

TECHNICAL UNIVERSITY OF CRETE
SCHOOL OF ELECTRICAL & COMPUTER ENGINEERING



DIPLOMA THESIS

Data Analytics with Differential Privacy

Author:
Vassilis V. Digalakis, Jr.

Thesis Committee:
Prof. Minos Garofalakis
(*Advisor*)
Prof. George Karystinos
Prof. Aggelos Bletsas

A thesis submitted in partial fulfillment of the requirements for the degree of
Diploma in Electrical and Computer Engineering

July 14, 2018

Data Analytics with Differential Privacy

by Vassilis V. Digalakis, Jr.

Thesis Advisor: Prof. Minos Garofalakis

Abstract

Differential privacy is the state-of-the-art definition for privacy, that “addresses the paradox of learning nothing about an individual while learning useful information about a population” (Dwork and Roth). In other words, differential privacy guarantees that any analysis performed on a sensitive dataset leaks no information about the individuals whose data are contained therein. In this thesis, we develop differentially private algorithms to analyze distributed and streaming data.

In the distributed model, we consider the particular problem of learning -in a distributed fashion- a global model of the data, that can subsequently be used for arbitrary analyses. We build upon PrivBayes, a differentially private method that approximates the high-dimensional distribution of a centralized dataset as a product of low-order distributions, utilizing a Bayesian Network model. Specifically, we examine three novel approaches to learning a global Bayesian Network from distributed data, while offering the differential privacy guarantee to all local datasets. Our work includes a detailed theoretical analysis of the distributed, differentially private entropy estimator which we use in one of our algorithms, as well as a detailed experimental evaluation, using both synthetic and real-world data.

In the streaming model, we focus on the problem of estimating the density of a stream of users (or, more generally, elements), which expresses the fraction of all users that actually appear in the stream. We offer one of the strongest privacy guarantees for the streaming model, namely user-level pan-privacy, which ensures that the privacy of any user is protected, even against an adversary that observes, on rare occasions, the internal state of the algorithm. We provide a detailed analysis of an existing, sampling-based algorithm for the problem, and propose two novel modifications that significantly improve it, both theoretically and experimentally, by optimally using all the allocated “privacy budget”.

Acknowledgments

After five years of truly hard yet exciting work, and before leaving Greece to pursue my PhD degree in the US, I cannot but express my gratitude towards the people who were by my side and made everything I achieved possible.

First and foremost, I would like to thank my advisor, *Prof. Minos Garofalakis*, for exposing me to research and pointing me towards exciting topics. I am also grateful towards the other two members of the committee. I had numerous intriguing discussions with *Prof. George Karystinos*, both about technical issues and about Olympiacos. I have to admit that the ones I enjoyed the most were -surprisingly!- the former. *Prof. Aggelos Bletsas* always encouraged and supported me, and taught me my first two MIT courses. Hopefully the ones that will follow at MIT will be equally interesting.

I owe a lot to many, among my other Professors at TUC, especially *Prof. Stavros Christodoulakis*, for his wise advice when I had to make critical decisions, *Prof. Michail Lagoudakis*, for his kindness and support, and *Prof. Daphne Manoussaki*, for offering me the opportunity to serve as a teaching assistant.

Nothing would have been possible without *my family and my close friends*. I want to thank *my mother*, for her constant and unconditional support, and *Korina*, whose achievements inspired me as well, and motivated me to work even harder. And of course, no words are enough to thank *Areti*; she has been by my side both in good and in bad times, for many many years now. Last but not least, I acknowledge the person to whom I owe the most, *my father*, who managed -through his continuous guidance- to convert one of the world's most impatient persons, to a more patient, persistent and determined one.

Contents

1	Introduction	1
1.1	The Promise of Differential Privacy	2
1.2	Distributed and Streaming Data	2
1.3	Thesis Organization & Contributions	2
2	Differential Privacy	5
2.1	Basic Definitions & Theorems	5
2.2	Achieving Differential Privacy	7
2.2.1	Randomized Response	8
2.2.2	Laplace Mechanism	8
2.2.3	Exponential Mechanism	9
2.3	Data Analytics with Differential Privacy	10
3	Distributed Bayesian Network Learning with Differential Privacy	11
3.1	Differential Privacy in the Distributed Model	13
3.1.1	Distributed Data Mining with Differential Privacy	14
3.1.2	Data Publishing with Differential Privacy	15
3.1.3	Our Approach: Distributed Data Publishing with Differential Privacy	16
3.2	Bayesian Networks	16
3.2.1	Basic Definitions	16
3.2.2	Learning Bayesian Networks	18
3.3	PrivBayes	20
3.3.1	Structure Learning Phase	21
3.3.2	Parameter Learning Phase	23
3.3.3	PrivBayes in the Distributed Model	23
3.4	Sharing the Noisy Sufficient Statistics	25
3.4.1	Privacy Analysis	26
3.4.2	Accuracy Analysis	26
3.5	Noisy Majority Voting	37
3.6	Sharing the Noisy Model	38
3.7	Experimental Evaluation	39
3.7.1	Evaluation Metric	39
3.7.2	Synthetic Data	40
3.7.3	Real-World Data	43
3.8	Discussion	45
4	Pan-Private Stream Density Estimation	47
4.1	Differential Privacy in the Streaming Model	48
4.1.1	Pan-Privacy	50
4.1.2	Continual Observation	51
4.2	Dwork’s Density Estimator	52

4.2.1	Privacy Analysis	53
4.2.2	Accuracy Analysis	54
4.2.3	Picking the Optimal Sample Size	59
4.3	Optimal Bernoulli Density Estimator	59
4.3.1	On the Use of the Allocated Privacy Budget	59
4.3.2	Optimally Tuning the Bernoulli Distributions	60
4.3.3	Estimator & Analysis	62
4.4	Laplace Density Estimator	66
4.4.1	Using Continuous Distributions	66
4.4.2	Estimator & Analysis	68
4.4.3	Quantized Laplace Density Estimator	73
4.4.4	Comments on the Use of Other Continuous Distributions	74
4.5	Experimental Evaluation	75
4.6	Discussion	78
5	Conclusions & Future Work	81
5.1	Distributed Model	81
5.2	Streaming Model	81
	References	83

Chapter 1

Introduction

“Historically, privacy was almost implicit, because it was hard to find and gather information. But in the digital world, whether it’s digital cameras or satellites or just what you click on, we need to have more explicit rules - not just for governments but for private companies.”

Bill Gates

The importance of privacy in the era of Big Data is well-understood. The recent Facebook-Cambridge Analytica scandal is only the last in a series of major privacy breaches, that point out the need to reconsider our perspective on the significance of protecting the data that we generate. At the same time, however, it would be a huge mistake to miss the opportunities that the massive availability of data offers.

Until recently, the need for privacy was much less established, as our data mostly lived within our personal computers and were protected by our passwords. The need to securely store our data within our computers, combined with the needs to perform secure electronic transactions and communicate securely (to mention but a few), were addressed by cryptography and formal approaches were developed early on in the history of computer science.

In contrast, nowadays most of our data live in the Cloud, or in the servers of Google, Facebook, Amazon etc., and only a tiny fraction of them still lives within our personal computers. Furthermore, taking the explosion of analytics into account, we may even not be able to prevent our data from being included in analyses, both for advertisement and for research purposes. How will our privacy be protected?

Cryptography may still be the answer in some cases, but even if we manage to develop methods to analyze encrypted data sufficiently well, there are still some unsurmountable problems. Once an individual’s data enter a database, they may even stay there forever; nobody can guarantee that the cryptographic techniques used will still protect the individual’s privacy in, say, 25 years from now. Moreover, the effect of an individual’s data on the result of an analysis, which will essentially be published, cannot be protected by cryptography.

That being said, we have to distinguish between security and privacy. The fact that our data are securely stored today does not mean that our privacy is protected; neither today nor in the future. Classical cryptography has provided decent solutions in guaranteeing security. In the case of privacy, however, a totally different calculus applies. It is clear that a formal definition for privacy is also essential, together with a set of tools (e.g. algorithms,

software), that will allow us guarantee that the privacy (according to this definition) of individuals who contribute their data to datasets and analyses is protected.

1.1 The Promise of Differential Privacy

“You will not be affected, adversely or otherwise, by allowing your data to be used in any study or analysis, no matter what other studies, datasets or information sources are available.”

The promise of differential privacy.
Cynthia Dwork & Aaron Roth

Early approaches in privacy-preserving data analysis suffered from a series of attacks (e.g. linkage attacks, differencing attacks, reconstruction attacks), and commonly used privacy definitions, like k-anonymity [57], were shown to be flawed. On the contrary, *differential privacy* meets the requirements in being a formal privacy definition. Introduced in 2006 by Cynthia Dwork, Frank McSherry, Kobi Nissim, and Adam Smith [21], differential privacy is a strong and mathematically rigorous guarantee, that describes a promise, made by a data holder, to an individual that contributes its data to a dataset, that the individual’s privacy will be protected.

As Dwork and Roth [24] argue, differential privacy addresses the paradox of learning nothing about an individual while learning useful information about a population. In addition to being a strong privacy guarantee, differential privacy as a definition has mathematical properties that facilitate the design of algorithms which satisfy it. As a result, numerous methods that realize the differential privacy guarantee have been developed in the past few years, and now we are definitely one step closer to privacy-preserving data analysis.

1.2 Distributed and Streaming Data

The paradigm of our data living in a static database that is protected by a curator (e.g. database administrator) belongs to the past to a large extent. Nowadays, our data are everywhere and in various forms:

- They are often distributed among several databases. This is common, for example, in medical applications, where different hospitals possess different parts of a joint database of the clinical records of individuals.
- Or, they are dynamically created and arrive continuously in a stream. Being able to monitor such a stream and extract statistics is important for many disciplines, like -for instance- epidemiology.

We mention these two particular applications due to their obvious connections with privacy; medical data are by definition sensitive.

1.3 Thesis Organization & Contributions

In this section we jointly outline the organization of this thesis and its key contributions.

In **Chapter 2**, we formalize the notion of differential privacy. We give the basic definitions and theorems, and present the most common mechanisms that are used as building

blocks to provide the differential privacy guarantee. In addition, we describe in detail the *traditional differential privacy model*, and emphasize some points that often cause confusion in the literature.

In **Chapter 3**, we develop differentially private algorithms to analyze *distributed data*, and we consider a model where a large dataset is horizontally distributed among mutually distrustful parties. We address the particular problem of distributed learning of Bayesian Networks with differential privacy. The key contributions of the chapter are the following:

- We formally describe differential privacy in the distributed model, we examine several alternative approaches to solving our problem, and we provide a detailed survey of the related work that addresses ours, or similar models.
- We present PrivBayes [64],[65], the state-of-the-art method in learning Bayesian Networks with differential privacy, we identify the challenges that arise when moving PrivBayes from a centralized to a distributed environment, and we examine three novel approaches to learning a global Bayesian Network from distributed data, while offering the differential privacy guarantee to all local datasets.
- We provide a detailed theoretical analysis of the distributed, differentially private entropy estimator which we use in one of our algorithms.
- We experimentally evaluate our algorithms using both synthetic and real-world data.

In **Chapter 4**, we develop differentially private algorithms to analyze *streaming data*, and we consider the cashier-register streaming model. We address the particular problem of estimating the density of a stream of users, and we offer one of the strongest privacy guarantees for the streaming model, namely user-level pan-privacy, which ensures that the privacy of any user is protected, even against an adversary that observes, on rare occasions, the internal state of the algorithm. The key contributions of the chapter are the following:

- We formally describe differential privacy in the streaming model, we analyze in-depth the existing definitions and approaches, and we provide an extensive survey of the related work.
- We provide for the first time a detailed analysis of the sampling-based, pan-private density estimator, proposed by Dwork et al. [23], we identify its main limitation, in that it does not use all the allocated privacy budget.
- We examine two different novel approaches to modifying the original estimator, based on optimally tuning the Bernoulli distributions it uses, as well as on using continuous distributions (Laplace, Gaussian), and we analyze the theoretic guarantees that our modified estimators offer.
- We experimentally compare our algorithms.

In **Chapter 5**, we conclude this thesis, and we provide directions for future work and possible extensions to our solutions.

Chapter 2

Differential Privacy

In this chapter, we formalize the notion of differential privacy. In addition, we present some key theorems that follow from the definitions which we introduce, and describe the most common mechanisms that are used as building blocks to provide the differential privacy guarantee.

We consider a *model of computation*, where a database, or (more generally) a dataset D contains the sensitive data of individuals, called the data owners; each record/tuple/data point (the terms are used interchangeably) in D corresponds to a single individual. We assume the existence of a trusted entity by the data owners, called the data holder or curator. The data holder has direct access to the sensitive dataset and analyzes it, ensuring that any output produced by the analysis does not violate the owners' privacy. A key thing to notice is that the data holder and the data analyst is a common (and trusted) entity. A line of work considers the data analyst to be a separate, untrusted entity that observes the result of the analysis, which is performed by the (trusted) data holder. We do not follow this approach, and we refer to the untrusted entity against which we want to protect the data as the adversary. We call this the *traditional differential privacy model*.

We also need to distinguish between the *interactive* and the *non-interactive* case.

- In the interactive case, the data holder is adaptively asked queries on the sensitive dataset and still has to protect the owners' privacy. In other words, the adversary observes the output of every analysis performed by the data holder, and may pose a new query based on this output.
- In the non-interactive case, the data holder analyzes the sensitive dataset only once. The output of the analysis is published, so the adversary has access to it, but the original dataset is not accessed again and may even be destroyed.

In our work, we focus on the non-interactive case, so that the analysis that is to be performed on the sensitive data is decided and known in advance.

2.1 Basic Definitions & Theorems

In defining differential privacy, it is useful to think of the dataset D as a finite collection of records from a universe \mathcal{U} . The histogram representation of D is a function $\text{hist}(D) : \mathcal{U} \rightarrow \mathbb{N}^{|\mathcal{U}|}$, where the i -th entry $\text{hist}(D)_i$ (for some $i \in \{1, \dots, |\mathcal{U}|\}$) represents the number of tuples in D of type $\mathcal{U}[i]$ (assuming that an arbitrary ordering is defined over \mathcal{U}).

We proceed by introducing the notion of adjacency between datasets, which (informally) refers to a pair of datasets that differ on a single record. Depending on the interpretation of the word “differ”, two different definitions have been used in the literature.

