Using Noisy Binary Search for Differentially Private Anomaly Detection

Daniel M. Bittner[1(✉)], Anand D. Sarwate[2], and Rebecca N. Wright[3]

[1] Department of Computer Science, Rutgers University, Piscataway, NJ, USA
dbittner@cs.rutgers.edu
[2] Department of Electrical and Computer Engineering, Rutgers University, Piscataway, NJ, USA
anand.sarwate@rutgers.edu
[3] Department of Computer Science and DIMACS, Rutgers University, Piscataway, NJ, USA
rebecca.wright@rutgers.edu

Abstract. In this paper, we study differential privacy in noisy search. This problem is connected to noisy group testing: the goal is to find a defective or anomalous item within a group using only aggregate group queries, not individual queries. Differentially private noisy group testing has the potential to be used for anomaly detection in a way that provides differential privacy to the non-anomalous individuals while still helping to allow the anomalous individuals to be located. To do this, we introduce the notion of anomaly-restricted differential privacy. We then show that noisy group testing can be used to satisfy anomaly-restricted differential privacy while still narrowing down the location of the anomalous samples, and evaluate our approach experimentally.

1 Introduction

We consider the problem of privacy-sensitive anomaly detection—screening to detect individuals, behaviors, areas, or data samples of high interest. What defines an anomaly is context-specific: examples include a spoofed rather than genuine user attempting to log in to a web site, a fraudulent credit card transaction, or a suspicious traveler in an airport. The unifying assumption is that the number of truly anomalous points is quite small with respect to the population, so that deep screening of all individual data points would potentially be time-intensive, costly, and unnecessarily invasive of privacy. Anomaly detection is well studied (see the survey of Chandola et al. [11]), but methods to provide anomaly detection along with privacy are less well studied. In this paper we provide a framework for identifying anomalous data while guaranteeing quantifiable privacy in a rigorous sense. Once identified, such anomalies could warrant further data collection and investigation, depending on the context and relevant policies.

I. Dinur et al. (Eds.): CSCML 2018, LNCS 10879, pp. 20–37, 2018.
https://doi.org/10.1007/978-3-319-94147-9_3

While anomaly detection is important for many applications, it can also raise privacy concerns when the underlying data is sensitive. Search algorithms on private data can violate data use agreements and can make people uncomfortable with potential anomaly detection methods. In this paper, we focus on guaranteeing privacy during the deployment of anomaly detection. To achieve this, we take as our starting point the notion of *group testing* [14], which was most famously proposed for screening US military draftees for syphilis during World War II. In group testing, individuals are tested in groups to limit the number of tests. Using multiple rounds of screenings, a small number of positive individuals can be detected very efficiently. Group testing has the added benefit of providing privacy to individuals through plausible deniability—since the group tests use aggregate data, individual contributions to the test are masked by the group.

Our work takes the first steps toward strengthening and formalizing these privacy guarantees to achieve differential privacy. Differential privacy is a statistical measure of disclosure risk that was introduced in 2006 [18] and captures the intuition that an individual's privacy is protected if the results of a computation have at most a very small and quantifiable dependence on that individual's data. In the last decade, there has been an explosion of research in differential privacy, with many techniques and algorithms poised for practical application [20,27,31] and adoption underway by high-profile companies such as Apple [21] and Google [20].

Potential anomaly detection applications for group testing would rely on existing or new sensing technologies that can perform (reasonably accurate) queries in aggregate to reveal and isolate anomalous outliers. Applications might include privacy-sensitive methods for searching for outlying cell phone activity patterns or Internet activity patterns in a geographic location. These techniques are also in line with the US Department of Homeland Security's visionary goal of "screening at speed" [13]—unobtrusive screening of people, baggage, or cargo.

Our main contribution is a differentially private access mechanism for narrowing down the location of anomalies in a set of samples using noisy group testing. Our goal is to guarantee privacy for non-anomalous individuals while identifying anomalous samples. To formalize this we introduce the notion of *anomaly-restricted differential privacy*. By adding noise to group query results, we can guarantee differential privacy while allowing efficient and accurate detection of non-anomalous individuals. The adaptive sequential query design is an active learning algorithm for noisy binary search that is connected to information-theoretic models of communication with feedback.

A summary of our contributions is as follows:

- We introduce a new notion of anomaly-restriction differential privacy, which may be of independent interest.
- We provide a noisy group-based search algorithm that satisfies the anomaly-restricted differential privacy definition.
- We provide both theoretical and empirical analysis of our noisy search algorithm, showing that it performs well in some cases and exhibits the usual privacy/accuracy tradeoff of differentially private mechanisms.

2 Related Work

Machine learning methods have found widespread use in anomaly detection due to their ability to analyze and extract patterns from large amounts of data. Several surveys cover the wide variety of anomaly detection techniques and applications. For example, Hodge and Austin [23] and Agyemang et al. [2] survey anomaly detection techniques in the context of outlier detection via proximity and statistical approaches. Chandola et al. [11] provide a comprehensive survey addressing techniques in these categories as well as covering information theoretic and spectral approaches and techniques used in range of applicable fields including popular applications such as intrusion detection and fraud detection, as well as medical, industrial, image, and text anomalies.

Group testing describes a set of techniques for detection of anomalies from sets primarily containing non-anomalous items by performing testing on groups rather than querying individual items. Group testing was initially conceived during World War II as a cost-efficient method to test for syphilis by grouping multiple individuals' blood into a single sample [14]. A negative result for the single sample would imply all the individuals were negative, while a less-common positive result would require further follow up. The technique was not put into practice due to the limited number of individuals that could be tested at any one time, and group testing languished for several years before eventually being revived for industrial testing purposes [15].

Group testing has received more recent interest in the statistics and information theory communities. In particular, classical connections between group testing and error control coding have led to relaxations of the group testing problem, as surveyed in a recent paper by Mazumdar [29]. Group testing has also been used for multiaccess communications [5,37], data mining [28], molecular biology [12], and DNA screening [32]. Related concepts have been explored in constructing compressed sensing matrices [9,30].

Introduced by Dwork et al. in 2006 [18], differential privacy has become a widely studied framework for providing privacy-sensitive results from data analyses. Differential privacy for anomaly detection has been studied previously in the context of training classifiers using machine learning [22]. In contrast, our work addresses differential privacy during the deployment of an anomaly search algorithm by using differentially private group testing.

Our method of differentially private group testing makes use of noisy group testing [3,8,10], which provides methods that successfully identify anomalies using group queries among a set of items even if the answers to the group queries are not completely accurate. Specifically, we use a probabilistic binary search [4, 25,33–35], which is intimately connected to the problem of communication over noisy channels with feedback. The classical scheme by Horstein [24] uses what we would now call a Bayesian active learning approach to learn a threshold with noisy labels. In our case, the noise is used (and may even be deliberately introduced) to provide differential privacy.

3 Problem Formulation

The main idea behind our approach is to query individuals in groups and use noise to provide differential privacy. For this to work, we must have a group query which can detect the presence of an anomalous sample. As in active learning algorithms, we use multiple adaptive queries to locate the anomalies. In particular, we use a Bayesian formulation in which the algorithm maintains a probability distribution, or posterior belief, over the point representing its belief about where the anomaly lies. The number of queries can be controlled by either a stopping rule based on the belief or limits on overall privacy risk.

