# PCOR: Private Contextual Outlier Release
# via Differentially Private Search

Masoumeh Shafieinejad
University of Waterloo
Waterloo, Canada
masoumeh@uwaterloo.ca

Florian Kerschbaum
University of Waterloo
Waterloo, Canada
fkerschb@uwaterloo.ca

Ihab F. Ilyas
University of Waterloo
Waterloo, Canada
ilyas@uwaterloo.ca

## ABSTRACT

Outlier detection plays a significant role in various real world applications such as intrusion, malfunction, and fraud detection. Traditionally, outlier detection techniques are applied to find outliers in the context of the whole dataset. However, this practice neglects the data points, namely contextual outliers, that are not outliers in the whole dataset but in some specific neighborhoods. Contextual outliers are particularly important in data exploration and targeted anomaly explanation and diagnosis. In these scenarios, the data owner computes the following information: i) The attributes that contribute to the abnormality of an outlier (metric), ii) Contextual description of the outlier's neighborhoods (context), and iii) The utility score of the context, e.g. its strength in showing the outlier's significance, or in relation to a particular explanation for the outlier. However, revealing the outlier's context leaks information about the other individuals in the population as well, violating their privacy. We address the issue of population privacy violations in this paper. There are two main challenges in defining and applying privacy in contextual outlier release. In this setting, the data owner is required to release a *valid context* for the queried record, i.e. a context in which the record is an outlier. Hence, the first major challenge is that the privacy technique must preserve the validity of the context for each record. We propose techniques to protect the privacy of individuals through a relaxed notion of differential privacy to satisfy this requirement. The second major challenge is applying the proposed techniques efficiently, as they impose intensive computation to the base algorithm. To overcome this challenge, we propose a graph structure to map the contexts to, and introduce differentially private graph search algorithms as efficient solutions for the computation problem caused by differential privacy techniques.

## CCS CONCEPTS

• **Security and privacy → Data anonymization and sanitization**.

## KEYWORDS

differential privacy, contextual outlier detection, graph search, private sampling

## 1 INTRODUCTION

We borrow the contextual outliers example of Han et al. [24]: "The temperature today is $28°C$. Is it exceptional (i.e., an outlier)? It depends, for example, on the time and location! If it is in winter in Toronto, yes, it is an outlier. If it is a summer day in Toronto, then it is normal. Unlike traditional outlier detection, in this case, whether or not today's temperature value is an outlier depends on the context — the date, the location, and possibly some other factors". Contextual outlier detection helps with the problem of *hidden* outliers — the data points that are considered *normal* compared to the whole population, but are outliers when compared to a subset of the dataset. Consider an analyst performing market research who wishes to determine the companies that are unusually profitable. While Company $V$ might have normal reported profit compared to all companies, it may have outlying profit when compared to media companies with less than 2000 employees. In this case, the released context for Company $V$ is [Business = media] $\wedge$ [$|Employees|$ < 2000]. For the same data point $V$ there might be multiple contexts (explanations) in which $V$ is an outlier, each of which might have a different *utility* value. One example for utility is the size of the population that is covered by the context, which indicates the significance of the outlier. Another example is the context's overlap with a desired context, suitable for scenarios where the private explanation is required to be in a close relation with a chosen context. The data owner releases a context for $V$ that ideally achieves the highest utility. We assume reporting the outlier data point $V$ is allowed when covered by the business purpose of performing the analysis. However, releasing the context — although useful for the data analyst — leaks information about other individuals in the dataset. This might not be permitted without the users' consent, according to privacy laws (e.g. GDPR [4]).

### 1.1 Technical Challenges

The privacy for statistical queries of various types, such as average, sum, variance, histogram, median, and maximum likelihood estimator, have been investigated [14, 15, 37] in the literature. However,

the studies of privacy in outlier analysis are few, and mostly – described in more detail in Section 7 – addressing the privacy of the outlier [7, 20, 33, 39]. Consequently, the privacy of individuals in the remaining population has remained neglected, leaving the private contextual outlier reporting (PCOR) design as an open problem.

Differential privacy [13, 15, 17] is a popular privacy notion in data analysis [19] that has shown success in providing privacy when applied to products made by organizations such as the US Census Bureau [25, 34], Google [18], Apple [22] and Uber [29]. Since differential privacy i) guarantees privacy for individual records, ii) ensures privacy even in the presence of side information about the data, and iii) bounds the information leakage to a total privacy budget ($\epsilon$) across multiple data releases; it thoroughly meets the privacy needs in PCOR. However, the challenges of utilizing differential privacy in PCOR are many, including defining a meaningful notion of differential privacy for contextual outliers, finding a proper technique that fits PCOR's requirements, and proving that the privacy budget is reasonably bounded. When queried about a record, PCOR must select a *private valid* context with ideally high utility among all candidates. To achieve this goal, we propose a variation of the Exponential mechanism [1] in PCOR, as this mechanism provides differential privacy when the nature of the candidates is discrete. However, the conventional definition of privacy and neighboring datasets in the Exponential mechanism does not guarantee the validity of the context in the outlier release case. This motivates us to propose a variation of the Exponential mechanism that guarantees the query response is valid and provides a relaxed notion of differential privacy. Our proposed technique, similar to the conventional Exponential mechanism, in its formulaic application inevitably enumerates all possible contexts to select one. This requirement heavily affects the performance of the scheme, impeding its practicality. To summarize, our challenge is to provide a design for private contextual outlier reporting that

(1) Defines a notion of differential privacy for contextual outliers, and provides a technique for achieving it. As a result, in most cases, the output of PCOR is approximately the same ($\leq e^\epsilon$), even if any single record in the input database is arbitrarily added or removed;

(2) Reduces the complexity from exponential time in the formulaic (direct) approach to polynomial time. This reduction however, should not affect whether the design supports differential privacy;

(3) Achieves high utility. Supporting privacy and scalability, should not prevent the design from meeting its purpose in providing *useful* data for the analyst, and;

(4) Is compatible with any utility function (e.g. population size) as well as any outlier detection algorithm, instead of being designed for a restricted number of specific cases.

## 1.2 Our Contributions

We propose PCOR, a framework for privately releasing contextual outliers, that addresses the aforementioned challenges:

(1) PCOR utilizes a technique for a relaxed notion of differential privacy that prevents an adversary from learning information about the individuals in the context other than the reported outlier.

(2) We propose deploying a sampling layer – based on a graph structure – prior to applying the Exponential mechanism, to reduce the complexity of PCOR to polynomial time. In our efficient sampling, we analyze the Breadth-First and the Depth-First search algorithms. We modify these algorithms to support differential privacy for the outcome. We contribute the first presentation of differentially private graph search algorithms in the literature.

(3) We prove our proposed solution is private and compare our algorithm in privacy and computational complexity to alternative solutions. The performance and utility of our algorithm is also supported by our experimental results.

(4) We show that our framework can accommodate any utility function for the contexts. We provide results for two examples of utility functions that correspond to two separate classes of outlier release. One, evaluates the context based on the size of the population it covers while the other, evaluates the contexts based on its relation (overlap) with a chosen context. Ultimately, we claim our algorithm works for the generic case of reporting the contextual outliers and is not limited to any particular outlier detection algorithm. Our claims are reinforced by experimental results over three algorithms for different outlier detection categories.

PCOR solves the challenge of simultaneously providing privacy, utility, and performance in reporting contextual outliers. Applied on a dataset of 50,000 records, PCOR reduces the runtime from three days in the direct differentially private approach to 37 minutes; while it maintains 90% of the maximum utility[2] and guarantees relaxed (output constrained, defined in Section 2.4) differential privacy with $\epsilon = 0.2$.

## 1.3 Paper Organization

The rest of this paper is organized as follows: In Section 2 we describe the outlier detection algorithms and define differential privacy. In Section 3 we provide the formal model and problem definitions for our private contextual outlier release (PCOR). We discuss the direct approach of embedding a differential privacy mechanism in contextual outlier detection in Section 4 and address the performance challenge. In Section 5, we introduce a graph representation of contexts used in our efficient sampling algorithms. Privacy proofs of the candidate algorithms are also covered in Section 5. Their performance and accuracy are evaluated through the experiments described in Section 6. Our results in Section 6 confirm that our final candidate succeeds in achieving privacy, accuracy and efficiency. In Section 7, we situate our work in the current body of research, and in Section 8 we conclude our work.

---

[1]We introduce the Exponential mechanism in Section 2.3.

[2]These numbers correspond to the experiments where the utility is "population size". The experiments for the "overlap" utility run even faster (Section 6.4).

## 2 PRELIMINARIES

### 2.1 Outlier Detection Algorithms

Outlier detection algorithms fall into three main categories: i) Statistics-based, ii) Distance-based and iii) Model-based methods [28]. In statistics-based outlier detection techniques [3, 9, 23, 27, 35, 40, 41, 43] it is assumed that normal data points appear in high probability regions of a stochastic model, while outliers appear in its low probability regions. There are two different types of statistics-based approaches: a) Hypothesis testing, and b) Distribution fitting approaches. The hypothesis testing [23, 41, 43] approach calculates a test statistic based on the data points, then it determines whether the null hypothesis claiming no outlier exists in the dataset, is valid or not. We evaluate Grubbs' test [23] as a hypothesis testing outlier detection. The distribution fitting approach [3, 27, 35, 40] first fits a probability density function to the data. Next, it labels the points with low probability as outliers. We select the Histogram approach [3] for the fitting distribution category. We also investigate Local Outlier Factor [8] as the distance-based outlier detection method in our experiments. The local outlier factor (LOF) method scores data points based on the density of their neighboring data points and detects the outlier points based on the scores.

Contrary to the outlier detection algorithms described above, model-based techniques detect the outliers in a supervised approach. However, in this paper we focus on the unsupervised outlier detection algorithms as they do not require access to labeled data.

### 2.2 Differential Privacy

Differential privacy (DP), was introduced in 2006 by Dwork et al. [16]. To protect the individual's data, DP mechanisms require that changing one record in the database should not change the output distribution much. The formal definition is as follows.

**Definition 2.1.** (Differential Privacy [16]) A randomized function $M : \mathcal{D} \to O$ : gives $\epsilon$-differential privacy if for any set $S \subseteq O$ and any pair of neighboring datasets $D_1$ and $D_2$ from $\mathcal{D}$ denoted by $(D_1, D_2) \in \mathcal{N}$, such that $D_1$ and $D_2$ differ by adding/removing a record:

$$Pr[M(D_1) \in S] \leq e^\epsilon \times Pr[M(D_2) \in S] \qquad (1)$$

### 2.3 The Exponential Mechanism

Exponential mechanism is a differential privacy mechanism used for non-numeric values. The goal of this mechanism is to randomly map a set of $n$ inputs each from a domain $\mathcal{D}$ to some output in a range $\mathcal{R}$ [36]. There are no specific assumptions about the nature of $\mathcal{D}$ or $\mathcal{R}$ other than a base measure $\mu$ on $\mathcal{R}$ exists. We proceed with the common assumption that the base measure $\mu$ is uniform. The Exponential mechanism is driven by a utility function $u : \mathcal{D} \times \mathcal{R} \to \mathbb{R}$ that assigns a real valued score to any pair $(D, r)$ from $\mathcal{D} \times \mathcal{R}$, higher scores indicate more appealing $r$'s for $D$. Given a $D \in \mathcal{D}$ the goal of the mechanism is to return an $r \in \mathcal{R}$ such that $u(D, r)$ is (approximately) maximized while guaranteeing differential privacy. To provide the definition for the Exponential mechanism, we need to define *sensitivity* first.