Definition 2.1 (Adjacency). *Datasets D and D' are adjacent, denoted $\text{adj}(D, D')$, if ...*

A: ... D can be obtained from D' by either adding or removing a single record, so:

$$\|hist(D) - hist(D')\|_1 = \sum_{i=1}^{|\mathcal{U}|} |hist(D)_i - hist(D')_i| \leq 1$$

B: ... D can be obtained from D' by changing the value of a single record, so:

$$\|hist(D) - hist(D')\|_1 = \sum_{i=1}^{|\mathcal{U}|} |hist(D)_i - hist(D')_i| \leq 2$$

As we will shortly see, the notion of adjacency plays a key role in the definition of differential privacy, and depending on the approach we follow to defining adjacency, we end up with a (slightly) different flavor of differential privacy (Kifer and Machanavajjhala [43]). If Definition 2.1(A) is used, we speak of *unbounded* differential privacy, whereas if Definition 2.1(B) is adopted, we speak of *bounded* differential privacy. Bounded differential privacy derives its name from the fact that the adjacent datasets involved have essentially the same size; in unbounded differential privacy there is no such restriction. In the literature, both approaches are assumed to be fine and one is free to choose whichever is more convenient.

We are now ready to formally define differential privacy, which intuitively guarantees that a randomized algorithm accessing a sensitive dataset produces similar outputs on similar (adjacent) inputs. As a result, the impact of any single record (individual) on the algorithm's output is negligible, and hence no information is leaked about the individuals whose data are in the dataset.

Definition 2.2 (Differential Privacy). *A randomized algorithm $\mathbf{Alg} : \mathcal{U} \rightarrow \mathcal{O}$ is ε -differentially private if for all $O \subseteq \mathcal{O}$, and for all pairs of adjacent datasets D, D' ,*

$$\mathbb{P}[\mathbf{Alg}(D) \in O] \leq e^\varepsilon \mathbb{P}[\mathbf{Alg}(D') \in O]$$

where the probability space is over the coin flips of \mathbf{Alg} .

The parameter ε , called the *privacy budget*, quantifies the privacy risk. In general, smaller values of ε imply more privacy, as the distributions of outputs of the algorithm for adjacent inputs tend to “come closer”. How we should pick ε in practice is an open question; Hsu et al. [38] address this question and propose a method for choosing ε .

An additional remark is that ε -differential privacy is a statistical in nature, *worst-case guarantee*, in that it ensures that for every run of algorithm \mathbf{Alg} (on every possible input), the output observed is almost equally likely to be observed on every adjacent input. Therefore, several *relaxations* of the original Definition 2.2 have been considered; for example, (ε, δ) -differential privacy guarantees that for every pair of adjacent datasets, it is extremely unlikely that the observed output will be more likely to be observed with the one input of the pair than with the other. In other words, (ε, δ) -differential privacy allows a privacy leakage with some small probability δ .

We next present some key theorems; their proofs can be found in [24]. Theorem 2.1 is known as the Post-Processing Theorem, and illustrates that differential privacy is immune to post-processing; an adversary cannot compute a function of the output of a differentially private algorithm and make it less private.

Theorem 2.1 (Post-Processing). *Let $\mathbf{Alg} : \mathcal{U} \rightarrow \mathcal{O}$ be a randomized algorithm that satisfies ε -differential privacy. Let $f : \mathcal{O} \rightarrow \mathcal{O}'$ be an arbitrary randomized mapping. Then $f(\mathbf{Alg}) : \mathcal{U} \rightarrow \mathcal{O}'$ also satisfies ε -differential privacy.*

Theorem 2.2, known as the Sequential Composition Theorem, shows that if we compose multiple differentially private mechanisms, the privacy budget adds up. For simplicity, we only present the theorem for the case of two mechanisms, but we note that it holds in general.

Theorem 2.2 (Sequential Composition). *Let $\mathbf{Alg}_1 : \mathcal{U} \rightarrow \mathcal{O}_1$ be a randomized algorithm that satisfies ε_1 -differential privacy, and let $\mathbf{Alg}_2 : \mathcal{U} \rightarrow \mathcal{O}_2$ be a randomized algorithm that satisfies ε_2 -differential privacy. Then their combination $\mathbf{Alg}_{1,2} = (\mathbf{Alg}_1, \mathbf{Alg}_2) : \mathcal{U} \rightarrow \mathcal{O}_1 \times \mathcal{O}_2$ satisfies $(\varepsilon_1 + \varepsilon_2)$ -differential privacy.*

Differential privacy also composes when a sequence of differentially private mechanisms is applied in parallel on non-intersecting subsets of the entire dataset. Theorem 2.3 is known as the Parallel Composition Theorem.

Theorem 2.3 (Parallel Composition). *Let $\mathbf{Alg}_1 : \mathcal{U} \rightarrow \mathcal{O}_1$ be a randomized algorithm that satisfies ε_1 -differential privacy, and let $\mathbf{Alg}_2 : \mathcal{U} \rightarrow \mathcal{O}_2$ be a randomized algorithm that satisfies ε_2 -differential privacy. Let D_1, D_2 be two partitions of the dataset D , such that $D_1 \cup D_2 = D$ and $D_1 \cap D_2 = \emptyset$. Then their combination $\mathbf{Alg}_{1,2}(D) = (\mathbf{Alg}_1(D_1), \mathbf{Alg}_2(D_2)) : \mathcal{U} \rightarrow \mathcal{O}_1 \times \mathcal{O}_2$ satisfies $(\max\{\varepsilon_1, \varepsilon_2\})$ -differential privacy.*

An important thing to notice is that parallel composition does not hold for bounded differential privacy; Definition 2.2 is violated when removing a record from D_1 and adding a record to D_2 (or vice-versa). Nevertheless, if our goal is to separately offer the differential privacy guarantee to each subset of D , so that the adjacency relation is considered separately for D_1 and D_2 , then parallel composition does hold. This is the case in the approach we follow to applying differential privacy in the distributed model, in Chapter 3.

Finally, the next Theorem 2.4, protects the privacy of groups of size k . It also addresses the case that multiple (k) records in the dataset refer to the same individual.

Theorem 2.4 (Group Privacy). *Any ε -differentially private algorithm $\mathbf{Alg} : \mathcal{U} \rightarrow \mathcal{O}$ is $(k\varepsilon)$ -differentially private for groups of size k . That is, for all datasets D, D' such that $\|hist(D) - hist(D')\|_1 \leq k$ and for all $O \subseteq \mathcal{O}$,*

$$\mathbb{P}[\mathbf{Alg}(D) \in O] \leq e^{k\varepsilon} \mathbb{P}[\mathbf{Alg}(D') \in O]$$

where the probability space is over the coin flips of \mathbf{Alg} .

Now that we have developed a better understanding of differential privacy, we make an additional remark on definition 2.1. Any algorithm that satisfies ε -unbounded differential privacy also satisfies 2ε -bounded differential privacy, since changing the value of one record is equivalent to first removing the old version of the record and then adding the new one.

2.2 Achieving Differential Privacy

Differential privacy is a definition, and not an algorithm. In practice, we are interested in developing algorithms that satisfy Definition 2.2 and hence offer the differential privacy guarantee to their input datasets.

Definition 2.3. *A randomized algorithm that satisfies Definition 2.2 is called a privacy mechanism.*

As the reader may have noticed, a privacy mechanism is *essentially* a randomized algorithm, i.e. an algorithm that employs a degree of randomness as part of its logic and produces an output that is a random variable (or vector). We do not formally introduce

the notion of randomized algorithms here and we refer the interested reader to the book by Mitzenmacher and Upfal [47].

In this section, we present three primitive differentially private mechanisms, which we use throughout our work. More sophisticated mechanisms have been developed (e.g. sparse vector technique, multiplicative weights mechanism, subsample and aggregate framework), that achieve much better results by reconsidering the computational goal of specific tasks. We refer the interested reader to the monograph by Dwork and Roth [24] for a detailed presentation of such mechanisms.

2.2.1 Randomized Response

Randomized response is a research method proposed by Warner [59] that allows respondents to a survey on a sensitive issue to protect their privacy against the interviewer, while still providing credible answers. We next present a simple version of randomized response, based on an example Dwork and Roth [24]. We slightly modify the original example, in that the data holder perturbs the sensitive dataset D before publishing or analyzing it. We assume that D consists of a binary record b (a bit) per individual, which indicates whether the individual does or does not have a particular property.

Algorithm 2.1: Simple Randomized Response	
Input:	Dataset D
Output:	Differentially private dataset \tilde{D}
1	Initialize $\tilde{D} = \emptyset$
2	for each $b \in D$ do
3	Flip a fair coin
4	if <i>tails</i> then
5	$\tilde{b} = b$
6	else
7	Flip a second fair coin
8	if <i>heads</i> then
9	$\tilde{b} = 1$
10	else
11	$\tilde{b} = 0$
12	Add \tilde{b} to \tilde{D}
13	Return \tilde{D}

Theorem 2.5 examines the privacy guarantees of 2.1. We remark that, throughout this work, by log we refer to the natural logarithm.

Theorem 2.5 (Randomized response). *The version of randomized response described in Algorithm 2.1 satisfies $\log 3$ -differential privacy.*

The power of randomized response is that it provides plausible deniability, and it directly perturbs the sensitive dataset (privacy by process). As a result, even if an individual's record indicates that it has the property in question, the individual may still credibly argue that it does not.

2.2.2 Laplace Mechanism

The Laplace mechanism [21] provides a way to transform a numeric function $f : \mathcal{U} \rightarrow \mathbb{R}^N$ (that inputs a dataset D and outputs a vector $f(D) \in \mathbb{R}^N$) into a differentially private

mechanism. We first introduce the notion of the sensitivity of f , which intuitively captures the effect of a single record on the output of f .

Definition 2.4 (ℓ_1 -sensitivity). *The ℓ_1 -sensitivity of a function $f : \mathcal{U} \rightarrow \mathbb{R}^N$ is:*

$$\Delta f = \max_{adj(D,D')} \|f(D) - f(D')\|_1$$

We next introduce the Laplace distribution, which is a symmetric, double-sided version of the exponential distribution. Throughout our work, we slightly abuse notation and use $\text{Laplace}(\mu, b)$ (instead of $X \sim \text{Laplace}(\mu, b)$) to refer to a random variable that follows the Laplace Distribution. Similarly, we use $\text{Bernoulli}(p)$ to refer to a Bernoulli random variable with parameter p , and so forth.

Definition 2.5 (Laplace distribution). *A random variable $X \sim \text{Laplace}(\mu, b)$ has probability density function:*

$$f_X(x|\mu, b) = \frac{1}{2b} e^{-\frac{|x-\mu|}{b}}, \quad x \in \mathbb{R}$$

where μ is a location parameter, such that $\mathbb{E}[X] = \mu$, and $b > 0$ is a scale parameter, such that $\text{var}(X) = 2b^2$.

Theorem 2.6 (Laplace mechanism). *Given any function $f : \mathcal{U} \rightarrow \mathbb{R}^N$, the Laplace mechanism, that on input $D \in \mathcal{U}$ outputs:*

$$\tilde{f}(D) = f(D) + [X_1 \dots X_N]^T$$

where X_i are i.i.d. $\text{Laplace}(0, \frac{\Delta f}{\epsilon})$ random variables, satisfies ϵ -differential privacy.

2.2.3 Exponential Mechanism

The main limitation of the Laplace mechanism is that it can only handle numeric functions of the dataset. In contrast, the Exponential mechanism [45] provides a way to transform arbitrary (e.g. categorical) functions of the dataset into differentially private mechanisms. We now consider an arbitrary function $f : \mathcal{U} \rightarrow \mathcal{O}$ that maps an input dataset D to an arbitrary object $O \in \mathcal{O}$. The Exponential mechanism is based upon a scoring/quality/utility function (the terms are used interchangeably) $q : \mathcal{U} \times \mathcal{O} \rightarrow \mathbb{R}$, that measures the quality $q(D, O)$ of the output O when the input is D .

Theorem 2.7 (Exponential mechanism). *Given any function $f : \mathcal{U} \rightarrow \mathcal{O}$, the Exponential mechanism, that on input $D \in \mathcal{U}$ outputs an element $O \in \mathcal{O}$ with probability $\propto e^{\frac{\epsilon q(D, O)}{2\Delta q}}$, where $q : \mathcal{U} \times \mathcal{O} \rightarrow \mathbb{R}$ and $\Delta q = \max_{O \in \mathcal{O}} \max_{adj(D, D')} |q(D, O) - q(D', O)|$, satisfies ϵ -differential privacy.*

The Exponential mechanism defines a distribution \mathbf{p}_{EM} over the set of possible outputs, and then samples from \mathbf{p}_{EM} . Intuitively high-quality outputs are favored, as they are more likely to be sampled. \mathbf{p}_{EM} can be arbitrarily complex, so the Exponential mechanism can be a double-edged sword:

- On the one hand, the Exponential mechanism is general; for example, the Laplace mechanism can be viewed as an instance of the Laplace mechanism.
- On the other hand, in many occasions it may not even be possible to sample efficiently from \mathbf{p}_{EM} .

2.3 Data Analytics with Differential Privacy

Differential privacy has been established as the state-of-the-art model in privacy-preserving data analytics. Therefore, there now exists a vast literature in the field of data analytics with differential privacy; a survey of the related work is beyond the scope of this thesis. We refer the interested reader to the survey by Chaudhuri and Sarwate [52] (and the many references therein), who provide an overview of the work that connects differential privacy with many areas that relate to data analytics, namely statistics and robust statistics, signal processing and machine learning (classification, regression, dimensionality reduction, filtering). Nevertheless, in Chapters 3 and 4 we do extensively review efforts that relate to the models that we examine, namely the distributed and the streaming model.

Chapter 3

Distributed Bayesian Network Learning with Differential Privacy

We consider a model where a large dataset is *horizontally distributed among mutually distrustful parties* (data holders) that are not able or willing to share their part, forming a distributed database. Our goal is to perform privacy-preserving data mining, and in particular, to make inferences about the population, without compromising the privacy of the individuals whose data are used. The model we described applies, for instance, in biomedical data analysis, and constitutes a major limitation in biomedical research. Hospitals and other trustworthy entities maintain the clinical records of individuals, but are unable to share and accurately analyze them, due to the risk of privacy breaches (Figure 3.1).

We aim to achieve something more general compared to (most) previous approaches in *privacy-preserving data mining over distributed data*. In this direction, we examine three approaches in learning -in a distributed and privacy-preserving fashion- a model that approximates the (high-dimensional) data distribution as a product of low-order marginal and conditional distributions. This is achieved by exploiting the attribute dependencies that are present in the data distribution; the theory of probabilistic graphical models offers tools to efficiently examine these dependencies. Our first approach is an exact approach, in that it requires each data holder to share noisy versions of the algorithm's sufficient statistics; we also provide a detailed theoretical analysis of this approach. The other two approaches are based on heuristic techniques, inspired from popular ideas from the distributed machine learning literature. Finally, once we have the privacy-preserving approximation of the data distribution in hand, we are able to perform arbitrary analyses on our data (e.g. classification).

