cover). To formalize this, let the universe be  $\mathcal{U} = \mathcal{P}$  and let the set system be  $\mathcal{S}_R = \{S_j(R) : j \in \mathcal{C}\}$ , where

$S_j(R) = \{p \in \mathcal{P} : d(S_p, j) \leq R\}$ . It is easy to see that this indeed reduces to a Partial Set Cover problem. In the following formalization of our algorithm, let  $\alpha(\epsilon, \delta)$  denote the approximation guarantee of Algorithm 1:

---

**Algorithm 3** DPCLIENTCOVER:

---

**Input:** MOBILEVACCCLINIC instance  $(\mathcal{C}, \mathcal{S}, \mathcal{P}, k, \rho)$ , additive error  $\gamma$ , and privacy parameters  $(\epsilon, \delta)$   
**let**  $\epsilon' \leftarrow \epsilon/\log_2(1/\gamma), \delta' \leftarrow \delta/\log_2(1/\gamma)$   
**let**  $low \leftarrow 0, high \leftarrow 1$   
**while**  $high - low > \gamma$  **do**  
     $\mathcal{R} \leftarrow (high + low)/2$   
    Calculate  $\mathcal{S}_{\mathcal{R}}$  as described above  
     $F_{\mathcal{R}} \leftarrow \mathcal{M}_{PartialSetCover}(\mathcal{C}, \mathcal{S}_{\mathcal{R}}, \rho, \epsilon', \delta')$   
    **if**  $|F_{\mathcal{R}}| > \alpha(\epsilon', \delta') \cdot k$  **then**  $low \leftarrow \mathcal{R}$   
    **otherwise**  $high \leftarrow \mathcal{R}$   
**end while**  
**Output**  $F_{\mathcal{R}}$  for minimum  $R$  such that  $|F_{\mathcal{R}}| \leq \alpha(\epsilon', \delta') \cdot k$ .

---

**Theorem 3.2.** For any  $\gamma > 0$ , Algorithm 3 is  $(2\epsilon, \delta)$ -DP and an  $(O(\frac{\log^2 |\mathcal{C}| \log(1/\gamma) \log(1/\delta)}{\epsilon}), \gamma)$ -approximation algorithm for MOBILEVACCCLINIC with outliers, with probability at least  $1 - \tilde{O}(\frac{1}{|\mathcal{C}|})$ , when  $|\mathcal{C}| = O(n)$ .

*Proof.* The privacy guarantee is easy: since the guesses  $R$  don't depend on the private information,  $(2\epsilon, \delta)$ -differential privacy follows directly by basic composition. Note that we are allowed to find the minimum  $R$  such that  $|F_{\mathcal{R}}| \leq \alpha \cdot k$  by post-processing, since the  $F_{\mathcal{R}}$ 's are private already.

Next, we turn to the utility guarantee. Let us analyze the algorithm for the iterations of binary search where  $R \geq R^*$ . We will show that for these iterations, we necessarily have  $|F_{\mathcal{R}}| \leq \alpha(\epsilon', \delta') \cdot k$ . This implies that the binary search converges to some  $R$  satisfying  $R \leq R^* + \gamma$ , so the minimum  $R$  satisfying  $|F_{\mathcal{R}}| \leq \alpha(\epsilon', \delta') \cdot k$  also will satisfy  $R \leq R^* + \gamma$ .

To complete the proof, let's prove our claim. Suppose  $R \geq R^*$  and consider the corresponding Partial Set Cover instance. We know there exists some Partial Set Cover of size  $k$  (by the definition of  $R^*$ ), so our approximation guarantees in Theorem 2.4 imply that  $|F_{\mathcal{R}}| \leq \alpha(\epsilon', \delta') \cdot k$  with probability  $1 - O(\frac{1}{n})$ . Apply the union bound over all iterations of the binary search completes the proof.  $\square$

**Remark 3.** Our algorithm needs to violate the budget  $k$  by a poly-logarithmic multiplicative factor, which is often not possible in the real world. To circumvent this, we note that it has been observed experimentally that set cover algorithms obtain near optimal solutions; thus, we can implement Algorithm 3 with  $\alpha$  set to 1 and still obtain a near-optimal radius  $R$  in practice. Even this contribution is non-trivial; before our work, even good heuristics were not known for differentially private  $k$ -supplier. We use this heuristic in the experiments in Section 4.

