

Detecting Violations of Differential Privacy

Zeyu Ding
 Pennsylvania State University
 dxd437@psu.edu

Yuxin Wang
 Pennsylvania State University
 ywang@cse.psu.edu

Guanhong Wang
 Pennsylvania State University
 gpw5092@psu.edu

Danfeng Zhang
 Pennsylvania State University
 zhang@cse.psu.edu

Daniel Kifer
 Pennsylvania State University
 dkifer@cse.psu.edu

ABSTRACT

The widespread acceptance of differential privacy has led to the publication of many sophisticated algorithms for protecting privacy. However, due to the subtle nature of this privacy definition, many such algorithms have bugs that make them violate their claimed privacy. In this paper, we consider the problem of producing counterexamples for such incorrect algorithms. The counterexamples are designed to be short and human-understandable so that the counterexample generator can be used in the development process – a developer could quickly explore variations of an algorithm and investigate where they break down. Our approach is statistical in nature. It runs a candidate algorithm many times and uses statistical tests to try to detect violations of differential privacy. An evaluation on a variety of incorrect published algorithms validates the usefulness of our approach: it correctly rejects incorrect algorithms and provides counterexamples for them within a few seconds.

CCS CONCEPTS

• Security and privacy → Privacy protections;

KEYWORDS

Differential privacy; counterexample detection; statistical testing

ACM Reference Format:

Zeyu Ding, Yuxin Wang, Guanhong Wang, Danfeng Zhang, and Daniel Kifer. 2018. Detecting Violations of Differential Privacy. In *2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18)*, October 15–19, 2018, Toronto, ON, Canada. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3243734.3243818>

1 INTRODUCTION

Differential privacy has become a de facto standard for extracting information from a dataset (e.g., answering queries, building machine learning models, etc.) while protecting the confidentiality of individuals whose data are collected. Implemented correctly, it guarantees that any individual's record has very little influence on the output of the algorithm.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '18, October 15–19, 2018, Toronto, ON, Canada

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5693-0/18/10...\$15.00

<https://doi.org/10.1145/3243734.3243818>

However, the design of differentially private algorithms is very subtle and error-prone – it is well-known that a large number of published algorithms are incorrect (i.e. they violate differential privacy). A sign of this problem is the existence of papers that are solely designed to point out errors in other papers [12, 28]. The problem is not limited to novices who may not understand the subtleties of differential privacy; it even affects experts whose goal is to design sophisticated algorithms for accurately releasing statistics about data while preserving privacy.

There are two main approaches to tackling this prevalence of bugs: programming platforms and verification. Programming platforms, such as PINQ [29], Airavat [37], and GUPT [32] provide a small set of primitive operations that can be used as building blocks of algorithms for differential privacy. They make it easy to create correct differentially private algorithms at the cost of accuracy (the resulting privacy-preserving query answers and models can become less accurate). Verification techniques, on the other hand, allow programmers to implement a wider variety of algorithms and verify proofs of correctness (written by the developers) [2–7] or synthesize most (or all) of the proofs [1, 24, 35, 41].

In this paper, we take a different approach: finding bugs that cause algorithms to violate differential privacy, and generating counterexamples that illustrate these violations. We envision that such a counterexample generator would be useful in the development cycle – variations of an algorithm can be quickly evaluated and buggy versions could be discarded (without wasting the developer's time in a manual search for counterexamples or a doomed search for a correctness proof). Furthermore, counterexamples can help developers understand why their algorithms fail to satisfy differential privacy and thus can help them fix the problems. This feature is absent in all existing programming platforms and verification tools. To the best of our knowledge, this is the first paper that treats the problem of detecting counterexamples in incorrect implementations of differential privacy.

Although recent work on relational symbolic execution [22] aims for simpler versions of this task (like detecting incorrect calculations of sensitivity), it is not yet powerful enough to reason about probabilistic computations. Hence, it cannot detect counterexamples in sophisticated algorithms like the *sparse vector technique* [19], which satisfies differential privacy but is notorious for having many incorrect published variations [12, 28].

Our counterexample generator is designed to function in black-box mode as much as possible. That is, it executes code with a variety of inputs and analyzes the (distribution of) outputs of the code. This allows developers to use their preferred languages and libraries as much as possible; in contrast, most language-based

tools will restrict developers to specific programming languages and a very small set of libraries. In some instances, the code may include some tuning parameters. In those cases, we can use an optional symbolic execution model (our current implementation analyzes python code) to find values of those parameters that make it easier to detect counterexamples. Thus, we refer to our method as a *semi-black-box* approach.

Our contributions are as follows:

- We present the first counterexample generator for differential privacy. It treats programs as semi-black-boxes and uses statistical tests to detect violations of differential privacy.
- We evaluate our counterexample generator on a variety of sophisticated differentially private algorithms and their common incorrect variations. These include the sparse vector method and noisy max [19], which are cited as the most challenging algorithms that have been formally verified so far [1, 5]. In particular, the sparse vector technique is notorious for having many incorrect published variations [12, 28]. We also evaluate the counterexample generator on some simpler algorithms such as the histogram algorithm [14], which are also easy for novices to get wrong (by accidentally using too little noise). In all cases, our counterexample generator produces counterexamples for incorrect versions of the algorithms, thus showing its usefulness to both experts and novices.
- The false positive error (i.e. generating "counterexamples" for correct code) of our algorithm is controllable because it is based on statistical testing. The false positive rate can be made arbitrarily small just by giving the algorithm more time to run.

Limitations: it is impossible to create counterexample/bug detector that works for all programs. For this reason, our counterexample generator is not intended to be used in an adversarial setting (where a rogue developer wants to add an algorithm that appears to satisfy differential privacy but has a back door). In particular, if a program satisfies differential privacy except with an extremely small probability (a setting known as *approximate differential privacy* [16]) then our counterexample generator may not detect it. Solving this issue is an area for future work.

The rest of the paper is organized as follows. Related work is discussed in Section 2. Background on differential privacy and statistical testing is discussed in Section 3. The counterexample generator is presented in Section 4. Experiments are presented in Section 5. Conclusions and future work are discussed in Section 6.

2 RELATED WORK

Differential privacy. The term differential privacy covers a family of privacy definitions that include pure ϵ -differential privacy (the topic of this paper) [17] and its relaxations: approximate (ϵ, δ) -differential privacy [16], concentrated differential privacy [8, 20], and Renyi differential privacy [31]. The pure and approximate versions have received the most attention from algorithm designers (e.g., see the book [19]). However, due to the lack of availability of easy-to-use debugging and verification tools, a considerable fraction of published algorithms are incorrect. In this paper, we focus on algorithms for which there is a public record of an error (e.g., variants of the sparse vector method [12, 28]) or where a seemingly small change to an algorithm breaks an important component of

the algorithm (e.g., variants of the noisy max algorithm [5, 19] and the histogram algorithm [14]).

Programming platforms and verification tools. Several dynamic tools [21, 29, 37, 39, 40] exist for enforcing differential privacy. Those tools track the privacy budget consumption at runtime, and terminates a program when the intended privacy budget is exhausted. On the other hand, static methods exist for verifying that a program obeys differential privacy during any execution, based on relational program logic [1–7] and relational type system [24, 35, 41]. We note that those methods are largely orthogonal to this paper: their goal is to verify a correct program or to terminate an incorrect one, while our goal is to detect an incorrect program and generate counterexamples for it. The counterexamples provide valuable guidance for fixing incorrect algorithms for algorithm designers. Moreover, we believe our tool fills in the currently missing piece in the development of differentially private algorithms: with our tool, immature designs can first be tested for counterexamples, before being fed into those dynamic and static tools.

Counterexample generation. Symbolic execution [9, 10, 26] is widely used for program testing and bug finding. One attractive feature of symbolic execution is that when a property is being violated, it generates counterexamples (i.e., program inputs) that lead to violations. More relevant to this paper is work on testing relational properties based on symbolic execution [22, 30, 33]. However, those work only apply to deterministic programs, but the differential privacy property inherently involves probabilistic programs, which is beyond the scope of those work.

3 BACKGROUND

In this section, we discuss relevant background on differential privacy and hypothesis testing.

3.1 Differential Privacy

We view a database as a finite multiset of records from some domain. It is sometimes convenient to represent a database by a histogram, where each cell is the count of times a specific record is present.

Differential privacy relies on the notion of *adjacent* databases. The two most common definitions of adjacency are: (1) two databases D_1 and D_2 are *adjacent* if D_2 can be obtained from D_1 by *adding or removing* a single record. (2) two databases D_1 and D_2 are *adjacent* if D_2 can be obtained from D_1 by *modifying* one record. The notion of adjacency used by an algorithm must be provided to the counterexample generator. We write $D_1 \sim D_2$ to mean that D_1 is adjacent to D_2 (under whichever definition of adjacency is relevant in the context of a given algorithm).

We use the term *mechanism* to refer to an algorithm M that tries to protect the privacy of its input. In our case, a mechanism is an algorithm that is intended to satisfy ϵ -differential privacy:

Definition 3.1 (Differential Privacy [17]). Let $\epsilon \geq 0$. A mechanism M is said to be ϵ -differentially private if for every pair of adjacent databases D_1 and D_2 , and every $E \subseteq \text{Range}(M)$, we have

$$P(M(D_1) \in E) \leq e^\epsilon \cdot P(M(D_2) \in E).$$

The value of ϵ , called the privacy budget, controls the level of the privacy: the smaller ϵ is, the more privacy is guaranteed.

One of the most common building blocks of differentially private algorithms is the *Laplace mechanism* [17], which is used to answer numerical queries. Let \mathcal{D} be the set of possible databases. A numerical query is a function $q : \mathcal{D} \rightarrow \mathbb{R}^k$ (i.e. it outputs a k -dimensional vector of numbers). The Laplace mechanism is based on a concept called *global sensitivity*, which measures the worst-case effect one record can have on a numerical query:

Definition 3.2 (Global Sensitivity [17]). The ℓ_1 -global sensitivity of a numerical query q is

$$\Delta_q = \max_{D_1 \sim D_2} \|q(D_1) - q(D_2)\|_1.$$

The Laplace mechanism works by adding Laplace noise (having density $f(x | b) = \frac{1}{2b} \exp(-|x|/b)$ and variance $2b^2$) to query answers. The chosen variance depends on ϵ and the global sensitivity. We use the notation $\text{Lap}(b)$ to refer to the Laplace noise.