We assume that our distributed database consists of homogeneous records, and each record consists of N attributes X_1, \dots, X_N , from a set \mathcal{A} . In statistical machine learning, each attribute is viewed as a random variable, and, thus, a record \mathbf{x} (or data point - the terms are used interchangeably) can be viewed as a realization of the random vector $\mathbf{X} = [X_1 \dots X_N]^T$. All attributes are observed, so there exist no hidden variables. A (relatively small) number of missing values may exist for some attributes; this can be faced using elementary techniques, like mean imputation.

We also assume that we deal with discrete data, so $\forall i \in \{1, \dots, N\}$, X_i takes values from the discrete alphabet $\text{domain}(X_i)$. If continuous attributes exist, we propose to discretize them using any discretization technique; in doing so, we are able to accurately estimate quantities like the mutual information, that play a major role in our solution. On the contrary, estimating such quantities for continuous data is a much harder problem.

We next introduce some notation, that will allow us to formally describe our model.

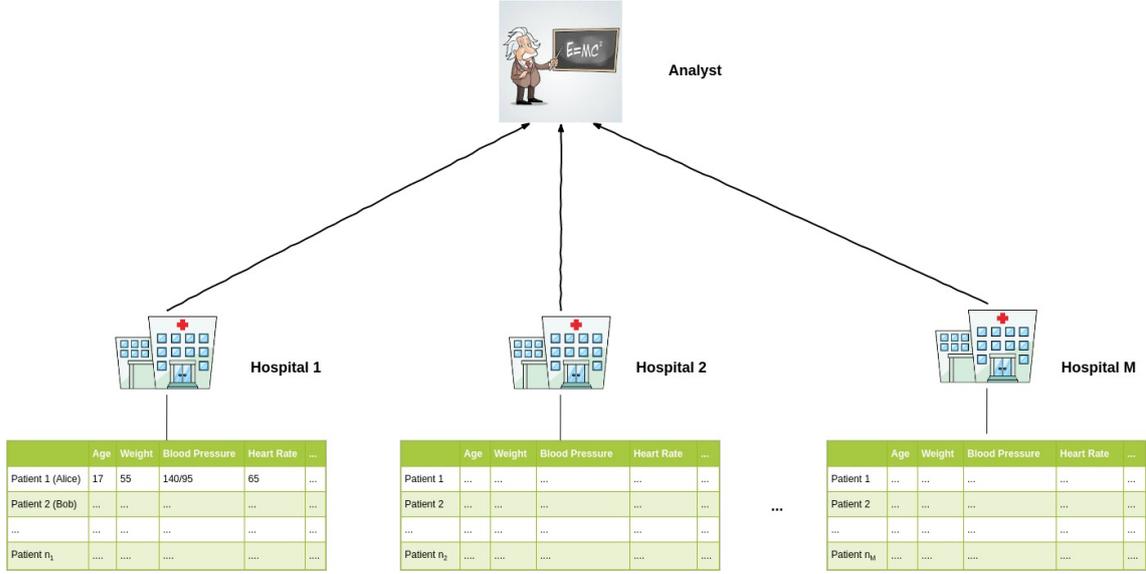


Figure 3.1: Motivating application

Table 3.1: Table of notations

Notation	Description
N	Data dimension - number of attributes per data point
X_i	Attribute i ($i \in \{1, \dots, N\}$)
\mathcal{A}	Set of all attributes: $\mathcal{A} = \{X_1, \dots, X_N\}$
\mathbf{x}	Data point: $\mathbf{x} = (x_1, \dots, x_N)$ ($\forall i \in \{1, \dots, N\} : x_i \in \text{domain}(X_i)$)
M	Number of data holders (degree of distribution)
D_j	Dataset of holder j ($j \in \{1, \dots, M\}$)
n_j	Size of dataset of holder j ($j \in \{1, \dots, M\}$): $n_j = D_j $
n	Size of total dataset: $n = \sum_{j=1}^M n_j$

Therefore, for some $i \in \{1, \dots, N\}$, by X_i we refer to the i -th attribute (random variable) in \mathcal{A} , while x_i denotes a realization of X_i . Accordingly, by \mathbf{X} we refer to the random vector that consists of all N attributes in \mathcal{A} , while \mathbf{x} denotes a realization of \mathbf{X} (a data point). In addition, for some $j \in \{1, \dots, N\}$ and $j \neq i$, consider the set $\mathcal{A}' = \{X_i, X_j, \dots\} \subseteq \mathcal{A}$. By $\mathbf{X}_{\mathcal{A}'}$ we refer to the random vector that consists of all attributes in \mathcal{A}' , while $\mathbf{x}_{\mathcal{A}'}$ denotes a realization of $\mathbf{X}_{\mathcal{A}'}$, which is in fact a tuple that consists of the values that the attributes in \mathcal{A}' take in data point \mathbf{x} . So, for instance, both \mathbf{X} and $\mathbf{X}_{\mathcal{A}}$ refer to the same random vector, and \mathbf{x} , $\mathbf{x}_{\mathcal{A}}$ both refer to realizations of this random vector.

Finally, we introduce the notation that we use for frequency and probability distributions. We denote by $d_{\mathcal{A}'}$ the Cartesian product of the domains of the attributes in \mathcal{A}' :

$$d_{\mathcal{A}'} = \text{domain}(X_i) \times \text{domain}(X_j) \times \dots$$

Also, if $|d_{\mathcal{A}'}| = d$, and we define an arbitrary ordering over the set $d_{\mathcal{A}'}$, then for some $k \in \{1, \dots, d\}$, $d_{\mathcal{A}'}[k] = (x_i, x_j, \dots)$ represents the k -th element in $d_{\mathcal{A}'}$, and is, in fact, one of the d possible realizations of the random vector $\mathbf{X}_{\mathcal{A}'}$. Given a dataset D , the joint

frequency distribution of the attributes in \mathcal{A}' is defined (and computed) as:

$$\mathbf{c}_{\mathbf{X}_{\mathcal{A}'}} = \mathbf{c}_{\mathcal{A}'} = [c_1 \dots c_d]^T, \quad \text{where } c_k = \sum_{\mathbf{x} \in D} \mathbf{1}(\mathbf{x}_{\mathcal{A}'} = d_{\mathcal{A}'}[k])$$

Since we view the attributes as random variables, we make the assumption that there exists an underlying probability distribution (that generated D); we denote the joint probability distribution of the attributes in \mathcal{A}' by:

$$\mathbf{p}_{\mathbf{X}_{\mathcal{A}'}} = \mathbf{p}_{\mathcal{A}'} = [p_1 \dots p_d]^T, \quad \text{where } p_k = \mathbb{P}(\mathbf{X}_{\mathcal{A}'} = d_{\mathcal{A}'}[k])$$

Of course, in practice we only have the data in D ; $\mathbf{p}_{\mathcal{A}'}$ is just a model that lives in our head. The maximum likelihood estimate of $\mathbf{p}_{\mathcal{A}'}$ is computed from $\mathbf{c}_{\mathcal{A}'}$ as:

$$\hat{\mathbf{p}}_{\mathcal{A}'} = \frac{1}{|D|} \mathbf{c}_{\mathcal{A}'}$$

In general, using $\hat{\cdot}$ we refer to the maximum likelihood (empirical) estimate of the underlying quantity, and using $\tilde{\cdot}$ we refer to a perturbed (noisy) estimate.

Finally, we assume that there exist no zero-probability events, so $p_i > 0, \forall i \in \{1, \dots, d\}$, although it is possible not to observe an event, so that $c_i = 0$, for some $i \in \{1, \dots, d\}$.

3.1 Differential Privacy in the Distributed Model

Recall that, in the *traditional differential privacy model*, it is assumed that there exists a trusted entity by the data owners, whose sensitive data are in the dataset, called the data holder or curator. This entity has direct access to the private dataset and analyzes it, ensuring that any output produced by the analysis satisfies differential privacy. Therefore, the data holder and the data analyst is a common (and trusted) entity.

This is not the case in several real-world applications; the data owners often do not trust the entity that collects and analyzes their data, so the traditional model needs to be modified. The *local model*, introduced by Kasiviswanathan et al. [41], is motivated by such situations, and is, in fact, a generalization of randomized response. Each data owner maintains its own database (of size 1 - a single data point), and provides the analyst with noisy answers to queries about it (satisfying differential privacy).

In our model, there are two additional changes. Firstly, the data holder and the data analyst are *separate entities*. The former is trusted by the data owners (e.g. hospitals), whereas the latter is not. Secondly, we assume that there exist *multiple data holders* (all trusted), each maintaining the sensitive data that belong to a subset of the data owners. We also assume that these subsets are non-overlapping, so each data owner's record is stored by exactly one data holder. Therefore, although the local model could apply to our model, we can avoid having each data owner perturb its own record, and make the data holders responsible for the perturbation. In particular, each trusted data holder collects its subset of the sensitive data from the data owners and either responds to queries, or performs arbitrary analyses on them; critically, the answers given must satisfy differential privacy (using the standard definition), and the overall privacy budget consumed must meet the privacy requirements. Then, the analyst only gets to see these answers, and since differential privacy is immune to post-processing, the owners' privacy is preserved.

Several efforts in the differential privacy literature have addressed this, or closely related models. A line of work attempts to combine differential privacy with *secure multiparty computation*. For example, Pathak et al. [48] propose a method for the analyst to learn a global, differentially private classifier from locally trained classifiers, which are

perturbed and aggregated using a secure protocol. Alhadidi et al. [2] address the problem of privacy-preserving data publishing, in the special case that the data are horizontally partitioned among two parties; they develop a two-party protocol for the exponential mechanism, and their solution relies on generalization. Goryczka et al. [33] consider the problem of secure sum aggregation in the distributed model (using various secure multiparty computation protocols and encryption schemes), while preserving differential privacy for the aggregated data (using distributed versions of the Laplace mechanism).

The main limitation of these hybrid solutions is that they sacrifice some of the most significant advantages of differential privacy; for example, they introduce the need of having a secret key, and they require making assumptions about the computational power of the adversary. To avoid such limitations, we focus on achieving *pure differential privacy*. We identify two paths we could follow to address our model:

- **Distributed data mining with differential privacy.** The (untrusted) analyst develops a distributed algorithm, so it communicates with the (trusted) data holders during its execution. The data holders respond to the analyst’s queries, ensuring that their answers satisfy differential privacy.

This approach can be viewed as a collaborative approach, since a global model is jointly learned, and its major advantage is that the entire dataset is directly utilized.

- **Data publishing with differential privacy.** Each (trusted) data holder constructs a model (e.g. a probabilistic graphical model) or synopsis (e.g. a histogram) of its dataset, which it then publishes, or uses to generate and publish a synthetic dataset. If the published model/synopsis/dataset satisfies differential privacy, then any analysis performed on it will also guarantee differential privacy, since differential privacy is immune to post-processing. The (untrusted) analyst collects all models/synopses/datasets, merges them, and runs a centralized, non-private algorithm on the merged result.

Clearly, the major advantage of this approach is that the published datasets can be used for arbitrary analyses. Zheng [68] investigates (among others) the implications of publishing the model versus publishing a synthetic dataset (generated using the model).

3.1.1 Distributed Data Mining with Differential Privacy

Most efforts on distributed data mining with differential privacy are task-specific, so the result of the distributed, differentially private algorithm is, for instance, a privacy-preserving *classifier*.

The parallelizable nature of *stochastic gradient descent* facilitates the design of distributed machine learning algorithms that are based on this optimization method. Rajkumar et al. [51] build a global classifier, not by combining locally trained classifiers (like Pathak et al. [48]), but instead, by directly optimizing the overall multiparty objective; their solution is based on Gaussian objective perturbation and guarantees (ϵ, δ) -differential privacy. Shokri et al. [54] propose a method that allows multiple data owners to jointly learn a neural network model, without even sharing their sensitive datasets; each data owner only shares small subsets of its models parameters during training. Direct privacy leakage is prevented, since no sensitive dataset is shared, and indirect privacy leakage, which could occur by sharing the parameters, is also prevented by using the sparse vector technique (a primitive differentially private mechanism).

Some authors consider a slightly modified model; they assume the existence of a *public dataset*, originating from data owners who are willing to share their data, The public dataset is used to enhance the utility of their analyses, which are all classification-related. Ji et al. [40] develop a method to train a differentially private logistic regression model

from distributed data. Their solution is based on properly modifying the Newton-Raphson method, which they use to solve the underlying optimization problem. Xie et al. [62] propose an ensemble learning method to aggregate locally trained binary classifiers and regressors. The local models are trained in a differentially private manner (using the expected risk minimization technique of Chaudhuri et al. [10]), and the aggregation method is based on using the public data. Hamm et al. [34] also investigate the problem of aggregating locally trained classifiers (that guarantee differential privacy), but they assume that the public dataset consists of auxiliary unlabeled data. Their main contribution is in the aggregation method; they demonstrate the limitations of using majority voting, and therefore propose a new risk, weighted by class probabilities (estimated from the ensemble).

Wahab et al. [1] examine a different data mining task, and in particular, they propose a distributed *association rules mining* framework. Their solution is based on having the data owners anonymize their data using a simple differentially private scheme.

3.1.2 Data Publishing with Differential Privacy

As we already noted, the main advantage of this approach is that it produces a general and query/task independent result that can be used for arbitrary analyses. Hence, a vast literature has been developed on data publishing with differential privacy. An important line of work is based on constructing and publishing differentially private *synopses* of the input dataset. Barak et al. [3] examine the release of contingency tables (frequency distributions) via Fourier decompositions. Hay et al. [36] and Xu et al. [63] investigate the similar problem of releasing histograms, whereas Ding et al. [19] work with data cubes. Other synopses have also been used, like spatial decompositions (Cormode et al. [16]), and wavelet transforms, which allow to more accurately answer range queries (Xiao et al. [61]). Cormode et al. [16] examine more sophisticated summarization techniques, based on sampling, filtering and sketching.

The first connection between differential privacy and *probabilistic inference* is due to Williams et al. [60]; they apply probabilistic inference to the noisy data, and, taking into account that the perturbation process is known, they attempt to estimate the parameters of the model that generated the data. Dimitrakakis et al. [18] examine the connections between differential privacy and Bayesian inference by introducing a differentially private mechanism based on posterior sampling. Several subsequent efforts build upon their ideas (e.g. Zhang et al. [67], Foulds et al. [27], Bernstein et al. [5]).