**Definition 2.2.** (Sensitivity) The sensitivity of a function $u : \mathcal{D} \times \mathcal{R} \to \mathbb{R}$, is defined as

$$\Delta u = max_{(D_1, D_2) \in \mathcal{N}, r \in \mathcal{R}} = |u(D_1, r) - u(D_2, r)|.$$

The sensitivity of $u$ is the maximum change in the function output resulted from replacing any individual's data with a different one. In other words, sensitivity is the maximum change in $u$, if we switch from any input dataset $D_1$ to any neighbor $D_2$ of $D_1$.

**Definition 2.3.** (The Exponential mechanism) [36] For any function $u : \mathcal{D} \times \mathcal{R} \to \mathbb{R}$, the Exponential mechanism is defined as follows:

$Exp_u^\epsilon(D, \mathcal{R}) :=$ Choose $r$,

$$Pr[r] = \frac{exp(\frac{\epsilon u(D, r)}{2\Delta u})}{\sum_{r' \in \mathcal{R}} exp(\frac{\epsilon u(D, r')}{2\Delta u})} \qquad (2)$$

We include Theorem 2.1 for the privacy of the Exponential mechanism and refer to McSherry et al.'s work [36] for the its proof. For PCOR's accuracy, we present empirical evaluations in Section 6.

**Theorem 2.1.** Theorem 1. (Privacy [36]). The Exponential mechanism $Exp_u^\epsilon(d)$ gives ($2\epsilon\Delta u$)-differential privacy.

As shown in Theorem 2.1, for the Exponential mechanisms to be the most useful, sensitivity ($\Delta u$) must be limited. Commonly the utility function is chosen such that $\Delta u \leq 1$, so that $Exp_u^\epsilon$ ensures ($2\epsilon$)-differential privacy.

### 2.4 Output Constrained Differential Privacy

To provide privacy in our PCOR, we utilize a relaxed notion of differential privacy, namely Output Constrained Differential Privacy [26]. We postpone the reasoning for this choice to Section 3.1, after introducing the notations and defining the problem at the beginning of Section 3.

**Definition 2.4.** ($f$-Neighbors) Given function $f : \mathcal{D} \to O$, for any pair of datasets $D_1, D_2 \in \mathcal{D}$, the datasets $D_1$ and $D_2$ are neighbors w.r.t. $f$, denoted by $\mathcal{N}(f(.))$, if

(1) $f(D_1) = f(D_2) \neq \emptyset$
(2) $(D_1, D_2) \in \mathcal{N}$

The $f$-Neighbors $D_1$ and $D_2$, not only have to satisfy the general neighboring condition and differ from each other in only one record, but also need to result in the same (non-empty) output result/set when the function $f$ is applied to them. In our outlier release use case, the function $f$ outputs the set of all valid contexts in the dataset for a particular outlier. We elaborate outlier release in Section 3 and define $f$ in Definition 3.1.

**Definition 2.5.** (Output Constrained Differential Privacy – OCDP) A randomized mechanism $M$ satisfies $(\epsilon, f)$−OCDP, if for any $(D_1, D_2) \in \mathcal{N}(f(.))$, and every set of outputs $S \subseteq range(M)$, we have:

$$Pr[M(D_1) \in S] \leq e^\epsilon Pr[M(D_2) \in S] \qquad (3)$$

## 3 PROBLEM DEFINITION

The outlier release encourages the anomaly or discourages it, due to approval or disapproval incentives respectively. In the approval

**Table 1: A sample of income data set $D$**

| Record | Job title | City | District | Salary |
|--------|-----------|------|----------|--------|
| 1 | Medical Doctor | Montreal | Business | $S_1$ |
| 2 | Lawyer | Toronto | Business | $S_2$ |
| 3 | CEO | Ottawa | Diplomatic | $S_3$ |
| 4 | Lawyer | Toronto | Business | $S_4$ |
| 5 | Lawyer | Ottawa | Diplomatic | $S_5$ |
| 6 | Medical Doctor | Toronto | Historic | $S_6$ |
| 7 | Lawyer | Ottawa | Business | $S_7$ |
| 8 | Lawyer | Ottawa | Diplomatic | $S_8$ |
| 9 | CEO | Montreal | Historic | $S_9$ |
| 10 | Medical Doctor | Toronto | Diplomatic | $S_{10}$ |

incentive, the outlier benefits from the announcement as their out-lying is desired. However, in the latter, the anomaly is undesired and needs to be proven to penalize the outlier. For approval incentive, imagine a funding agency that awards people with outstanding achievement in underrepresented groups. In this example, the agency is allowed to announce the outlier and the privacy-enhanced context, while any individual in that context is able to deny their participation in the funding competition. For disapproval incentive, imagine an insurance company increasing the annual payment of a client due to the client's abnormal behavior in a group of clients. While the insurance company is required to provide justification for this increase, they are not allowed to reveal information about any other client. To clarify how contextual outlier release violates the privacy of the participants, we provide an example of income analysis with concrete attribute values in Table 1. The goal in this example is to reveal the context that covers the largest population for an outlier. $D$ is a dataset with categorical attributes of $Jobtitle = \{CEO, MedicalDoctor, Lawyer\}$, $City$ taking values from $\{Montreal, Ottawa, Toronto\}$, and $District$ with the domain $\{Business, Historic, Diplomatic\}$, and a numerical attribute Salary. Assume that the data owner reveals the record $V$, e.g. record 8 in Table 1, has its highest anomaly significance in the context $Jobtitle \in \{CEO, Lawyer\}$, in $Diplomatic$ district of $Ottawa$. By having a side information about the individuals in the maximum context, e.g. by knowing there is only one $CEO$ living in $Diplomatic$ district of $Ottawa$, the revealed deterministic statement about the outlier $V$, leaks information about the presence of the $CEO$ in the dataset as well.

In our setup, we consider a dataset instance, $D$, of a relational schema, $R$. Attributes of $R$ are presented by the set $attr(R) = \{A_1, \cdots, A_m, M\}$; where $M$ is the *metric* attribute we define the outlier with regard to. We denote by $P_{ij}$, a predicate that is defined over the $j^{th}$ value in the domain of a categorical or numerical attribute $A_i$ derived from $attr(R)$. For instance, in our running example, $P_{23}$ is the third attribute value in the second attribute in $attr(R) = \{Jobtitle, City, District, Salary\}$, which is $City = Toronto$. The domain of an attribute $A_i$ includes all possible values for $A_i$, regardless of them being covered in the particular instance $D$; we refer to the domain size of $A_i$ by $|A_i|$. In reporting private contextual outliers, it is necessary to take all the values in the domain of the attributes into account; the reason for this

requirement is discussed in Section 4. A predicate $P_{ij}$ filters the tuples from $D$ to a subset that satisfies the predicate. This filtering can be represented as conjunction of disjunction predicates; $[P_{11} \vee P_{12} \vee \cdots \vee P_{1|A_1|}] \wedge \cdots \wedge [P_{m1} \vee P_{m2} \vee \cdots \vee P_{m|A_m|}]$; this format covers many common types of SQL queries.

Now, we can define a *context*, $C$. A context is represented as a binary vector of the form $\langle c_{11}, \cdots, c_{1|A_1|}, \cdots, c_{m1}, \cdots, c_{m|A_m|} \rangle$ with length $t = \sum_{i=1}^{m} |A_i|$. The bit $c_{ij}$ in $C$ is set to 1, if the predicate $P_{ij}$ is covered in the context. $C$ filters the dataset $D$ to the population $D_C$. A tuple $V$ belongs to $D_C$, denoted as $V \in D_C$, if $V$ is selected by a set of predicates $P_{ij}, \forall i \in [1, m]$ in $C$, where $1 \leq j \leq |A_i|$. We denote the total number of tuples in $C$'s population as $|D_C|$. In our example of the dataset with categorical attributes of $attr(R) = \{Jobtitle, City, District, Salary\}$ with the attribute domains of $\{CEO, MedicalDoctor, Lawyer\}$ for $Jobtitle$, $\{Montreal, Ottawa, Toronto\}$ for $City$, as well as $District$ attributes $\{Business, Historic, Diplomatic\}$, this combination of predicates: $[P_{11} \vee P_{13}] \wedge [P_{23}] \wedge [P_{32}]$, filters the dataset to the $CEO$s and $Lawyer$s in $Toronto$ who live in the $Historic$ district. The corresponding context to this subset of the dataset is represented by $C = \langle 101001010 \rangle$. Evidently, any non-empty context should include at least one predicate of each attribute, i.e. the context vector has a minimum Hamming weight[3], $m$. We call context $C'$ connected to $C$, if the Hamming distance of $C'$ and $C$ is 1, i.e. $C'$ is constructable from $C$ by adding or removing (not both) only one predicate. For our running example of $C = \langle 101001010 \rangle$, being $CEO$s and $Lawyer$s in $Toronto$ who live in the $Historic$ district, the context $C' = \langle 100001010 \rangle$ which is $CEO$s in $Toronto$ who live in the $Historic$ district, is connected to $C$. Now, we can define the contextual outlier release. Given an outlier detection algorithm that takes as input any population $D_C$ and the metric $M$, and outputs all outlier in $D_C$, we can implement the outlier verification method $f_M(D_C, V)$ to determine if $V$ is an outlier in $C$ with regard to the metric $M$; formally

$$f_M(D_C, V) = \begin{cases} true, & \text{V is an outlier in } D_C \text{ w.r.t. M} \\ false, & \text{Otherwise} \end{cases}$$

We call a context $C$, a *matching context* for an outlier $V$ with regard to metric $M$ if and only if $f_M(D_C, V) = true$.

**Definition 3.1.** (Contextual Outlier Enumeration, $COE_M$) Given a dataset $D$ with schema $R$, a metric attribute M, and an outlier verification function $f_M$ for $M$, $COE_M(D, V)$ produces all matching contexts of $V$; i.e. all $C$'s from the attributes $attr(R)$ such that $V \in D_C$, and $f_M(D_C, V) = true$.

To the best of our knowledge, the problem $COE_M$ defined above is a hard problem with no existing efficient solution [32, 42]. Recall from Section 2.3, that each output of $COE_M$ is assigned a score by a utility function, where higher scores indicate more appealing outputs. We focus on two types of utility, $u_V(D, C)$, of a context $C$ for an outlier $V$: i) a quantitative evaluation of the context independent from other contexts. The population size of $C$ is an very common example for this utility type, since it indicates the outlier's significance. ii) an evaluation of the context that may score the context based on its relation with a chosen context. As an example of this utility type, we investigate the overlap of the candidate context

---

[3]The Hamming weight of a binary vector is the number of 1's in the vector.

with a fixed starting context – the results are shown in Section 6.4. In either case, the data owner desires to report $V$ with a context that is both private and achieves high utility. In this work, as implied in the definition, we only consider deterministic outlier detection algorithms embedded in $COE_M$.

**Definition 3.2.** (Private Contextual Outlier Reporting) Given a dataset $D$ with schema $R$, a metric attribute $M$, a local outlier $V$, and a contextual outlier enumeration $COE_M$ for $M$, private contextual outlier reporting ($PCOR$) produces a context $C$ for $V$ such that

- (a) $f_M(D_C, V) = true$
- (b) $C$ is produced by a differentially private mechanism
- (c) $C$ is expected to maximize utility within $COE_M(D, V)$
- (d) $C$ is calculated in $O(p(n))$

Note that the context $C$ is the only context that is released to explain the outlying of $V$. All the valid contexts calculated for generating $COE_M$ are solely known by the data owner and are used to compute $C$.