Definition 3.3 (The Laplace mechanism [17]). For any numerical query $q : \mathcal{D} \rightarrow \mathbb{R}^n$, the Laplace mechanism outputs

$$M(D, q, \epsilon) = q(D) + (\eta_1, \dots, \eta_n)$$

where η_i are independent random variables sampled from $\text{Lap}(\Delta_q/\epsilon)$.

THEOREM 3.4 ([19]). *The Laplace mechanism is ϵ -differentially private.*

3.2 Hypothesis Testing

A statistical *hypothesis* is a claim about the parameters of the distribution that generated the data. The *null hypothesis*, denoted by H_0 is a statistical hypothesis that we are trying to disprove. For example, if we have two samples, X and Y where X was generated by a $\text{Binomial}(n, p_1)$ distribution and Y was generated by a $\text{Binomial}(n, p_2)$ distribution, one null hypothesis could be $p_1 = p_2$ (that is, we would like to know if the data supports the conclusion that X and Y came from different distributions). The *alternative hypothesis*, denoted by H_1 , is the complement of the null hypothesis (e.g., $p_1 \neq p_2$).

A hypothesis test is a procedure that takes in a data sample Z and either rejects the null hypothesis or fails to reject the null hypothesis. A hypothesis test can have two types of errors: type I and type II. A type I error occurs if the test incorrectly rejects H_0 when it is in fact true. A type II error occurs if the test fails to reject H_0 when the alternative hypothesis is true. Type I and type II errors are analogous to false positives and false negatives, respectively.

In most problems, controlling type I error is the most important. In such cases, one specifies a *significance level* α and requires that the probability of a type I error be at most α . Commonly used values for α are 0.05 and 0.01. In order to allow users to control the type I error, the hypothesis test also returns a number p – known as the p -value – which is a probabilistic estimate of how unlikely it is that the null hypothesis is true. The user rejects the null hypothesis if $p \leq \alpha$. In order for this to work (i.e. in order for the Type I error to be below α), the p -value must satisfy certain technical conditions: (1) a p -value is a function of a data sample Z , (2) $0 \leq p(Z) \leq 1$, (3) if the null hypothesis is true, then $P(p(Z) \leq \alpha | H_0) \leq \alpha$.

A relevant example of a hypothesis test is Fisher's exact test [23] for two binomial populations. Let c_1 be a sample from a $\text{Binomial}(n_1, p_1)$

distribution and let c_2 be a sample from a $\text{Binomial}(n_2, p_2)$ distribution. Here p_1 and p_2 are unknown. Using these values of c_1 and c_2 , the goal is to test the null hypothesis $H_0 : p_1 \leq p_2$ against the alternative $H_1 : p_1 > p_2$. Let $s = c_1 + c_2$. The key insight behind Fisher's test is that if $C_1 \sim \text{Binomial}(n_1, p_1)$ ¹ and $C_2 \sim \text{Binomial}(n_2, p_2)$ and if $p_1 = p_2$, then the value $P(C_1 > c_1 | C_1 + C_2 = s)$ does not depend on the unknown parameters p_1 or p_2 and can be computed from the cumulative distribution function of the hypergeometric distribution; specifically, it is equal to $1 - \text{Hypergeometric.cdf}(c_1 | n_1 + n_2, n_1, s)$. When $p_1 > p_2$, then $P(C_1 > c_1 | C_1 + C_2 = s)$ cannot be computed without knowing p_1 and p_2 . However, it is less than $1 - \text{Hypergeometric.cdf}(c_1 | n_1 + n_2, n_1, s)$. Thus it can be shown that $1 - \text{Hypergeometric.cdf}(c_1 | n_1 + n_2, n_1, s)$ is a valid p -value and so the Fisher's exact test rejects the null hypothesis when this quantity is $\leq \alpha$.

4 COUNTEREXAMPLE DETECTION

For a mechanism M that does not satisfy ϵ -differential privacy, the goal is to prove this failure. By Definition 3.1, this involves finding a pair of adjacent databases D_1, D_2 and an output event E such that $P(M(D_1) \in E) > e^\epsilon P(M(D_2) \in E)$. Thus a counterexample involves finding these two adjacent inputs D_1 and D_2 , the bad output set E , and to show that for these choices, $P(M(D_1) \in E) > e^\epsilon P(M(D_2) \in E)$.

Ideally, one would compute the probabilities $P(M(D_1) \in E)$ and $P(M(D_2) \in E)$. Unfortunately, for sophisticated mechanisms, it is not always possible to compute these quantities exactly. However, we can sample from these distributions many times by repeatedly running $M(D_1)$ and $M(D_2)$ and counting the number of times that the outputs fall into E . Then, we need a statistical test to reject the null hypothesis $P(M(D_1) \in E) \leq e^\epsilon P(M(D_2) \in E)$ (or fail to reject it if the algorithm is ϵ -differentially private).

We will be using the following conventions:

- The input to most mechanisms is actually a list of queries $Q = (q_1, \dots, q_l)$ rather than a database directly. For example, algorithms to release differentially private histograms operate on a histogram of the data; the sparse vector mechanism operates on a sequence of queries that each have global sensitivity equal to 1. Thus, we require the user to specify how the input query answers can differ on two adjacent databases. For example, in a histogram, exactly one cell count changes by at most 1. In the sparse vector technique [19], every query answer changes by at most 1. To simplify the discussion, we abuse notation and use D_1, D_2 to also denote the answers of Q on the input adjacent databases. For example, when discussing the sparse vector technique, we write $D_1 = [0, 0]$ and $D_2 = [1, 1]$. This means there are adjacent databases and a list of queries $Q = [q_1, q_2]$ such that they evaluate to $[0, 0]$ on the first database and $[1, 1]$ on the second database.
- We use ϵ_0 to indicate the privacy level that a mechanism claims to achieve.
- We use Ω for the set of all possible outputs (i.e., range) of the mechanism M . We use ω for a single output of M .

¹This is read as " C_1 is a random variable having the $\text{Binomial}(n_1, p_1)$ distribution".

- We call a subset $E \subseteq \Omega$ an event. We use p_1 (respectively, p_2) to denote $P(M(D_1) \in E)$, the probability that the output of M falls into E when executing on database D_1 (respectively, D_2).
- Some mechanisms take additional inputs, e.g., the sparse vector mechanism. We collectively refer to them as *args*.

Our discussion is organized as follows. We provide an overview of the counterexample generator in Section 4.1. Then we incrementally explain our approach. In Section 4.2 we present the hypothesis test. That is, suppose we already have query sequences D_1 and D_2 that are generated from adjacent databases and an output set E , how do we test if $P(M(D_1) \in E) \leq e^\epsilon P(M(D_2) \in E)$ or $P(M(D_1) \in E) > e^\epsilon P(M(D_2) \in E)$? Next, in Section 4.3, we consider the question of output selection. That is, suppose we already have query answers D_1 and D_2 that are generated from adjacent databases, how do we decide which E should be used in the hypothesis test? Finally, in Section 4.4, we consider the problem of generating the adjacent query sequences D_1 and D_2 as well as additional inputs *args*.

The details of specific mechanisms we test for violations of differential privacy will be given in the experiments in Section 5.

4.1 Overview

At a high level, the counterexample generator can be summarized in the pseudocode in Algorithm 1. First, it generates an `InputList`, a set of candidate tuples of the form (D_1, D_2, args) . That is, instead of returning a single pair of adjacent inputs D_1, D_2 and any auxiliary arguments the mechanism may need, we return multiple candidates which will be filtered later. Each adjacent pair (D_1, D_2) is designed to be short so that a developer can understand the problematic inputs and trace them through the code of the mechanism M . For this reason, the code of M will also run fast, so that it will be possible to later evaluate $M(D_1, \text{args})$ and $M(D_2, \text{args})$ multiple times very quickly.

Algorithm 1: Overview of Counterexample Generator

```

1 function CounterExampleDetection( $M, \epsilon$ ):
  input:  $M$ : mechanism
            $\epsilon$ : desired privacy (is  $M$   $\epsilon$ -differentially private?)
           // get set of possible inputs:  $(D_1, D_2, \text{args})$ 
2  InputList  $\leftarrow$  InputGenerator( $M, \epsilon$ )
3   $E, D_1, D_2, \text{args} \leftarrow$  EventSelector( $M, \epsilon, \text{InputList}$ )
4   $p_\top, p_\perp \leftarrow$  HypothesisTest( $M, \epsilon, D_1, D_2, \text{args}, E$ )
5  return  $p_\top, p_\perp$ 

```

The next step is the `EventSelector`. It takes each tuple (D_1, D_2, args) from `InputList` and runs $M(D_1, \text{args})$ and $M(D_2, \text{args})$ multiple times. Based on the type of the outputs, it generates a set of candidates for E . For example, if the output is a real number, then the set of candidates is the set of intervals (a, b) . For each candidate E and each tuple (D_1, D_2, args) , it counts how many times $M(D_1, \text{args})$ produced an output $\omega \in E$ and how many times $M(D_2, \text{args})$ produced an output in E . Based on these results, it picks one specific E and one tuple (D_1, D_2, args) which it believes is most likely to show a violation of ϵ_0 -differential privacy.

Finally, the `HypothesisTest` takes the selected E, D_1, D_2 , and *args* and checks if it can detect statistical evidence that $P(M(D_1, \text{args}) \in E) > e^\epsilon P(M(D_2, \text{args}) \in E)$ – which corresponds to the p -value p_\top – or $e^\epsilon P(M(D_1, \text{args}) \in E) < P(M(D_2, \text{args}) \in E)$ – which corresponds to the p -value p_\perp .