Notation: We generally use calligraphic script to denote sets. For any positive integer K, we denote the set $\{1, 2, \ldots, K\}$ by $[K]$.

3.1 Data Model

In this paper, we analyze a simplified version of the full problem with a single anomaly: for this setting, we can characterize the performance theoretically.

The data is a vector $\mathcal{X} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{n+1})$ of $n+1$ individuals, where $\mathbf{x}_i \in \mathbb{R}^+$. With some abuse of notation we write this as an ordered multiset $\{\mathbf{x}_i : i \in [n+1]\}$. The i-th element $\mathbf{x}_i$ represents the output of some anomaly score function applied to individual i: larger $\mathbf{x}$ denotes a higher anomaly level. One of the data points is an anomaly $\mathbf{x}^*$. Let i^* be the index of the anomaly, so that $\mathbf{x}_{i^*} = \mathbf{x}^*$. Two thresholds t_ℓ and t_h separate the anomaly value of the anomalous points from the other points such that

$$\mathbf{x} \in \begin{cases} [0, t_\ell] & \mathbf{x} \neq \mathbf{x}^* \\ [t_h, \infty) & \mathbf{x} = \mathbf{x}^* \end{cases} \tag{1}$$

for a set of two thresholds $t_\ell, t_h \in \mathbb{R}^+$ where $t_\ell < t_h$. This corresponds to a scenario where there is some measurement that can distinguish the anomaly from the non-anomalous values.

The data is held by an oracle that has access to $\mathcal{X}$ and can answer queries about $\mathcal{X}$. The search algorithm knows the number of points $n+1$ and the index set $[n+1]$, the levels t_ℓ and t_h separating anomalous from non-anomalous values, and that $\mathcal{X}$ contains a single anomalous point. However, it does not know the actual values $\{\mathbf{x}_1, \ldots, \mathbf{x}_{n+1}\}$. We wish to model a situation in which the oracle can only query groups of points. This could correspond to a situation where there is a measurement or sensor which can access aggregates (for example, all items in a given area) but not individual records.

3.2 Differential Privacy

The search algorithm queries the oracle, which provides *differentially private* responses. Traditional differential privacy protects privacy for every individual in the database [18]. The key difference in our model is that we only require that the oracle provide differential privacy for the non-anomalous points: we define a new notion of *anomaly-restricted* neighbors.

Definition 1 (Anomaly-Restricted Neighbors). *We say that two data sets* $\mathcal{D}$ *and* $\mathcal{D}'$ *are anomaly-restricted neighbors (and write* $\mathcal{D} \sim \mathcal{D}'$*) if* $\mathbf{x}^* \in \mathcal{D} \cap \mathcal{D}'$ *and* $|\mathcal{D} \cap \mathcal{D}'| = n$.

Definition 2 (Differential Privacy [18]). *A randomized mechanism* $\mathcal{A}(\cdot)$ *is* ϵ*-differentially private if for any set of measurable outputs* $\mathcal{Y}$ *and any two databases* $\mathcal{D}$ *and* $\mathcal{D}'$ *with* $\mathcal{D} \sim \mathcal{D}'$,

$$\Pr\left[\mathcal{A}(\mathcal{D}) \in \mathcal{Y}\right] \leq e^{\epsilon} \Pr\left[\mathcal{A}\left(\mathcal{D}'\right) \in \mathcal{Y}\right]. \tag{2}$$

A differentially private algorithm $\mathcal{A}(\cdot)$ guarantees that neighboring databases create similar outputs: for anomaly-restricted neighbors this means that adding or removing a single non-anomalous individual does not significantly alter the output of the mechanism. The privacy parameter ϵ is the privacy risk: larger values of ϵ allow larger differences between the distributions of $\mathcal{A}(\mathcal{D})$ and $\mathcal{A}(\mathcal{D}')$ [16–18]. Differential privacy controls the error probabilities in the hypothesis test between $\mathcal{D}$ and $\mathcal{D}'$ given the output of the mechanism [26,36].

The Laplace mechanism [18] is a common approach to making differentially private approximation to scalar functions $H(\cdot)$. This approach adds Laplace noise with a parameter that is a function of the privacy risk ϵ and the global sensitivity Δ_g of $H(\cdot)$. Corresponding to our new neighbor definition, we also need a model for anomaly-restricted global sensitivity.

Definition 3 (Anomaly-Restricted Global Sensitivity). *Let* $H(\cdot)$ *be a scalar-valued function. The anomaly-restricted global sensitivity of* $H(\cdot)$ *is*

$$\Delta_g = \max_{\mathcal{D},\mathcal{D}':\mathcal{D}\sim\mathcal{D}'} \left|H\left(\mathcal{D}\right) - H\left(\mathcal{D}'\right)\right|. \tag{3}$$

Given ϵ and $H(\cdot)$, the Laplace mechanism computes $\mathcal{A}(\mathcal{D}) = H(\mathcal{D}) + Z$ where $Z \sim \mathsf{Lap}\left(\Delta_g/\epsilon\right)$ where the Laplace distribution $\mathsf{Lap}(\lambda)$ has density

$$p\left(z;\lambda\right) = \frac{1}{2\lambda} \exp\left(-\frac{z}{\lambda}\right). \tag{4}$$

Differential privacy satisfies several *composition properties.*

Definition 4 (Simple Sequential Composition [18]). *Given a series of* n *independent differentially private mechanisms* $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ *with privacy parameters* $\epsilon_1, \epsilon_2, \ldots, \epsilon_n$ *computed on* $\mathcal{D}$*, the resulting function is differentially private with privacy parameter* $\sum_{i=1}^{n} \epsilon_i$.

Definition 5 (Parallel Composition [18]). *Given a series of* n *independent differentially-private mechanisms* $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ *with privacy parameters* $\epsilon_1, \epsilon_2, \ldots, \epsilon_n$ *computed on disjoint subsets of* $\mathcal{D}$*, then the resulting function is differentially private with privacy parameter* $\max_i \epsilon_i$.

In this paper, we restrict our attention to ϵ-differentially private methods. For approximate (ϵ, δ)-differential privacy there are stronger composition results in which the total privacy risk for sequential composition grows sublinearly with the number of terms [6,19,26], including the so-called "moments accountant" [1].

4 Algorithms

At each time t the search algorithm issues a query $\mathcal{Q}_t \subset [n+1]$ to the oracle that depends on the responses to past queries. A search algorithm consists of rules for sequentially selecting sets $\mathcal{Q}_1, \mathcal{Q}_2, \ldots$ with privacy risks $\epsilon_1, \epsilon_2, \ldots$ where $\mathcal{Q}_t \subset [n+1]$. A standard (noiseless) bisection search algorithm receives accurate queries and can then discard non-anomalous data points with certainty. When the oracle responses are noisy, we cannot fully discard any data points. We use a discretized version [4,7] of a probabilistic bisection algorithm [24] to adaptively determine the location of the anomaly. In particular, the algorithm uses a Bayesian inference step to update a probability mass function on $[n+1]$ that represents the *belief* about i^*.