In the previous section, we elaborated on the outlier detection algorithms to deploy in $f_M(D_C, V)$, affecting the results of $COE_M$. We also defined differential privacy's [13] requirement for PCOR and presented the Exponential mechanism for implementing it. In Section 3.1 we explain why we choose Output Constrained Differential Privacy for PCOR. Subsequently, in Section 3.2, we discuss PCOR's compatibility with any utility function through the two examples mentioned above.

## 3.1 Outlier-Preserving Privacy in PCOR

As shown in Equation 1, a mechanism $\mathcal{M}$ provides differential privacy, if the probability of obtaining any of its possible outcomes $S$ changes with a maximum ratio of $e^\epsilon$, if an input dataset $D_1$ is changed to any neighboring dataset $D_2$. Hence, if there is an $S$ in range of $\mathcal{M}$ that has a zero probability of occurrence for a dataset $D_1$ but a non-zero probability for a neighboring dataset $D_2$, the differential privacy guarantee does not hold. The mechanism must range over all $S \subseteq Range(\mathcal{M})$ with non-zero probability. Since PCOR must output a valid context for $V$ as the final answer, it assigns zero probability to all non-valid contexts. This zero probability is a violation of differential privacy, if adding or removing a record in a dataset $D_1$ changes a valid context $S$ for $V$ to a non-valid context for $V$ in $D_2$. In PCOR, differential privacy is guaranteed by applying the Exponential mechanism to $COE_M(D, V)$. $COE_M(D, V)$, as introduced in Definition 3.1, outputs the set of all contexts in which $V$ is an outlier so that $f_M(D_C, V) = true$ – property (a) in Definition 3.2 – is satisfied. Therefore, to guarantee differential privacy as introduced in its original form as represented in Equation 1, for any outlier $V$ and any neighboring datasets $D_1$ and $D_2$, the equality $COE_M(D_1, V) = COE_M(D_2, V)$ must hold. In other words, adding or removing a record, should not change the set of valid contexts for a particular outlier $V$. This condition is too strict for outlier detection, and as we show in our experiments in Section 6.7, there exist several neighboring datasets for various outlier detection algorithms that violate this constraint. Hence, we need some notion of differential privacy with satisfiable requirements for outlier detection algorithms, and we need to show that the shift to this notion, does not sacrifice the privacy guarantee that we are looking for.

Since violating the constraint in the original differential privacy notion is algorithm-dependent, we propose a relaxed notion of (Output Constrained) differential privacy that is algorithm-dependent as well[4]. This notion guarantees $\epsilon-$differential privacy when the equality $COE_M(D_1, V) = COE_M(D_2, V)$ for the outlier detection algorithm holds. We ran several experiments for the next two observations: i) to what extent the $COE_M(D_1, V) = COE_M(D_2, V)$ assumption in OCDP matches the results of outlier detection algorithms in practice, ii) what are the effects if this assumption does not hold and whether it results in a privacy sacrifice. Our results include evaluations of this assumption in group privacy, where the datasets $D_1$ and $D_2$ differ in more than one record. We ran experiments on the outlier detection algorithms introduced earlier in Section 2.1, and provide the results in Section 6.7.

## 3.2 Utility Functions

As described in Section 2.3, the differential privacy level in a mechanism is directly related to the sensitivity of the utility function. We discussed in Section 2.3 that in order to support reasonable levels of privacy, we aim for utility functions with sensitivity close to 1. We introduce two such utility functions here: i) Maximizing the context population size, and ii) Maximizing the overlap with the starting context.

*3.2.1 Context Population Size.* The outlier verification function, $f$, can return *true* on various contexts for a tuple $V$. Intuitively, a larger context population indicates a higher significance of the outlier.

**Definition 3.3.** (Maximum Context) $C$ is the Maximum context for outlier $V$ with regard to $f_M$ if and only if

- (i) $f_M(D_C, V)$ returns true
- (ii) $\forall C'$ s.t. $f_M(D_{C'}, V) = True$: $|D_C| \geq |D_{C'}|$

By choosing the population size as utility score, the mechanism guarantees the output to be *close enough* to the highest score answer, i.e. the maximum context. Formally,

$$u_V(D, C) = \begin{cases} -\infty, & f_M(D_C, V) = false \\ |D_C|, & \text{Otherwise} \end{cases}$$

The utility of $-\infty$ is assigned to non-valid contexts, so that it has zero probability to be picked by the Exponential mechanism. Furthermore, since the population size would differ by at most 1, if replacing the dataset with a neighboring one, the utility function has sensitivity of 1.

*3.2.2 Overlap with the Starting Context.* Another subject of interest in outlier analysis is a utility function that scores a context based on its relation to a chosen context. For this category of utility functions, we focus on one that scores a context $C$ based on its population's intersection with the population of a chosen/starting context for $V$, $C_V$. Formally,

$$u_V(D, C) = \begin{cases} -\infty, & f_M(D_C, V) = false \\ |D_C \cap D_{C_V}|, & \text{Otherwise} \end{cases}$$

Note that, changing a record in the dataset $D$ alters the score, i.e. the intersection size, by at most 1 (sensitivity). As mentioned

---

[4]Note that the common notion of differential privacy, $(\epsilon, \delta) - DP$, is not algorithm dependent.

earlier, PCOR utilizes the Exponential mechanism to provide differential privacy. We start with the direct application of the mechanism and argue why it imposes a substantial computational complexity to PCOR. To resolve the issue, we propose applying a sampling layer prior to deploying the the Exponential mechanism. However, as we explain in Section 5.1, sampling in its basic form does not reduce the complexity of the direct approach. To improve the sampling method, we map the contexts and their connections to a graph. Searching over the graph, we observe "locality" in the matching contexts and use this property to propose efficient sampling methods for our PCOR. A random walk on the graph shows significant superiority over the basic sampling approach. To improve the sampling even further, we use the utility to direct the search and introduce differentially private versions of Depth-First and Breadth-First search algorithms for sampling. This results in a private scalable design for PCOR that achieves high utility.

## 4 DIRECT APPROACH FOR PCOR

In the direct application of the Exponential mechanism in OCDP, we apply the mechanism over the outlier enumeration algorithm, as described in Algorithm 1. As we mentioned in Section 3, in order to provide more privacy protection for the individuals it is important that the enumeration algorithm considers all the values in the domain of attributes in $attr(R)$, not just the values covered in the dataset $D$. We show the rationality behind this requirement through the example in Section 3. The data owner reports the record $V$ is an outlier in the context $Jobtitle \in \{CEO, Lawyer\}$, in $Diplomatic$ district of $Ottawa$. By enumerating just over the values in the dataset, this report reveals that the individuals in the context are either CEOs or Lawyers. However, by enumerating over all possible values in the domain, the context description will be larger, e.g. $Jobtitle \in \{CEO, Lawyer, CFO, Diplomat\}$, in $Diplomatic$ district of $Ottawa$. This report leaves it unclear that which of the attribute values CEO, Lawyer, CFO, Diplomat is in the dataset.

---

**Algorithm 1:** PCOR - Direct Approach

**Input** : $D, attr(R), V, u, \epsilon_1$
**Output**: $C_p$

1 $C_M = \emptyset$
2         ▷ Ranging over all possible contexts, to find matches
3 **for** $C \subseteq attr(R)$ **do**
4     **if** $f_M(D_C, V) = true$ **then**
5        $C_M \leftarrow C_M \cup C$
6     **end**
7 **end**
8                   ▷ The Exponential mechanism
9 return $C_p \leftarrow Exp_u^{\epsilon_1}(D, C_M)$;
10 **Comment:** A context $C$ is a subset of the values of the attributes in $attr(R)$. All $C$'s included in $COE_M(D, V)$ are candidates for the private output, but with a probability determined by the utility function $u$. The mechanism drawing from the candidates set terminates the algorithm.

---

**Theorem 4.1.** Algorithm 1 satisfies $(\epsilon = 2\epsilon_1, COE_M(., V)) - OCDP$, according to Definition 2.5.

PROOF. Consider the neighboring pair $(D_1, D_2) \in N(f(.))$, for $COE_M$ being the outlier enumeration used in Algorithm 1. According to Definition 2.4, $D_1$ and $D_2$ differ in one record and result in the same range of outputs, $C_M$. To prove the algorithm satisfies the claim in the Theorem 4.1, we need to show that the probability of algorithm outputting a particular $C_p^*$ for $D_1$ is $2\epsilon_1$-different from $D_2$.

$$\frac{Pr[Alg1(D_1) = C_p^*]}{Pr[Alg1(D_2) = C_p^*]} = \frac{\frac{e^{\epsilon_1 u(C_p^*, D_1)}}{\sum_{c \in C_M} e^{\epsilon_1 u(c, D_1)}}}{\frac{e^{\epsilon_1 u(C_p^*, D_2)}}{\sum_{c \in C_M'} e^{\epsilon_1 u(c, D_2)}}}$$

$$= \frac{e^{\epsilon_1 u(C_p^*, D_1)}}{e^{\epsilon_1 u(C_p^*, D_2)}} \times \frac{\sum_{c \in C_M'} e^{\epsilon_1 u(c, D_2)}}{\sum_{c \in C_M} e^{\epsilon_1 u(c, D_1)}} \quad (4)$$

Since $|D_1| - |D_2| = 1$, any context selected for $D_1$ loses in maximum one record when selected for $D_2$. The utility for each context is proportional to its population size, hence its sensitivity is at most 1, i.e. $\Delta u \leq 1$. This means that that for any context in R (including $C_p$) the relation $u(c, D_1) - u(c, D_2) \leq \Delta u \leq 1$ holds. Hence the Equation 4 simplifies to

$$\frac{Pr[Alg1(D_1) = C_p^*]}{Pr[Alg1(D_2) = C_p^*]} \leq e^{\epsilon_1 \Delta u} \times e^{\epsilon_1 \Delta u} = e^{2\epsilon_1 \Delta u} \leq e^{2\epsilon_1} \quad (5)$$

Note that the number of elements in $C_M$ is the same as $C_M'$, due to the neighboring definition for output constrained differential privacy. However, the population size for the contexts in $C_M$ and $C_M'$ might differ in one record. □

**Theorem 4.2.** The computation complexity of Algorithm 1 is $O(2^{t+1})$, where $t$ is the total number of attribute values.

PROOF. Direct application of the Exponential mechanism requires brute forcing over all possible contexts for $V$ (lines 3-7 in Algorithm 1) and find the matching ones. With a context being a binary vector $\langle c_{11}, \cdots, c_{1|A_1|}, \cdots, c_{m1}, \cdots, c_{m|A_m|} \rangle$ of length $t = \sum_{i=1}^{m} |A_i|$, this phase is $O(2^t)$. It also calculates the weight of each candidate (lines 3-7), which is $O(2^t)$ as well. Hence, Algorithm 1 is $O(2^{t+1})$. □

## 5 SAMPLING APPROACH FOR PCOR

To resolve the efficiency problem in PCOR, we propose deploying a sampling layer, prior to applying the Exponential mechanism. We start with Uniform Sampling, and analyze its privacy and complexity of the algorithm.