It is important to note that the `EventSelector` also uses `HypothesisTest` internally as a sub-routine to filter out candidates. That is, for every candidate E and every candidate (D_1, D_2, args) , it runs `HypothesisTest` and treats the returned value as a score. The combination of E and (D_1, D_2, args) with the best score is returned by the `EventSelector`. Note that `EventSelector` is using the `HypothesisTest` in an exploratory way – it evaluates many hypotheses and returns the best one it finds. This is why the E and (D_1, D_2, args) that are finally chosen need to be evaluated again on Line 4 using fresh samples from M .

Interpreting the results. One of the best ways of understanding the behavior of the counterexample generator is to look at the p -values it outputs. That is, we take a mechanism M that claims to satisfy ϵ_0 -differential privacy and, for each ϵ close to ϵ_0 , we test whether it satisfies ϵ -differential privacy (that is, even though M claims to satisfy ϵ_0 -differential privacy, we may want to test if it satisfies ϵ -differential privacy for some other value of ϵ that is close to ϵ_0). The hypothesis tester returns two p -values:

- p_\top . Small values indicate that probably $P(M(D_1, \text{args}) \in E) > e^\epsilon P(M(D_2, \text{args}) \in E)$.
- p_\perp . Small values indicate that probably $P(M(D_2, \text{args}) \in E) > e^\epsilon P(M(D_1, \text{args}) \in E)$.

For each ϵ , we plot the minimum of p_\top and p_\perp . Figure 1 shows typical results that would appear when the counterexample detector is run with real mechanisms M as input.

In Figure 1a, M correctly satisfies the claimed $\epsilon_0 = 0.7$ differential privacy. In that plot, we see that the p -values corresponding to $\epsilon = 0.2, 0.4, 0.6$ are very low, meaning that the counterexample generator can prove that the algorithm does not satisfy differential privacy for those smaller values of ϵ . Near 0.7 it becomes difficult to find counterexamples; that is, if an algorithm satisfies 0.7-differential privacy, it is very hard to statistically prove that it does not satisfy 0.65 differential privacy. This is a typical feature of hypothesis tests as it becomes difficult to reject the null hypothesis when it is only slightly incorrect (e.g., when the true privacy parameter is only slightly different from the ϵ we are testing). Now, any algorithm that satisfies 0.7-differential privacy also satisfies ϵ -differential privacy for all $\epsilon \geq 0.7$. This behavior is seen in Figure 1a as the p -values are large for all larger values of ϵ .

Figure 1b shows a graph that can arise from two distinct scenarios. One of the situations is when the mechanism M claims to provide 0.7-differential privacy but actually provides more privacy (i.e. ϵ -differential privacy for $\epsilon < 0.7$). In this figure, the counterexample generator could prove, for example, that M does not satisfy 0.4-differential privacy, but leaves open the possibility that it satisfies 0.5-differential privacy. The other situation is when our tool has failed to find good counterexamples. Thus when a mechanism is correct, good precision by the counterexample generator means that the line starts rising close to (but before the dotted line), and worse precision means that the line starts rising much earlier.




Figure 1: Interpreting experimental results on hypothesis tests. A hypothetical algorithm M claims to achieve ϵ_0 -differential privacy. For each ϵ from 0.2 to 2.4 we would evaluate if M satisfies ϵ -differential privacy. We show a typical graph when M does satisfy ϵ_0 -differential privacy (left), a graph where M possibly provides more privacy (center) and a graph where M provides less privacy than advertised.

Figure 1c shows a typical situation in which an algorithm claims to satisfy 0.7-differential privacy but actually provides less privacy than advertised. In this case, the counterexample generator can generate good counterexamples at $\epsilon = 0.7$ (the dotted line) and even at much higher values of ϵ . When a mechanism is incorrect, such a graph indicates good precision by the counterexample generator.

Limitations. In some cases, finding counterexamples requires a large input datasets. In those cases, searching for the right inputs and running algorithms on them many times will impact the ability of our counterexample generator to find counterexamples. This is a limitation of all techniques based on statistical tests.

Another important case where our counterexample generator is not expected to perform well is when violations of differential privacy happen very rarely. For example, consider a mechanism M that checks if its input is $D_1 = [1]$. If so, with probability e^{-9} it outputs 1 and otherwise it outputs 0 (if the input is not 1, M always outputs 0). M does not satisfy ϵ -differential privacy for any value of ϵ . However, showing it statistically is very difficult. Supposing $D_1 = [1]$ and $D_2 = [0]$ are adjacent databases, it requires running $M(D_1)$ and $M(D_2)$ billions of times to observe that an output of 1 is possible under D_1 but is at least e^ϵ times less likely under D_2 .

Addressing both of these problems will likely involve incorporation of program analysis, such as symbolic execution, into our statistical framework and is a direction for future work.

4.2 Hypothesis Testing

Suppose we have a mechanism M , inputs D_1, D_2 and an output set E (we discuss the generation of D_1, D_2 in Section 4.4 and E in Section 4.3). We would like to check if $P(M(D_1) \in E) > e^\epsilon P(M(D_2) \in E)$ or if $P(M(D_2) \in E) > e^\epsilon P(M(D_1) \in E)$, as that would demonstrate a violation of ϵ -differential privacy. We treat the $P(M(D_1) \in E) > e^\epsilon P(M(D_2) \in E)$ case in this section, as the other case is symmetric.

To do this, the high level idea is to:

- Define $p_1 = P(M(D_1) \in E)$ and $p_2 = P(M(D_2) \in E)$

Algorithm 2: Hypothesis Test. Parameter n : # of iterations

```

1 function pvalue( $c_1, c_2, n, \epsilon$ ):
2    $\tilde{c}_1 \leftarrow \text{Binomial}(c_1, 1/e^\epsilon)$ 
3    $s \leftarrow \tilde{c}_1 + c_2$ 
4   return  $1 - \text{Hypergeom.cdf}(\tilde{c}_1 | 2n, n, s)$ 
5 function HypothesisTest( $M, args, \epsilon, D_1, D_2, E$ ):
   input:  $M$ : mechanism
          $args$ : additional arguments for  $M$ 
          $\epsilon$ : privacy budget to test
          $D_1, D_2$ : adjacent databases
          $E$ : Event
6    $O_1 \leftarrow$  results of running  $M(D_1, args)$  for  $n$  times
7    $O_2 \leftarrow$  results of running  $M(D_2, args)$  for  $n$  times
8    $c_1 \leftarrow |\{i \mid O_1[i] \in E\}|$ 
9    $c_2 \leftarrow |\{i \mid O_2[i] \in E\}|$ 
10   $p_\top \leftarrow \text{pvalue}(c_1, c_2, n, \epsilon)$ 
11   $p_\perp \leftarrow \text{pvalue}(c_2, c_1, n, \epsilon)$ 
12  return  $p_\top, p_\perp$ 

```

- Formulate the null hypothesis as $H_0 : p_1 \leq e^\epsilon \cdot p_2$ and the alternative as $H_1 : p_1 > e^\epsilon \cdot p_2$.
- Run M with inputs D_1 and D_2 independently n times each. Record the results as O_1 and O_2 .
- Count the number of times the result falls in E in each case. Let $c_1 = |\{i \mid O_1[i] \in E\}|$ and $c_2 = |\{i \mid O_2[i] \in E\}|$. Intuitively, $c_1 \gg e^\epsilon c_2$ provides strong evidence against the null hypothesis.
- Calculate a p -value based on c_1, c_2 to determine how unlikely the null hypothesis is.

The challenge is, of course, in the last step as we don't know what p_1 and p_2 are. One direction is to estimate them from c_1 and c_2 . However, it is also challenging to estimate the variance of our estimates \hat{p}_1 and \hat{p}_2 (the higher the variance, the less the test should trust the estimates).

Instead, we take a different approach that allows us to conduct the test without knowing what p_1 and p_2 are. First, we note that c_1 and c_2 are equivalent to samples from a Binomial(n, p_1) distribution and a Binomial(n, p_2) distribution respectively. We first consider the border case where $p_1 = e^\epsilon p_2$. Consider sample \tilde{c}_1 from a Binomial($c_1, 1/e^\epsilon$) distribution. We note that this sample enjoys the following property (which implies that in the border case, \tilde{c}_1 has the same distribution as c_2):

LEMMA 4.1. *Let $X \sim \text{Binomial}(n, p_1)$ and Z be generated from X by sampling from the Binomial($X, 1/e^\epsilon$) distribution. The marginal distribution of Z is Binomial($n, p_1/e^\epsilon$).*

PROOF. The relationship between Binomial and Bernoulli random variables means that $X = \sum_{i=1}^n X_i$, where X_i is a Bernoulli(p_1) random variable. Generating Z from X is the same as doing the following: set $Z_i = 0$ if $X_i = 0$. If $X_i = 1$, set $Z_i = 1$ with probability $1/e^\epsilon$ (and set $Z_i = 0$ otherwise). Then set $Z = \sum_{i=1}^n Z_i$. Hence, the marginal distribution of Z_i is a Bernoulli(p_1/e^ϵ) random variable:

$$\begin{aligned} P(Z_i = 1) &= P(Z_i = 1 \mid X_i = 1)P(X_i = 1) + P(Z_i = 1 \mid X_i = 0)P(X_i = 0) \\ &= (1/e^\epsilon) \cdot p_1 + 0 \cdot (1 - p_1) = p_1/e^\epsilon \end{aligned}$$

This means that the marginal distribution of Z is Binomial($n, p_1/e^\epsilon$). \square

Thus we have the following facts that follow immediately from the lemma:

- If $p_1 > e^\epsilon p_2$ then the distribution of \tilde{c}_1 is Binomial(n, \tilde{p}_1) with $\tilde{p}_1 = p_1/e^\epsilon$ and so has a larger Binomial parameter than c_2 (which is Binomial(n, p_2)). We want our test to be able to reject the null hypothesis in this case.
- If $p_1 = e^\epsilon p_2$ then the distribution of \tilde{c}_1 is Binomial(n, \tilde{p}_1) with $\tilde{p}_1 = p_2$ and so has the same Binomial parameter as c_2 . We do not want our test to reject the null hypothesis in this case.
- If $p_1 < e^\epsilon p_2$ then the distribution of \tilde{c}_1 is Binomial(n, \tilde{p}_1) with $\tilde{p}_1 = p_1/e^\epsilon$ and so has a smaller Binomial parameter than c_2 (which is Binomial(n, p_2)). We do not want to reject the null hypothesis in this case.