4.1 Warmup: Randomized Response

A baseline algorithm for privacy binary search is noisy binary search using randomized response. At each time t the algorithm chooses a query $\mathcal{Q}_t$ and sends it to the oracle, which responds with

$$\mathcal{Y}_t = \mathbf{1}\left(i^* \in \mathcal{Q}_t\right) \oplus z_t \tag{5}$$

where $\oplus$ is addition modulo 2 and $z_t \sim \mathsf{Bernoulli}(p)$.

Proposition 1. *The response in* (5) *guarantees* $\log \frac{1-p}{p}$*-differential privacy.*

Given a response $\mathcal{Y}_t$ and noise parameter p, the algorithm can compute a posterior distribution on the location of the anomaly. Given $\bar{\mathcal{Q}}_t = [n+1] \setminus \mathcal{Q}_t$, let $\mathcal{R}_t = \mathcal{Q}_t$ if $\mathcal{Y}_t = 1$ and $\mathcal{R}_t = \bar{\mathcal{Q}}_t$ if $\mathcal{Y}_t = 0$. Given an initial estimate $\mathbf{f}_{t-1}$ on $[n+1]$, the Bayesian update is given by

$$\mathbf{f}_t\left(i\right) = \begin{cases} \frac{\mathbf{f}_{t-1}(i)(1-p)}{\sum_{j\in\mathcal{R}_t}\mathbf{f}_{t-1}(j)(1-p)+\sum_{k\notin\mathcal{R}_t}\mathbf{f}_{t-1}(k)p} & i \in \mathcal{R}_t \\ \frac{\mathbf{f}_{t-1}(i)p}{\sum_{j\in\mathcal{R}_t}\mathbf{f}_{t-1}(j)(1-p)+\sum_{k\notin\mathcal{R}_t}\mathbf{f}_{t-1}(k)p} & i \notin \mathcal{R}_t \end{cases}. \tag{6}$$

Because $p < \frac{1}{2}$, this rule increases $f_{t-1}(i)$ for $i \in \mathcal{R}_t$ and decreases $f_{t-1}(i)$ for $i \notin \mathcal{R}_t$ and eventually concentrates the posterior on i^*. If at each iteration the algorithm chooses a query $\mathcal{Q}_t$ with posterior probability close to 1/2 (i.e. a median split) this is a classic algorithm first analyzed by Burnashev and Zigangirov [7] (see also Horstein [24]) for i^* chosen uniformly in $[n+1]$; we can initialize by uniformly permuting the indices to use their result.

4.2 Proposed Algorithm: Differentially Private Binary Search

Before presenting the search algorithm, we introduce a modified oracle. Randomized response forces the oracle to determine whether $i^* \in \mathcal{Q}_t$ or $i^* \in \bar{\mathcal{Q}}_t$ and then obfuscates that value. In some cases, the oracle may simply be a noisy

privacy-preserving sensor that instead returns noisy estimates $\mathcal{A}(\mathcal{Q}_t; \mathcal{X})$ of some function $H(\mathcal{Q}_t; \mathcal{X})$. Consider an oracle that computes

$$\mathcal{Y}_t = \mathcal{A}\left(\mathcal{Q}_t; \mathcal{X}\right) \qquad \bar{\mathcal{Y}}_t = \mathcal{A}\left(\bar{\mathcal{Q}}_t; \mathcal{X}\right), \tag{7}$$

where the oracle splits the data set into $\mathcal{Q}_t$ and $\bar{\mathcal{Q}}_t = [n+1] \setminus \mathcal{Q}_t$ and returns anomaly-restricted differentially private approximation to both components. Notationally, we suppress $\mathcal{X}$ from $H(\mathcal{Q}_t; \mathcal{X})$ when it is clear from context.

There are many choices for the aggregation function $H(\cdot)$ used to calculate $\mathcal{A}$. For example, we could take the average $H(\mathcal{Q}) = \frac{1}{|\mathcal{Q}|}\sum_{i\in\mathcal{Q}} \mathbf{x}_i$. The anomaly-restricted global sensitivity is $\Delta_g = \frac{t_\ell}{|\mathcal{Q}|}$, so we can hypothetically add Laplace noise $Z \sim \mathsf{Lap}(\frac{t_\ell}{|\mathcal{Q}|\epsilon})$, $\bar{Z} \sim \mathsf{Lap}(\frac{t_\ell}{|\bar{\mathcal{Q}}|\epsilon})$ to form $\mathcal{Y} = H(\mathcal{Q}) + Z$ and $\bar{\mathcal{Y}} = H(\bar{\mathcal{Q}}) + \bar{Z}$, respectively.

In this work, we consider instead the max function:

$$H\left(\mathcal{Q}\right) = \max\{\mathbf{x}_i : i \in \mathcal{Q}\}. \tag{8}$$

Due to our definition of anomaly-restricted sensitivity, averages that include the anomaly can "dilute" the effect of the anomaly level. The max function can show the difference between $\mathcal{Y}$ and $\bar{\mathcal{Y}}$ in a way the depends less strongly on the distribution of the non-anomalous population. It has a higher sensitivity than the average function but we demonstrate its effectiveness empirically.

Lemma 1. *The anomaly-restricted global sensitivity of the aggregation function* $H(\mathcal{Q}; \mathcal{X}) = \max\{\mathbf{x}_i : i \in \mathcal{Q}\}$ *in* (8) *is* $\Delta_g(H) = t_\ell$.

Proof. Let $\mathcal{Q}$ be any query. Consider two anomaly-restricted neighboring data sets $\mathcal{X}$ and $\mathcal{X}'$ and let i^* be the index of the anomalous point. If $i^* \in \mathcal{Q}$ then $|H(\mathcal{Q}; \mathcal{X}) - H(\mathcal{Q}; \mathcal{X}')| = 0$ and $|H(\bar{\mathcal{Q}}; \mathcal{X}) - H(\bar{\mathcal{Q}}; \mathcal{X}')| \leq t_\ell$. If $i^* \in \bar{\mathcal{Q}}$ then $|H(\mathcal{Q}; \mathcal{X}) - H(\mathcal{Q}; \mathcal{X}')| \leq t_\ell$ and $|H(\bar{\mathcal{Q}}; \mathcal{X}) - H(\bar{\mathcal{Q}}; \mathcal{X}')| = 0$. Thus $\max|H(\mathcal{Q}; \mathcal{X}) - H(\mathcal{Q}; \mathcal{X}')| = t_\ell$. □

The oracle can then provide a differentially private query mechanism $\mathcal{A}$ for $H(\mathcal{Q}) = \max\{\mathbf{x}_i : i \in \mathcal{Q}\}$ by generating

$$\mathcal{A}\left(\mathcal{Q}\right) = \max\{\mathbf{x}_i : i \in \mathcal{Q}\} + Z \qquad \text{and} \qquad \mathcal{A}\left(\bar{\mathcal{Q}}\right) = \max\{\mathbf{x}_i : j \notin \mathcal{Q}\} + \bar{Z}, \tag{9}$$

where Z and $\bar{Z}$ are independent random variables with distribution $\mathsf{Lap}(t_\ell/\epsilon)$.