### 5.1 Uniform Sampling

The first sampling algorithm we evaluate, picks contexts uniformly from the set of all valid contexts. We apply the Exponential mechanism after sampling. The idea of utilizing Uniform Sampling prior to the Exponential mechanism was investigated by Lantz et al. [30].

**Theorem 5.1.** The Uniform Sampling described in Algorithm 2 satisfies $(\epsilon = 2\epsilon_1, COE_M(., V)) - OCDP$.

**Algorithm 2:** PCOR - Uniform Sampling

---

**Input** : $D$ , $attr(R)$, $V$, $u$, $\epsilon_1$, $p$ ($p = \frac{1}{2}$ here)

**Output:** $C_p$

1                     ▷ Finding $n$ matching contexts

2   $C_M = \emptyset$

3 **while** $|C_M| \le n$ **do**

4     **for** $i \le t$ **do**

5        $C[i] = \begin{cases} 1, & \text{w.p. p} \\ 0, & \text{w.p. 1-p} \end{cases}$

6     **end**

7     **if** $f_M(D_C, V) = true$ **then**

8        $C_M \leftarrow C_M \cup C$

9     **end**

10 **end**

11                 ▷ The Exponential mechanism on the samples

12 return $C_p \leftarrow Exp_u^{\epsilon_1}(D, C_M)$;

13 **Comment:** We form the context vector, by setting one's and zero's randomly. In the general case the probabilities are correspondingly $p$ and $1 - p$. We consider $p = \frac{1}{2}$ to achieve Uniform Sampling. We obtain the $C_M$ from the sampling, then apply Exp. mechanism as in Algorithm 1.

---

PROOF.

$$\frac{Pr[Alg2(D_1) = C_p^*]}{Pr[Alg2(D_2) = C_p^*]}$$

$$= \frac{\frac{e^{\epsilon_1 u(C_p^*, D_1)}}{\sum_{c \in C_M} e^{\epsilon_1 u(c, D_1)}} \times Pr_{D_1}[C_p^* \in C_M]}{\frac{e^{\epsilon_1 u(C_p^*, D_2)}}{\sum_{c \in C_M'} e^{\epsilon_1 u(c, D_2)}} \times Pr_{D_2}[C_p^* \in C_M]}$$

$$\le e^{\epsilon_1 \Delta u} \times e^{\epsilon_1 \Delta u} \times 1 = e^{2\epsilon_1 \Delta u} \quad (6)$$

The first two items in the inequality in Equation 6 follows the same reasoning as the proof in Theorem 4.1. The $Pr[C_p^* \in C_M]$ is independent of the database being $D_1$ or $D_2$, as the contexts are chosen based on attribute values not the records, and the attribute values for $D_1$ and $D_2$ are the same. □

**Theorem 5.2.** The computation complexity of Algorithm 2 is $O(2^t)$, where $t$ is the total number of attribute values.

PROOF. We calculate the complexity of Algorithm 2 in lines 3-10 first. The algorithm keeps sampling among all $2^t$ contexts until it finds $n$ matching contexts for $V$. We want to find out how many contexts the algorithm should sample on average to achieve the goal. Assume the total number of matching contexts for $V$ is $N$; i.e. the probability of a sample being a matching context is $\frac{N}{2^t}$. Hence, the number of success in sampling follows Binomial distribution $B(x, \frac{N}{2^t})$, with $x$ being the total number of samples [11]. Since the expected value of the number of success in Binomial distribution in this case is $x \times \frac{N}{2^t}$, we need $n \times \frac{2^t}{N}$ samples to obtain $n$ matching contexts on average. Furthermore, Algorithm 2 runs the Exponential mechanism on the matching contexts, which adds $O(n)$ complexity. Therefore, Uniform Sampling is $O(\frac{n \times 2^t}{N} + n)$.

Considering $n$ and $N$ being constant values, Uniform Sampling does not effectively change the complexity of the direct approach. □

## 5.2 Graph-based Sampling

We assume the data owner knows a valid starting context[5] $C_V$ for a local outlier $V$ and desires to find the private maximum one through PCOR. We map this problem to a search over a graph $G$ initiating form $V$'s starting context $C_V$, aiming to reach the maximum context. We define the context graph with the set of vertices $Vtx = \{C\}$, where $\{C\}$ is all possible contexts defined over the attribute values from $attr(R)$. There is an edge between two contexts $C$ and $C'$ if they are *connected* to each other, i.e. they are different in the presence of just one attribute value. Every context $C$ can be represented as a binary vector of the form $\langle c_{11}, \cdots, c_{1|A_1|}, \cdots, c_{m1}, \cdots, c_{m|A_m|} \rangle$ with length $t = \sum_{i=1}^m |A_i|$. As a result, flipping a bit in $C$ forms a binary vector for context $C'$ which is different from $C$ in just one (added/removed) attribute value. Hence, there are $t$ number of $C'$s connected to $C$; i.e. each vertex is of degree $t$. Given that the connected vertices are different in just one attribute value, we hypothesize that if $V$ is an outlier in $C$, then it is more probable to be an outlier in a connected vertex than some randomly chosen vertex among Vtx. In other words, we hypothesize the existence of *locality* in outlier detection algorithms. We show in this paper the hypothesis holds for algorithms from any outlier detection category. We start by our first sampling algorithm on the context graph: Random Walk.

*5.2.1 Random Walk Sampling.* The data owner knows a valid context for a local outlier $V$, namely starting context $C_V$, and desires to find the private maximum one. In random walk sampling, we initiate from the starting context and change the attribute values to obtain the next context to continue the chain. The change in the context in each iteration includes adding/removing an attribute value to/from the context.

**Theorem 5.3.** The random walk sampling described in Algorithm 3 satisfies $(\epsilon = 2\epsilon_1, COE_M(., V))$–OCDP, where $\epsilon_1$ is the privacy parameter in the Exponential mechanism.

PROOF. To calculate the privacy provided by an algorithm, we evaluate the probability change of the privacy mechanism outputting a particular value when switching from a database $D_1$ to a neighbor $D_2$. These probabilities are shown by $Pr[Alg3(D) = C_p^*]$ and $Pr[Alg3(D_2) = C_p^*]$ in this case respectively. A closer look to the algorithm, reveals that for $C_p^*$ to be an output for $D_1$, it must be selected by the Exponential mechanism from $C_M$. Hence,

$$Pr_{D_1}[Alg3(D_1) = C_p^*] = Pr_{D_1}[C_p^*|C_M] \times Pr_{D_1}[C_M]$$
$$= Pr_{D_1}[C_p^*|C_M] \times \Pi_{i=1}^n Pr_{D_1}[C_i|C_{i-1}] \times Pr_{D_1}[C_V] \quad (7)$$

Now, we calculate the probability of Algorithm 3 resulting in the same final answer ($C_p^*$) for the two databases $D_1$, $D_2$. Let $C_M = \{c_1, c_2, \cdots, c_c, \cdots, c_p\}$ be the path of contexts sampled by Algorithm 3 for database $D_1$. To update the $C_M$ from any path of contexts, e.g. $\{c_1, c_2, \cdots, c_i\}$, to another path that is different from the former in only the last context, i.e. $\{c_1, c_2, \cdots, c_{i+1}\}$, two events need to take place. First, the predecessor context of $c_{i+1}$ be selected

---

**Algorithm 3:** PCOR - Random Walk Sampling

---

**Input** : $D$, $attr(R)$, $V$, $C_V$, $u$, $\epsilon_1$

**Output:** $C_p$

1   $C_M = [C_V]$

2                    ▷ $n$ is the total number of samples

3 **while** $|C_M| \leq n$ && $C_{conn} \neq \emptyset$ **do**

4     $C \leftarrow C_V$

5     $C_{conn} = \{$connected vertices to $C\}$

6                ▷ $C_i$ is selected randomly from $C_{conn}$

7     $C_i \xleftarrow{r} C_{conn}$

8     **if** $f_M(D_{C_i}, V) = true$ **then**

9        $C_M \leftarrow C_M \cup C_i$

10       $C_V \leftarrow C_i$

11     **end**

12     **else**

13       $C_{conn}.remove(C_i)$

14     **end**

15 **end**

16           ▷ The Exponential mechanism on the samples

17 return $C_p \leftarrow Exp_u^{\epsilon_1}(D, C_M)$;

18 **Comment:** The random walk adds/removes a random attribute value to/from a starting context to obtain the next context, the algorithm continues to do so until the next context is a matching one for $V$, this context is the second sample. We continue random walking from the obtained matching context, until we reach the desired number of sampled contexts in the multiset $C_M$. The the Exponential mechanism is applied afterwards as in Alg. 1 to the samples to select the final differentially private output.

---

from the $i$ elements in the $C_M$, we show this context by $c_j$. Second, among all nodes connected to $c_j$, $c_{i+1}$ be selected. By changing from $D_1$ to $D_2$, the former's probability reduces by $2\epsilon\Delta u$; the latter however maintains the same probability, as the selection among the connected nodes is random. Hence, as we also showed in proving Theorem 5.1, the probability of selecting the same $C_p^*$ from a set of $n$ samples such as $C_M$, is bounded by $e^{2\epsilon\Delta u}$. Therefore,

$$\frac{Pr[Alg3(D_1) = C_p^*]}{Pr[Alg3(D_2) = C_p^*]} \leq e^{2\epsilon_1 \Delta u} \tag{8}$$

Note that multiple paths could lead Algorithm 3 to achieve the same output for $D_2$ as $D_1$. In our proof however, we consider the worst case where there is only one path, therefore Algorithm 3 has to follow the same path for $D_2$ that it does for $D_1$. □

**Theorem 5.4.** The computation complexity of Algorithm 3 is $O(t)$.

PROOF. Random walk starts from a starting context for $V$, $C_V$ and iteratively changes a random attribute value of $C_V$ until it finds a matching context among the $t$ connected contexts. Afterwards, it forms a path by repeating the same process for the last context on the chain. As searching through connected contexts to a context $C$ is done without replacement, its complexity is $O(t)$, the process is repeated for a constant number of times $n$ to form $C_M$. Ultimately,

the Exponential mechanism is applied to select the final candidate among $n$ (constant number of) samples. Hence, the algorithm's complexity is $O(nt + n)$. □

Exploiting the context graph in the Random Walk algorithm improves the sampling complexity from exponential in Uniform Sampling to linear. Our experimental results in Section 6 also affirms the superiority of the Random Walk in utility and performance over the Uniform Sampling. To improve the utility and performance even more, we alter the walk over the connected contexts from a random method to a one directed by the utility function; in other words, we can improve sampling further by taking the utility already during selection of the next vertex into account. Therefore, we explore the Depth-First and Breadth-First search algorithms for sampling.

*5.2.2 Depth-First Search Sampling.* We modify the Depth-First search (DFS) algorithm [10] to provide a differentially private version of the algorithm. We emphasize that differential privacy cannot be supported with a non-modified DFS algorithm, e.g. by perturbing its output. As described in Section 2.2, differential privacy requires the algorithm to generate any output with approximately the same ($e^\epsilon$ different) probability for neighboring datasets. However, as DFS is a deterministic algorithm, the probability of it generating any output is zero except for one particular value. Hence, applying DFS on a dataset $D_1$ might result in an output $r$ that has probability of zero when DFS is applied on a neighboring dataset $D_2$; this property prevents the original DFS from supporting differential privacy. To overcome the challenge, we modify DFS as presented in the Algorithm 4. The sampling method initiates a stack with a starting context for $V$, $C_V$. Next, it searches over all connected contexts to $C_V$ and selects a matching context by applying the Exponential mechanism to the candidates and pushes the result onto the stack. The last context on the stack is the starting point for the next iteration. The iterations continue until the stack contains the desired number of samples, $n$.