Thus, by randomly generating \tilde{c}_1 from c_1 , we have (randomly) reduced the problem of testing $p_1 > e^\epsilon p_2$ vs. $p_1 \leq e^\epsilon p_2$ (on the basis of c_1 and c_2) to the problem of testing $\tilde{p}_1 > p_2$ vs. $\tilde{p}_1 \leq p_2$ (on the basis of \tilde{c}_1 and c_2). Now, checking whether \tilde{c}_1 and c_2 come from the same distribution can be done with the Fisher's exact test (see Section 3): the p -value is $1 - \text{Hypergeom.cdf}(\tilde{c}_1 \mid 2n, n, \tilde{c}_1 + c_2)$.² This is done in the function `pvalue` in Algorithm 2.

To summarize, given c_1 and c_2 , we first sample \tilde{c}_1 from the Binomial($c_1, 1/e^\epsilon$) distribution and then return the p -value of $(1 - \text{Hypergeom.cdf}(\tilde{c}_1 \mid 2n, n, \tilde{c}_1 + c_2))$. Since this is a random reduction, we reduce its variance by sampling \tilde{c}_1 multiple times and averaging the p -values. That is, we run the p -value function (Algorithm 2) multiple times with the same inputs and average the p -values it returns.

²Here we use a notation from SciPy [25] package where `Hypergeom.cdf` means the cumulative distribution function of hypergeometric distribution.

4.3 Event Selection

Having discussed how to test if $P(M(D_1) \in E) > e^\epsilon P(M(D_2) \in E)$ or if $P(M(D_2) \in E) > e^\epsilon P(M(D_1) \in E)$ when D_1, D_2 , and E were pre-specified, we now discuss how to select the event E that is most likely to show violations of ϵ -differential privacy.

Algorithm 3: Event Selector. Parameter n : # of iterations

```

1 function EventSelector( $n, M, \epsilon, InputList$ ):
   input:  $M$ : mechanism
          $InputList$ : possible inputs
          $\epsilon$ : privacy budget to test
2   SearchSpace  $\leftarrow$  search space based on return type
3   pvalues  $\leftarrow$  [ ]
4   results  $\leftarrow$  [ ]
5   foreach ( $D_1, D_2, args$ )  $\in$   $InputList$  do
6      $O_1 \leftarrow$  results of running  $M(D_1, args)$  for  $n$  times
7      $O_2 \leftarrow$  results of running  $M(D_2, args)$  for  $n$  times
8     foreach  $E \in$  SearchSpace do
9        $c_1 \leftarrow$   $|\{i \mid O_1[i] \in E\}|$ 
10       $c_2 \leftarrow$   $|\{i \mid O_2[i] \in E\}|$ 
11       $p_\top \leftarrow$  pvalue( $c_1, c_2, n, \epsilon$ )
12       $p_\perp \leftarrow$  pvalue( $c_2, c_1, n, \epsilon$ )
13      pvalues.append( $\min(p_\top, p_\perp)$ )
14      results.append( $D_1, D_2, args, E$ )
15   end
16 end
17 return results[ $\text{argmin}(pvalues)$ ]

```

One of the challenges is that different mechanisms could have different output types (e.g., a discrete number, a vector of numbers, a vector of categorical values, etc.). To address this problem, we define a search space S of possible events to look at. The search space depends on the type of the output ω of M , which can be determined by running $M(D_1)$ and $M(D_2)$ multiple times.

(1) **The output ω is a fixed length list of categorical values.**

We first run $M(D_1)$ once and ask it to not use any noise (i.e. tell it to satisfy ϵ -differential privacy with $\epsilon = \infty$). Denote this output as ω_0 . Now, when M runs with its preferred privacy settings to produce an output ω , we define $t(\omega)$ be the Hamming distance between the output ω and ω_0 . The search space is

$$S = \{\{\omega \mid t(\omega) = k\} : k = 0, 1, \dots, l\}$$

where l is the fixed length of output of M . Another set of events relate to the count of a categorical value in the output. If there are m values, then define

$$S_i = \{\{\omega \mid \omega.count(value_i) = k\} : k = 0, 1, \dots, l\},$$

$1 \leq i \leq m$. The overall search space is the union of S and all S_i .

(2) **The output ω is a variable length list of categorical values.** In this case, one extra set of events E we look at correspond to the length of the output. For example, we may check if $P(M(D_1)$ has length k) $>$ $P(M(D_2)$ has length k). Hence, we define

$$S_0 = \{\{\omega \mid \omega.length = k\} : k = 0, 1, \dots\}$$

For the search space S , we use this S_0 unioned with the search space from the previous case.

(3) **The output ω is a fixed length list of numeric values.**

In this case, the output is of the form $\omega = (a_1, \dots, a_m)$. Our search space is the union of the following:

$$\begin{aligned} & \{\{\omega \mid \omega[i] \in (a, b)\} : i = 1, \dots, m \text{ and } a < b\}, \\ & \{\{\omega \mid \text{avg}(\omega) \in (a, b)\} : a < b\}, \\ & \{\{\omega \mid \min(\omega) \in (a, b)\} : a < b\}, \\ & \{\{\omega \mid \max(\omega) \in (a, b)\} : a < b\}. \end{aligned}$$

That is, we would end up checking if $P(\text{avg}(M(D_1)) \in (a, b)) > P(\text{avg}(M(D_2)) \in (a, b))$, etc.. To save time, we often restrict a and b to be multiples of a small number like ± 0.2 , or $\pm \infty$. In the case that the output ω is always an integer array, we replace the condition “ $\in (a, b)$ ” with “ $= k$ ” for each integer k .

(4) **M outputs a variable length list of numeric values.**

The search space is the union of Case 3 and S_0 in Case 2.

(5) **M outputs a variable length list of mixed categorical and numeric values.** In this case, we separate out the categorical values from numeric values and use the cross product of the search spaces for numeric and categorical values. For instance, events would be of the form “ ω has k categorical components equal to ℓ and the average of the numerical components of ω is in (a, b) ”

The EventSelector is designed to return one event E for use in the hypothesis test in Algorithm 1. The way EventSelector works is it receives an InputList, which is a set of tuples $(D_1, D_2, args)$ where D_1, D_2 are adjacent databases and $args$ is a set of values for any other parameters M needs. For each such tuple, it runs $M(D_1)$ and $M(D_2)$ for n times each. Then for each possible event in the search space, it runs the hypothesis test (as an exploratory tool) to get a p-value. The combination of $(D_1, D_2, args)$ and E that produces the lowest p-value is then returned to Algorithm 1. Algorithm 1 uses those choices to run the real hypothesis test on fresh executions of M on D_1 and D_2 .

The pseudocode for the EventSelector is shown in Algorithm 3.³

4.4 Input Generation

In this section we discuss our approaches for generating candidate tuples $(D_1, D_2, args)$ where D_1, D_2 are adjacent databases and $args$ is a set of auxiliary parameters that a mechanism M may need.

4.4.1 Database Generation. To find the adjacent databases that are likely to form the basis of counterexamples that illustrate violations of differential privacy, we adopt a simple and generic approach that works surprisingly well. Recalling that the inputs to mechanisms are best modeled as a vector of query answers, we use the type of patterns shown in Table 1.

The “One Above” and “One Below” categories are suitable for algorithms whose input is a histogram (i.e. in adjacent databases, at most one query can change, and it will change by at most 1). The rest of the categories are suitable when in adjacent databases every

³In practice, to avoid choosing bad E , we let c_E be the total number of times $M(D_1)$ and/or $M(D_2)$ produced an output in E . Then it only executes Line 11-14 in Algorithm 3 if $c_E \geq 0.001 \cdot n \cdot e^\epsilon$, otherwise the selection of E is too noisy.

Algorithm 4: Input Generator.

```

1 function ArgumentGenerator( $M, D_1, D_2$ ):
2    $args_0 \leftarrow$  Arguments used in noise generation with values
   that minimize the noises
3    $constraints \leftarrow$  Traverse the source code of  $M$  and generate
   constraints to force  $D_1$  and  $D_2$  to diverge on branches
4    $args_1 \leftarrow$  MaxSMT( $constraints$ )
5   return  $args_0 + args_1$ 
6 function InputGenerator( $M, len$ ):
   input:  $M$ : mechanism
    $len$ : length of input to generate
7    $candidates \leftarrow$  Empirical pairs of databases of length  $len$ 
8    $InputList \leftarrow [ ]$ 
9   foreach  $(D_1, D_2) \in candidates$  do
10     $args \leftarrow$  ArgumentGenerator( $M, D_1, D_2$ )
11     $InputList.append(D_1, D_2, args)$ 
12  end
13  return  $InputList$ 

```

Table 1: Database categories and samples

Category	Sample D1	Sample D2
One Above	[1, 1, 1, 1, 1]	[2, 1, 1, 1, 1]
One Below	[1, 1, 1, 1, 1]	[0, 1, 1, 1, 1]
One Above Rest Below	[1, 1, 1, 1, 1]	[2, 0, 0, 0, 0]
One Below Rest Above	[1, 1, 1, 1, 1]	[0, 2, 2, 2, 2]
Half Half	[1, 1, 1, 1, 1]	[0, 0, 0, 2, 2]
All Above & All Below	[1, 1, 1, 1, 1]	[2, 2, 2, 2, 2]
X Shape	[1, 1, 0, 0, 0]	[0, 0, 1, 1, 1]

query can change by at most one (i.e. the queries have sensitivity⁴ $\Delta_q = 1$).

The design of the categories is based on the wide variety of changes in query answers that are possible when evaluated on one database and on an adjacent database. For example, it could be the case that a few of the queries increase (by 1, if their sensitivity is 1, or by Δ_q in the general case) but most of them decrease. A simple representative of this situation is “One Above Rest Below” in which one query increases and the rest decrease. The category “One Below Rest Above” is the reverse.