Given this revised oracle, we can turn to the search algorithm. The search is greedy: the searcher picks a query set which yields the most information (measured with respect to its belief) about the location of the anomaly. To represent our relative certainty about whether a given point is the anomaly, our search procedure updates a probability mass function $\mathbf{f}_t$ on $[n+1]$ where $\mathbf{f}_t(i) = \Pr(i^* = i)$. At each iteration we treat the previous posterior as a new prior and use $\mathbf{f}_{t-1}$ to determine the new query $\mathcal{Q}_t$. Since we do not have any prior knowledge about what element of $\mathcal{X}$ is the anomaly, at $t = 0$, we assume that each point is equally likely to be the anomaly: the initial prior distribution $\mathbf{f}_0$ is uniformly distributed on $[n+1]$, so $\mathbf{f}_0(i) = \frac{1}{n+1}$.

The algorithm uses the probability mass function $\mathbf{f}_{t-1}$ in order to select a query at each iteration $\mathcal{Q}_t$. First, the algorithm chooses a uniformly chosen random permutation σ on $[n+1]$. The corresponding permutation of the prior distribution is $\tilde{\mathbf{f}}_{t-1}(\sigma(i)) = \mathbf{f}_{t-1}(i)$. For a probability mass function on $[n+1]$ define the median $\mathcal{M}(\mathbf{f}) = \max\{m : \sum_{i=1}^{m} \mathbf{f}(i) < \sum_{i=m+1}^{n+1} \mathbf{f}(i)\}$.

The algorithm selects a query that maximizes information gain by dividing each query along the median of the permuted probability mass function.

At each iteration t the algorithm queries the oracle with

$$\mathcal{Q}_t = \left\{ i : \sigma(i) \leq \mathcal{M}\left(\tilde{\mathbf{f}}_{t-1}\right) \right\}. \tag{10}$$

Let $q_{t-1} = \sum_{i=0}^{\mathcal{M}(\tilde{\mathbf{f}}_{t-1})} \tilde{\mathbf{f}}_{t-1}(i)$ be the probability mass of the query set $\mathcal{Q}_t$. Note that $q \leq \frac{1}{2}$. Correspondingly, randomly choosing σ prevents reductions in information gain when q deviates significantly from $\frac{1}{2}$.

The oracle returns noisy values $\mathcal{Y}_t$ and $\bar{\mathcal{Y}}_t$ using (7) and (9) and the algorithm updates using a Bayesian update step similar to the case of randomized response. Given a prior belief $\mathbf{f}_{t-1}(i)$ that $i^* = i$, the likelihood of observing $(\mathcal{Y}_t, \bar{\mathcal{Y}}_t)$ is approximated by

$$\phi\left(\mathcal{Y}_t, \bar{\mathcal{Y}}_t \mid i^* = i\right) = \begin{cases} \frac{\epsilon^2}{4t_\ell{}^2} \exp\left(-\frac{\epsilon}{t_\ell}|\mathcal{Y}_t - t_h|\right) \exp\left(-\frac{\epsilon}{t_\ell}|\bar{\mathcal{Y}}_t - t_\ell|\right) & i \in \mathcal{Q}_t \\ \frac{\epsilon^2}{4t_\ell{}^2} \exp\left(-\frac{\epsilon}{t_\ell}|\mathcal{Y}_t - t_\ell|\right) \exp\left(-\frac{\epsilon}{t_\ell}|\bar{\mathcal{Y}}_t - t_h|\right) & i \in \bar{\mathcal{Q}}_t \end{cases}. \tag{11}$$

We can use this approximation in the Bayes update:

$$\mathbf{f}_t(i) = \frac{\mathbf{f}_{t-1}(i)\,\phi\left(\mathcal{Y}_t, \bar{\mathcal{Y}}_t \mid i^* = i\right)}{\sum_{j\in[n+1]} \mathbf{f}_{t-1}(j)\,\phi\left(\mathcal{Y}_t, \bar{\mathcal{Y}}_t \mid i^* = j\right)}. \tag{12}$$

There are two ways in which this procedure can halt. The first is if the algorithm expends the *privacy budget*. From the composition results, after T queries with ϵ-differentially private responses, the algorithm has incurred privacy risk $T\epsilon$. Given a total privacy budget b, we therefore halt the algorithm when $(T+1)\epsilon > b$.

The second halting condition is on the estimated posterior distribution $\mathbf{f}_t$. If the posterior has concentrated around a single point or small interval, we can halt the procedure and output the posterior distribution. This is characterized by computing some stopping time $\tau(\mathbf{f}_t)$. For example, Ben-Or and Hassidim [4] proposed a multi-epoch recursive search strategy and suggest taking $\tau(\mathbf{f}) = \mathbf{1}(\max_i \mathbf{f}_t(i) > \epsilon_{par})$ for $\epsilon_{par} = (24 \log n)^{-1/2}$ to prune the initial set $[n+1]$ into a smaller set of indices with larger posterior probability. In the approach studied by Burnashev and Zigangirov [7], the algorithm terminates when $\max_i \log \frac{\mathbf{f}_t(i)}{1-\mathbf{f}_t(i)} > \log(1/\delta)$ for a target error probability δ. In this case, the goal is to guarantee that the largest posterior probability is $\mathbf{f}_t(i^*)$ with probability $1-\delta$.

Pseudocode for the algorithm is shown in Algorithm 1.

Algorithm 1. PrivateBinarySearch$(\mathcal{X}, \epsilon, b, t_\ell, t_h, \epsilon_{par})$

$\mathbf{f}_0 \leftarrow \frac{1}{|\mathcal{X}|}$ for $i = 1, 2, \ldots, |\mathcal{X}|$, $t = 1$
while $\tau(\mathbf{f}_{t-1}) \neq 1$ and $t\epsilon < b$ **do**
  Draw σ uniformly at random from permutations on $[n+1]$.
  $\mathcal{Q}_t \leftarrow \{i : \sigma(i) \leq \mathcal{M}(\tilde{\mathbf{f}}_{t-1})\}$
  $\mathcal{Y}_{\mathcal{Q}_t} \leftarrow \mathcal{A}(\mathcal{Q}_t)$ and $\mathcal{Y}_{\bar{\mathcal{Q}}_t} \leftarrow \mathcal{A}(\bar{\mathcal{Q}}_t)$ from (9)
  Update $\mathbf{f}_t$ using (12)
  $t \leftarrow t + 1$
end while
return $\mathbf{f}_{t-1}$

4.3 Finding the Output

The search algorithm uses a halting condition based on $\mathbf{f}_{t-1}$ and then outputs $\mathbf{f}_{t-1}$, leaving open the question of how to determine the location of the anomaly i^*. If the algorithm waits for $\mathbf{f}_{t-1}$ to concentrate significantly, then with high probability the largest value in $\mathbf{f}_{t-1}$ corresponds to i^*. If instead it prioritizes the privacy budget, then it could pass a list of the largest entries of $\mathbf{f}_{t-1}$ for further processing. More issues regarding practical deployment of this algorithm are discussed in Sect. 7.

5 Analysis

The sensitivity of the max query in Lemma 1 immediately implies that each iteration guarantees ϵ-differential privacy.