The state of the art solution in data publishing with differential privacy is *PrivBayes*, introduced by Zhang et al. [64],[65]. The authors identify that the main problem in publishing high-dimensional data (that consist of a relatively large number of attributes $N = |\mathcal{A}|$) with differential privacy is that the perturbation required inevitably overlaps the signal in the data. By high-dimensional, we refer to data whose domain size $|d_{\mathcal{A}}|$ is comparable with the total number of data points n . The proposed solution, namely PrivBayes, is inspired from the theory of probabilistic graphical models, and is based on learning the Bayesian Network (directed graphical model) that best fits the data, while satisfying differential privacy. The learned Bayesian Network provides an approximation of the high-dimensional data distribution as a product of low-order conditional and marginal distributions; these low-order distributions contain much more compact signal that is not severely damaged by the required perturbation. Finally, a synthetic dataset is published by sampling tuples (data points) from the approximate distribution. As we already noted in the introduction of this chapter, our solution is also based on this idea, and is strongly inspired by PrivBayes.

Following PrivBayes, Chen et al. [11] develop a sampling-based framework to explore the dependencies among all attributes and build a dependency graph, and then approxi-

mate the data distribution based on the junction tree algorithm. However, the version of the sparse vector technique they use to satisfy differential privacy is shown to be flawed. Ping et al. [50] implement PrivBayes as part of their proposed DataSynthesizer, a software tool that takes a sensitive dataset as input and outputs a synthetic, statistically similar dataset.

3.1.3 Our Approach: Distributed Data Publishing with Differential Privacy

We aim to combine the two approaches, in that:

- we jointly learn a global, differentially private model (specifically a Bayesian Network) utilizing the entire dataset, and
- we publish either the model itself, or a dataset generated from the model, so that it can be used for arbitrary analyses.

Thus, our approach can be characterized as a *distributed data publishing with differential privacy* approach. One of the solutions that we offer is based on a purely distributed algorithm, in that the data holders incrementally respond to the analyst’s queries during the execution of the algorithm, and a single (global) model is learned in a distributed fashion. The other two solutions are hybrid, as the information that each data holder shares is a synopsis/model of its local dataset, which the analyst utilizes to learn a global model.

In a recent work, Su et al. [56] also develop a distributed version of PrivBayes, in order to privately learn (at once) a (global) Bayesian Network from distributed data, and then use it to publish a synthetic dataset. However, their solution does not meet the requirements we have posed, in two ways. Firstly, they assume the existence of a semi-trusted curator, an intermediate entity that assists the data holders to collectively learn the global model. In particular, the data holders and the curator collaboratively identify the Bayesian Network that best fits the integrated dataset D in a sequential manner; their key contribution is in the construction of the search frontier, which consists of the set of candidate edges to add to the Bayesian Network in the next update step. Secondly, their solution is based on the distributed (multi-party) version of the Laplace mechanism (Pathak et al. [48]), which requires the use of cryptography.

3.2 Bayesian Networks

In this section, we provide the necessary background for Bayesian Networks, which constitute the basic building block of our solution.

3.2.1 Basic Definitions

A Bayesian Network (Pearl [49]) is a probabilistic graphical model, and in particular a directed graphical model, that defines a family of joint probability distributions over a set of attributes/random variables. More formally, we give the following definitions, following Koller and Friedman [44]

Definition 3.1 (Bayesian Network structure). *A Bayesian Network structure \mathcal{G} is a directed acyclic graph whose nodes represent attributes X_1, \dots, X_N from a set \mathcal{A} . Let $\mathbf{Pa}(X_i)$ denote the parents of X_i in \mathcal{G} and $\mathbf{NonDesc}(X_i)$ denote the set of attributes that are not descendants of X_i in \mathcal{G} . Then \mathcal{G} encodes the following set of conditional independence assumptions:*

$$(X_i \perp \mathbf{NonDesc}(X_i)) \mid \mathbf{Pa}(X_i) , \quad \forall i \in \{1, \dots, N\}$$

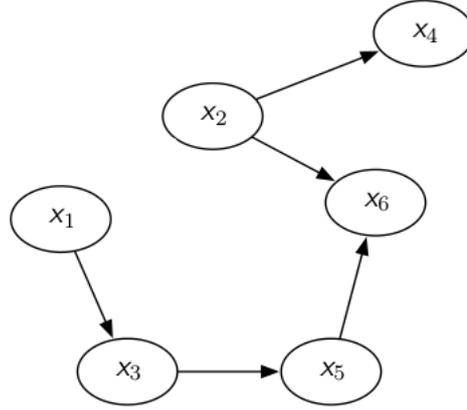


Figure 3.2: Example Bayesian Network structure

An example Bayesian network structure is illustrated in Figure 3.2. Using this example, we identify the types of connections between nodes, and the (in)dependence relations they imply for the corresponding attributes:

- $X_1 \rightarrow X_3$: direct dependence, $X_3 \not\perp X_1$.
- $X_1 \rightarrow X_3 \rightarrow X_5$: indirect causal effect, $X_5 \not\perp X_1$, but $(X_5 \perp X_1)|X_3$.
- $X_5 \leftarrow X_3 \leftarrow X_1$: indirect evidential effect, $X_1 \not\perp X_5$, but $(X_1 \perp X_5)|X_3$.
- $X_6 \leftarrow X_2 \rightarrow X_4$: common cause, $X_4 \not\perp X_6$, but $(X_4 \perp X_6)|X_2$.
- $X_5 \rightarrow X_6 \leftarrow X_2$: common effect, $X_5 \perp X_2$, but $(X_5 \not\perp X_2)|X_6$.

Definition 3.2 (Factorization). Let \mathcal{G} be a Bayesian Network structure over a set of attributes $\mathcal{A} = \{X_1, \dots, X_N\}$. The distribution $\mathbf{p}_{\mathcal{A}}$ factorizes according to \mathcal{G} if $\mathbf{p}_{\mathcal{A}}$ can be expressed as a product:

$$\mathbf{p}_{\mathcal{A}} = \prod_{i=1}^N \mathbf{p}_{X_i | \mathbf{Pa}(X_i)}$$

Each individual factor $\mathbf{p}_{X_i | \mathbf{Pa}(X_i)}$ represents the set of all conditional probability distributions of X_i , one for each realization of the attributes in $\mathbf{Pa}(X_i)$, and is called a conditional probability table.

We denote the set of all conditional probability tables by: $\Theta = \{\Theta_1, \dots, \Theta_N\}$, where $\Theta_i = \mathbf{p}_{X_i | \mathbf{Pa}(X_i)}$, $\forall i \in \{1, \dots, N\}$.

Definition 3.3 (Bayesian Network). A Bayesian Network over a set of attributes \mathcal{A} is a pair $\mathcal{B} = (\mathcal{G}, \Theta)$, where $\mathbf{p}_{\mathcal{A}}$ factorizes according to \mathcal{G} and is specified based on the conditional probability distributions in Θ .

Definition 3.4 (Bayesian Network degree). Let $\mathcal{B} = (\mathcal{G}, \Theta)$ be a Bayesian Network over a set of attributes $\mathcal{A} = \{X_1, \dots, X_N\}$. We define the degree k of \mathcal{B} as the in-degree \mathcal{G} :

$$k = \max_{i \in \{1, \dots, N\}} |\mathbf{Pa}(X_i)|$$

Based on Definitions 3.1, 3.2, 3.3, we see that *any* joint probability distribution over a set of attributes \mathcal{A} can be represented using a Bayesian Network (by the *chain rule*). In practice, we are particularly interested in distributions that factorize according to *sparse*

graphs, that lead to compact representations and efficient inference.

Example. Consider a probability distribution \mathbf{p} over N attributes $\{X_1, \dots, X_N\}$, such that $\text{domain}(X_i) = d$, $\forall i \in \{1, \dots, N\}$. Then, \mathbf{p} is originally represented using d^N parameters, each corresponding to the probability of a specific realization of the random vector \mathbf{X} . A Bayesian Network allows us to represent \mathbf{p} using only

$$\sum_{i=1}^N |\Theta_i| = \sum_{i=1}^N d^{|\text{Pa}(X_i)|+1} = \mathcal{O}(Nd^k)$$

parameters, which can be substantially smaller (usually $k \ll N$).

The time complexity of both making *inferences* using \mathbf{p} , and *sampling* from \mathbf{p} , can significantly decrease. More specifically,

- By inference, we basically refer to two tasks. Assume that we observe a subset \mathcal{A}_1 of the attributes in \mathcal{A} , and let $\mathcal{A}_2 = \mathcal{A} \setminus \mathcal{A}_1$. The first task is the calculation of posterior probabilities, namely $\mathbf{p}_{\mathcal{A}_2|\mathcal{A}_1}$. The second task is the calculation of the most probable configuration (realization) for the unobserved attributes, namely $\text{argmax}_{\mathbf{x}_{\mathcal{A}_2} \in d_{\mathcal{A}_2}} \mathbf{p}_{\mathcal{A}_2|\mathcal{A}_1}$. As an example, in case the unobserved attributes correspond to a set of class attributes, calculating the most probable configuration provides a solution for the corresponding classification problem.
- As far as sampling is concerned, efficient algorithms exist, and probably the simplest among them is *prior sampling* (which we utilize throughout this work). Instead of sampling from the N -dimensional joint distribution \mathbf{p} , we separately draw a sample for each attribute X_i ($\forall i \in \{1, \dots, N\}$) from the proper conditional distribution in Θ_i . To be able to do so, we need to ensure that, by the time we draw a sample for X_i , we have already sampled for all attributes in $\text{Pa}(X_i)$; this is achieved by topologically ordering the attributes based on \mathcal{G} .

The power of Bayesian Networks is twofold. Besides the *computational gain* they offer, by allowing us to approximate a high-dimensional distribution as a product of low-order conditional and marginal distributions, Bayesian Networks are also highly *interpretable* models, and can hence be used for knowledge discovery.

3.2.2 Learning Bayesian Networks

In real-world applications, we only have access to data, which we assume that follow a distribution \mathbf{p} over a set $\mathcal{A} = \{X_1, \dots, X_N\}$ of attributes, that can be encoded by a Bayesian Network. Sometimes the Bayesian Network structure is known; for example, it may be given by an expert, or it may be determined by the physical properties of the application. However, our focus is on the case that the structure is unknown, and our objective is to find both the structure and the parameters of the Bayesian Network that best fits our dataset D . As we already noted in the introduction of this chapter, we make the assumption that our data are fully-observed, in the sense that no hidden variables exist.

In general, there are three approaches to learning Bayesian Networks of unknown structure from data, namely:

- score-based, where the Bayesian Network learning problem is viewed as a model selection problem and is solved using optimization methods,
- constraint-based, where the Bayesian Network is viewed as a representation of independences and the learning process is based on independence tests, and

- Bayesian model averaging, where an ensemble of possible structures is learned and aggregated.

We focus on the score-based approach, which we adopt in our work.

The first step in *score-based Bayesian Network learning* is to assign each possible Bayesian Network \mathcal{B} a score. A natural choice for *scoring function* is the likelihood function \mathcal{L} , which measures the probability of observing the data in D (assuming that individual data points are i.i.d.), given a model. In our case, the model is a Bayesian Network $\mathcal{B} = (\mathcal{G}, \Theta)$, so, denoting by $\mathbf{p}_{\mathcal{B}}$ the distribution encoded by \mathcal{B} , we conclude that the score of \mathcal{B} is:

$$\mathcal{L}(\mathcal{B}; D) = \prod_{\mathbf{x} \in D} \mathbf{p}_{\mathcal{B}}(\mathbf{x}) = \prod_{\mathbf{x} \in D} \prod_{i=1}^N \mathbf{p}_{X_i | \mathbf{Pa}(X_i)}(x_i, \mathbf{x}_{\mathbf{Pa}(X_i)}) = \prod_{i=1}^N \prod_{\mathbf{x} \in D} \Theta_i(x_i, \mathbf{x}_{\mathbf{Pa}(X_i)})$$

which illustrates the decomposability of the global likelihood into local likelihoods (one for each parameter Θ_i), based on the Bayesian Network structure. If we instead use the logarithm of \mathcal{L} as scoring function, it turns out that:

$$\ell(\mathcal{B}; D) = \log \mathcal{L}(\mathcal{B}; D) = n \sum_{i=1}^N \hat{I}(X_i; \mathbf{Pa}(X_i)) - n \sum_{i=1}^N \hat{H}(X_i)$$

where by \hat{H} and \hat{I} we refer to the empirical entropy and mutual information respectively.

The next step is to find the Bayesian Network structure \mathcal{G} that achieves the highest score. Noting that the only term in the log-likelihood score that depends on \mathcal{G} is the empirical mutual information, gives that:

$$\max_{\mathcal{B}} \ell(\mathcal{B}; D) = \max_{\mathcal{G}} \underbrace{\max_{\Theta} \ell(\mathcal{G}, \Theta; D)}_{\text{MLE of } \Theta \text{ given } \mathcal{G}} = \max_{\mathcal{G}} \ell(\mathcal{G}, \hat{\Theta}; D) = \max_{\mathcal{G}} \sum_{i=1}^N \hat{I}(X_i; \mathbf{Pa}(X_i))$$

which demonstrates that the structure \mathcal{G}^* that maximizes the log-likelihood score is the one that has the maximum sum over all attributes, of the mutual information between each attribute and its parents in \mathcal{G}^* , when the maximum likelihood parameters $\hat{\Theta}$ are used for \mathcal{G}^* . From an information theoretic point of view, it can be shown that \mathcal{G}^* minimizes the Kullback-Leibler divergence between the actual data distribution, and the distribution encoded by the Bayesian Network structure \mathcal{G}^* .

Solving the aforementioned optimization problem turns out to be hard in general. Specifically, taking into account that our graph consists of N nodes, there are $\mathcal{O}(2^{N^2})$ potential structures that form our search space (super-exponential in the number of attributes). Consequently, heuristic (local-search) algorithms are employed in practice, like hill-climbing and simulated annealing. Nevertheless, for the special case that $k = 1$ (Bayesian Networks with degree 1), the well-known Chow-Liu algorithm [12] (Algorithm 3.1) allows us to greedily find the optimal structure.

An important limitation of the likelihood score is that it favors more complex structures over simpler ones. In fact, if we do not artificially constrain the number of parents allowed for each attribute (as in the case of the Chow-Liu algorithm), then any algorithm using the likelihood score will almost always return a fully connected structure. This results from the fact that adding an additional parent to any attribute will almost always increase the score, since for any random variables X, Y, Z , we have $\hat{I}(X; Y, Z) \geq \hat{I}(X; Y)$ (information never hurts). Therefore, in practice, we either learn fixed-degree Bayesian Networks (this is the case in PrivBayes), or we penalize more complex structures depending on the dataset size (e.g. BIC score).