**Theorem 5.5.** The differentially private DFS described in Algorithm 4 satisfies $(\epsilon = (2n + 2)\epsilon_1, COE_M(., V))-$OCDP, where $n$ is the total number of samples, $\epsilon_1$ is the privacy parameter in the the Exponential mechanism.

PROOF. A similar reasoning to the one in Theorem 5.3 holds for DFS. The idea is calculating the probability of forming the same stack of nodes for neighboring datasets $D_1$ and $D_2$. However, the probability of obtaining the same stack, i.e. the corresponding probability to $\Pi_{i=1}^n Pr_{D_1}[C_i | C_{i-1}] \times Pr_{D_1}[C_V]$ in Equation 7, is not the same value for $D_1$ and $D_2$ anymore. As in differential private DFS, we apply the Exponential mechanism before adding any child to the stack, each conditional probability in Equation 7 is $\epsilon_1 \Delta u$ different for $D_1$ and $D_2$. Since there are $n$ nodes in the stack, the upper bound is: $\frac{Pr_{D_1}[Stack]}{Pr_{D_2}[Stack]} \leq e^{2n\epsilon_1\Delta u}$. We also know from Equation 6, that the probability of selecting the same $C_p^*$ from a set of $n$ samples ($Stack$), is bounded by $e^{2\epsilon_1\Delta u}$. This results in

$$\frac{Pr[Alg4(D_1) = C_p^*]}{Pr[Alg4(D_2) = C_p^*]} \leq e^{(2n+2)\epsilon_1\Delta u} \tag{9}$$

□

**Algorithm 4:** PCOR - Depth-First Search Sampling

---

**Input** : $D$ , $attr(R)$, $V$, $C_V$, $u$, $\epsilon$

**Output:** $C_p$

1   $Stack \leftarrow \{C_V\}$, $Samples \leftarrow 0$, $Visited \leftarrow \emptyset$

2             $\triangleright$ $n$ is the total number of samples

3 **while** $|Visited| \leq n$ && $|Stack| > 0$ **do**

4     $C \leftarrow Stack.top()$, $Visited \leftarrow Visited \cup C$, $C_{chldn} \leftarrow \emptyset$

5                 $\triangleright$ Generating all children of C

6     **for** $i \leq t$ **do**

7        $C[i] \leftarrow \begin{cases} 0, & \text{If } C[i] = 1 \\ 1, & \text{If } C[i] = 0 \end{cases}$

8        **if** $f_M(D_C, V) = true$ $and$ $C \notin Visited$ **then**

9           $C_{chldn} \leftarrow C_{chldn} \cup C$

10        **end**

11     **end**

12     **if** $not$ $C_{chldn}$ **then**

13        $Stack.pop()$

14     **end**

15     **else**

16           $\triangleright$ The Exponential mechanism on children of $C$

17        $C \leftarrow Exp_u^{\epsilon_1}(D, C_{chldn})$

18        push $C$ onto $Stack$

19     **end**

20 **end**

21         $\triangleright$ The Exp. mechanism for the final private answer

22 return $C_p \leftarrow Exp_u^{\epsilon_1}(D, Visited)$;

23 **Comment:** The algorithm starts from a starting context, $C_V$, adds it to the stack. It then calculates all the connected nodes to $C_V$ that are matching contexts for $V$. Next, the algorithm applies the Exponential mechanism to select one of the connected nodes, $c_c$. Afterwards, the $c_c$ is pushed onto the stack, to be the starting point for the next search as well. This procedure continues until $n$ samples are pushed onto the stack. The last sample in the stack is the final privacy preserving answer.

---

**Theorem 5.6.** The computational complexity of Algorithm 4 is $O(t)$, where $t$ is the total number of attribute values.

PROOF. The differential private DFS algorithm starts a path from a starting context for $V$, $C_V$, selects the next matching context the $t$ connected contexts by applying the Exponential mechanism to them and pushes the result onto the stack; which forms an $O(2t)$ step. Then repeats the process for the last context on the stack until the stack size reaches the constant desired value, $n$, times. Hence, the algorithm's complexity is $O(2nt)$. □

*5.2.3 Breadth-First Search Sampling.* We introduce a modified and differentially private version of the Breadth-First search (BFS) algorithm [31] as our last sampling method. The reason for necessity of modifying the algorithm to support differential privacy is the same as the one provided for Depth-First search algorithm in Section 5.2.2. Our sampling method starts from a starting context for $V$, $C_V$. It initiates the set $C_M$ with that context, then searches over all

connected contexts to $C_V$ and adds all the matching contexts to $C_M$. To decide the starting point for the next iteration, it applies the Exponential mechanism to the contexts in the $C_M$ (treating $C_M$ as a priority queue) and proceeds similarly with the selected context. The process continues until $n$ number of samples are added to $C_M$.

---

**Algorithm 5:** PCOR - Breadth-First Search Sampling

---

**Input** : $D$ , $attr(R)$, $V$, $C_V$, $u$, $\epsilon$

**Output:** $C_p$

1   $C_M = \{C_V\}$, $Samples \leftarrow 0$, $Visited \leftarrow \emptyset$

2             $\triangleright$ $n$ is the total number of samples

3 **while** $|Visited| \leq n$ && $|C_M| > 0$ **do**

4             $\triangleright$ The Exponential mechanism on $C_M$

5     $C \leftarrow Exp_u^{\epsilon_1}(D, C_M)$

6     $Visited \leftarrow Visited \cup C$, $C_M.remove(C)$

7             $\triangleright$ Adding all children of C to $C_M$

8     **for** $i \leq t$ **do**

9        $C[i] \leftarrow \begin{cases} 0, & \text{If } C[i] = 1 \\ 1, & \text{If } C[i] = 0 \end{cases}$

10        **if** $f_M(D_C, V) = true$ $and$ $C \notin Visited$ **then**

11           $C_M \leftarrow C_M \cup C$

12        **end**

13     **end**

14 **end**

15         $\triangleright$ The Exp. mechanism for the final private answer

16 return $C_p \leftarrow Exp_u^{\epsilon_1}(D, Visited)$;

17 **Comment:** The algorithm starts from a starting context $C_V$, and calculates all the connected contexts to $C_V$, adds the matching ones to the priority queue $C_M$. It continues by applying the the Exponential mechanism to the contexts in $C_M$, and iterates the procedure until it finds all $n$ samples.

---

**Theorem 5.7.** The differentially private BFS described in Algorithm 5 satisfies $((2n + 2)\epsilon_1, COE_M(., V))$–OCDP, where $n$ is the total number of samples and $\epsilon_1$ is the privacy parameter in the Exponential mechanism.

PROOF. Similar to DFS, the privacy proof of BFS follows the proof for Theorem 5.3. Which results in

$$\frac{Pr[Alg5(D_1) = C_p^*]}{Pr[Alg5(D_2) = C_p^*]} \leq e^{(2n+2)\epsilon_1 \Delta u} \tag{10}$$

□

**Theorem 5.8.** The computation complexity of Algorithm 5 is $O(t)$, where $t$ is the total number of attribute values.

PROOF. BFS initiates the search from a starting context for $V$, $C_V$, searches over its connected contexts to find the matching ones to add to $C_M$ which forms an $O(t)$ step. This step is repeated for a context chosen by the Exponential mechanism from $C_M$ ($O(nt)$) for the constant desired number of samples, $n$, times. Hence, the algorithm is $O(n^2t + nt)$. □

# 6 EXPERIMENTS

We run five groups of experiments to support our choices for the parameters in our PCOR design varying: i) Sampling algorithm, ii) Utility function, iii) Outlier detection algorithms, iv) Privacy parameter $\epsilon$, and v) Group privacy limit in OCDP.

## 6.1 Datasets

We evaluate PCOR over two datasets. The first one is the Ontario's public sector salary dataset. The province of Ontario (Canada) annually publishes a list of public employees who earn $100,000 and above [2]. We filter the dataset to obtain the information of 51000 employees with income higher than 100K in Ontario, the attributes are $attr(R) = \{Jobtitle, Employer, Year, Salary\}$, with domain sizes: $|Jobtitle| = 9$, $|Employer| = 8$, and $|Year| = 8$. The second dataset we use is homicide reports in the United States [1], compiled and made available by the Murder Accountability Project. The data is gathered from multiple agencies and includes the age, sex, ethnicity of victims and perpetrators, their relationship, and the weapon used. We filter the dataset to obtain the information of 110,000 records. The attributes of the data are $attr(R) = \{AgencyType, State, Weapon, VictimAge\}$. The categorical attributes have domains of size 4, 6, and 6 respectively. Our results for both datasets follow the same pattern. Due to space limit, we mainly provide the experiment results on the salary dataset, and refer to our full paper for the detailed experiment results on the homicide dataset.

## 6.2 Evaluation Setup

We ran our experiments on a machine with 1 TB RAM, 132 cores, Intel(R) Xeon(R) CPU E7-8870 2.40GHz specifications. We repeated each experiment 200 times to achieve reliable results for PCORs and fair comparison of the algorithms in utility and performance. The performance is measured as the runtime of the algorithm, in two metrics: i) The range of shortest to longest runtime, and ii) The average runtime. The utility is measured as the proportion of the utility of the PCOR's output to the maximum utility. The maximum utility for each data point (outlier), is gained from a reference file. The reference file contains all possible contexts in $attr(R)$ accompanied with their associated utility, and the list of outliers for each context. Generating this file, for our data set of 51,000 records utilizing an optimized program took three days to complete on the machine mentioned earlier. This is also the running time of the direct/formulaic application of Exponential mechanism in OCDP (described in Section 4). From the reference file, we find all matching contexts for outlier $V$ and find the one with maximum utility to divide the utility results of the PCOR by. For the 200 utility results, we provide the mean and the 90% confidence interval (CI) for each PCOR. We run five groups of experiments to answer the following questions: 1) what is the best (private, accurate, fast) sampling algorithm for PCOR? 2) Is PCOR flexible with embedding other utility functions? 3) Is PCOR flexible with embedding other outlier detection algorithms? We evaluate PCOR over three deterministic outlier detection algorithms, each from a main category described in Section 2.1, namely: Grubbs [23], Local Outlier Factor (LOF) [8], and Histogram [3]. 4) How does changing each one of privacy, utility and performance parameters affect the other two. 5)

How does changing a (group of) record affect the set of possible outcomes, $COE_M$?