Another situation is where roughly half of the queries increase and half decrease (when evaluated on a database compared to when evaluated on an adjacent database). This scenario is captured by the “Half Half” category. Another situation is where all of the queries increase. This is captured by the “All Above & All Below” category. Finally, the “X Shape” category captures the setting where the query answers are not all the same and some increase and others decrease when evaluated on one database compared to an adjacent database.

These categories were chosen from our desire to allow counterexamples to be easily understood by mechanism designers (and to make it easier for them to manually trace the code to understand

⁴For queries with larger sensitivity, the extension is obvious. For example $D_1 = [1, 1, 1, 1, 1]$ and $D_2 = [1 + \Delta_q, 1 + \Delta_q, \dots, 1 + \Delta_q]$

the problems). Thus the samples are short and simple. We consider inputs of length 5 (as in Table 1) and also versions of length 10.

4.4.2 Argument Generation. Some differentially-private algorithms require extra parameters beyond the database. For example, the sparse vector technique [19], shown in Algorithm 11, takes as inputs a threshold T and a bound N . It tries to output numerical queries that are larger than T . However, for privacy reasons, it will stop after it returns N noisy queries whose values are greater than T . These two arguments are specific to the algorithm and their proper values depend on the desired privacy level as well as algorithm precision.

To find values of auxiliary parameters (such as N and T in Sparse Vector), we build argument generator based on *Symbolic Execution* [26], which is typically used for bug finding: it generates concrete inputs that violate assertions in a program. In general, a symbolic executor assigns symbolic values, rather than concrete values as normal execution would do, for inputs. As the execution goes, the executor maintains a symbolic program state at each assertion and generates constraints that will violate the assertion. When those constraints are satisfiable, concrete inputs (i.e., a solution of the constraints) are generated.

Compared with standard symbolic execution, a major difference in our argument generation is that we are interested in algorithm arguments that will likely maximize the privacy cost of an algorithm. In other words, there is no obvious assertion to be checked in our argument generation. To proceed, we use two heuristics that likely will cause large privacy cost of an algorithm:

- The first heuristic applies to parameters that affect noise generation. For example in Sparse Vector, the algorithm adds $\text{Lap}(2 \cdot N \cdot \Delta_q / \epsilon_0)$ noise. For such a variable, we use the value that results in small amount of noise (i.e., $N = 1$). Small amount of noise is favorable since it reduces the variance in the hypothesis testing (Section 4.2).
- The second heuristic (for variables that do not affect noise) prefers arguments that make two program executions using two different databases (as described in Section 4.4.1) to take as many diverging branches as possible. The reason is that diverging branches will likely use more privacy budget.

Next, we give a more detailed overview of our customized symbolic executor. The symbolic executor takes a pair of *concrete* databases as inputs (as described in Section 4.4.1) and uses *symbolic* values for other input parameters. Random samples in the program (e.g., a sample from Laplace distribution) are set to value 0 in the symbolic execution. Then, the symbolic executor tracks symbolic program states along program execution in the standard way [26]. For example, the executor will generate a constraint⁵ $x = y + 1$ after an assignment ($x \leftarrow y + 1$), assuming that variable y has a symbolic value y before the assignment. Also, the executor will unroll loops in the source code, which is standard in most symbolic executors.

Unlike standard symbolic executors, the executor conceptually tracks a pair of symbolic program states along program execution (one on concrete database D_1 , and one on concrete database D_2). Moreover, it also generate extra constraints, according to the two heuristics above, in the hope of maximizing the privacy cost of an

⁵For simplicity, we use a simple representation for constraints; Z3 has an internal format and a user can either use Z3's APIs or SMT2 [34] format to represent constraints.

algorithm. In particular, it handles two kinds of statements in the following way:

- **Sampling.** The executor generates two constraints for a sampling statement: a constraint that eliminates randomness in symbolic execution by assigning sample to value 0, and a constraint that ensures a small amount of noise. Consider a statement ($\eta \leftarrow \text{Lap}(e)$). The executor generates two constraints: $\eta = 0$ as well as a constraint that minimizes expression e .
- **Branch.** The executor generates a constraint that makes the two executions diverge on branches. Consider a branch statement (**if** e **then** \dots). Assume that the executor has symbolic values e_1 and e_2 for the value of expression e on databases D_1 and D_2 respectively; it will generate a constraint $(e_1 \wedge \neg e_2) \vee (\neg e_1 \wedge e_2)$ to make the executions diverge. Note that unlike other constraints, a diverging constraint might be unsatisfiable (e.g., if the query answers under D_1 and D_2 are the same). However, our goal is to *maximize* the number of satisfiable diverging constraints, which can be achieved by a MaxSMT solver.

The executor then uses an external MaxSMT solver such as Z3 [13] on all generated constraints to find arguments that maximizes the number of diverged branches.

For example, the correct version of the Sparse Vector algorithm (see the complete algorithm in Algorithm 11) has the parameter T (a threshold). It has a branch that tests whether the noisy query answer is above the threshold T :

$$q + \eta_2 \geq \hat{T}$$

Here, η_2 is a noise variable, q is one query answer (i.e. one of the components of the input D_1 of the algorithm) and \hat{T} is a noisy threshold ($\hat{T} = T + \eta_1$). Suppose we start from a database candidate ($[1, 1, 1, 1, 1], [2, 2, 2, 2, 2]$). The symbolic executor assigns symbolic values to the parameters T and unrolls the loop in the algorithm, where each iteration handles one noisy query. Along the execution, it updates program states. For example, statement $\hat{T} \leftarrow T + \eta_1$ results in $\hat{T} = T + \eta_1$. For the first execution of the branch of interest, the executor tracks the following symbolic program state:

$$q_1 = 1 \wedge q_2 = 2 \wedge \eta_1 = 0 \wedge \eta_2 = 0 \wedge \hat{T}_1 = T + \eta_1 \wedge \hat{T}_2 = T + \eta_2$$

as well as the following constraint for diverging branches:

$$(q_1 + \eta_1 \geq \hat{T}_1 \wedge q_2 + \eta_2 < \hat{T}_2) \vee (q_1 + \eta_1 < \hat{T}_1 \wedge q_2 + \eta_2 \geq \hat{T}_2)$$

Similarly, the executor generates constraints from other iterations. In this example, the MaxSMT solver returns a value in between of 1 and 2 so that constraints from all iterations are satisfied. This value of T is used as arg in the candidate tuple (D_1, D_2, arg) .

5 EXPERIMENTS

We implemented our counterexample detection framework with all components, including hypothesis test, event selector and input generator. The implementation is publicly available⁶. The tool takes in an algorithm implementation and the desired privacy bound ϵ_0 , and generates counterexamples if the algorithm does not satisfy ϵ_0 -differential privacy.

⁶ <https://github.com/cmla-psu/statdp>.

In this section we evaluate our detection framework on some of the popular privacy mechanisms and their variations. We demonstrate the power of our tool: for mechanisms that falsely claim to be differentially private, our tool produces convincing evidence that this is not the case in just a few seconds.

5.1 Noisy Max

Report Noisy Max reports which one among a list of counting queries has the largest value. It adds $Lap(2/\epsilon_0)$ noise to each answer and returns the *index* of the query with the largest noisy answer. The correct versions have been proven to satisfy ϵ_0 -differential privacy [19] no matter how long the input list is. A naive proof would show that it satisfies $(\epsilon_0 \cdot |Q|/2)$ -differential privacy (where $|Q|$ is the length of the input query list), but a clever proof shows that it actually satisfies ϵ_0 -differential privacy.

Algorithm 5: Correct Noisy Max with Laplace noise

```

1 function NoisyMax( $Q, \epsilon_0$ ):
  input:  $Q$ : queries to the database,  $\epsilon_0$ : privacy budget.
2   NoisyVector  $\leftarrow []$ 
3   for  $i = 1 \dots \text{len}(Q)$  do
4     | NoisyVector[ $i$ ]  $\leftarrow Q[i] + Lap(2/\epsilon_0)$ 
5   end
6   return  $\text{argmax}(\text{NoisyVector})$ 

```

5.1.1 Adding Laplace Noise. The correct Noisy Max algorithm (Algorithm 5) adds independent $Lap(2/\epsilon_0)$ noise to each query answer and returns the index of the maximum value. As Figure 2a shows, we test this algorithm for different privacy budget ϵ_0 at 0.2, 0.7, 1.5. All lines rise when the test ϵ is slightly less than the claimed privacy level ϵ_0 of the algorithm. This demonstrates the precision of our tool: before ϵ_0 , there is almost 0 chance to falsely claim that this algorithm is not private; after ϵ_0 , the p -value is too large to conclude that the algorithm is incorrect. We note that the test result is very close to the ideal cases, illustrated by the vertical dashed lines.

Algorithm 6: Correct Noisy Max with Exponential noise

```

1 function NoisyMax( $Q, \epsilon_0$ ):
  input:  $Q$ : queries to the database,  $\epsilon_0$ : privacy budget.
2   NoisyVector  $\leftarrow []$ 
3   for  $i = 1 \dots \text{len}(Q)$  do
4     | NoisyVector[ $i$ ]  $\leftarrow Q[i] + \text{Exponential}(2/\epsilon_0)$ 
5   end
6   return  $\text{argmax}(\text{NoisyVector})$ 

```

5.1.2 Adding Exponential Noise. One correct variant of Noisy Max adds $\text{Exponential}(2/\epsilon_0)$ noise, rather than Laplace noise, to each query answer (Algorithm 6). This mechanism has also been proven to be ϵ_0 -differential private [19]. Figure 2b shows the corresponding test result, which is similar to that of Figure 2a. The result indicates that this correct variant likely satisfies ϵ_0 -differential privacy for the claimed privacy budget.