**Remark 4.** We state Algorithm 3 and its guarantees in terms of guessing the optimal radius  $R^*$  up to an additive error of  $\gamma$ . We find that this converges faster in practice and thus obtains a better privacy/utility tradeoff. As mentioned before, binary

searching on the interpoint distances can save us this additive error of  $\gamma$  but will lose a logarithmic factor in the number of facilities placed. The proof of this is very similar, and thus omitted.

### 3.1 Lower Bounds

To give our lower bound for MOBILEVACCCLINIC, we will first show that Partial Set Cover cannot be solved exactly under  $(\epsilon, \delta)$ -differential privacy and that an approximation is necessary (deferred to Appendix).

**Lemma 3.3.** Under  $(\epsilon, \delta)$ -differential privacy, Partial Set Cover cannot be solved exactly (with high probability).

We will use this fact to derive an information theoretic lower bound for MOBILEVACCCLINIC, stating that even computationally inefficient algorithms cannot give a finite approximation factor for this problem while simultaneously satisfying approximate differential privacy. Overall, this lower bound justifies our use of bicriteria approximation algorithms for our problem. The reduction we give is similar to the one found in [Anegg *et al.*, 2022] for a separate problem,  $\gamma$ -colorful  $k$ -center.

**Theorem 3.4.** There can be no finite approximation algorithm for MOBILEVACCCLINIC which satisfies  $(\epsilon, \delta)$ -differential privacy, even when the metric space is the Euclidean line.

*Proof.* Suppose for contradiction some differentially private approximation algorithm did exist with approximation ratio  $\alpha = \alpha(n, m, \rho, \epsilon, \delta)$ . We will show that it can be used to solve any instance  $(\mathcal{U}, \mathcal{S}, \rho, \epsilon, \delta)$  of Differentially Private Partial Set Cover to optimality. By contradiction, we will conclude such an approximation algorithm cannot exist.

Let the metric space be the Euclidean line, let  $\mathcal{C} = \{1, \dots, |\mathcal{S}|\}$  and let the covering requirement be  $\rho = \rho$ . For each element  $u_\ell \in \mathcal{U}$ , define a person  $p \in \mathcal{P}$  and let their travel locations be  $S_p = \{i \in [|\mathcal{S}|] : u_\ell \in S_i\}$ . By construction, it is clear that the optimal solution for this MOBILEVACCCLINIC instance has radius 0 and gives a solution (of size  $k$ ) to the Partial Set Cover instance. Furthermore, any  $\alpha$ -approximation will output a solution with radius  $\alpha \cdot 0 = 0$  as well. Consequently, we can binary search over  $k \in [m]$  to use any  $\alpha$ -approximation for MOBILEVACCCLINIC to obtain an optimal solution for Partial Set Cover.  $\square$

## 4 Experiments

	Clients	Locations	Diameter (km)
Charlottesville	33156	5660	8.12
Albemarle	74253	9619	61.62

Table 1: Network Information

For our experiments in our main paper, we focus on the vaccine distribution problem we studied in Section 3. We run all of our experiments on synthetic mobility data from Charlottesville city and Albemarle county, which are part of a synthetic U.S. population (see [Chen *et al.*, 2021; Machi *et*

al., 2021] for details). The data keeps track of the locations each person travels to throughout a week, so each person is associated with a set of activity locations. We use this data in our experiments to study the following:

- **Privacy/Accuracy Tradeoff:** how does the quality of our algorithm vary as the privacy parameters ( $\epsilon, \delta$ ) vary?
- **Cost of Privacy:** how does our private algorithm compare with the non-private one given in [Li et al., 2022]?

Due to space constraints, we defer some of these experiments and analyses to the full version. There, we also investigate the above questions for our Partial Set Cover algorithms.

**Baseline:** We use a non-private greedy algorithm to solve the partial set cover problem as our *Baseline*. Starting with an empty set of locations, the greedy algorithm at each step chooses one location that covers the most of the uncovered clients. Then it adds the location to its set and updates the current state of covered clients. The baseline algorithm utilizes the same binary search routines with DPCLIENTCOVER to find the optimal radius.