Algorithm 3.1: Chow-Liu Algorithm

<p>Input: Dataset D Output: BN structure \mathcal{G}</p> <ol style="list-style-type: none"> 1 Initialize \mathcal{G}_0 to a fully connected (undirected) graph 2 for $i=1$ to N do 3 Estimate (and store) $\hat{\mathbf{p}}_{X_i}$ from D 4 for $j=1$ to $i-1$ do 5 Estimate (and store) $\hat{\mathbf{p}}_{X_j, X_i}$ from D 6 Compute weight $\hat{I}(X_j; X_i)$ of edge (X_j, X_i) in \mathcal{G}_0 7 $\mathcal{G} = \text{MaximumWeightSpanningTree}(\mathcal{G}_0)$ 8 Give directions to edges in \mathcal{G} 9 Return \mathcal{G}
--

Once the Bayesian Network structure \mathcal{G} is known, we are left with a simple parameter estimation task; we have to estimate the parameters in Θ . A common approach to achieve this is to (once again) use the maximum likelihood principle. The decomposability of the likelihood function, which we demonstrated earlier, allows us to maximize each local likelihood $\mathcal{L}_i(\Theta_i; D) = \prod_{\mathbf{x} \in D} \Theta_i(x_i, \mathbf{x}_{\text{Pa}(X_i)})$ separately, and then combine the solutions to get the (global) maximum likelihood estimate $\hat{\Theta}$ for Θ :

$$\hat{\Theta} = \{\hat{\Theta}_1, \dots, \hat{\Theta}_N\}, \text{ where } \hat{\Theta}_i = \text{argmax}_{\Theta_i} \mathcal{L}_i(\Theta_i; D)$$

We described a purely frequentist approach to learning Bayesian Networks from data, in that we utilize the likelihood score to learn the structure and then estimate the parameters that maximize the likelihood function. An alternative path would be to adopt a Bayesian approach, and treat the candidate structures (in the structure learning phase) or the parameters (in the parameter learning phase) as random; we would assign them prior distributions, and we would maximize the posterior distribution of the data, instead of the likelihood. We do not further discuss the Bayesian approach here, but we mention that its main advantage over the frequentist approach is that it avoids overfitting.

3.3 PrivBayes

In this section, we provide an overview of PrivBayes. We also highlight the changes that need to be performed in our model, as well as the assumptions we make. Apparently, a first major change originates from the fact that PrivBayes addresses the traditional differential privacy model, where the data holder and the analyst are the same, trusted entity. Therefore, the original solution depends on a centralized algorithm (the single dataset D is analyzed by a trusted entity), and the only requirement is that the resulting Bayesian Network satisfies differential privacy (both the structure and the parameters).

PrivBayes consists of three phases, which we briefly present here in order to be able argue about the privacy guarantees of the overall algorithm. In the next subsections, we present each of the first two phases separately; the third phase is trivial, so we fully develop it here.

- **Structure learning phase.** During this phase, the analyst accesses the sensitive data in D . It uses an ϵ_1 -differentially private algorithm to learn the structure of a k -degree Bayesian Network that accurately encodes the conditional independencies that are present in the underlying data distribution.

- **Parameter learning phase.** During this phase, the analyst utilizes the learned structure, and again accesses the sensitive data in D . It uses an ε_2 -differentially private algorithm to estimate the parameters of the learned k -degree Bayesian Network.
- **Synthetic data generation phase.** During this phase, no access to the sensitive data is performed. The analyst utilizes the learned model to sample n' tuples (data points) from the approximate distribution (encoded by the Bayesian Network). The sampling step is performed using prior sampling, and does not access the sensitive data - only the output of the first two phases is used. Since (as we argue in theorem 3.1) the output of the first two phases satisfies differential privacy, according to the post-processing theorem, so does the result of the third phase, and no additional perturbation is required.

We note that the value of n' depends on the application, but a reasonable choice would be to set $n' = n$. Also, in case we are interested in the model, and not in publishing a synthetic dataset, the data generation phase may even be skipped.

Theorem 3.1. *Let ε_1 and ε_2 be the privacy budget consumed during each of the first two phases of PrivBayes respectively. Then, the overall algorithm satisfies $(\varepsilon_1 + \varepsilon_2)$ -differential privacy.*

Theorem 3.1 directly follows from the composition theorem of differential privacy. The choice of ε_1 and ε_2 determines the balance between the quality of the learned structure and the learned parameters. Assuming that the total available privacy budget is ε , and by setting $\varepsilon_1 = \beta\varepsilon$ and $\varepsilon_2 = (1 - \beta)\varepsilon$, the authors of PrivBayes experimentally conclude that the optimal split is achieved for some $\beta \in [0.2, 0.5]$. For simplicity, in our work we assume that, whenever the privacy budget has to be split between the two phases, $\varepsilon_1 = \varepsilon_2 = \frac{\varepsilon}{2}$, so the overall algorithm satisfies ε -differential privacy.

3.3.1 Structure Learning Phase

The structure learning phase of PrivBayes is based on a greedy extension of the Chow-Liu algorithm (Algorithm 3.1) for higher degree Bayesian Networks, which we present in Algorithm 3.2. The authors of PrivBayes propose the aforementioned approach, instead of the local-search algorithms that are commonly used for learning the structure of higher degree Bayesian Networks, because these algorithms perform too many data accesses and incur a high cost in terms of sensitivity; as a result, a prohibitive amount of noise is added.

Lemma 3.1 quantifies the sensitivity of the empirical mutual information; the proof can be found in the full PrivBayes paper [65]. We remark that the sensitivity is computed using the definition of adjacent datasets that leads to *bounded* differential privacy, so $\Delta\hat{I}$ expresses the maximum change in the empirical mutual information that is caused by changing the value of one data point.

Lemma 3.1. *For any random variables X and Y , the sensitivity of $\hat{I}(X; Y)$ is:*

$$\Delta\hat{I} = \begin{cases} \frac{1}{n} \log n + \frac{n-1}{n} \log \frac{n}{n-1} & \text{if } X \text{ or } Y \text{ is binary} \\ \frac{2}{n} \log \frac{n+1}{2} + \frac{n-1}{n} \log \frac{n+1}{n-1} & \text{otherwise} \end{cases}$$

The output \mathcal{G} of Algorithm 3.2 consists of the $N - 1$ attribute-parent pairs each of which was added using the exponential mechanism (line 12) with the empirical mutual information as scoring function and with privacy budget $\frac{\varepsilon_1}{N-1}$. Thus, Theorem 3.2 directly follows (by the composition theorem).

Theorem 3.2. *Algorithm 3.2 satisfies ε_1 -differential privacy.*

Algorithm 3.2: Greedy BN Structure Learning	
Input:	Dataset D , BN degree k , Privacy budget ϵ_1
Output:	BN structure \mathcal{G}
1	Initialize $\mathcal{G} = \emptyset$ and $V = \emptyset$
2	for each $\mathcal{A}' \subseteq \mathcal{A}$ such that $ \mathcal{A}' \leq k + 1$ do
3	Estimate (and store) $\hat{\mathbf{p}}_{\mathcal{A}'}$ from D
4	Arbitrarily select an attribute X from \mathcal{A}
5	Add (X, \emptyset) to \mathcal{G} and X to V
6	for $i=1$ to $N-1$ do
7	Initialize $\Omega = \emptyset$
8	for each $X \in \mathcal{A} \setminus V$ do
9	for each $\mathbf{Pa}(X) \subseteq V$ such that $ \mathbf{Pa}(X) \leq k$ do
10	Compute $\hat{I}(X; \mathbf{Pa}(X))$ using $\hat{\mathbf{p}}_X$, $\hat{\mathbf{p}}_{\mathbf{Pa}(X)}$, and $\hat{\mathbf{p}}_{X, \mathbf{Pa}(X)}$
11	Add $(X, \mathbf{Pa}(X), \hat{I}(X; \mathbf{Pa}(X)))$ to Ω
12	Sample a tuple $(X, \mathbf{Pa}(X), \hat{I}(X; \mathbf{Pa}(X)))$ from Ω w/ prob. $\propto \frac{\hat{I}(X; \mathbf{Pa}(X)) \epsilon_1}{2(N-1)\Delta \hat{I}}$
13	Add $(X, \mathbf{Pa}(X))$ to \mathcal{G} and X to V
14	Return \mathcal{G}

How to pick k . Based on our earlier discussion on Bayesian Networks, it is not hard to notice that, excluding the first k attributes that are inserted in \mathcal{G} , all other $N - k$ attributes will have exactly k parents (see Section 3.2.2 - information never hurts). A natural question is how to pick k in practice. To answer this question, the authors of PrivBayes introduce the notion of θ -*usefulness*, which enables PrivBayes to automatically select k . In particular, the user-specified parameter θ expresses the minimum-allowed ratio of average scale of signal to average scale of noise in each entry in the conditional probability tables of the resulting k -degree Bayesian Network. So, for example, when the privacy budget is small, PrivBayes will favor smaller values for k , so that the dimension of the resulting conditional probability distributions is small enough and the signal contained is not dominated by the noise. Although this notion could apply to our model as well, we focus on different aspects of PrivBayes, and, hence, assume that k is known in advance.

Improving the scoring function. A key contribution of PrivBayes is that the authors identify a significant limitation in using the empirical mutual information as scoring function for the exponential mechanism. Specifically, although $\lim_{n \rightarrow \infty} \Delta \hat{I} = 0$ (for both binary and non-binary domains), $\Delta \hat{I} > \frac{\log n}{n}$, so the empirical mutual information can be quite large compared to the range of \hat{I} , and, as a result, the exponential mechanism cannot distinguish between high-scoring and low-scoring attribute-parent pairs. Recall that the probability of sampling a particular pair is $\propto \frac{\hat{I}(X; \mathbf{Pa}(X)) \epsilon_1}{2(N-1)\Delta \hat{I}}$, so when $\hat{I}(X; \mathbf{Pa}(X))$ and $\Delta \hat{I}$ are comparable for many candidate pairs $(X; \mathbf{Pa}(X))$, their sampling probabilities tend to become uniform. PrivBayes tackles this limitation by introducing two alternative scoring functions, that have small sensitivities compared to their ranges, but, at the same time, their behavior is similar with that of the empirical mutual information.

- The first, namely F , relies on the notion of maximum joint distributions and has sensitivity $\Delta F = \frac{1}{n}$. It can be efficiently computed only for binary domains.
- The second, namely R , relies on the L_1 distance between $\hat{\mathbf{p}}_{X, \mathbf{Pa}(X)}$ and $\hat{\mathbf{p}}_X \hat{\mathbf{p}}_{\mathbf{Pa}(X)}$ (the distribution that minimizes the empirical mutual information), and has sensitivity $\Delta R = \frac{3}{n} + \frac{2}{n^2}$. It can be efficiently computed for arbitrary domains.

This optimization directly applies to all our proposed solutions. Nevertheless, we choose to maintain the empirical mutual information as scoring function (in all the solutions we examine), because, firstly, it is a well-understood function, and secondly, it facilitates our theoretical analysis.

3.3.2 Parameter Learning Phase

The parameter learning phase of PrivBayes also follows the frequentist approach we described in Section 3.2.2.

Algorithm 3.3: BN Parameter Learning

Input: Dataset D , BN structure \mathcal{G} , Privacy budget ε_2
Result: BN parameters Θ

- 1 Initialize $\Theta_i = \emptyset, \forall i \in \{1, \dots, N\}$
- 2 **for** $i = 1$ **to** N **do**
- 3 Estimate $\hat{\mathbf{p}}_{X_i, \mathbf{Pa}(X_i)}$
- 4 Compute noisy $\tilde{\mathbf{p}}_{X_i, \mathbf{Pa}(X_i)} = \hat{\mathbf{p}}_{X_i, \mathbf{Pa}(X_i)} + \text{Laplace}(0, \frac{2N}{n\varepsilon_2})$
- 5 Set negative values in $\tilde{\mathbf{p}}_{X_i, \mathbf{Pa}(X_i)}$ to 0 and normalize
- 6 Compute $\tilde{\Theta}_i = \tilde{\mathbf{p}}_{X_i | \mathbf{Pa}(X_i)}$ (by marginalizing $\tilde{\mathbf{p}}_{X_i, \mathbf{Pa}(X_i)}$)
- 7 Add $\tilde{\Theta}_i$ to $\tilde{\Theta}$
- 8 **Return** $\tilde{\Theta}$

The output $\tilde{\Theta}$ of Algorithm 3.3 consists of the N conditional probability tables $\tilde{\Theta}_i$, $i \in \{1, \dots, N\}$, each of which represents the set of conditional distributions of an attribute, given all realizations of its parents. Notice that each $\tilde{\Theta}_i$ is constructed using a noisy version of the maximum likelihood estimate of the joint probability distribution of attribute X_i and its parents. The sensitivity of the maximum likelihood estimate (using n data points) is $\Delta \hat{\mathbf{p}} = \frac{2}{n}$; the proof is almost identical with that of Theorem 3.4, where we argue about the sensitivity of the frequency distribution over a set of attributes (which is the basis of the maximum likelihood estimate). In total, we estimate N noisy distributions using the Laplace mechanism (line 4), each with privacy budget $\frac{\varepsilon_2}{N}$, so (again) by the composition theorem:

Theorem 3.3. *Algorithm 3.3 satisfies ε_2 -differential privacy.*

Consistency. We remark that, although the distributions in $\tilde{\Theta}$ have no consistency issues (since each of them is the conditional distribution of a different attribute, given a different realization of its parents), the intermediate joint distributions $\tilde{\mathbf{p}}$ computed by Algorithm 3.3 are (generally) not consistent. Specifically, an arbitrary attribute X is likely to be involved in several joint distributions, and we have no guarantees that the marginals of X derived using different joint distributions will be the same. Following PrivBayes, we also ignore these consistency issues, as they have no impact on the final result.