Algorithm 7: Incorrect Noisy Max with Laplace noise, returning the maximum value

```

1 function NoisyMax( $Q, \epsilon_0$ ):
  input:  $Q$ : queries to the database,  $\epsilon_0$ : privacy budget.
2   NoisyVector  $\leftarrow []$ 
3   for  $i = 1 \dots \text{len}(Q)$  do
4     | NoisyVector[ $i$ ]  $\leftarrow Q[i] + Laplace(2/\epsilon_0)$ 
5   end
6   // returns maximum value instead of index
7   return  $\text{max}(\text{NoisyVector})$ 

```

Algorithm 8: Incorrect Noisy Max with Exponential noise, returning the maximum value

```

1 function NoisyMax( $Q, \epsilon_0$ ):
  input:  $Q$ : queries to the database,  $\epsilon_0$ : privacy budget.
2   NoisyVector  $\leftarrow []$ 
3   for  $i = 1 \dots \text{len}(Q)$  do
4     | NoisyVector[ $i$ ]  $\leftarrow Q[i] + \text{Exponential}(2/\epsilon_0)$ 
5   end
6   // returns maximum value instead of index
7   return  $\text{max}(\text{NoisyVector})$ 

```

Algorithm 9: Histogram

```

1 function Histogram( $Q, \epsilon_0$ ):
  input:  $Q$ : queries to the database,  $\epsilon_0$ : privacy budget.
2   NoisyVector  $\leftarrow []$ 
3   for  $i = 1 \dots \text{len}(Q)$  do
4     | NoisyVector[ $i$ ]  $\leftarrow Q[i] + Lap(1/\epsilon_0)$ 
5   end
6   return NoisyVector

```

Algorithm 10: Histogram with wrong scale

```

1 function Histogram( $Q, \epsilon_0$ ):
  input:  $Q$ : queries to the database,  $\epsilon_0$ : privacy budget.
2   NoisyVector  $\leftarrow []$ 
3   for  $i = 1 \dots \text{len}(Q)$  do
4     | // wrong scale of noise is added
5     |  $\text{NoisyVector}[i] \leftarrow Q[i] + Lap(\epsilon_0)$ 
6   end
7   return NoisyVector

```

5.1.3 Incorrect Variants of Exponential Noise. An incorrect variant of *NoisyMax* has the same setup but instead of returning the *index* of maximum value, it directly returns the maximum value. We evaluate on two variants that report the maximum value instead of the index (Algorithm 7 and 8) and show the test result in Figure 2c and 2d.

For the variant using Laplace noise (Figure 2c), we can see that for $\epsilon_0 = 0.2$, the line rises at around test ϵ of 0.4, indicating that this algorithm is incorrect for the claimed privacy budget of 0.2.

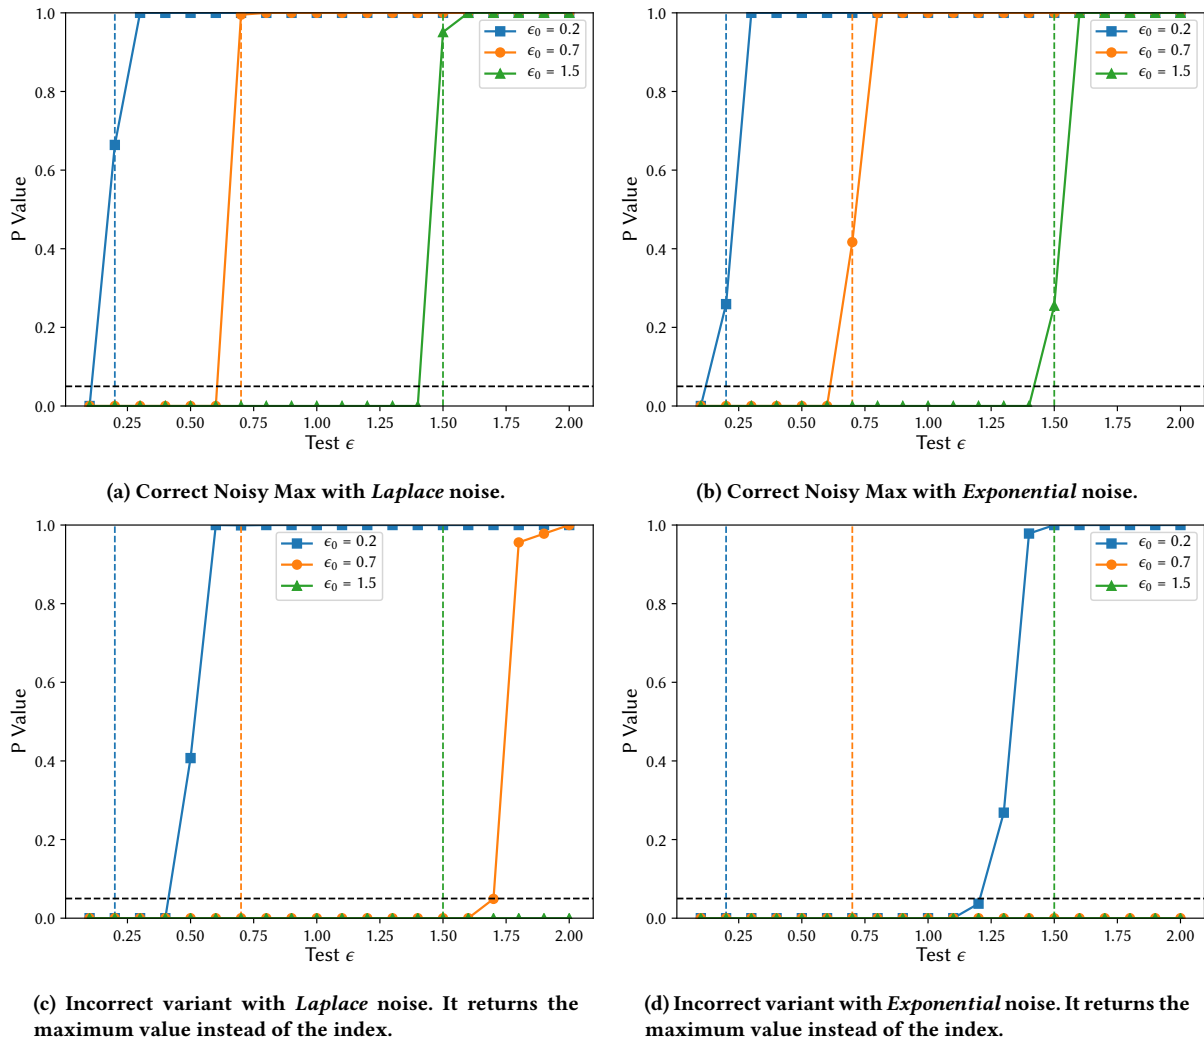


Figure 2: Results of Noisy Max algorithm and its variants.

The same pattern happens when we set privacy budget to be 0.7 and 1.5: all lines rise much later than their claimed privacy budget. In this incorrect version, returning the maximum value (instead of its index) causes the algorithm to actually satisfy $\epsilon_0 \cdot |Q|/2$ differential privacy instead of ϵ_0 -differential privacy. For the variant using Exponential noise (Figure 2d), the lines rise much later than the claimed privacy budgets, indicating strong evidence that this variant is indeed incorrect. Also, we can hardly see the lines for privacy budgets 0.7 and 1.5, since their p-values remain 0 for all the test ϵ ranging from 0 to 2.2 in the experiment.

5.2 Histogram


The Histogram algorithm [14] is a very simple algorithm for publishing an approximate histogram of the data. The input is a histogram and the output is a noisy histogram with the same dimensions. The Histogram algorithm requires input queries to differ in at most one

element. Here we evaluate with different scale parameters for the added Laplace noise.

The correct Histogram algorithm adds independent $\text{Lap}(1/\epsilon_0)$ noise to each query answer, as shown in Algorithm 9. Since at most one query answer may differ by at most 1, returning the maximum value is ϵ_0 -differentially private [14].

To mimic common mistakes made by novices of differential privacy, we also evaluate on an incorrect variant where $\text{Lap}(\epsilon_0)$ noise is used in the algorithm (Algorithm 10). We note that the incorrect variant here satisfies $1/\epsilon_0$ -differential privacy, rather the claimed ϵ_0 -differential privacy.

Figures 3a and 3b show the test results for the correct and incorrect variants respectively. Here, Figure 3a indicates that the correct implementation satisfies the claimed privacy budgets. For the incorrect variant, the claimed budgets of 0.2 and 0.7 are correctly rejected; this is expected since the true privacy budgets are $1/0.2$ and $1/0.7$ respectively for this incorrect version. Interestingly,



(a) Correct Histogram algorithm with Lap($1/\epsilon_0$) noise.

(b) Incorrect Histogram algorithm with Lap(ϵ_0) noise. It provides more privacy than advertised when $\epsilon_0 \geq 1$ and less privacy than advertised when $\epsilon_0 < 1$.

Figure 3: Results of Histogram algorithm and its variants

the result indicates that for $\epsilon_0 = 1.5$, this algorithm is likely to be more private than claimed (the line rise around 0.6 rather than 1.5). Again, this is expected since in this case, the variant is indeed $1/1.5 = 0.67$ -differentially private.

5.3 Sparse Vector

The Sparse Vector Technique (SVT) [18] (see Algorithm 11) is a powerful mechanism for answering numerical queries. It takes a list of numerical queries and simply reports whether their answers are above or below a preset threshold T . It allows the program to output some noisy query answers without any privacy cost. In particular, arbitrarily many “below threshold” answers can be returned, but only at most N “above threshold” answers can be returned. Because of this remarkable property, there are many variants proposed in both published papers and practical use. However, most of them turn out to be actually not differentially private [28]. We test our tool on a correct implementation of SVT and the major incorrect variants summarized in [28]. In the following, we describe what the variants do and list their pseudocodes.

5.3.1 SVT [28]. Lyu et al. have proposed an implementation of SVT and proved that it satisfies ϵ_0 -differential privacy. This algorithm (Algorithm 11) tries to allocate the global privacy budget ϵ_0 into two parts: half of the privacy budget goes to the threshold, and the other half goes to values which are above the threshold. There will not be any privacy cost if the noisy value is below the noisy threshold, in which case the program will output a *False*. If the noisy value is above the noisy threshold, the program will output a *True*. After outputting a certain amount (N) of *True*'s, the program will halt.

Figure 4a shows the test result for this correct implementation. All lines rise around the true privacy budget, indicating that our tool correctly conclude that this algorithm is correct.