**Metric:** All of our comparisons are made based on the objective value of MOBILEVACCCLINIC with outliers. Informally, by calculating the service cost of each person, our objective is to minimize the distance the  $\lceil \rho n \rceil^{th}$  person needs to travel in order to get vaccinated. Formally, our objective is  $\min_{P_1 \subseteq \mathcal{P}: |P_1| = \lceil \rho n \rceil} \max_{p \in P_1} d(S_p, F)$ .

### 4.1 Privacy/Accuracy/Budget Tradeoff

In this subsection, we investigate the two questions regarding the privacy/accuracy trade-off and the cost of privacy. We set  $\delta = 10^{-6}$  and run our private vaccine distribution algorithm for different budgets  $k \in \{4, \dots, 16\}$  and different privacy parameters  $\epsilon \in \{0.25, 0.5, 1, 2, 4\}$ . A careful reader may notice that Theorem 2.4 only applies for  $\epsilon \in (0, 1)$ , and may worry that our experimental regime doesn't have privacy guarantees. However, DPCLIENTCOVER makes calls to the Partial Set Cover algorithm with parameter  $\epsilon' = \epsilon / \log_2(1/\gamma)$ , so we still have differential privacy.

The experimental results show that, generally, with higher values of budget, all algorithms perform better, i.e., they achieve lower objective values. Our private algorithm DPCLIENTCOVER usually achieves higher accuracy with higher  $\epsilon$  at every level of budget, while the non-private algorithm (*Baseline*) always has the best performance. These results are expected. Figure 2 also shows that our private algorithm performs well, obtaining near optimal choices of facility locations for  $\epsilon \geq 1$ . For example, DPCLIENTCOVER with  $\epsilon = 4$  matches the performance of *Baseline* on the Albemarle dataset and almost matches the *Baseline* at several values of budget on the Charlottesville dataset. However, we can observe some trade-offs between privacy and utility, especially in high-privacy settings. In particular, for  $\epsilon = 0.25$  and  $k = 4$ , there is a high cost of guaranteeing privacy in both datasets. In these cases, the objectives of the private settings are 4–7 times higher than those of the *Baseline*.

A more subtle result that Figure 2 gives is to help us understand the trade-off between the budget and the privacy loss  $\epsilon$ . From a policy-making perspective, this is interesting since we can allocate more mobile vaccine distribution sites in return

for a stronger privacy guarantee. For example, the accuracy from  $k = 8$  facilities with  $\epsilon = 0.25$  is similar to using  $k = 4$  facilities with  $\epsilon = 0.5$ . Figure 1 shows the vaccine clinics chosen by the private ( $\epsilon = 0.5$ , white spots) and non-private (black spots) algorithms for the Charlottesville dataset with  $k = 8$ . The two algorithms tend to choose two totally different distributions of vaccine clinics, although they do not differ much in term of objective values (see Figure 2, where  $k = 8$ ).




Figure 2: Utilities of DPCLIENTCOVER at  $\rho = 0.8$  with different values of privacy parameter ( $\epsilon$ ) and budget ( $k$ ) on both datasets. The  $y$ -axis indicates the objectives measured at the 80<sup>th</sup> percentiles (in meters). The  $x$ -axis shows different budget values, i.e., the number of selected locations. Generally, DPCLIENTCOVER performs closely to the *Baseline* with large values of  $\epsilon$  ( $\epsilon \geq 2$ ) on both datasets. For the Albemarle county dataset, curves for the *Baseline* and DPCLIENTCOVER with  $\epsilon = 4$  overlap.

## 5 Discussion

In this paper, we consider the Set Cover problem and the  $k$ -supplier problem, which have strong hardness results under differential privacy. We observe that partial coverage of elements/clients suffices to avoid the impossibility results, and give the first non-trivial approximation algorithms for both problems. Overall, our work is an important step in tackling the impossibility results in differentially private algorithm design and leaves many interesting problems open:

- Both of our algorithms for Partial Set Cover require some (relatively loose) assumption on the set system. An interesting question is whether we can remove these assumptions: can we obtain a general approximation algorithm for Partial Set Cover under differential privacy?
- Our algorithm for  $k$ -supplier violates the budget  $k$  by a poly-logarithmic factor, which is impractical in some settings. It is interesting to see what guarantees are possible without violating the budget: can we obtain true approximations for  $k$ -supplier with outliers with privacy?
- As mentioned before, the facility location problem has a  $\Omega(\sqrt{n})$  approximation hardness result under differential privacy. It is interesting to see if our ideas can help avoid this: does allowing partial coverage circumvent hardness results of the uniform facility location problem?