Proposition 2. *Each query in Algorithm 1 is ϵ-differentially private. After t iterations of the loop, the overall privacy risk is $t\epsilon$.*

Proof. The result follows from the fact that the noisy computation in (9) guarantees ϵ-differential privacy for $Z, \bar{Z} \sim \mathsf{Lap}(t_\ell/\epsilon)$. Fix neighboring anomaly-restricted datasets $\mathcal{X}$ and $\mathcal{X}'$ and queries $\mathcal{Q} \subset [|\mathcal{X}|]$ and $\mathcal{Q}' \subset [|\mathcal{X}'|]$. Since each iteration of the algorithm splits the dataset into disjoint subsets and applies $\mathcal{A}$ to each independently, by demonstrating that each $\mathcal{A}$ is ϵ-differentially private, we can apply the parallel composition theorem of differential privacy in Definition 5.

If $\mathcal{Q} = \mathcal{Q}'$, then clearly

$$\Pr\left[\mathcal{A}\left(\mathcal{Q}\right) = \mathcal{Y}\right] = \Pr\left[\mathcal{A}\left(\mathcal{Q}'\right) = \mathcal{Y}\right], \tag{13}$$

so the application of $\mathcal{A}$ is ϵ-differentially private. We are therefore left with the case where $\mathcal{Q}$ and $\mathcal{Q}'$ differ in a single non-anomalous point. By the post-processing invariance of differential privacy [18], it is sufficient to show that $\mathcal{Y} = \mathcal{A}(\mathcal{Q})$ is ϵ-differentially private. This follows from Lemma 1 and the differential privacy of the Laplace mechanism. □

Analyzing the convergence of Algorithm 1 is challenging because using Laplace noise means the amount of "progress" made by the algorithm using (12) varies from iteration to iteration. Furthermore, because we only know bounds on the non-anomalous and anomalous values, the update rule is performing an approximation to a Bayes update.

To understand the convergence of the method, we show that a modified version of the update reduces the problem to a noisy binary search. There are two changes: firstly, we do away with the random permutation and secondly, we compute a binary response from $(\mathcal{Y}_t, \bar{\mathcal{Y}}_t)$ and then apply the same Bayes update as randomized response update in (6). More specifically, the algorithm computes $\mathcal{Z}_t = \mathbf{1}\left(\mathcal{Y}_t > \bar{\mathcal{Y}}_t\right)$ and performs a Bayesian update of the prior distribution $\mathbf{f}_{t-1}$ to form the posterior $\mathbf{f}_t$. Because the determination of the subset containing the anomaly $\mathcal{Z}_t$ may be inaccurate, in order to perform the update, we must determine $p = \Pr(i^* \in \mathcal{Z}_t)$.

Lemma 2.

$$\Pr\left(i^* \in \mathcal{Z}_t\right) \geq 1 - \left(\frac{1}{2} + \frac{t_h - t_\ell}{t_\ell} \cdot \frac{\epsilon}{4}\right) \exp\left(-\epsilon \frac{t_h - t_\ell}{t_\ell}\right). \tag{14}$$

Proof. Without loss of generality, let us assume $i^* \in \mathcal{Q}$. We want to find the probability that the following difference is positive:

$$\mathcal{Y} - \bar{\mathcal{Y}} = \max\{\mathbf{x}_i : i \in \mathcal{Q}\} + Z - \max\{\mathbf{x}_i : i \in \bar{\mathcal{Q}}\} - Z'. \tag{15}$$

By assumption, $H(\mathcal{Q}) \geq t_h$ and $H(\bar{\mathcal{Q}})\} \leq t_\ell$, thus $H(\mathcal{Q}) - H(\bar{\mathcal{Q}}) \geq t_h - t_\ell$. Therefore $\Pr(Z' - Z) > t_h - t_\ell$ serves as a lower bound on the probability of that the query will return an erroneous result due to noise.

Since the Z and Z' both have zero mean, the distribution of $W = Z' - Z$ is the same as that of $Z + Z'$, which can be found by convolving the two Laplace densities given by (4) with parameter $\lambda = t_\ell/\epsilon$. By assumption, $t_h - t_\ell > 0$, so the probability density function for $w > 0$ is

$$\begin{aligned}
\mathbf{f}(w) &= \int_{-\infty}^{\infty} \frac{1}{2\lambda} \exp\left(-|z|/\lambda\right) \frac{1}{2\lambda} \exp\left(-|z - w|/\lambda\right)\, dz && (16)\\
&= \int_{-\infty}^{0} \frac{1}{4\lambda^2} \exp\left((2z - w)/\lambda\right)\, dz + \int_{0}^{w} \frac{1}{4\lambda^2} \exp\left(-w/\lambda\right)\, dz \\
&\quad + \int_{w}^{\infty} \frac{1}{4\lambda^2} \exp\left(-(2z - w)/\lambda\right)\, dz && (17)\\
&= \frac{1}{8\lambda} \exp\left(-w/\lambda\right) + \frac{w}{4\lambda^2} \exp\left(-w/\lambda\right) + \frac{1}{8\lambda} \exp\left(-w/\lambda\right) && (18)\\
&= \frac{\lambda + w}{4\lambda^2} \exp\left(-w/\lambda\right). && (19)
\end{aligned}$$

The cumulative distribution function for $w > 0$ is

$$\mathcal{F}(W \leq w) = \frac{1}{2} + \int_0^w \frac{\lambda + u}{4\lambda^2} \exp(-u/\lambda)\ du \tag{20}$$

$$= \frac{1}{2} + \left[-\frac{1}{4}\exp(-u/\lambda)\right]_{u=0}^{w} + \left[-\frac{u}{4\lambda}\exp(-u/\lambda)\right]_{u=0}^{w} - \int_0^w -\frac{1}{4\lambda}\exp(-u/\lambda)\ du \tag{21}$$

$$= \frac{1}{2} - \frac{1}{4}\exp(-w/\lambda) + \frac{1}{4} - \frac{w}{4\lambda}\exp(-w/\lambda) - \frac{1}{4}\exp(-w/\lambda) + \frac{1}{4} \tag{22}$$

$$= 1 - \left(\frac{1}{2} + \frac{w}{4\lambda}\right)\exp(-w/\lambda)\,. \tag{23}$$

Now, plugging in $\lambda = \frac{t_\ell}{\epsilon}$ and $w = t_h - t_\ell$ we have (14). □

Thus, we define

$$p = \left(\frac{1}{2} + \frac{t_h - t_\ell}{t_\ell} \cdot \frac{\epsilon}{4}\right)\exp\left(-\epsilon\frac{t_h - t_\ell}{t_\ell}\right). \tag{24}$$

from (14) and apply the Bayes update in (6).

Proposition 3. *Suppose the anomaly i^* is uniformly distributed in $[n+1]$. For any $\delta \in (0,1)$, let*

$$T = \min\left\{t : \max_i \log\frac{\mathbf{f}_t(i)}{1 - \mathbf{f}_t(i)} > \log\frac{1}{\delta}\right\}. \tag{25}$$

Set the stopping time $\tau(\mathbf{f}_{t-1}) = \mathbf{1}(t = T)$. Then the modified version of Algorithm 1 using $\mathcal{Z}_t = \mathbf{1}\left(\mathcal{Y}_t > \bar{\mathcal{Y}}_t\right)$ and (24) with update (6) satisfies

$$\mathbb{E}[T] \leq \frac{\log(n+1) + \log(1/\delta) + \epsilon}{1 - h_{\mathrm{b}}(p)} \tag{26}$$

where $h_{\mathrm{b}}(p) = -p\log p - (1-p)\log(1-p)$ is the binary entropy function.