Algorithm 11: SVT [28].

input: Q : queries to the database, ϵ_0 : privacy budget
 T : threshold, N : bound of outputting *True*'s
 Δ : sensitivity

```

1 function SVT( $Q, T, \epsilon_0, \Delta, N$ ):
2    $out \leftarrow []$ 
3    $\eta_1 \leftarrow \text{Lap}(2 * \Delta / \epsilon_0)$ 
4    $\tilde{T} \leftarrow T + \eta_1$ 
5    $count \leftarrow 0$ 
6   foreach  $q$  in  $Q$  do
7      $\eta_2 \leftarrow \text{Lap}(2 * N * \Delta / \epsilon_0)$ 
8     if  $q + \eta_2 \geq \tilde{T}$  then
9        $out \leftarrow True :: out$ 
10       $count \leftarrow count + 1$ 
11      if  $count \geq N$  then
12        Break
13      end
14    else
15       $out \leftarrow False :: out$ 
16    end
17  end
18  return ( $out$ )

```

5.3.2 *iSVT 1* [38]. One incorrect variant (Algorithm 12) adds no noise to the query answers, and has no bound on the number of *True*'s that the algorithm can output. This implementation does not satisfy ϵ_0 -differential privacy for any finite ϵ_0 .

This expectation is consistent with the test result shown in Figure 4b: the p-value never rises at any test ϵ . This result strongly indicates that this implementation with claimed privacy budget 0.2, 0.7, 1.5 is not private for at least any $\epsilon \leq 2.2$.

Algorithm 12: iSVT 1 [38]. This does not add noise to the query answers, and has no bound on number of *True*'s to output (i.e., N). This is not private for any privacy budget ϵ_0 .

input: Q : queries to the database, ϵ_0 : privacy budget
 T : threshold, Δ : sensitivity

```

1 function iSVT1( $Q, T, \epsilon_0, \Delta$ ):
2    $out \leftarrow []$ 
3    $\eta_1 \leftarrow \text{Lap}(\Delta/\epsilon_0)$ 
4    $\tilde{T} \leftarrow T + \eta_1$ 
5   // no bounds on number of outputs
6   foreach  $q$  in  $Q$  do
7     // adds no noise to query answers
8      $\eta_2 \leftarrow 0$ 
9     if  $q + \eta_2 \geq \tilde{T}$  then
10      |  $out \leftarrow \text{True} :: out$ 
11     else
12      |  $out \leftarrow \text{False} :: out$ 
13     end
14   end
15   return ( $out$ )

```

Algorithm 13: iSVT 3 [27]. The noise added to queries doesn't scale with N . The actual privacy cost is $\frac{1+6N}{4}\epsilon_0$.

input: Q : queries to the database, ϵ_0 : privacy budget
 T : threshold, N : bound of outputting *True*'s
 Δ : sensitivity

```

1 function iSVT3( $Q, T, \epsilon_0, \Delta, N$ ):
2    $out \leftarrow []$ 
3    $\eta_1 \leftarrow \text{Lap}(4 * \Delta/\epsilon_0)$ 
4    $\tilde{T} \leftarrow T + \eta_1$ 
5    $count \leftarrow 0$ 
6   foreach  $q$  in  $Q$  do
7     // noise added doesn't scale with N
8      $\eta_2 \leftarrow \text{Lap}(4 * \Delta/(3 * \epsilon_0))$ 
9     if  $q + \eta_2 \geq \tilde{T}$  then
10      |  $out \leftarrow \text{True} :: out$ 
11      |  $count \leftarrow count + 1$ 
12      | if  $count \geq N$  then
13        | Break
14      | end
15     else
16      |  $out \leftarrow \text{False} :: out$ 
17     end
18   end
19   return ( $out$ )

```

5.3.3 *iSVT 2* [11]. Another incorrect variant (Algorithm 14) has no bounds on the number of *True*'s the algorithm can output. Without the bounds, the algorithm will still output *True* even if it has

Algorithm 14: iSVT 2 [11]. This one has no bounds on number of *True*'s (i.e., N) to output. This is not private for any finite privacy budget ϵ_0 .

input: Q : queries to the database, ϵ_0 : privacy budget
 T : threshold, Δ : sensitivity

```

1 function iSVT2( $Q, T, \epsilon_0, \Delta$ ):
2    $out \leftarrow []$ 
3    $\eta_1 \leftarrow \text{Lap}(2 * \Delta/\epsilon_0)$ 
4    $\tilde{T} \leftarrow T + \eta_1$ 
5   // no bounds on number of outputs
6   foreach  $q$  in  $Q$  do
7      $\eta_2 \leftarrow \text{Lap}(2 * \Delta/\epsilon_0)$ 
8     if  $q + \eta_2 \geq \tilde{T}$  then
9       |  $out \leftarrow \text{True} :: out$ 
10     else
11       |  $out \leftarrow \text{False} :: out$ 
12     end
13   end
14   return ( $out$ )

```

Algorithm 15: iSVT 4 [36]. When the noisy query answer is above the threshold, output the actual value of noisy query answer.

input: Q : queries to the database, ϵ_0 : privacy budget
 T : threshold, N : bound of outputting *True*'s
 Δ : sensitivity


```

1 function iSVT4( $Q, T, \epsilon_0, \Delta, N$ ):
2    $out \leftarrow []$ 
3    $\eta_1 \leftarrow \text{Lap}(2 * \Delta/\epsilon_0)$ 
4    $\tilde{T} \leftarrow T + \eta_1$ 
5    $count \leftarrow 0$ 
6   foreach  $q$  in  $Q$  do
7      $\eta_2 \leftarrow \text{Lap}(2 * N * \Delta/\epsilon_0)$ 
8     if  $q + \eta_2 \geq \tilde{T}$  then
9       // output numerical value instead of boolean value
10      |  $out \leftarrow (q + \eta_2) :: out$ 
11      |  $count \leftarrow count + 1$ 
12      | if  $count \geq N$  then
13        | Break
14      | end
15     else
16      |  $out \leftarrow \text{False} :: out$ 
17     end
18   end
19   return ( $out$ )

```



exhausted its privacy budget. So this variant is not private for any finite ϵ_0 .

Figure 4c indicates this implementation with privacy budget $\epsilon_0 = 0.2$ is most likely *not* private for any $\epsilon \leq 0.5$. When $\epsilon_0 = 0.7$,




(a) Correct implementation of SVT [28].

(b) iSVT 1 [38] adds no noise to query and threshold.



(c) iSVT 2 [11] no bounds on outputting True's.

(d) iSVT 3 [27] query noise does not scale with N.



(e) iSVT 4 [36] outputs the actual query answer when it is above the threshold.

Figure 4: Results for variants of Sparse Vector Technique

Table 2: Counterexamples detected for incorrect privacy mechanisms

Mechanism ($\epsilon_0 = 1.5$)	Event E	D1	D2
Incorrect Noisy Max with Laplace Noise	$\omega \in (-\infty, 0.0)$	[1, 1, 1, 1, 1]	[0, 0, 0, 0, 0]
Incorrect Noisy Max with Exponential Noise	$\omega \in (-\infty, 1.0)$	[1, 1, 1, 1, 1]	[0, 0, 0, 0, 0]
Incorrect Histogram [17]	$\omega[0] \in (-\infty, 1.0)$	[1, 1, 1, 1, 1]	[2, 1, 1, 1, 1]
iSVT 1 [38]	$t(\omega) = 0$	[1, 1, 1, 1, 1, 1, 1, 1, 1, 1]	[0, 0, 0, 0, 0, 2, 2, 2, 2, 2]
iSVT 2 [11]	$t(\omega) = 9$	[1, 1, 1, 1, 1, 0, 0, 0, 0, 0]	[0, 0, 0, 0, 0, 1, 1, 1, 1, 1]
iSVT 3 [27]	$t(\omega) = 0$	[1, 1, 1, 1, 1, 0, 0, 0, 0, 0]	[0, 0, 0, 0, 0, 1, 1, 1, 1, 1]
iSVT 4 [36]	$(\omega.count(False), \omega[9]) \in \{9\} \times (-2.4, 2.4)$	[1, 1, 1, 1, 1, 1, 1, 1, 1, 1]	[0, 0, 0, 0, 0, 0, 0, 0, 0, 0]

we have detected counterexamples showing the algorithm is likely not private for any $\epsilon \in (0, 2.1]$. When $\epsilon_0 = 1.5$, we have detected counterexamples showing the algorithm is likely not private for any $\epsilon \in (0, 3.0]$.

5.3.4 iSVT 3 [27]. Another incorrect variant (Algorithm 13) adds noise to queries but the noise doesn't scale with the bound N . The actual privacy budget for this variant is $\frac{1+6N}{4}\epsilon_0$ where ϵ_0 is the input privacy budget.

We note that our tool detects the actual privacy cost, as shown in Figure 4d, for this incorrect algorithm. Consider privacy budget $\epsilon_0 = 0.2$. The corresponding line rises at 0.3, right before the actual budget $\frac{1+6N}{4}\epsilon_0 = 0.35$ ($N = 1$), suggesting the precision of our tool. The same happens for $\epsilon_0 = 0.7$ and 1.5. The two lines rise at 1.1 and 2.3, which are close to but before the actual budget 1.225 and 2.625, respectively.

5.3.5 iSVT 4 [36]. Another incorrect variant (Algorithm 15) outputs the actual value of noisy query answer when it is above the noisy threshold.

The interesting part of this algorithm is that, since it outputs heterogeneous list of booleans and values, our event selector chooses $\{9\} \times (-2.4, 2.4)$. This means we choose an event that consists of 9 booleans (in this case, Falses) followed by a number in $(-2.4, 2.4)$. Figure 4e shows much noise in it because this one is almost correct in the sense that violations of differential privacy happen with very low probability; thus it is hard to detect its incorrectness. But we can still see that the lines all rise later than the corresponding claimed privacy budget ϵ_0 . Hence, our tool correctly concludes that this algorithm does not satisfy ϵ_0 -differential privacy.