3.3.3 PrivBayes in the Distributed Model

As we already argued, PrivBayes follows a score-based approach to learning the target Bayesian Network, and, in particular, PrivBayes finds a greedy solution to the optimization problem which we described in Section 3.2.2. We re-formulate this optimization problem, taking into account that in our case, the data are distributed among M data holders.

$$\max_{\mathcal{B}} \ell(\mathcal{B}; D_1, \dots, D_M) = \max_{\mathcal{G}} \sum_{i=1}^N \hat{I}(X_i; \mathbf{Pa}(X_i))$$

$$\begin{aligned}
&= \max_{\mathcal{G}} \sum_{i=1}^N \sum_{\substack{x \in d_{X_i} \\ \mathbf{p} \in d_{\mathbf{Pa}(X_i)}}} \left(\frac{\sum_{j=1}^M \mathbf{c}_{X_i, \mathbf{Pa}(X_i)}^{(j)}(x, \mathbf{p})}{n} \right) \\
&\quad \times \log_2 \frac{n \sum_{j=1}^M \mathbf{c}_{X_i, \mathbf{Pa}(X_i)}^{(j)}(x, \mathbf{p})}{\sum_{j=1}^M \mathbf{c}_{X_i}^{(j)}(x) \sum_{j=1}^M \mathbf{c}_{\mathbf{Pa}(X_i)}^{(j)}(\mathbf{p})}
\end{aligned}$$

The main observation is that the scoring function which we optimize is *non-linear* with respect to the frequency distributions \mathbf{c} , and, consequently, we cannot compose the global value of the scoring function from its local values. Although here we examine the empirical mutual information, our argument is true for all scoring functions introduced in PrivBayes. Notice that the *sufficient statistic* to estimate the empirical mutual information between an attribute and its parents is their joint frequency distribution, and, hence, to solve our optimization problem we need to have all the $(k+1)$ -dimensional frequency distributions (where k is the degree of the target Bayesian Network).

The key question that needs to be answered is what information each data owner shares with the (untrusted) analyst; we refer to this question as the *What to share?* question. Furthermore, depending on the answer to this question, we need to figure out how to combine the information shared by different data holders. In the next sections, we examine different answers to the *What to share?* question.

Once the Bayesian Network structure \mathcal{G} is known, we again follow a frequentist approach and adjust Algorithm 3.3 to the distributed model. As we will see, depending on the structure learning approach used, the analyst may need or need not re-access the data to estimate the Bayesian Network parameters. The latter may be the case if the analyst has already retrieved the required local distributions, or is (somehow) able to estimate them. Therefore, we introduce the boolean parameter *retrieved* that indicates whether the analyst already possesses the distributions $\tilde{\mathbf{p}}_{X_i, \mathbf{Pa}(X_i)}$, $\forall i \in \{1, \dots, N\}$.

Algorithm 3.4: Distributed BN Parameter Learning

Input: Datasets D_1, \dots, D_M , BN structure \mathcal{G} , Boolean *retrieved*
Result: BN parameters Θ

- 1 Initialize $\Theta_i = \emptyset$, $\forall i \in \{1, \dots, N\}$
- 2 **for** $i = 1$ **to** N **do**
- 3 **if** *retrieved* == *False* **then**
- 4 **for** $j = 1$ **to** M **do**
- 5 **QUERY**(D_j): retrieve local $\tilde{\mathbf{c}}_{X_i, \mathbf{Pa}(X_i)}^{(j)}$
- 6 Estimate global $\tilde{\mathbf{p}}_{X_i, \mathbf{Pa}(X_i)} = \frac{1}{n} \sum_{j=1}^M \tilde{\mathbf{c}}_{X_i, \mathbf{Pa}(X_i)}^{(j)}$
- 7 Compute $\tilde{\Theta}_i = \tilde{\mathbf{p}}_{X_i | \mathbf{Pa}(X_i)}$ (by marginalizing $\tilde{\mathbf{p}}_{X_i, \mathbf{Pa}(X_i)}$)
- 8 Add $\tilde{\Theta}_i$ to $\tilde{\Theta}$
- 9 **Return** $\tilde{\Theta}$

Notice that, in contrast to Algorithm 3.3, no perturbation is performed in Algorithm 3.4. Since the analyst is not trusted, it is the data holders' responsibility to properly perturb their local frequency distributions, and handle any negative frequencies that may appear, prior to sharing them. Although it is a straightforward application of the Laplace mechanism, we explain in detail how each data holder properly perturbs its local frequency distribution in Section 3.4.1.

3.4 Sharing the Noisy Sufficient Statistics

The first approach to answering the *What to share?* question is based on asking each data holder to share its part of the sufficient statistics, that is, all $(k + 1)$ -dimensional frequency distributions. In doing so, the analyst is able to compose the global sufficient statistics (since the empirical frequency distribution is composable by simply summing the counts), and evaluate the scoring function for all candidate structures. In that sense, this approach, which we call *Sharing the Noisy Sufficient Statistics* (Algorithm 3.5), is an exact approach.

Algorithm 3.5: Sharing the Noisy Sufficient Statistics

Input: Datasets D_1, \dots, D_M , BN degree k
Output: BN structure \mathcal{G}

- 1 Initialize $\mathcal{G} = \emptyset$ and $V = \emptyset$
- 2 **for** each $\mathcal{A}' \subseteq \mathcal{A}$ such that $|\mathcal{A}'| = k + 1$ **do**
- 3 **for** $j = 1$ **to** M **do**
- 4 **QUERY**(D_j): retrieve local $\tilde{\mathbf{c}}_{\mathcal{A}'}^{(j)}$
- 5 Estimate global $\tilde{\mathbf{p}}_{\mathcal{A}'} = \frac{1}{n} \sum_{j=1}^M \tilde{\mathbf{c}}_{\mathcal{A}'}^{(j)}$
- 6 Arbitrarily select an attribute X from \mathcal{A}
- 7 Add (X, \emptyset) to \mathcal{G} and X to V
- 8 **for** $i=1$ **to** $N-1$ **do**
- 9 Initialize $\Omega = \emptyset$
- 10 **for** each $X \in \mathcal{A} \setminus V$ **do**
- 11 **for** each $\mathbf{Pa}(X) \subseteq V$ such that $|\mathbf{Pa}(X)| \leq k$ **do**
- 12 Compute $\tilde{\mathbf{p}}_{X, \mathbf{Pa}(X)}, \tilde{\mathbf{p}}_X, \tilde{\mathbf{p}}_{\mathbf{Pa}(X)}$ by marginalizing the proper distributions, and then $\tilde{I}(X; \mathbf{Pa}(X))$
- 13 Add $(X, \mathbf{Pa}(X), \tilde{I}(X; \mathbf{Pa}(X)))$ to Ω
- 14 Select $(X, \mathbf{Pa}(X))$ with the highest $\tilde{I}(X; \mathbf{Pa}(X))$
- 15 Add $(X, \mathbf{Pa}(X))$ to \mathcal{G} and X to V
- 16 **Return** \mathcal{G}

In order to satisfy differential privacy, each data holder responds to the queries on its dataset using the Laplace mechanism. We provide a more detailed analysis of the privacy guarantees of Algorithm 3.5 in Section 3.4.1.

Once the analyst collects the frequency distributions, it does not need to access the data again. This leads to the following two significant advantages:

- First, with the (noisy) sufficient statistics to estimate the empirical mutual information at hand, any algorithm (e.g. local search) can be used to address the structure learning optimization problem. Nevertheless, as shown in Algorithm 3.5, we keep using the greedy extension of the Chow-Liu Algorithm (to fairly compare the approaches we examine).
- Second, once the structure learning phase is completed, the parameter learning Algorithm 3.4 utilizes the already-retrieved distributions, and hence the entire privacy budget can be consumed in the structure learning phase.

On the downside, each data holder has to share $\binom{N}{k+1}$ frequency distributions, which may be prohibitive for high-degree Bayesian Networks, in terms of both perturbation and communication cost. To attack this problem, we could examine more sophisticated

techniques in sharing the required frequency distributions, like the ones we presented in Section 3.1.2. Nevertheless, as Xu et al. [63] demonstrate, even the naive application of the Laplace mechanism performs sufficiently well compared to such techniques for a variety of datasets.

In the next two subsections, we provide a detailed privacy and accuracy analysis of Algorithm 3.5.

3.4.1 Privacy Analysis

We first examine the privacy guarantees of Algorithm 3.5. The following theorem quantifies the amount of noise each data holder must add to each frequency distribution it shares in order to preserve differential privacy. Recall that the frequency distribution $\mathbf{c}_{\mathcal{A}'}$ of the attributes in \mathcal{A}' can be viewed as a d -dimensional vector, where $d = |d_{\mathcal{A}'}|$.

Theorem 3.4. *Let $b = \frac{2\binom{N}{k+1}}{\varepsilon}$. If, $\forall \mathcal{A}' \subseteq \mathcal{A}$ such that $|\mathcal{A}'| = k+1$, each data holder shares $\tilde{\mathbf{c}}_{\mathcal{A}'} = \mathbf{c}_{\mathcal{A}'} + \boldsymbol{\eta}$, where $\boldsymbol{\eta} = [\eta_1 \ \eta_2 \ \dots \ \eta_d]^T$ is a random vector of i.i.d. Laplace(0, b) entries, then Algorithm 3.5 preserves ε -differential privacy for any dataset D_j ($j \in \{1, \dots, M\}$).*

Proof. Let $D, D' = D \cup \{\mathbf{x}'\} \setminus \{\mathbf{x}\}$ be two adjacent datasets of the same size, possessed by an arbitrary data holder. Assume that $\mathbf{x}_{\mathcal{A}'} = d_{\mathcal{A}'}[j]$ for some $j \in \{1, \dots, d\}$, that is, the attributes \mathcal{A}' in \mathbf{x} take the j -th value from their joint domain $d_{\mathcal{A}'}$. Accordingly, assume that $\mathbf{x}'_{\mathcal{A}'} = d_{\mathcal{A}'}[j']$ for some $j' \in \{1, \dots, d\}$. The sensitivity of the frequency distribution \mathbf{c} is:

$$\begin{aligned} \Delta \mathbf{c} &= \max_{D, D'} \|\mathbf{c}(D) - \mathbf{c}(D')\|_1 = \max_{D, D'} \sum_{i=1}^d |c_i(D) - c_i(D')| \\ &= |c_j(D) - (c_j(D) - 1)| + |c_{j'}(D) - (c_{j'}(D) + 1)| = 2 \end{aligned}$$

since removing \mathbf{x} from the dataset will cause the count of $d_{\mathcal{A}'}[j]$ to decrease by one, and adding \mathbf{x}' to the dataset will cause the count of $d_{\mathcal{A}'}[j']$ to increase by one. By simple application of the Laplace mechanism, it follows that adding i.i.d. Laplace(0, $\frac{\Delta \mathbf{c}}{\varepsilon'}$) noise to each count will preserve ε' -differential privacy.

The untrusted analyst interacts with the data only by viewing the underlying frequency distributions. In total, each data holder shares $\binom{N}{k+1}$ frequency distributions, and to compute each of them, a new data access is required. Therefore, by the composition theorem, if $\varepsilon' = \frac{\varepsilon}{\binom{N}{k+1}}$, the overall algorithm preserves ε -differential privacy. \square

3.4.2 Accuracy Analysis

To argue about the accuracy of Algorithm 3.5, we examine in detail part of the pipeline it implements.

For each $\mathcal{A}' \subseteq \mathcal{A}$ such that $|\mathcal{A}'| = k+1$, each data holder computes the (local) joint frequency distribution of the attributes in \mathcal{A}' . Therefore, if $\mathbf{c}^{(j)}$ is the distribution that holder $j \in \{1, \dots, M\}$ computes, then, assuming that $|d_{\mathcal{A}'}| = d$, holder j shares the following d -dimensional vector:

$$\tilde{\mathbf{c}}^{(j)} = \mathbf{c}^{(j)} + \boldsymbol{\eta}^{(j)} = [c_1^{(j)} \ c_2^{(j)} \ \dots \ c_d^{(j)}]^T + [\eta_1^{(j)} \ \eta_2^{(j)} \ \dots \ \eta_d^{(j)}]^T = [\tilde{c}_1^{(j)} \ \tilde{c}_2^{(j)} \ \dots \ \tilde{c}_d^{(j)}]^T$$

where $\boldsymbol{\eta}^{(j)}$ is picked as described in Theorem 3.4. Then, the analyst collects the noisy vectors $\tilde{\mathbf{c}}^{(1)}, \dots, \tilde{\mathbf{c}}^{(M)}$, it merges them as:

$$\tilde{\mathbf{c}} = \sum_{j=1}^M \tilde{\mathbf{c}}^{(j)} = \sum_{j=1}^M \mathbf{c}^{(j)} + \boldsymbol{\eta}^{(j)} = \mathbf{c} + \sum_{j=1}^M \boldsymbol{\eta}^{(j)}$$

and it estimates the corresponding probability distribution $\tilde{\mathbf{p}} = \frac{1}{n'} \tilde{\mathbf{c}}$. Notice that:

$$n' = \sum_{i=1}^d \sum_{j=1}^M \tilde{c}_i^{(j)} = \sum_{i=1}^d \sum_{j=1}^M c_i^{(j)} + \eta_i^{(j)} = n + \sum_{i=1}^d \sum_{j=1}^M \eta_i^{(j)}$$

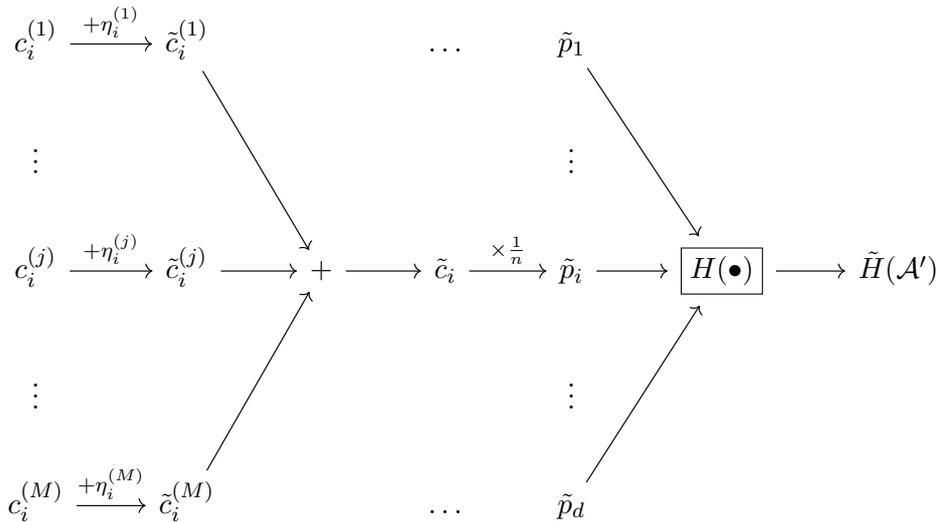
so n' is the sum of the actual total dataset size n , plus the sum of $d \times M$ zero-mean Laplace random variables, and, hence, it is also random with $\mathbb{E}[n'] = n$. To simplify our analysis, we assume that the analyst divides by n and we ignore the consistency issues this causes to the corresponding probability distribution. The resulting vector $\tilde{\mathbf{p}} = \tilde{\mathbf{p}}_{\mathcal{A}'}$ (we have so far skipped the subscript \mathcal{A}' for simplicity in presentation), is subsequently utilized in the following two ways:

- To compute the joint entropy of the attributes in \mathcal{A}' , $\tilde{H}(\mathcal{A}')$.
- To compute the joint probability distributions of attributes in subsets of \mathcal{A}' by marginalization, and then, their joint entropies.