Proof. The result follows by mapping the algorithm to the interval estimation problem studied by Burnashev and Zigangirov [7]. The main difference is that when using $\mathcal{Z}_t$, (24) is only an upper bound on the error probability of the oracle for randomized response. However, this means that the oracle is only potentially *less* noisy than the randomized response oracle. Using the stopping rule in (25), we get the upper bound on the expected number of queries [7, Theorem 3]. □

6 Experimental Results

We demonstrate the practical performance of our approach through experiments on a data set for anomaly detection. The experiments investigate how different

configurations of input parameters and constraints on the datasets can affect accuracy and total privacy risk. Specifically, we are interested in the impact of the thresholds t_h and t_ℓ, the oracle response configuration, and the halting conditions τ and privacy budget b.

6.1 Dataset

The experiments use the A1 Benchmark from the Yahoo Labeled Anomaly Detection Dataset, part of the Yahoo Webscope reference library [38]. Each dataset in the benchmark is preprocessed down to single anomaly by selecting the largest anomalous point in each dataset and selecting thresholds by letting $\mathbf{x}_j = \max\{\mathbf{x}_i : i \neq i^*\}$ and setting $t_h = \mathbf{x}_{i^*} - .1(\mathbf{x}_{i^*} - \mathbf{x}_j)$ and $t_\ell = \mathbf{x}_j + .1(\mathbf{x}_{i^*} - \mathbf{x}_j)$. Some experiments are run specifically on datasets 6 and 8 in order to explore the effects of the non-anomalous point distribution on the algorithm performance. These two datasets exemplify the two primary distributions for sets contained in benchmark: datasets that are a mixture of normal distributions, and datasets where points are heavily skewed toward 0.

6.2 Procedure

Because we are interested in approximate detection of the anomaly, we declare that the algorithm succeeds if it halts and can output a small set S of indices such that $i^* \in S$. In particular, we choose $|S| = 4$ and set S to be the indices with the 4 largest posterior probabilities. This selection is to capture the difference between $\mathbf{f}(i^*)$ being the close to the largest posterior probability and being much smaller. Cases where $\mathbf{f}(i) = \mathbf{f}(j)$ for $i \neq j$ are prevented in practice by the randomized permutation of the probability mass function after each iteration. For these experiments, $\tau = \mathbf{1}(\max\{\mathbf{f}(i) : i \in [n+1]\} > 0.5)$ is used as a halting condition when not otherwise specified. Each configuration of the algorithm parameters are run for a set number of cycles c. The approximate average error rate for the configuration is $(1 - \frac{\sum_{i=1}^{c} \mathbf{1}(i^* \in S_i)}{c})$ and the average total privacy risk is $\frac{\sum_{i=1}^{c} (t\epsilon)_i}{c}$. For these experiments, we take $c = 100$.

6.3 Results

We demonstrate the algorithm's performance as a function of the privacy parameter ϵ. Smaller ϵ values result in noisier responses from the oracle which require more iterations to reach the halting condition. Correspondingly, larger values of ϵ decrease noise which requires fewer total iterations, but at greater privacy cost per iteration. The tradeoff between error rate and total privacy risk forms a concave upward curve. Lower values of the privacy parameter are more costly in total privacy risk as the noise at each iteration strongly decreases $\Pr(i^* \in \mathcal{R})$. Increasing the privacy parameter increases $\Pr(i^* \in \mathcal{R})$ at a greater rate than the privacy cost per iteration increases, thus decreasing total privacy risk. However, these improvements have diminishing returns. Eventually, increasing the privacy

parameter no longer improves the error rate as $\Pr(i^* \in \mathcal{R}) \to 1$. At this point, increasing the privacy parameter doesn't improve the error rate, but continues to increase total privacy risk.

Threshold Ratios. Figure 1 demonstrates the effect of the thresholds on the algorithm's performance. Each point in the figure depicts the error rate as a function of that dataset's threshold ratio $\frac{t_h - t_\ell}{t_\ell}$ with privacy parameter set to $\epsilon = 1$. A dataset with a higher threshold ratio tends to perform better than an equivalent dataset with a lower threshold ratio for a given value of the privacy parameter. This is due to $\Delta_g = t_\ell$, which causes smaller differences between thresholds $t_h - t_\ell$ to be more likely to be overcome by noise. The steep improvement in error rate for small changes in the threshold ratio highlight the importance of tuning the privacy parameter to the thresholds of the dataset. (Note that datasets 6 and 8 were selected to have similar threshold ratios at 0.647 and 0.701 respectively).

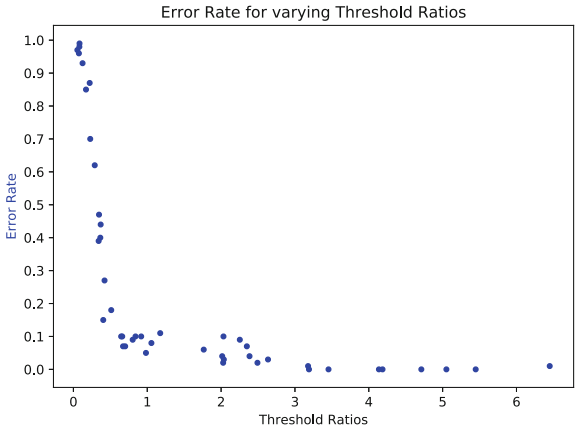

Fig. 1. Error rate for each dataset as a function of the threshold ratio $\frac{t_h - t_\ell}{t_\ell}$.

Oracle Response Constructions. Figure 2 demonstrates how different constructions of the oracle response and Bayesian update methods affect the error rate. The proposed oracle response approaches include the randomized response oracle (5), the binarized noisy response oracle (14) and the direct noisy result oracle (11). Despite all constructions achieving $t\epsilon$-differential privacy, there is a strong difference in effect on the error rate and total privacy risk.

Randomized response has the worst error rate because the oracle error probability is fixed. This contrasts with the oracle mechanisms that use the noisy aggregations: the actual noisy response depend on the values $(\mathcal{Y}, \bar{\mathcal{Y}})$, which can be more informative depending on the noise. For example, when the actual difference between $\mathcal{Y}$ or $\bar{\mathcal{Y}}$ exceeds the difference between t_h and t_ℓ, added noise is less likely to cause incorrect responses than in randomized response. Similarly, the oracle that directly uses the noisy response performs better than the

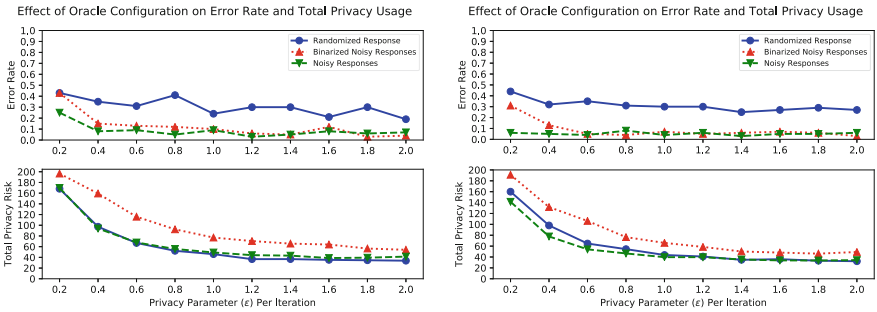

Fig. 2. The error rate and total privacy risk as a function of the privacy parameter ϵ for different oracle response constructions on data sets 6 and 8.

binary oracle construction as the likelihood for the binary oracle at each iteration is a lower bound given by (14) which gives up some information gain on each iteration. Because the binarized construction is a lower bound on the actual likelihood, more updates become required to achieve the same effect as the other constructions and thus ends up having greater total privacy risk.