## 6.3 Choosing a Sampling algorithm for PCOR

We show that using Breadth-First Search algorithm for private sampling results in the highest utility and performance for PCOR. We evaluate four sampling algorithms: i) Uniform Sampling, ii) Random Walk, iii) Breadth-First Search, and iv) Depth-First Search [6]. We generate sets of length 50 samples and set $\epsilon = 0.2$ as the total differential privacy budget. This translates to $\epsilon_1 \approx 0.002$ in Depth-First Search and Breadth-First Search as shown in Equations 9 and 10, and $\epsilon_1 = 0.1$ in Uniform Sampling and Random Walk as shown in Equations 6 and 8 respectively. The run times of the sampling algorithms [7] are summarized in Table 2. Utility is mea-

**Table 2: Sampling Methods Comparison - Performance**

| Algorithm | $T_{min}$ | $T_{max}$ | $T_{avg}$ | $\epsilon$ | Outlier |
|-----------|-----------|-----------|-----------|-----------|---------|
| Uniform | 7m | 24h | 97m | 0.2 | LOF |
| Random Walk | 15s | 109s | 51s | 0.2 | LOF |
| DFS | 8m | 80m | 40m | 0.2 | LOF |
| BFS | 6m | 61m | 37m | 0.2 | LOF |

sured as the ratio of the population size of the private answer to that of the maximum context's for an outlier $V$. Utility comparison results are summarized in Table 3. The Random Walk improves

**Table 3: Sampling Methods Comparison - Utility**

| Algorithm | Utility | CI | $\epsilon$ | Outlier |
|-----------|---------|-----|-----------|---------|
| Uniform | 0.65 | (0.64, 0.67) | 0.2 | LOF |
| Random Walk | 0.57 | (0.55, 0.60) | 0.2 | LOF |
| DFS | 0.88 | (0.85, 0.90) | 0.2 | LOF |
| BFS | 0.90 | (0.88, 0.93) | 0.2 | LOF |

the performance of the uniform sampling significantly, but it sacrifices utility. The BFS and DFS algorithms perform close to each other in running time and utility, they make up for the utility loss in Random Walk by a considerable amount, but require a longer runtime. Their advantage over Uniform sampling in both utility and performance is noticeable. The slight advantage of BFS over DFS in utility and runtime, also holds for PCOR with a different utility function (Section 6.3). Therefore, we select BFS as PCOR's final sampling method.

## 6.4 PCOR and Other Utility Functions

We show the adaptability of PCOR with any utility function by providing the results of embedding another utility function than the population size in Tables 4 and 5. The utility function in this experiment calculates the population intersection of the private context and the starting context, $C_V$. The utility and performance

---

[6]The results are depicted in Figure 1 in the appendix
[7]Figure 1[a-d] in the appendix

**Table 4: Intersection Overlap Utility - Performance**

| Algorithm | $T_{min}$ | $T_{max}$ | $T_{avg}$ | $\epsilon$ | Outlier |
|---|---|---|---|---|---|
| DFS | 3m | 47m | 19m | 0.2 | LOF |
| BFS | 5m | 48m | 20m | 0.2 | LOF |

**Table 5: Intersection Overlap Utility - Utility**

| Algorithm | Utility | CI | $\epsilon$ | Outlier |
|---|---|---|---|---|
| DFS | 0.88 | (0.86, 0.91) | 0.2 | LOF |
| BFS | 0.97 | (0.95, 0.98) | 0.2 | LOF |

evaluations [8] confirm that BFS is a superior candidate over the DFS algorithm.

## 6.5 PCOR and Outlier Detection Algorithms

We show that PCOR can be successfully used with various outlier detection algorithm. We explore two more outlier detection algorithms in our PCOR design: Grubbs, and Histogram methods. The results are [9] summarized in Table 6 and Table 7. We filter the original dataset to obtain a subset of 11000 records with 14 attributes values in total, and use BFS as the sampling algorithms in the PCOR design. In this set of experiments, the number of samples is $n = 50$ and the privacy parameter is $\epsilon = 0.2$. For Histogram method, we

**Table 6: Outlier Detection Algorithms - Performance**

| Algorithm | $T_{min}$ | $T_{max}$ | $T_{avg}$ | $\epsilon$ | Sampling |
|---|---|---|---|---|---|
| Grubbs | 0.5m | 1m | 0.8m | 0.2 | BFS |
| Histogram | 2m | 4m | 3.4m | 0.2 | BFS |

bin the samples corresponding to a context $C$ to $\sqrt{|D_C|}$ bins, where $|D_C|$ is the size of the population covered by $C$. The bins with less frequency than $2.5 \times 10^{-3}|D_C|$ are labeled as outliers. The results

**Table 7: Outlier Detection Algorithms - Utility**

| Algorithm | Utility | CI | $\epsilon$ | Sampling |
|---|---|---|---|---|
| Grubbs | 0.86 | (0.84, 0.89) | 0.2 | BFS |
| Histogram | 0.89 | (0.87, 0.91) | 0.2 | BFS |

confirm the compatibility of our PCOR design with various outlier detection algorithms. Furthermore, the success of applying BFS on the algorithms implies that the locality (discussed in Section 5.2) exists in all evaluated outlier detection algorithms.

---

[8] Figure 2 in the appendix
[9] Figure 3 in the appendix

## 6.6 Privacy, Utility, Performance Trade-off

We show the trade-off between privacy, utility and performance for our final candidate, BFS sampling algorithm, by changing the privacy parameter $\epsilon$ for $n = 50$ number of samples. The results are [10] summarized in Table 8 and Table 9. Increasing the $\epsilon$ from

**Table 8: Effect of privacy parameter on performance**

| $\epsilon$ | $T_{min}$ | $T_{max}$ | $T_{avg}$ | Sampling | Outlier |
|---|---|---|---|---|---|
| 0.05 | 2m | 29m | 15m | BFS | LOF |
| 0.1 | 2m | 29m | 16m | BFS | LOF |
| 0.2 | 3m | 30m | 17m | BFS | LOF |
| 0.4 | 3m | 30m | 17m | BFS | LOF |

**Table 9: Effect of privacy parameter on utility**

| $\epsilon$ | Utility | CI | Sampling | Outlier |
|---|---|---|---|---|
| 0.05 | 0.67 | (0.62, 0.71) | BFS | LOF |
| 0.1 | 0.82 | (0.78, 0.85) | BFS | LOF |
| 0.2 | 0.90 | (0.88, 0.93) | BFS | LOF |
| 0.4 | 0.92 | (0.90, 0.94) | BFS | LOF |

0.05 to 0.4 generally increases the utility, however the parameter $\epsilon = 0.2$ is the optimum value and the utility does not increase significantly afterwards. Moreover, changing the privacy parameter does not impose a notable effect on the algorithm runtime. We also

**Table 10: Effect of # of samples on performance**

| # Samples | $T_{min}$ | $T_{max}$ | $T_{avg}$ | Sampling | Outlier |
|---|---|---|---|---|---|
| 25 | 1m | 14m | 7m | BFS | LOF |
| 50 | 3m | 29m | 16m | BFS | LOF |
| 100 | 6m | 61m | 37m | BFS | LOF |
| 200 | 21m | 163m | 99m | BFS | LOF |

**Table 11: Effect of # of samples on utility**

| # Samples | Utility | CI | Sampling | Outlier |
|---|---|---|---|---|
| 25 | 0.85 | (0.81, 0.88) | BFS | LOF |
| 50 | 0.88 | (0.85, 0.91) | BFS | LOF |
| 100 | 0.90 | (0.88, 0.93) | BFS | LOF |
| 200 | 0.84 | (0.81, 0.87) | BFS | LOF |

investigate the effect of changing the number of samples while the privacy parameter is fixed $\epsilon = 0.2$. The results are summarized in Tables 10 and 11[11]. As we showed in the proof of Theorem 5.8, the only parameters in the computation complexity of BFS are: i) the total number of attribute values $t$, and ii) the number of samples $n$

---

[10] Figure 4[a-d] in the appendix
[11] Figure 5[a-h] in the appendix

– which we referred to as a constant number in the proof. Hence, when $t$ is fixed, $n$ determines the performance of PCOR. Increasing the number of samples from 25 to 100 decreases the performance, but increases the utility. This increment does not continue for for $n = 200$. Recall from Theorem 5.7, that keeping a fixed $\epsilon$ in PCOR while increasing $n$, requires using smaller $\epsilon_1$'s in algorithm design, this can cancel out the positive affect of higher $n$'s on utility.

## 6.7 Context Match and Group Privacy

Our reduced salary dataset consists of 11,000 records with 3 attributes that include 14 attribute values in total. The homicide dataset consists of 28,000 records with three attributes that include 12 attribute values in total. We chose small datasets to run several experiments in a reasonable amount of time. We are aware that our results do not benefit from this choice, as changing a single record in a small dataset, more strongly affects the set of outliers than in a large dataset. We repeated each experiment for 100 random outliers and measured the contexts set match of the original dataset and its neighboring datasets. Recall from Section 3.1 that the experiments in this section are to observe: i) to what extent the assumption $COE_M(D_1, V) = COE_M(D_2, V)$ in OCDP matches the results of outlier detection algorithms in practice, ii) what are the effects if this assumption does not hold and whether it results in a privacy sacrifice. For both mentioned datasets and all three outlier detection algorithm, we recorded the $COE_M$ results of the original dataset and its 50 randomly chosen neighboring datasets. The results for the objective i) are summarized in Tables 12 - 13. These tables show the results for neighboring datasets that are different from the original dataset ($\Delta D$) in 1, 5, 10, and 25 records, for the salary and the homicide dataset respectively. For objective ii) and

**Table 12: $COE$ Match - Salary dataset**

| Algorithm | $\Delta D = 1$ | $\Delta D = 5$ | $\Delta D = 10$ | $\Delta D = 25$ |
|-----------|------|------|------|------|
| Grubbs | 99.8% | 96.9% | 94.5% | 91.9% |
| LOF | 89% | 87.9% | 86.7% | 85.7% |
| Histogram | 95.5% | 82.1% | 70.8% | 58.8% |

**Table 13: $COE$ Match - Homicide dataset**

| Algorithm | $\Delta D = 1$ | $\Delta D = 5$ | $\Delta D = 10$ | $\Delta D = 25$ |
|-----------|------|------|------|------|
| Grubbs | 100% | 100% | 100% | 97.8% |
| LOF | 99.9% | 99.5% | 98.7% | 97.7% |
| Histogram | 98.5% | 85.2% | 69.3% | 53.3% |

to evaluate the privacy in the case of non-matching $COE$'s of the neighboring datasets, we ran another set of experiments. Consider the contexts in the intersection of non-equal sets $COE_M(D_1, V)$ and $COE_M(D_2, V)$, where $|D_1 - D_2| = 1$. We measured the maximum ratio of the probability of selecting any of these context for $D_1$ to the probability of selecting the same context for $D_2$. This experiment was repeated for 200 outlier samples for three outlier detection algorithm over the Salary dataset. All the experiments confirm this ratio

is below $e^\epsilon$ (recall the differential privacy requirement in Equation 1) for $\epsilon = 0.2$ that is the privacy budget in our experiments that we presented earlier in this section. In other words, for non-matching $COE$'s of neighboring datasets, we found no instance that violates the privacy bound in (unconstrained) $\epsilon-$differential privacy.

## 7 RELATED WORK

The attempts to provide privacy in outlier detection, initiated with using secure computation methods for distributed [12, 44] or non-distributed [5] data. In a different approach, Bhaduri et al. [6] use nonlinear data distortion transformation and show how it can be useful for privacy preserving anomaly detection from sensitive datasets. Gehrke et al. [20] introduce *crowd blending privacy*, in which for every individual in the database, either the individual blends in to a crowd of $k$ people in the database with respect to the privacy mechanism, or the mechanism ignores the individual's data. Inspired by crowd blending [20] and providing different levels of privacy [21], Lui and Pass [33] proposed *tailored differential privacy* for protecting outliers, in which the privacy parameter for an individual is determined by the "outlierness" of the individual's data in the dataset. Böhler et al. [7] evaluate the opposite by relaxing the differential privacy definition and granting outliers *less* protection, as they consider the outliers as "faulty systems or sensors one need to detect". In another attempt in private outlier investigation, Nissim et al. [38] use differential to locate a small outlier cluster privately. Okada et al. in [39] break the outlier analysis to two tasks: i) counting outliers in a given subspace and ii) discovering sub-spaces containing many outliers, under the constraints of differential privacy. They show their method achieves better utility compared to the original global sensitivity based methods. Nonetheless, our work is the first to investigate the privacy of the individuals in the outlier's context and not the outlier's privacy.