Once the analyst has computed all the required entropies, it may compute the mutual information between an attribute X and its candidate parent set $\mathbf{Pa}(X)$, utilizing the well-known formula:

$$\tilde{I}(X; \mathbf{Pa}(X)) = \tilde{H}(X) + \tilde{H}(\mathbf{Pa}(X)) - \tilde{H}(X, \mathbf{Pa}(X))$$

The pipeline we described is (partly) depicted in the diagram below.



The aforementioned process is executed for all $\binom{N}{k+1}$ sets \mathcal{A}' . Thus, when the analyst needs to compute the mutual information between X and $\mathbf{Pa}(X)$, it first selects a set \mathcal{A}' such that $\{X\} \cup \mathbf{Pa}(X) \subseteq \mathcal{A}'$. Then, it performs the necessary marginalizations, to extract $\tilde{\mathbf{p}}_X, \tilde{\mathbf{p}}_{\mathbf{Pa}(X)}, \tilde{\mathbf{p}}_{X, \mathbf{Pa}(X)}$; the last distribution needs to be computed only if $\{X\} \cup \mathbf{Pa}(X) \subset \mathcal{A}'$ - otherwise the analyst already has it. Finally, after computing the entropies of these three distributions ($\tilde{H}(X)$, $\tilde{H}(\mathbf{Pa}(X))$, and -potentially- $\tilde{H}(X, \mathbf{Pa}(X))$),

the analyst utilizes them to estimate $I(X; \mathbf{Pa}(X))$. Of course, in practice, the analyst would not need to compute the entropies, as the mutual information can be derived directly from the probability distributions. Nevertheless, we adopt this approach to facilitate our theoretical analysis of the algorithm.

Our main argument is that, if the noisy entropy estimate $\tilde{H}(\mathcal{A}')$ for an arbitrary set \mathcal{A}' is close to the true entropy $H(\mathcal{A}')$, then so will be the mutual information estimates computed over the course of the algorithm, and therefore, Algorithm 3.5 will construct a BN structure that accurately encodes the data distribution.

Before starting our analysis, we point out that we use the Big \mathcal{O} for multiple variables as described by Cormen et al. [13], despite the inconsistencies demonstrated by Howell [37], so for a vector $\mathbf{x} \in \mathbb{R}^n$ and two functions f, g defined on some subset of \mathbb{R}^n :

$$f(\mathbf{x}) = \mathcal{O}(g(\mathbf{x})) \text{ as } \mathbf{x} \rightarrow \infty$$

$$\Leftrightarrow \exists M, \exists C > 0 \text{ such that } \forall \mathbf{x} \text{ with } \|\mathbf{x}\|_\infty \geq M, |f(\mathbf{x})| \leq C|g(\mathbf{x})|$$

In practice, we may be interested in settings where some of the variables are fixed (for instance, we may not care about having an arbitrarily large number of data holders). If this is the case, we may view these variables as constants, which would simplify the resulting expressions.

We start our analysis with a few useful lemmas. Proving the first of them (Lemma 3.2) involves tedious computations; equivalent results were obtained by Harris [35].

Lemma 3.2. *Let $\mathbf{c} = [c_1 \dots c_d]^T$ be a random vector that follows a multinomial distribution with parameters n, \mathbf{p} , where $\mathbf{p} = [p_1 \dots p_d]^T$ and $\sum_{i=1}^d p_i = 1$. Then, $\forall i \in \{1, \dots, d\}$, the first 6 moments of c_i are:*

$$\begin{aligned} \mathbb{E}[c_i] &= np_i \\ \mathbb{E}[c_i^2] &= n^2 p_i^2 + n(-p_i^2 + p_i) \\ \mathbb{E}[c_i^3] &= n^3 p_i^3 + n^2(-3p_i^3 + 3p_i^2) + n(2p_i^3 - 3p_i^2 + p_i) \\ \mathbb{E}[c_i^4] &= n^4 p_i^4 + n^3(-6p_i^4 + 6p_i^3) + n^2(11p_i^4 - 18p_i^3 + 7p_i^2) \\ &\quad + n(-6p_i^4 + 12p_i^3 - 7p_i^2 + p_i) \\ \mathbb{E}[c_i^5] &= n^5 p_i^5 + n^4(-10p_i^5 + 10p_i^4) + n^3(35p_i^5 - 60p_i^4 + 25p_i^3) \\ &\quad + n^2(-50p_i^5 + 110p_i^4 - 75p_i^3 + 15p_i^2) \\ &\quad + n(24p_i^5 - 60p_i^4 + 50p_i^3 - 15p_i^2 + p_i) \\ \mathbb{E}[c_i^6] &= n^6 p_i^6 + n^5(-15p_i^6 + 15p_i^5) + n^4(85p_i^6 - 150p_i^5 + 65p_i^4) \\ &\quad + n^3(-225p_i^6 + 525p_i^5 - 390p_i^4 + 90p_i^3) \\ &\quad + n^2(274p_i^6 - 750p_i^5 + 715p_i^4 - 270p_i^3 + 31p_i^2) \\ &\quad + n(-120p_i^6 + 360p_i^5 - 390p_i^4 + 180p_i^3 - 31p_i^2 + p_i) \end{aligned}$$

In addition, for $j \neq i$:

$$\begin{aligned} \mathbb{E}[c_i c_j] &= n(n-1)p_i p_j \\ \mathbb{E}[c_i^2 c_j] &= n(n-1)p_i p_j + n(n-1)(n-2)p_i^2 p_j \end{aligned}$$

In proving the next Lemma 3.3 we utilize the multinomial theorem, the linearity of expectations and the independence assumption between the Laplace random variables.

Lemma 3.3. *Let η_1, \dots, η_M be i.i.d. Laplace(0, b) random variables. Then, $\forall i \in \{1, \dots, M\}$ and for some $\kappa \in \mathbb{N}$:*

$$\mathbb{E}[\eta_i^\kappa] = \begin{cases} 0 & \text{if } \kappa \text{ is odd} \\ b^\kappa \kappa! & \text{if } \kappa \text{ is even} \end{cases}$$

Also, taking into account that:

$$\mathbb{E}[(\sum_{i=1}^M \eta_i)^\kappa] = \sum_{\substack{(\kappa_1, \kappa_2, \dots, \kappa_M): \\ \kappa_1 + \kappa_2 + \dots + \kappa_M = \kappa}} \frac{\kappa!}{\kappa_1! \kappa_2! \dots \kappa_M!} \prod_{i=1}^M \mathbb{E}[\eta_i^{\kappa_i}]$$

it follows that:

$$\begin{aligned} \mathbb{E}[\sum_{i=1}^M \eta_i] &= 0 \quad , \quad \mathbb{E}[(\sum_{i=1}^M \eta_i)^2] = 2Mb^2 \\ \mathbb{E}[(\sum_{i=1}^M \eta_i)^3] &= 0 \quad , \quad \mathbb{E}[(\sum_{i=1}^M \eta_i)^4] = 12(M^2 + M)b^4 \\ \mathbb{E}[(\sum_{i=1}^M \eta_i)^5] &= 0 \quad , \quad \mathbb{E}[(\sum_{i=1}^M \eta_i)^6] = 240(M^3 + 2M)b^6 \end{aligned}$$

Lemma 3.4 is an application of the well-known Cauchy-Schwarz inequality on the expectations of random variables.

Lemma 3.4. For any two random variables X and Y :

$$|\mathbb{E}[XY]| \leq \sqrt{\mathbb{E}[X^2] \mathbb{E}[Y^2]}$$

where equality holds if and only if $X = cY$ for some constant $c \in \mathbb{R}$.

The next Lemma 3.5 will allow us to bound sums that appear in our analysis.

Lemma 3.5. Let $\mathbf{p} = [p_1 \dots p_d]^T$ be a probability distribution over a discrete alphabet of size d , and let $0 < p_{\min} \leq p_i, \forall i \in \{1, \dots, d\}$. Then:

$$(i) \quad \sum_{i=1}^d \frac{1}{(p_i)^\kappa} \leq \frac{d}{(p_{\min})^\kappa} = \mathcal{O}\left(\frac{d}{(p_{\min})^\kappa}\right) \quad (\text{for some } \kappa \in \mathbb{N})$$

If, in addition, $p_{\min} < \frac{1}{e^2}$, then:

$$\begin{aligned} (ii) \quad \sum_{i=1}^d (1 + \log p_i)^2 &\leq d(1 + \log p_{\min})^2 = \mathcal{O}(d \log^2 p_{\min}) \\ (iii) \quad \sum_{i=1}^d \frac{|1 + \log p_i|}{p_i} &\leq d \frac{|1 + \log p_{\min}|}{p_{\min}} = \mathcal{O}\left(\frac{d \log p_{\min}}{p_{\min}}\right) \\ (iv) \quad \sum_{i=1}^d \frac{|1 + \log p_i|}{(p_i)^2} &\leq d \frac{|1 + \log p_{\min}|}{(p_{\min})^2} = \mathcal{O}\left(\frac{d \log p_{\min}}{(p_{\min})^2}\right) \end{aligned}$$

The asymptotic expressions hold for $d \rightarrow \infty, p_{\min} \rightarrow 0$.

We make the following remarks concerning the proofs of the inequalities that appear in Lemma 3.5:

- Inequality (i): the proof is trivial.
- Inequalities (ii), (iii), (iv): the function $f(x) = |1 + \log(x)|, x \in (0, 1]$, has the following behavior:

$$\begin{aligned} \frac{1}{e^2} \leq x \leq 1 &\Leftrightarrow f(x) \leq 1 \\ 0 < x < \frac{1}{e^2} &\Leftrightarrow f(x) > 1 \text{ and } f \text{ is monotonically decreasing.} \end{aligned}$$

Thus, if $p_{\min} < \frac{1}{e^2}$, then $\max_i f(p_i) = f(p_{\min})$. Taking this into account, we derive the desired results by simple algebra.

For example, we apply inequalities (i) and (ii) for a uniform distribution (assuming $d \geq 8 \Leftrightarrow p_{\min} = \frac{1}{d} < \frac{1}{e^2}$):

$$\begin{aligned} \sum_{i=1}^d \frac{1}{(p_i)^\kappa} &= \sum_{i=1}^d \left(\frac{1}{d}\right)^\kappa = d^{\kappa+1} \\ \sum_{i=1}^d (1 + \log p_i)^2 &= d(1 + \log \frac{1}{d})^2 \end{aligned}$$

The last Lemma 3.6 is due to Basharin [4] who first examined the properties of the empirical entropy estimator, and Harris [35], who performed a much more detailed analysis.

Lemma 3.6. *Let $\mathbf{p} = [p_1 \dots p_d]^T$ be a probability distribution over a discrete alphabet of size d , and let $H(\mathbf{p}) = -\sum_{i=1}^d p_i \log_2 p_i$ be the entropy functional of \mathbf{p} . The empirical entropy estimator is computed as:*

$$\hat{H} = H(\hat{\mathbf{p}}) = -\sum_{i=1}^d \hat{p}_i \log_2 \hat{p}_i = -\log_2(e) \sum_{i=1}^d \hat{p}_i \log \hat{p}_i$$

where $\hat{\mathbf{p}}$ is the maximum likelihood estimate of \mathbf{p} using n independent data points. Ignoring the $\log_2(e)$ scale factor, the bias and mean squared error of \hat{H} are:

$$\begin{aligned} \mathbb{E}[\hat{H} - H] &= -\frac{d-1}{2n} + \frac{1 - \sum_{i=1}^d \frac{1}{p_i}}{12n^2} + \mathcal{O}\left(\frac{1}{n^3}\right) \\ \mathbb{E}[(\hat{H} - H)^2] &= \frac{\sum_{i=1}^d p_i \log_2^2 p_i - H^2}{n} + \frac{d^2 - 1}{4n^2} + \mathcal{O}\left(\frac{1}{n^3}\right) \end{aligned}$$

Following all previous authors on the topic, we also ignore the scale factor in our analysis, and work with the entropy defined using the natural logarithm. This has no essential effect on the computation, and at the end, we can multiply everything through by $\log_2(e)$ to change the base back to base two.

Taking these into account, we provide two theorems that give the conditions under which our distributed, differentially private entropy estimator accurately approximates the true entropy. In the first (Theorem 3.5), we examine the estimator's bias.

Theorem 3.5. *The absolute bias of the distributed, differentially private entropy estimator \tilde{H} used by Algorithm 3.5 is:*

$$|\mathbb{E}[\tilde{H} - H]| = \frac{d-1}{2n} + \mathcal{O}\left(\frac{dMb^2}{p_{\min}n^2} + \frac{dMb^2}{(p_{\min})^2n^3} + \frac{dM^2b^4}{(p_{\min})^3n^4}\right)$$

If $p_{\min} = \omega\left(\frac{1}{n}\right)$ and we pick $\varepsilon = \omega\left(\sqrt{\frac{M}{p_{\min}n}}N^{k+1}\right)$, then the noise terms get dominated by the $\mathcal{O}\left(\frac{d}{n}\right)$ term, and the bias of \tilde{H} converges asymptotically to that of the empirical entropy estimator \hat{H} (Lemma 3.6).

Proof. Let k be the degree of the target Bayesian Network. We fix a set of attributes $\mathcal{A}' \subseteq \mathcal{A}$ such that $|\mathcal{A}'| = k + 1$, and assume $|d_{\mathcal{A}'}| = d$. For simplicity in presentation, we skip the subscripts and denote the joint frequency and probability distributions of the attributes in \mathcal{A}' by \mathbf{c} and \mathbf{p} respectively.