Algorithm Halting Conditions. The algorithm's two termination conditions, τ and total privacy risk exceeding budget b, are explored in Figs. 3 and 4. Figure 3 depicts the algorithm's error rate with varying budget constraints where the halting constraint τ has been removed. When the total privacy risk passes pre-assigned budget checkpoints, S is checked for the presence of the anomaly and the algorithm continues. Similarly Fig. 4 depicts various halting constraints where the budget constraint has been removed and again checks S at pre-assigned halting checkpoints. When the algorithm is forced to preemptively halt because total privacy risk exceeds the budget, errors are excessively high. This is due to the increased chance that not enough iterations have been run to allow the algorithm to overcome noisy oracle responses. When the privacy parameter ϵ is larger, the

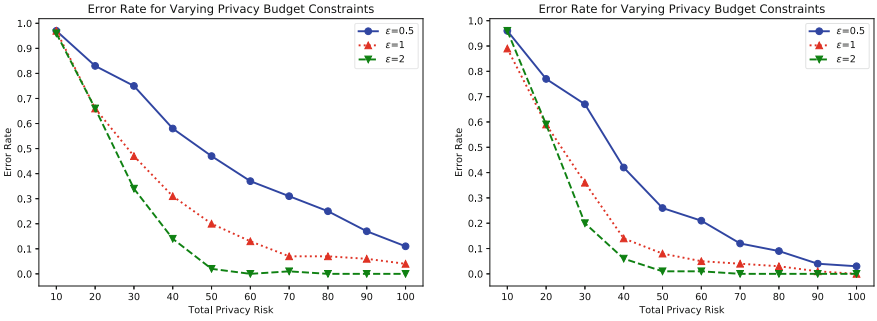

Fig. 3. Error rates for varying inputs of the privacy parameter ϵ with differing maximum budget constraints b for datasets 6 and 8.

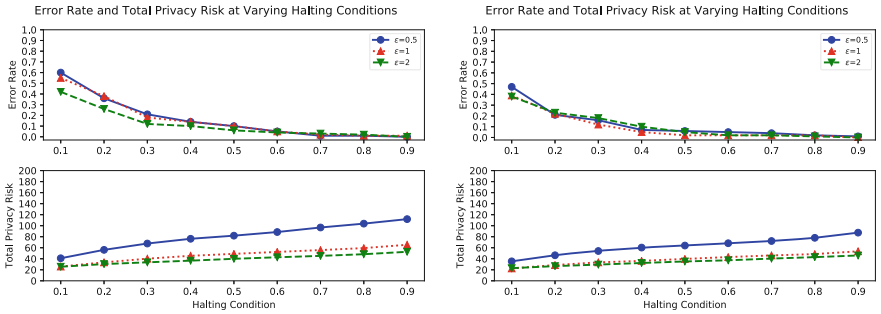

Fig. 4. Error rates and total privacy risk across varying halting constraints τ for data sets 6 and 8.

algorithm is more likely to suffer errors from the algorithm terminating early. Correspondingly, when total privacy risk does not prevent early termination due to budget b, larger values of ϵ result in fewer errors. Thus, a proper privacy budget should be allocated to perform enough iterations to prevent errors due to halting early.

Figure 4 demonstrates how different halting conditions τ affect the error rate with unlimited privacy budget. Specifically, the figure depicts the effect of altering α for $\tau = \mathbf{1}(\max\{\mathbf{f}(i) : i \in [n+1]\} > \alpha)$. As the halting condition serves as a requirement of convergence of the probability mass toward a single point, the algorithm can steadily improve the error rates by increasing α. This requires correspondingly more iterations to achieve, incurring greater total privacy risk for any run of the algorithm.

7 Discussion

We have described a differentially private search algorithm using noisy binary search with applications to anomaly detection. For this application, we defined a new notion of anomaly-restricted neighboring databases to capture the idea that anomalous points (which potentially merit scrutiny even if it is privacy-invasive) are not given privacy guarantees. The noise in the algorithm provides quantifiable privacy during the search. We showed theoretically and empirically that the greedy Bayesian search strategy can quickly narrow down a small set of samples that contain the anomaly.

There are a number of practical considerations that must be further addressed for our work to be useful in particular applications. For example, in most cases, it will be necessary to handle multiple anomalies rather than only a single anomaly. If a good upper bound is known on the expected maximum number of anomalous points, then one approach for using our method would be to first divide the set into disjoint subsets that with high probability contain only a single anomaly, and then proceeding to apply our method to each of those subsets individually.

In any particular application, it is also necessary to specify what points the algorithm should return. This depends on various factors, including what will be done with those points. We envision a scenario in which the points returned undergo some further screening, presumably after appropriate policies are followed. However, this creates a tradeoff between false positives and false negatives. To provide the most privacy, it would be desirable for the returned set to be as small as possible. However, narrowing down too far increases the chance of returning a set that misses the anomalous point. In very large search spaces or problems with many anomalies, one option would be to recursively prune out non-anomalous points: while this should work well in practice, theoretically analyzing the corresponding privacy-utility tradeoffs may be quite complex.

Our method uses a fixed privacy loss ϵ_t per iteration, not without loss of generality. Varying ϵ_t across iterations in a decaying manner could correspond to active learning or noisy search under the Tsybakov noise condition. Results from active learning can yield bounds on convergence to interpret the error/privacy tradeoff. A key difference between our search model and standard noisy search is that we can design the noise to optimize the privacy-utility tradeoff.

In order to provide privacy without relying on a trusted party, our method relies on the existence of a sensor or other measurement device that carries out the noisy aggregate queries directly, without carrying out individual queries and computing a noisy aggregate result from them. Practical use of our techniques therefore depends on the practical creation and deployment of such sensors.

Acknowledgements. This work was partially supported by NSF under award CCF-1453432, DARPA and SSC Pacific under contract N66001-15-C-4070, and DHS under award 2009-ST-061-CCI002 and contract HSHQDC-16-A-B0005/HSHQDC-16-J-00371.