## 8 CONCLUSION

The revealed context in contextual outlier release leaks information about the individual records in the dataset. We address this privacy violation and propose techniques for a relaxed notion of differential privacy to provide PCOR and resolve the issue. However, the differential privacy solution in its formulaic application suffers from weak performance, impeding its usage in practice. To achieve efficiency, we propose utilizing a sampling layer in the design. We present differentially private graph search algorithms, first time in the literature, and use them for sampling. We prove PCOR with these sampling methods supports worthwhile levels of differential privacy, while providing the desired utility and performance. We articulate and demonstrate empirically that PCOR is compatible with any utility function in outlier detection. Our results indicate that the relaxation required for privacy in PCOR is a not a strong requirement, and is satisfied in most cases in practice. We also show that PCOR is generic and fits any outlier detection algorithm.

## 9 ACKNOWLEDGEMENT

# REFERENCES

[1] [n.d.]. The Murder Accountability Project. https://www.kaggle.com/murderaccountability/homicide-reports#database.csv. Accessed: 2019-12-15.

[2] [n.d.]. Ontario's public sector salary dataset. https://www.ontario.ca/page/public-sector-salary-disclosure. Accessed: 2019-07-31.

[3] 2008. *Pearson, Karl*. Springer New York, New York, NY, 418–419. https://doi.org/10.1007/978-0-387-32833-1_313

[4] 2016-06-27. Regulation (eu) 2016/679 of the European parliament and of the council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data,and repealing directive 95/46/ec (general data protection regulation).

[5] Abdulatif Alabdulatif, Heshan Kumarage, Ibrahim Khalil, and Xun Yi. 2017. Privacy-preserving anomaly detection in cloud with lightweight homomorphic encryption. *J. Comput. System Sci.* 90 (2017), 28 – 45. https://doi.org/10.1016/j.jcss.2017.03.001

[6] K. Bhaduri, M. D. Stefanski, and A. N. Srivastava. 2011. Privacy-Preserving Outlier Detection Through Random Nonlinear Data Distortion. *Trans. Sys. Man Cyber. Part B* 41, 1 (Feb. 2011), 260–272. https://doi.org/10.1109/TSMCB.2010.2051540

[7] Jonas Böhler, Daniel Bernau, and Florian Kerschbaum. 2017. Privacy-Preserving Outlier Detection for Data Streams. In *Data and Applications Security and Privacy XXXI*, Giovanni Livraga and Sencun Zhu (Eds.). Springer International Publishing, Cham, 225–238.

[8] Markus Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. 2000. LOF: Identifying Density-Based Local Outliers. In *PROCEEDINGS OF THE 2000 ACM SIGMOD INTERNATIONAL CONFERENCE ON MANAGEMENT OF DATA*. ACM, 93–104.

[9] Varun Chandola, Arindam Banerjee, and Vipin Kumar. 2009. Anomaly Detection: A Survey. *ACM Comput. Surv.* 41, 3, Article 15 (July 2009), 58 pages. https://doi.org/10.1145/1541880.1541882

[10] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms, Third Edition* (3rd ed.). The MIT Press.

[11] Jay L. Devore. 2011. *Probability and Statistics for Engineering and the Sciences* (8th ed.). Brooks/Cole. ISBN-13: 978-0-538-73352-6.

[12] L. T. Dung and H. T. Bao. 2011. A Distributed Solution for Privacy Preserving Outlier Detection. In *2011 Third International Conference on Knowledge and Systems Engineering*. 26–31. https://doi.org/10.1109/KSE.2011.13

[13] Cynthia Dwork. 2006. Differential Privacy. In *Automata, Languages and Programming*, Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–12.

[14] Cynthia Dwork. 2008. Differential Privacy: A Survey of Results. In *Theory and Applications of Models of Computation*, Manindra Agrawal, Dingzhu Du, Zhenhua Duan, and Angsheng Li (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–19.

[15] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284.

[16] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Third Theory of Cryptography Conference (TCC 2006)* (third theory of cryptography conference (tcc 2006) ed.) *(Lecture Notes in Computer Science, Vol. 3876)*. Springer, 265–284. https://www.microsoft.com/en-us/research/publication/calibrating-noise-to-sensitivity-in-private-data-analysis/

[17] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.* 9, 3&#8211;4 (Aug. 2014), 211–407. https://doi.org/10.1561/0400000042

[18] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security* (Scottsdale, Arizona, USA) *(CCS '14)*. ACM, New York, NY, USA, 1054–1067. https://doi.org/10.1145/2660267.2660348

[19] Chang Ge, Xi He, Ihab F. Ilyas, and Ashwin Machanavajjhala. 2019. APEx: Accuracy-Aware Differentially Private Data Exploration. In *Proceedings of the 2019 International Conference on Management of Data* (Amsterdam, Netherlands) *(SIGMOD '19)*. ACM, New York, NY, USA, 177–194. https://doi.org/10.1145/3299869.3300092

[20] Johannes Gehrke, Michael Hay, Edward Lui, and Rafael Pass. 2012. Crowd-Blending Privacy. In *Advances in Cryptology – CRYPTO 2012*, Reihaneh Safavi-Naini and Ran Canetti (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 479–496.

[21] Arpita Ghosh and Aaron Roth. 2015. Selling privacy at auction. *Games and Economic Behavior* 91 (2015), 334 – 346. https://doi.org/10.1016/j.geb.2013.06.013

[22] A. Greenberg. 2016. Apple's 'differential privacy' is about collecting your data — but not your data. (2016).

[23] Frank E. Grubbs. 1969. Procedures for Detecting Outlying Observations in Samples. *Technometrics* 11, 1 (1969), 1–21. https://doi.org/10.1080/00401706.1969.10490657

[24] Jiawei Han. 2005. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

[25] Samuel Haney, Ashwin Machanavajjhala, John M. Abowd, Matthew Graham, Mark Kutzbach, and Lars Vilhuber. 2017. Utility Cost of Formal Privacy for Releasing National Employer-Employee Statistics. In *Proceedings of the 2017 ACM International Conference on Management of Data* (Chicago, Illinois, USA) *(SIGMOD '17)*. ACM, New York, NY, 1339–1354. https://doi.org/10.1145/3035918.3035940

[26] Xi He, Ashwin Machanavajjhala, Cheryl Flynn, and Divesh Srivastava. 2017. Composing Differential Privacy and Secure Computation: A Case Study on Scaling Private Record Linkage. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (Dallas, Texas, USA) *(CCS '17)*. ACM, New York, NY, USA, 1389–1406. https://doi.org/10.1145/3133956.3134030

[27] Rebecca Hellerstein. 2008. Who bears the cost of a change in the exchange rate? Pass-through accounting for the case of beer. *Journal of International Economics* 76, 1 (2008), 14–32. https://EconPapers.repec.org/RePEc:eee:inecon:v:76:y:2008:i:1:p:14-32

[28] Ihab F. Ilyas and Xu Chu. 2019-06-17. Data Cleaning.

[29] Noah Johnson, Joseph P. Near, and Dawn Song. 2018. Towards Practical Differential Privacy for SQL Queries. *Proc. VLDB Endow.* 11, 5 (Jan. 2018), 526–539. https://doi.org/10.1145/3187009.3177733

[30] Eric Lantz, Kendrick Boyd, and David Page. 2015. Subsampled Exponential Mechanism: Differential Privacy in Large Output Spaces. In *Proceedings of the 8th ACM Workshop on Artificial Intelligence and Security* (Denver, Colorado, USA) *(AISec '15)*. ACM, New York, NY, USA, 25–33. https://doi.org/10.1145/2808769.2808776

[31] C. Y. Lee. 1961. An Algorithm for Path Connections and Its Applications. In *IRE Transactions on Electronic Computers (EC-10 (2))*. IEEE, 346 – 365. https://doi.org/10.1109/TEC.1961.5219222

[32] Jiongqian Liang and Srinivasan Parthasarathy. 2016. Robust Contextual Outlier Detection: Where Context Meets Sparsity. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management* (Indianapolis, Indiana, USA) *(CIKM '16)*. ACM, New York, NY, USA, 2167–2172. https://doi.org/10.1145/2983323.2983660

[33] Edward Lui and Rafael Pass. 2015. Outlier Privacy. In *Theory of Cryptography*, Yevgeniy Dodis and Jesper Buus Nielsen (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 277–305.

[34] A. Machanavajjhala, D. Kifer, J. Abowd, J. Gehrke, and L. Vilhuber. 2008. Privacy: Theory meets Practice on the Map. In *2008 IEEE 24th International Conference on Data Engineering*. 277–286. https://doi.org/10.1109/ICDE.2008.4497436

[35] G Mclachlan. 1999. Mahalanobis Distance. *Resonance* 4 (06 1999), 20–26. https://doi.org/10.1007/BF02834632

[36] Frank McSherry and Kunal Talwar. 2007. Mechanism Design via Differential Privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS '07)*. IEEE Computer Society, Washington, DC, USA, 94–103. https://doi.org/10.1109/FOCS.2007.41

[37] Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. 2007. Smooth Sensitivity and Sampling in Private Data Analysis. In *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing* (San Diego, California, USA) *(STOC '07)*. ACM, New York, NY, USA, 75–84. https://doi.org/10.1145/1250790.1250803

[38] Kobbi Nissim, Uri Stemmer, and Salil Vadhan. 2016. Locating a Small Cluster Privately. (04 2016). https://doi.org/10.1145/2902251.2902296

[39] Rina Okada, Kazuto Fukuchi, and Jun Sakuma. 2015. Differentially Private Analysis of Outliers. In *Machine Learning and Knowledge Discovery in Databases*, Annalisa Appice, Pedro Pereira Rodrigues, Vítor Santos Costa, João Gama, Alípio Jorge, and Carlos Soares (Eds.). Springer International Publishing, Cham, 458–473.