## Acknowledgements

Dung Nguyen and Anil Vullikanti were partially supported by NSF grants CCF 1918656, IIS 1955797, NIH grant 2R01GM109718-07, and CDC cooperative agreement MIND U01CK000589. George Li was supported in part by NSF award number CCF-1918749.

## References

- [Alon *et al.*, 2003] Noga Alon, Baruch Awerbuch, Yossi Azar, Niv Buchbinder, and Joseph Naor. The online set cover problem. In Lawrence L. Larmore and Michel X. Goemans, editors, *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, June 9-11, 2003, San Diego, CA, USA*, pages 100–105. ACM, 2003.
- [Anegg *et al.*, 2022] Georg Anegg, Haris Angelidakis, Adam Kurpisz, and Rico Zenklusen. A technique for obtaining true approximations for k-center with covering constraints. *Math. Program.*, 192(1):3–27, 2022.
- [Balcan *et al.*, 2017] Maria-Florina Balcan, Travis Dick, Yingyu Liang, Wenlong Mou, and Hongyang Zhang. Differentially private clustering in high-dimensional euclidean spaces. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 322–331. PMLR, 2017.
- [Blocki *et al.*, 2021] Jeremiah Blocki, Elena Grigorescu, and Tamalika Mukherjee. Differentially-private sublinear-time clustering. In *IEEE International Symposium on Information Theory, ISIT 2021, Melbourne, Australia, July 12-20, 2021*, pages 332–337. IEEE, 2021.
- [Brown *et al.*, 2021] Gavin Brown, Marco Gaboardi, Adam D. Smith, Jonathan R. Ullman, and Lydia Zakyntinou. Covariance-aware private mean estimation without private covariance estimation. *CoRR*, abs/2106.13329, 2021.
- [Canonne *et al.*, 2020] Clément L. Canonne, Gautam Kamath, Audra McMillan, Jonathan R. Ullman, and Lydia Zakyntinou. Private identity testing for high-dimensional distributions. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [Chen *et al.*, 2021] Jiangzhuo Chen, Stefan Hoops, Achla Marathe, Henning Mortveit, Bryan Lewis, Srinivasan Venkatramanan, Arash Haddadan, Parantapa Bhattacharya, Abhijin Adiga, Anil Vullikanti, Aravind Srinivasan, Mandy L Wilson, Gal Ehrlich, Maier Fenster, Stephen Eubank, Christopher Barrett, and Madhav Marathe. Prioritizing allocation of covid-19 vaccines based on social contacts increases vaccination effectiveness. *medRxiv*, 2021.
- [Cohen-Addad *et al.*, 2022a] Vincent Cohen-Addad, Alessandro Epasto, Silvio Lattanzi, Vahab S. Mirrokni, Andres Munoz, David Saulpic, Chris Schwiegelshohn, and Sergei Vassilvitskii. Scalable differentially private clustering via hierarchically separated trees. *CoRR*, abs/2206.08646, 2022.
- [Cohen-Addad *et al.*, 2022b] Vincent Cohen-Addad, Yunus Esencayi, Chenglin Fan, Marco Gaboardi, Shi Li, and Di Wang. On facility location problem in the local differential privacy model. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *International Conference on Artificial Intelligence and Statistics, AISTATS 2022, 28-30 March 2022, Virtual Event*, volume 151 of *Proceedings of Machine Learning Research*, pages 3914–3929. PMLR, 2022.
- [Dwork and Roth, 2014] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4):211–407, aug 2014.
- [Dwork *et al.*, 2006] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.
- [Esencayi *et al.*, 2019] Yunus Esencayi, Marco Gaboardi, Shi Li, and Di Wang. Facility location problem in differential privacy model revisited. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8489–8498, 2019.
- [Eubank *et al.*, 2004] Stephen G. Eubank, V. S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, and Nan Wang. Structural and algorithmic aspects of massive social networks. In J. Ian Munro, editor, *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, New Orleans, Louisiana, USA, January 11-14, 2004*, pages 718–727. SIAM, 2004.
- [Ghazi *et al.*, 2020] Badih Ghazi, Ravi Kumar, and Pasin Manurangsi. Differentially private clustering: Tight approximation ratios. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4040–4054. Curran Associates, Inc., 2020.
- [Ghazi *et al.*, 2021] Badih Ghazi, Noah Golowich, Ravi Kumar, Pasin Manurangsi, and Chiyuan Zhang. On deep learning with label differential privacy. *CoRR*, abs/2102.06062, 2021.
- [Gupta *et al.*, 2010] Anupam Gupta, Katrina Ligett, Frank McSherry, Aaron Roth, and Kunal Talwar. Differentially private combinatorial optimization. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 1106–1125. SIAM, 2010.