References

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., Zhang, L.: Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communication Security (CCS 2016), Vienna, Austria, 24–28 October 2016, pp. 303–318. ACM (2016). https://doi.org/10.1145/2976749.2978318
2. Agyemang, M., Barker, K., Alhajj, R.: A comprehensive survey of numeric and symbolic outlier mining techniques. Intell. Data Anal. **10**(6), 521–538 (2006)
3. Atia, G.K., Saligrama, V.: Boolean compressed sensing and noisy group testing. IEEE Trans. Inf. Theory **58**(3), 1880–1901 (2012). https://doi.org/10.1109/TIT.2011.2178156
4. Ben-Or, M., Hassidim, A.: The Bayesian learner is optimal for noisy binary search (and pretty good for quantum as well). In: 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008), pp. 221–230 (2008). https://doi.org/10.1109/FOCS.2008.58
5. Berger, T., Mehravari, N., Towsley, D., Wolf, J.: Random multiple-access communication and group testing. IEEE Trans. Commun. **32**(7), 769–779 (1984). https://doi.org/10.1109/TCOM.1984.1096146

6. Bun, M., Steinke, T.: Concentrated differential privacy: simplifications, extensions, and lower bounds. In: Hirt, M., Smith, A. (eds.) TCC 2016. LNCS, vol. 9985, pp. 635–658. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53641-4_24
7. Burnashev, M.V., Zigangirov, K.S.: An interval estimation problem for controlled observations. Probl. Inf. Transm. **10**, 223–231 (1974)
8. Cai, S., Jahangoshahi, M., Bakshi, M., Jaggi, S.: GROTESQUE: noisy group testing (quick and efficient). Technical report arXiv:1307.2811 [cs.IT], ArXiV, July 2013. http://arxiv.org/abs/1307.2811
9. Calderbank, R., Howard, S., Jafarpour, S.: Construction of a large class of deterministic sensing matrices that satisfy a statistical isometry property. IEEE J. Sel. Topics Sig. Process. **4**(2), 358–743 (2010). https://doi.org/10.1109/JSTSP.2010.2043161
10. Chan, C.L., Jaggi, S., Saligrama, V., Agnihotri, S.: Non-adaptive group testing: explicit bounds and novel algorithms. IEEE Trans. Inf. Theory **60**(5), 3019–3035 (2014). https://doi.org/10.1109/TIT.2014.2310477
11. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: a survey. ACM Comput. Surv. (CSUR) **41**(3), 15 (2009)
12. Chen, H.B., Hwang, F.K.: A survey on nonadaptive group testing algorithms through the angle of decoding. J. Comb. Optim. **15**(1), 49–59 (2008). https://doi.org/10.1007/s10878-007-9083-3
13. Department of Homeland Security: Screening at speed (2017). https://www.dhs.gov/science-and-technology/apex-screening-speed. Accessed 3 Aug 2017
14. Dorfman, R.: The detection of defective members of large populations. Ann. Math. Stat. **14**(4), 436–440 (1943). http://www.jstor.org/stable/2235930
15. Du, D.Z., Hwang, F.K.: Combinatorial group testing and its applications, vol. 12, 2nd edn. World Scientific (1999). https://doi.org/10.1142/4252
16. Dwork, C.: Differential privacy. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4052, pp. 1–12. Springer, Heidelberg (2006). https://doi.org/10.1007/11787006_1. https://www.microsoft.com/en-us/research/publication/differential-privacy/
17. Dwork, C.: A firm foundation for private data analysis. Commun. ACM **54**(1), 86–95 (2011)
18. Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 265–284. Springer, Heidelberg (2006). https://doi.org/10.1007/11681878_14
19. Dwork, C., Rothblum, G., Vadhan, S.: Boosting and differential privacy. In: 2010 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS), Las Vegas, NV, pp. 51–60, October 2010. https://doi.org/10.1109/FOCS.2010.12
20. Erlingsson, Ú., Pihur, V., Korolova, A.: RAPPOR: randomized aggregatable privacy-preserving ordinal response. In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS 2014), pp. 1054–1067 (2014). https://doi.org/10.1145/2660267.2660348
21. Evans, J.: What Apple users need to know about differential privacy. IComputerWorld, June 2016. http://www.computerworld.com/article/3088179/apple-mac/what-apple-users-need-to-know-about-differential-privacy.html
22. Ghassemi, M., Sarwate, A.D., Wright, R.N.: Differentially private online active learning with applications to anomaly detection. In: Proceedings of the 9th ACM Workshop on Artificial Intelligence and Security (AISec), Vienna, Austria, pp. 117–128, October 2016

23. Hodge, V.J., Austin, J.: A survey of outlier detection methodologies. Artif. Intell. Rev. **22**(2), 85–126 (2004)
24. Horstein, M.: Sequential transmission using noiseless feedback. IEEE Trans. Inf. Theory **9**(3), 136–143 (1963). https://doi.org/10.1109/TIT.1963.1057832
25. Jedynak, B., Frazier, P.I., Sznitman, R.: Twenty questions with noise: Bayes optimal policies for entropy loss. J. Appl. Probab. **49**(1), 114–136 (2012). https://doi.org/10.1239/jap/1331216837
26. Kairouz, P., Oh, S., Viswanath, P.: The composition theorem for differential privacy. IEEE Trans. Inf. Theory **63**(6) (2017). https://doi.org/10.1109/TIT.2017.2685505
27. Machanavajjhala, A., Kifer, D., Abowd, J.M., Gehrke, J., Vilhuber, L.: Privacy: theory meets practice on the map. In: IEEE 24th International Conference on Data Engineering (ICDE), pp. 277–286 (2008). https://doi.org/10.1109/ICDE.2008.4497436
28. Macula, A.J., Popyack, L.J.: A group testing method for finding patterns in data. Discrete Appl. Math. **144**(1–2), 149–157 (2004). https://doi.org/10.1016/j.dam.2003.07.009
29. Mazumdar, A.: Nonadaptive group testing with random set of defectives. IEEE Trans. Inf. Theory **62**(12), 7522–7531 (2016). http://ieeexplore.ieee.org/document/7577749/
30. Mazumdar, A., Barg, A.: Sparse-recovery properties of statistical RIP matrices. In: Proceedings of the 49th Allerton Conference on Communication, Control and Computing, pp. 9–12, September 2011. https://doi.org/10.1109/Allerton.2011.6120142
31. Mir, D.J., Isaacman, S., Cáceres, R., Martonosi, M., Wright, R.N.: DP-WHERE: differentially private modeling of human mobility. In: Proceedings of the 2013 IEEE International Conference on Big Data, October 2013. https://doi.org/10.1109/BigData.2013.6691626
32. Ngo, H.Q., Du, D.Z.: A survey on combinatorial group testing algorithms with applications to DNA library screening. In: Du, D.Z., Pardalos, P.M., Wang, J. (eds.) Discrete Mathematical Problems with Medical Applications. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 55. AMS (2000)
33. Nowak, R.: Generalized binary search. In: 2008 46th Annual Allerton Conference on Communication, Control, and Computing, pp. 568–574. IEEE (2008)
34. Nowak, R.: Noisy generalized binary search. In: Advances in Neural Information Processing Systems, pp. 1366–1374 (2009)
35. Waeber, R., Frazier, P.I., Henderson, S.G.: Bisection search with noisy responses. SIAM J. Control Optim. **51**(3), 2261–2279 (2013)
36. Wasserman, L., Zhou, S.: A statistical framework for differential privacy. J. Am. Stat. Assoc. **105**(489), 375–389 (2010). https://doi.org/10.1198/jasa.2009.tm08651
37. Wolf, J.: Born again group testing: multiaccess communications. IEEE Trans. Inf. Theory **31**(2), 185–191 (1985). https://doi.org/10.1109/TIT.1985.1057026
38. Yahoo Labs: S5 - a labeled anomaly detection dataset, version 1.0 (2016). https://webscope.sandbox.yahoo.com/catalog.php?datatype=s&did=70