5.4 Performance

We performed all experiments on a double Intel® Xeon® E5-2620 v4 @ 2.10GHz CPU machine with 64 GB memory. Our tool is implemented in Anaconda distribution of Python 3 and optimized for running in parallel environment to fully utilize the 32 logical cores of the machine.

For each test ϵ , we set the samples of iteration n to be 500,000 for the hypothesis test and 100,000 for the event selector and query generator. Table 3 lists the average time spent on hypothesis test for a specific test ϵ (i.e., the average time spent on generating one single point in the figures) for each algorithm. The results suggest that it is very efficient to run a test for an algorithm against one privacy cost: all tests finish within 23 seconds.

The time difference between Noisy Max, Histogram and Sparse Vector Technique is due to the nature of the algorithms. For SVT, the

Table 3: Time spent on running tool for different algorithms

Mechanism	Time / Seconds
Correct Laplace Noisy Max [15]	4.32
Incorrect Laplace Noisy Max	9.49
Correct Exponential Noisy Max [15]	4.25
Incorrect Exponential Noisy Max	8.70
Histogram [14]	10.39
Incorrect Histogram	11.28
SVT [28]	1.99
iSVT 1 [38]	1.62
iSVT 2 [11]	4.56
iSVT 3 [27]	2.56
iSVT 4 [36]	22.97

parameter N is set to 1, meaning that the algorithm will halt once it hit a True branch. For Noisy Max and Histogram, all noise will be calculated and applied to each query answer, consuming more time to calculate p-values. Another factor that will also influence the test time is the search space of events. Correct *Noisy Max* returns an index which we would have a search space of only integers ranging from 1 to the length of queries. However, the incorrect *Noisy Max* will return a real number so the search space would be much larger than the correct one, thus taking more time to find a suitable event E . This also occurs in Sparse Vector Technique.

6 CONCLUSIONS AND FUTURE WORK

While it is invaluable to formally verify correct differentially-private algorithms, we believe that it is equally important to detect incorrect algorithms and provide counterexamples for them, due to the subtleties involved in algorithm development. We proposed a novel semi-black-box method of evaluating differentially private algorithms, and providing counterexamples for those incorrect ones. We show that within a few seconds, our tool correctly rejects incorrect algorithms (including published ones) and provides counterexamples for them.

Future work includes extensions that detect violations of differential privacy even if those violations occur with extremely small probabilities. This will require additional extensions such as a more refined use of program analysis techniques (including symbolic execution) that reason about what happens when a program is run on adjacent databases. Additional extensions include counterexample generation for other variants of differential privacy, such

as approximate differential privacy, zCDP, and renyi-differential privacy.

ACKNOWLEDGMENTS

We thank anonymous CCS reviewers for their helpful suggestions. This work was partially funded by NSF awards #1228669, #1702760 and #1566411.

REFERENCES

- [1] Aws Albarghouthi and Justin Hsu. 2017. Synthesizing coupling proofs of differential privacy. *Proceedings of the ACM on Programming Languages* 2, POPL (2017), 58.
- [2] Gilles Barthe, George Danezis, Benjamin Grégoire, César Kunz, and Santiago Zanella-Béguelin. 2013. Verified Computational Differential Privacy with Applications to Smart Metering. In *2013 IEEE 26th Computer Security Foundations Symposium*. 287–301.
- [3] Gilles Barthe, Noémie Fong, Marco Gaboardi, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. 2016. Advanced Probabilistic Couplings for Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 55–67.
- [4] Gilles Barthe, Marco Gaboardi, Emilio Jesús Gallego Arias, Justin Hsu, César Kunz, and Pierre-Yves Strub. 2014. Proving Differential Privacy in Hoare Logic. In *2014 IEEE 27th Computer Security Foundations Symposium*. 411–424.
- [5] Gilles Barthe, Marco Gaboardi, Benjamin Gregoire, Justin Hsu, and Pierre-Yves Strub. 2016. Proving differential privacy via probabilistic couplings. In *IEEE Symposium on Logic in Computer Science (LICS)*.
- [6] Gilles Barthe, Boris Köpf, Federico Olmedo, and Santiago Zanella Béguelin. 2012. Probabilistic Relational Reasoning for Differential Privacy. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 97–110.
- [7] Gilles Barthe and Federico Olmedo. 2013. Beyond differential privacy: Composition theorems and relational logic for f-divergences between probabilistic programs. In *ICALP*. 49–60.
- [8] Mark Bun and Thomas Steinke. 2016. Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds. In *Proceedings of the 14th International Conference on Theory of Cryptography - Volume 9985*.
- [9] Cristian Cadar, Daniel Dunbar, and Dawson R. Engler. 2008. KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs. In *8th USENIX Symposium on Operating Systems Design and Implementation, (OSDI)*. 209–224.
- [10] Cristian Cadar, Vijay Ganesh, Peter M. Pawlowski, David L. Dill, and Dawson R. Engler. 2006. EXE: Automatically Generating Inputs of Death. In *Conference on Computer and Communications Security (CCS)*. 322–335.
- [11] Rui Chen, Qian Xiao, Yu Zhang, and Jianliang Xu. 2015. Differentially private high-dimensional data publication via sampling-based inference. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 129–138.
- [12] Yan Chen and Ashwin Machanavajjhala. 2015. On the Privacy Properties of Variants on the Sparse Vector Technique. <http://arxiv.org/abs/1508.07306>. (2015).
- [13] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 337–340.
- [14] Cynthia Dwork. 2006. Differential Privacy. In *Proceedings of the 33rd International Conference on Automata, Languages and Programming - Volume Part II (ICALP'06)*. Springer-Verlag, Berlin, Heidelberg, 1–12. https://doi.org/10.1007/11787006_1
- [15] Cynthia Dwork. 2008. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*. Springer, 1–19.
- [16] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. 2006. Our Data, Ourselves: Privacy via Distributed Noise Generation. In *EUROCRYPT*. 486–503.
- [17] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *TCC*.
- [18] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil Vadhan. 2009. On the Complexity of Differentially Private Data Release: Efficient Algorithms and Hardness Results. In *STOC*. 381–390.
- [19] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* 9, 3–4 (2014), 211–407.
- [20] Cynthia Dwork and Guy N. Rothblum. 2016. Concentrated Differential Privacy. <http://arxiv.org/abs/1603.01887>. (2016).
- [21] Hamid Ebad, David Sands, and Gerardo Schneider. 2015. Differential privacy: Now it's getting personal. In *ACM Symposium on Principles of Programming Languages (POPL)*.
- [22] Gian Pietro Farina, Stephen Chong, and Marco Gaboardi. 2017. Relational Symbolic Execution. <http://arxiv.org/abs/1711.08349>. (2017).
- [23] R.A. Fisher. 1935. *The design of experiments*. 1935. Oliver and Boyd, Edinburgh.
- [24] Marco Gaboardi, Andreas Haeberlen, Justin Hsu, Arjun Narayan, and Benjamin C. Pierce. 2013. Linear Dependent Types for Differential Privacy. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (ACM Symposium on Principles of Programming Languages (POPL))*. 357–370.
- [25] Eric Jones, Travis Oliphant, Pearu Peterson, et al. 2001–. SciPy: Open source scientific tools for Python. (2001–). <http://www.scipy.org/> [Online; accessed 2018-08-17].
- [26] James C King. 1976. Symbolic execution and program testing. *Commun. ACM* 19, 7 (1976), 385–394.
- [27] Jaewoo Lee and Christopher W Clifton. 2014. Top-k frequent itemsets via differentially private fp-trees. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 931–940.
- [28] Min Lyu, Dong Su, and Ninghui Li. 2017. Understanding the sparse vector technique for differential privacy. *Proceedings of the VLDB Endowment* 10, 6 (2017), 637–648.
- [29] Frank D. McSherry. 2009. Privacy Integrated Queries: An Extensible Platform for Privacy-preserving Data Analysis. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data*. 19–30.
- [30] Dimitar Milushev, Wim Beck, and Dave Clarke. 2012. Noninterference via symbolic execution. In *Formal Techniques for Distributed Systems*. Springer, 152–168.
- [31] Ilya Mironov. 2017. Rényi Differential Privacy. In *30th IEEE Computer Security Foundations Symposium, CSF*.
- [32] Prashanth Mohan, Abhradeep Thakurta, Elaine Shi, Dawn Song, and David Culler. 2012. GUPT: Privacy Preserving Data Analysis Made Easy. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*.
- [33] Suzette Person, Matthew B Dwyer, Sebastian Elbaum, and Corina S Păsăreanu. 2008. Differential symbolic execution. In *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of software engineering*. ACM, 226–237.
- [34] Silvio Ranise and Cesare Tinelli. 2006. *The smt-lib standard: Version 1.2*. Technical Report. Technical report, Department of Computer Science, The University of Iowa, 2006. Available at www.SMT-LIB.org.
- [35] Jason Reed and Benjamin C. Pierce. 2010. Distance Makes the Types Grow Stronger: A Calculus for Differential Privacy. In *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming (ICFP '10)*. 157–168.
- [36] A. Roth. 2011. Sparse Vector Technique, Lecture notes for “The Algorithmic Foundations of Data Privacy”. (2011).
- [37] Indrajit Roy, Srinath Setty, Ann Kilzer, Vitaly Shmatikov, and Emmett Witchel. 2010. Airavat: Security and Privacy for MapReduce. In *NSDI*.
- [38] Ben Stoddard, Yan Chen, and Ashwin Machanavajjhala. 2014. Differentially private algorithms for empirical machine learning. *arXiv preprint arXiv:1411.5428* (2014).
- [39] Michael Carl Tschantz, Dilsun Kaynar, and Anupam Datta. 2011. Formal Verification of Differential Privacy for Interactive Systems (Extended Abstract). *Electron. Notes Theor. Comput. Sci.* 276 (Sept. 2011), 61–79.
- [40] Lili Xu, Konstantinos Chatzikokolakis, and Huimin Lin. 2014. *Metrics for Differential Privacy in Concurrent Systems*. 199–215.
- [41] Danfeng Zhang and Daniel Kifer. 2017. LightDP: Towards Automating Differential Privacy Proofs. In *ACM Symposium on Principles of Programming Languages (POPL)*. 888–901.