- [Hay *et al.*, 2009] Michael Hay, Chao Li, Gerome Miklau, and David D. Jensen. Accurate estimation of the degree distribution of private networks. In Wei Wang, Hillol Kargupta, Sanjay Ranka, Philip S. Yu, and Xindong Wu, editors, *ICDM 2009, The Ninth IEEE International Conference on Data Mining, Miami, Florida, USA, 6-9 December 2009*, pages 169–178. IEEE Computer Society, 2009.
- [Hsu *et al.*, 2014] Justin Hsu, Aaron Roth, Tim Roughgarden, and Jonathan R. Ullman. Privately solving linear programs. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 612–624. Springer, 2014.
- [Jones *et al.*, 2021] Matthew Jones, Huy L. Nguyen, and Thy D. Nguyen. Differentially private clustering via maximum coverage. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 11555–11563. AAAI Press, 2021.
- [Karwa *et al.*, 2011] Vishesh Karwa, Sofya Raskhodnikova, Adam D. Smith, and Grigory Yaroslavtsev. Private analysis of graph structure. *Proc. VLDB Endow.*, 4(11):1146–1157, 2011.
- [Lee and Kifer, 2018] Jaewoo Lee and Daniel Kifer. Concentrated differentially private gradient descent with adaptive per-iteration privacy budget. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '18*, page 1656–1665, New York, NY, USA, 2018. Association for Computing Machinery.
- [Li *et al.*, 2022] George Z. Li, Ann Li, Madhav Marathe, Aravind Srinivasan, Leonidas Tsepenekas, and Anil Vullikanti. Deploying vaccine distribution sites for improved accessibility and equity to support pandemic response. In Piotr Faliszewski, Viviana Mascardi, Catherine Pelachaud, and Matthew E. Taylor, editors, *21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2022, Auckland, New Zealand, May 9-13, 2022*, pages 789–797. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 2022.
- [Machi *et al.*, 2021] D. Machi, P. Bhattacharya, S. Hoops, J. Chen, H. Mortveit, S. Venkatramanan, B. Lewis, M. Wilson, A. Fadikar, T. Maiden, C. L. Barrett, and M. V. Marathe. Scalable epidemiological workflows to support covid-19 planning and response. In *2021 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 639–650, Los Alamitos, CA, USA, May 2021. IEEE Computer Society.
- [Mitrovic *et al.*, 2017] Marko Mitrovic, Mark Bun, Andreas Krause, and Amin Karbasi. Differentially private submodular maximization: Data summarization in disguise. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 2478–2487. PMLR, 2017.
- [Moshkovitz, 2015] Dana Moshkovitz. The projection games conjecture and the np-hardness of  $\ln n$ -approximating set-cover. *Theory of Computing*, 11(7):221–235, 2015.
- [Nguyen and Vullikanti, 2021] Dung Nguyen and Anil Vullikanti. Differentially private densest subgraph detection. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8140–8151. PMLR, 2021.
- [Nissim *et al.*, 2007] Kobbi Nissim, Sofya Raskhodnikova, and Adam D. Smith. Smooth sensitivity and sampling in private data analysis. In David S. Johnson and Uriel Feige, editors, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 75–84. ACM, 2007.
- [Vadhan, 2017] Salil P. Vadhan. The complexity of differential privacy. In Yehuda Lindell, editor, *Tutorials on the Foundations of Cryptography*, pages 347–450. Springer International Publishing, 2017.
- [Williamson and Shmoys, 2011] David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011.
- [Wolsey, 1982] Laurence A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Comb.*, 2(4):385–393, 1982.