[40] Emanuel Parzen. 1962. On Estimation of a Probability Density Function and Mode. *Ann. Math. Statist.* 33, 3 (09 1962), 1065–1076. https://doi.org/10.1214/aoms/1177704472

[41] Bernard Rosner. 1983. Percentage Points for a Generalized ESD Many-Outlier Procedure. *Technometrics* 25, 2 (1983), 165–172. http://www.jstor.org/stable/1268549

[42] Guanting Tang, James Bailey, Jian Pei, and Guozhu Dong. 2013. Mining multidimensional contextual outliers from categorical relational data. *Intelligent Data Analysis* 19 (07 2013). https://doi.org/10.1145/2484838.2484883

[43] Gary L. Tietjen and Roger H. Moore. 1972. Some Grubbs-Type Statistics for the Detection of Several Outliers. *Technometrics* 14, 3 (1972), 583–597. https://doi.org/10.1080/00401706.1972.10488948 arXiv:https://www.tandfonline.com/doi/pdf/10.1080/00401706.1972.10488948

[44] J. Vaidya and C. Clifton. 2004. Privacy-preserving outlier detection. In *Fourth IEEE International Conference on Data Mining (ICDM'04)*. 233–240. https://doi.org/10.1109/ICDM.2004.10081

arXiv:https://www.tandfonline.com/doi/pdf/10.1080/00401706.1969.10490657

# 10 APPENDIX

We present our experiment results in two types of figures: i) histogram of accuracy distribution (Utility), which ranges from 0 to 1, where 1 is the accuracy of the direct approach introduced in Section 4 and ii) histogram of runtime distribution (Performance). The range for runtime varies according to the represented experiment.
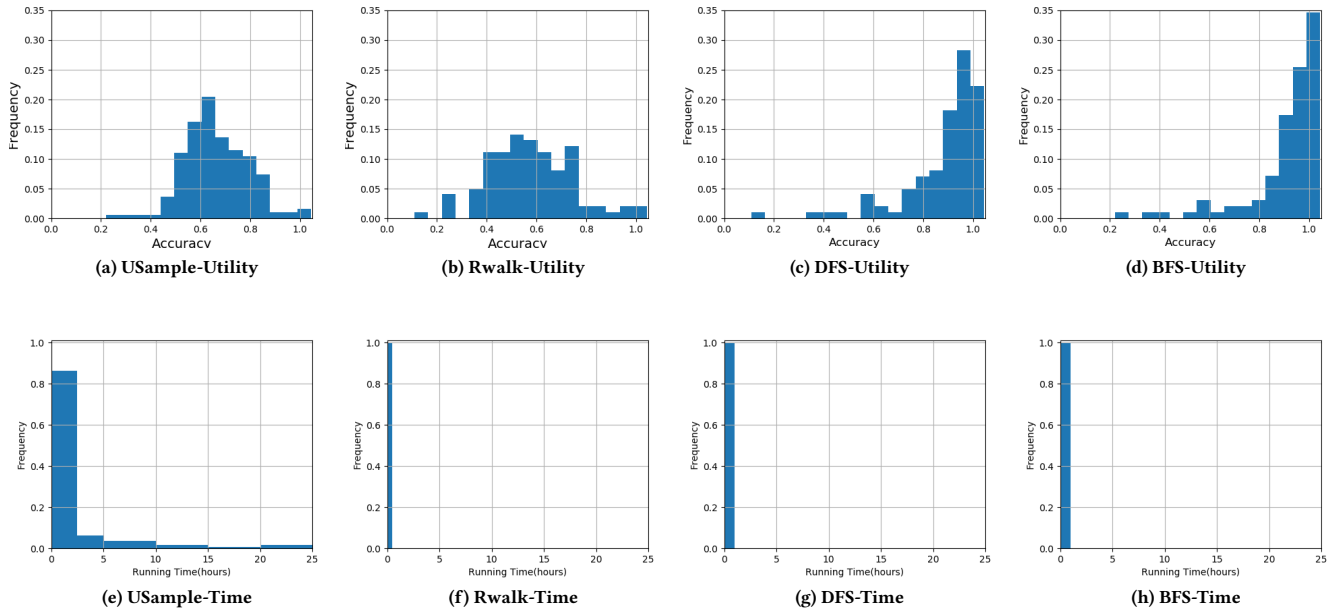
**Figure 1: Utility and Performance of PCORs for different sampling candidates, utility function outputs context population size, outlier detection algorithm is LOF, and $\epsilon = 0.2$ (a,e) Uniform Sampling, (b,f) Random walk, (c,g) DFS, (d,h) BFS; where (a), (b), (c) , (d) represent utility and (e), (f), (g) ,(h) demonstrate performance in running time**
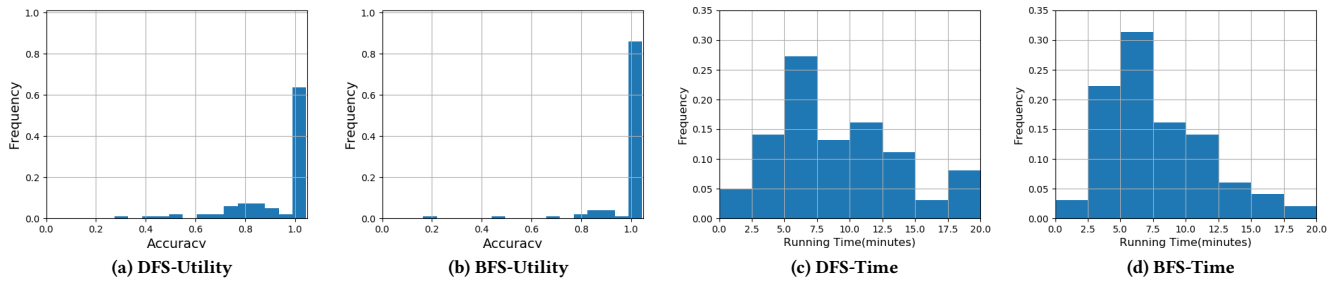


**Figure 2: Utility and Performance PCORS with different sampling candidates, utility is measured by the overlap with $C_V$, LOF is the outlier detection, and $\epsilon = 0.1$ (a,c) DFS, (b,d) BFS; where (a), (b) represent utility and (c), (d) demonstrate performance in running time**

**(a) Grubbs-Utility**  **(b) Histogram-Utility**  **(c) Grubbs-Time**  **(d) Histogram-Time**
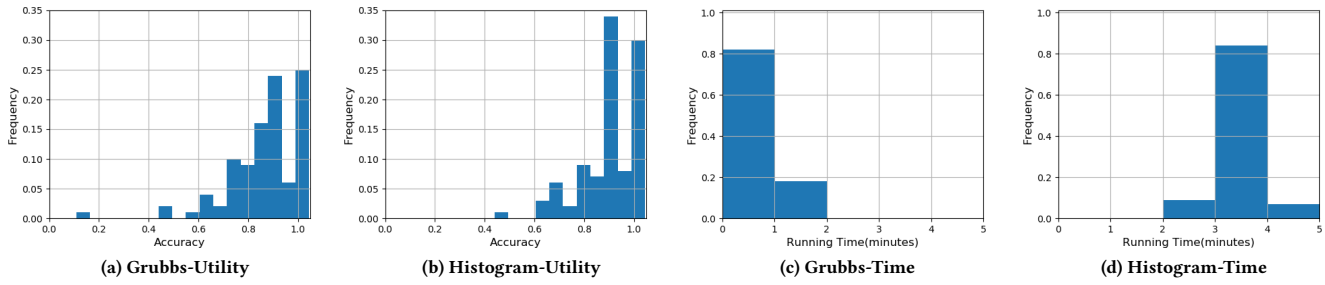
**Figure 3: Utility and Performance of PCORs with Grubbs and Histogram outlier detection algorithms using BFS sampling over utility of context population size for $\epsilon = 0.1$ (a,c) Grubbs, (b,d) Histogram; where (a), (b) represent utility and (c), (d) demonstrate performance in running time**
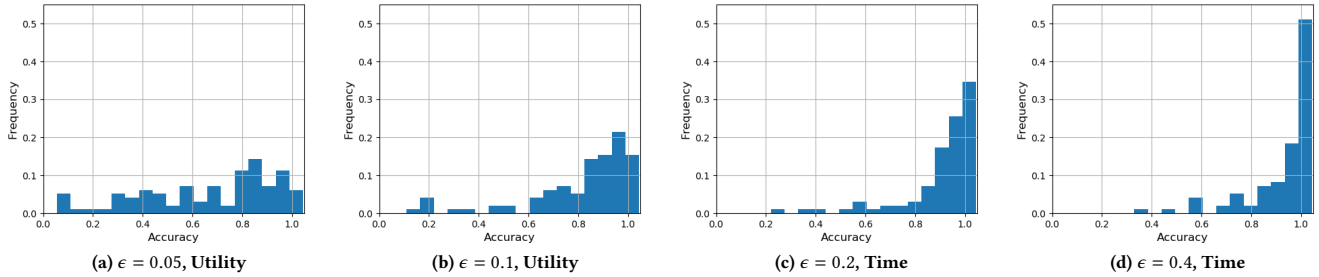


**(a) $\epsilon = 0.05$, Utility**  **(b) $\epsilon = 0.1$, Utility**  **(c) $\epsilon = 0.2$, Time**  **(d) $\epsilon = 0.4$, Time**

**Figure 4: The effect of privacy parameter over Utility in PCOR with BFS sampling and LOF outlier detection(a) $\epsilon = 0.05$, (b) $\epsilon = 0.1$, (c) $\epsilon = 0.2$, (d) $\epsilon = 0.4$**



**(a) $n = 25$,, Utility**  **(b) $n = 50$, Utility**  **(c) $n = 100$, Utility**  **(d) $n = 200$, Utility**

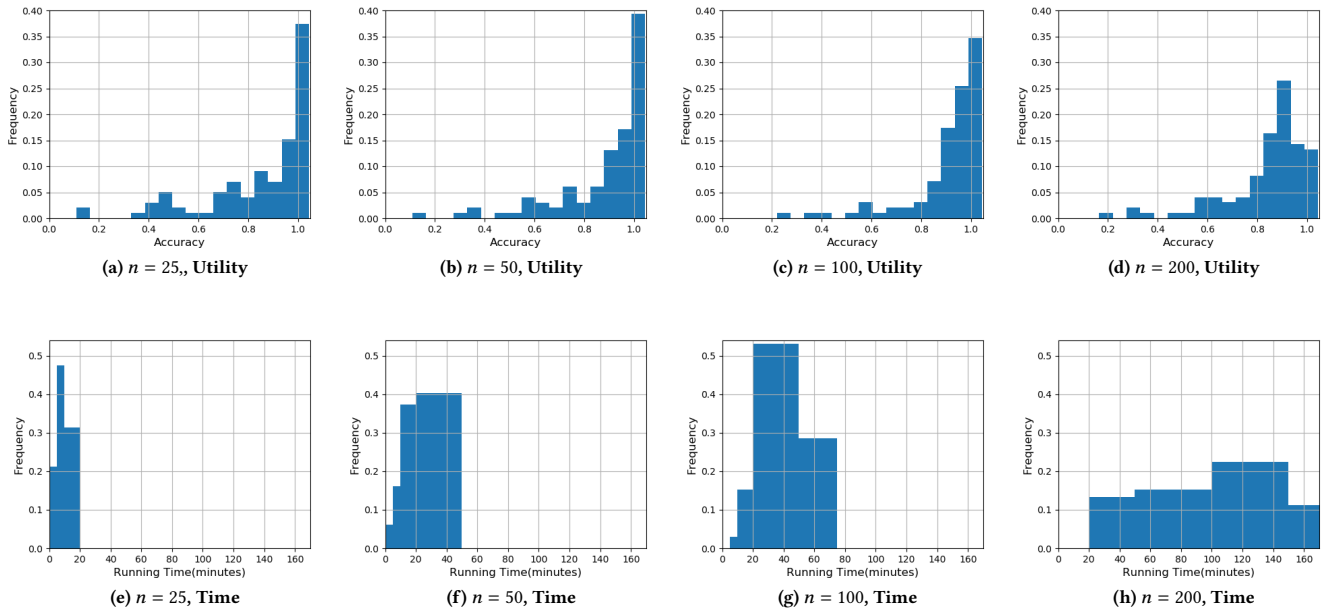**(e) $n = 25$, Time**  **(f) $n = 50$, Time**  **(g) $n = 100$, Time**  **(h) $n = 200$, Time**

**Figure 5: The effect of number of samples over utility and performance in PCOR with BFS sampling and LOF outlier detection for $\epsilon = 0.2$ (a,e) $n = 25$, (b,f) $n = 50$, (c,g) $n = 100$, (d,h) $n = 200$**