Recall that, based on the algorithm's pipeline and our assumptions:

$$\tilde{\mathbf{c}} = \mathbf{c} + \sum_{j=1}^M \boldsymbol{\eta}^{(j)}$$

so $\tilde{\mathbf{c}}$ is the sum of vector $\mathbf{c} \sim \text{Multinomial}(n, \mathbf{p})$ and M i.i.d. Laplace(0, b) random vectors (with independent entries). Each entry in $\tilde{\mathbf{c}}$ is computed as $\tilde{c}_i = c_i + \sum_{j=1}^M \eta_i^{(j)}$, $\forall i \in \{1, \dots, d\}$. We utilize Lemmas 3.2 and 3.3, and we use red fonts to emphasize the additive terms that emerge due to the perturbation:

$$\begin{aligned} \mathbb{E}[\tilde{c}_i] &= np_i \\ \mathbb{E}[\tilde{c}_i^2] &= n^2 p_i^2 + n(-p_i^2 + p_i) + 2Mb^2 \\ \mathbb{E}[\tilde{c}_i^3] &= n^3 p_i^3 + n^2(-3p_i^3 + 3p_i^2) + n(2p_i^3 - 3p_i^2 + p_i + 6p_i Mb^2) \\ \mathbb{E}[\tilde{c}_i^4] &= n^4 p_i^4 + n^3(-6p_i^4 + 6p_i^3) + n^2(11p_i^4 - 18p_i^3 + 7p_i^2 + 12p_i^2 Mb^2) \\ &\quad + n[-6p_i^4 + 12p_i^3 - 7p_i^2 + p_i + 12(-p_i^2 + p_i)Mb^2] \\ &\quad + 12(M^2 + M)b^4 \\ \mathbb{E}[\tilde{c}_i^5] &= n^5 p_i^5 + n^4(-10p_i^5 + 10p_i^4) + n^3(35p_i^5 - 60p_i^4 + 25p_i^3 + 20p_i^3 Mb^2) \\ &\quad + n^2[-50p_i^5 + 110p_i^4 - 75p_i^3 + 15p_i^2 + 60(-p_i^3 + p_i^2)Mb^2] \\ &\quad + n[24p_i^5 - 60p_i^4 + 50p_i^3 - 15p_i^2 + p_i \\ &\quad + (40p_i^3 - 60p_i^2 + 20p_i)Mb^2 + 60p_i(M^2 + M)b^4] \\ \mathbb{E}[\tilde{c}_i^6] &= n^6 p_i^6 + n^5(-15p_i^6 + 15p_i^5) + n^4(85p_i^6 - 150p_i^5 + 65p_i^4 + 30p_i^4 Mb^2) \\ &\quad + n^3[-225p_i^6 + 525p_i^5 - 390p_i^4 + 90p_i^3 + 180(-p_i^4 + p_i^3)Mb^2] \\ &\quad + n^2[274p_i^6 - 750p_i^5 + 715p_i^4 - 270p_i^3 + 31p_i^2 \\ &\quad + (330p_i^4 - 540p_i^3 + 210p_i^2)Mb^2 + 180p_i^2(M^2 + M)b^4] \\ &\quad + n[-120p_i^6 + 360p_i^5 - 390p_i^4 + 180p_i^3 - 31p_i^2 + p_i \\ &\quad + 30(-p_i^2 + p_i)Mb^2 + 180(-p_i^2 + p_i)(M^2 + M)b^4] \\ &\quad + 240(M^3 + 2M)b^6 \end{aligned}$$

In addition, for $j \neq i$:

$$\begin{aligned} \mathbb{E}[\tilde{c}_i \tilde{c}_j] &= n(n-1)p_i p_j \\ \mathbb{E}[\tilde{c}_i^2 \tilde{c}_j] &= n^3 p_i^2 p_j + n^2(-3p_i^2 p_j + p_i p_j) + n(2p_i^2 p_j - p_i p_j + 2p_j Mb^2) \end{aligned}$$

Noting that $\tilde{p}_i = \frac{\tilde{c}_i}{n} \Rightarrow \mathbb{E}[\tilde{p}_i^\kappa] = \frac{\mathbb{E}[\tilde{c}_i^\kappa]}{n^\kappa}$ ($\forall \kappa \in \mathbb{N}$), and thus, $\mathbb{E}[\tilde{p}_i] = \frac{np_i}{n} = p_i$, we compute the central moments of each \tilde{p}_i :

$$\begin{aligned} \mathbb{E}[\tilde{p}_i - p_i] &= 0 \\ \mathbb{E}[(\tilde{p}_i - p_i)^2] &= \frac{1}{n}(-p_i^2 + p_i) + \frac{1}{n^2}2Mb^2 \\ \mathbb{E}[(\tilde{p}_i - p_i)^3] &= \frac{1}{n^2}(2p_i^3 - 3p_i^2 + p_i) \\ \mathbb{E}[(\tilde{p}_i - p_i)^4] &= \frac{1}{n^2}(3p_i^4 - 6p_i^3 + 3p_i^2) \\ &\quad + \frac{1}{n^3}[-6p_i^4 + 12p_i^3 - 7p_i^2 + p_i + 12(-p_i^2 + p_i)Mb^2] \\ &\quad + \frac{1}{n^4}12(M^2 + M)b^4 \\ \mathbb{E}[(\tilde{p}_i - p_i)^5] &= \frac{1}{n^3}(-20p_i^5 + 50p_i^4 - 40p_i^3 + 10p_i^2) \end{aligned}$$

$$\begin{aligned}
\mathbb{E}[(\tilde{p}_i - p_i)^6] &= \frac{1}{n^4} [24p_i^5 - 60p_i^4 + 50p_i^3 - 15p_i^2 + p_i \\
&\quad + (40p_i^3 - 60p_i^2 + 20p_i)Mb^2] \\
&+ \frac{1}{n^3} (-15p_i^6 + 45p_i^5 - 45p_i^4 + 15p_i^3) \\
&+ \frac{1}{n^4} [130p_i^6 - 390p_i^5 + 415p_i^4 - 180p_i^3 + 25p_i^2 \\
&\quad + (90p_i^4 - 180p_i^3 + 90p_i^2)Mb^2] \\
&+ \frac{1}{n^5} [-120p_i^6 + 360p_i^5 - 390p_i^4 + 180p_i^3 - 31p_i^2 + p_i \\
&\quad + 30(-p_i^2 + p_i)Mb^2 + 180(-p_i^2 + p_i)(M^2 + M)b^4] \\
&+ \frac{1}{n^6} 240(M^3 + 2M)b^6
\end{aligned}$$

In addition, for $j \neq i$:

$$\begin{aligned}
\mathbb{E}[(\tilde{p}_i - p_i)(\tilde{p}_j - p_j)] &= \frac{1}{n} p_i p_j \\
\mathbb{E}[(\tilde{p}_i - p_i)^2(\tilde{p}_j - p_j)] &= \frac{1}{n^2} (2p_i^2 p_j - p_i p_j)
\end{aligned}$$

We take a third order Taylor expansion (with remainder) of \tilde{H} around the true entropy H :

$$\begin{aligned}
\tilde{H} &= H - \sum_{i=1}^d (1 + \log p_i)(\tilde{p}_i - p_i) - \frac{1}{2} \sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^2}{p_i} + \frac{1}{6} \sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{p_i^2} \\
&\quad - \frac{1}{12} \sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^4}{[(1-c)p_i + c\tilde{p}_i]^3}
\end{aligned}$$

where $c \in [0, 1]$. We next take $\mathbb{E}[\tilde{H}]$, and by linearity of expectations, we examine the expectation of each term in the Taylor expansion separately:

$$\begin{aligned}
\sum_{i=1}^d (1 + \log p_i) \mathbb{E}[\tilde{p}_i - p_i] &= 0 \\
\sum_{i=1}^d \frac{\mathbb{E}[(\tilde{p}_i - p_i)^2]}{p_i} &= \sum_{i=1}^d \frac{1}{n} \frac{-p_i^2 + p_i}{p_i} + \frac{1}{n^2} 2Mb^2 \frac{1}{p_i} \\
&= \frac{d-1}{n} + \frac{2Mb^2 \sum_{i=1}^d \frac{1}{p_i}}{n^2} \\
&= \frac{d-1}{n} + \mathcal{O}\left(\frac{dMb^2}{p_{\min} n^2}\right) \\
\sum_{i=1}^d \frac{\mathbb{E}[(\tilde{p}_i - p_i)^3]}{p_i^2} &= \sum_{i=1}^d \frac{1}{n^2} \frac{2p_i^3 - 3p_i^2 + p_i}{p_i^2} \\
&= \frac{2 - 3d + \sum_{i=1}^d \frac{1}{p_i}}{n^2} \\
&= \mathcal{O}\left(\frac{d}{p_{\min} n^2}\right) \\
\sum_{i=1}^d \mathbb{E}\left[\frac{(\tilde{p}_i - p_i)^4}{[(1-c)p_i + c\tilde{p}_i]^3}\right] &\leq \sum_{i=1}^d \frac{\mathbb{E}[(\tilde{p}_i - p_i)^4]}{(1-c)^3 p_i^3}
\end{aligned}$$

$$\begin{aligned}
&= \frac{3 - 6d + 3 \sum_{i=1}^d \frac{1}{p_i}}{(1-c)^3 n^2} \\
&\quad + \frac{-6 + 12d - 7 \sum_{i=1}^d \frac{1}{p_i} + \sum_{i=1}^d \frac{1}{p_i^2}}{(1-c)^3 n^3} \\
&\quad + \frac{12Mb^2(-\sum_{i=1}^d \frac{1}{p_i} + \sum_{i=1}^d \frac{1}{p_i^2})}{(1-c)^3 n^3} \\
&\quad + \frac{12(M^2 + M)b^4 \sum_{i=1}^d \frac{1}{p_i^3}}{(1-c)^3 n^4} \\
&= \mathcal{O}\left(\frac{d}{p_{\min} n^2} + \frac{dMb^2}{(p_{\min})^2 n^3} + \frac{dM^2 b^4}{(p_{\min})^3 n^4}\right)
\end{aligned}$$

To derive the asymptotic expressions, we utilize lemma 3.5.

Let us review the variables that appear in our expression:

- n : total number of data points ($n \rightarrow \infty$).
- d : joint domain size of attributes whose entropy we estimate ($d \rightarrow \infty$).
- M : number of data holders ($M \rightarrow \infty$).
- b : noise scale ($b \rightarrow \infty$). Recall that $b = \frac{\binom{N}{k+1}}{\varepsilon} = \mathcal{O}\left(\frac{N^{k+1}}{\varepsilon}\right)$, so:
 - N : total number of attributes ($N \rightarrow \infty$).
 - ε : is the privacy budget ($\varepsilon \rightarrow 0$).
- p_{\min} : minimum entry in joint probability distribution of attributes whose entropy we estimate ($p_{\min} \rightarrow 0$).

So, for example, the $\mathcal{O}\left(\frac{d}{p_{\min} n^2}\right)$ terms get dominated by the $\mathcal{O}\left(\frac{dMb^2}{p_{\min} n^2}\right)$ term. Combining the above yields:

$$|\mathbb{E}[\tilde{H} - H]| = \frac{d-1}{2n} + \mathcal{O}\left(\frac{dMb^2}{p_{\min} n^2} + \frac{dMb^2}{(p_{\min})^2 n^3} + \frac{dM^2 b^4}{(p_{\min})^3 n^4}\right)$$

Therefore, to ensure that the bias of \tilde{H} converges asymptotically to that of \hat{H} , the following must hold:

$$\begin{cases} \frac{dMb^2}{p_{\min} n^2} = o\left(\frac{d}{n}\right) \\ \frac{dMb^2}{(p_{\min})^2 n^3} = o\left(\frac{d}{n}\right) \\ \frac{dM^2 b^4}{(p_{\min})^3 n^4} = o\left(\frac{d}{n}\right) \end{cases} \Leftrightarrow \begin{cases} \varepsilon = \omega\left(\sqrt{\frac{M}{p_{\min} n}} N^{k+1}\right) \\ p_{\min} = \omega\left(\frac{1}{n}\right) \end{cases}$$

which completes the proof. □

We next examine the estimator's mean squared error (Theorem 3.6).

Theorem 3.6. *Assume that the data distribution \mathbf{p} satisfies $p_{\min} < \frac{1}{e^2}$, $p_{\min} = \omega\left(\frac{1}{\sqrt{n}}\right)$ and $\sigma^2 = \mathcal{O}(d^2)$, where $\sigma^2 = (\sum_{i=1}^d p_i \log^2 p_i) - H^2$.*

Then, if we pick $\varepsilon = \omega\left(\sqrt{\frac{M}{n}} N^{k+1} |\log(p_{\min})|\right)$, the mean squared error of the distributed, differentially private entropy estimator \tilde{H} used by Algorithm 3.5 is:

$$\mathbb{E}[(\tilde{H} - H)^2] = \frac{\sigma^2}{n} + o\left(\frac{\sigma^2}{n}\right)$$

and it converges asymptotically to that of the empirical entropy estimator \hat{H} (Lemma 3.6).

Proof. The proof of Theorem 3.6 follows that of Theorem 3.5, but involves much more complicated computations. We therefore focus on the asymptotic behavior of the quantities involved in our analysis, and we analytically compute only the dominant ones. We begin by deriving asymptotic expressions for the central moments of each \tilde{p}_i (as $n, M, b \rightarrow \infty$), introducing the notation $\theta = \frac{Mb^2}{n^2}$:

$$\begin{aligned}\mathbb{E}[\tilde{p}_i - p_i] &= 0 \\ \mathbb{E}[(\tilde{p}_i - p_i)^2] &= \mathcal{O}\left(\frac{1}{n} + \frac{Mb^2}{n^2}\right) = \mathcal{O}\left(\frac{1}{n} + \theta\right) \\ \mathbb{E}[(\tilde{p}_i - p_i)^3] &= \mathcal{O}\left(\frac{1}{n^2}\right) \\ \mathbb{E}[(\tilde{p}_i - p_i)^4] &= \mathcal{O}\left(\frac{1}{n^2} + \frac{Mb^2}{n^3} + \frac{M^2b^4}{n^4}\right) = \mathcal{O}\left(\frac{1}{n^2} + \frac{\theta}{n} + \theta^2\right) \\ \mathbb{E}[(\tilde{p}_i - p_i)^5] &= \mathcal{O}\left(\frac{1}{n^3} + \frac{Mb^2}{n^4}\right) = \mathcal{O}\left(\frac{1}{n^3} + \frac{\theta}{n^2}\right) \\ \mathbb{E}[(\tilde{p}_i - p_i)^6] &= \mathcal{O}\left(\frac{1}{n^3} + \frac{Mb^2}{n^4} + \frac{M^2b^4}{n^5} + \frac{M^3b^6}{n^6}\right) \\ &= \mathcal{O}\left(\frac{1}{n^3} + \frac{\theta}{n^2} + \frac{\theta^2}{n} + \theta^3\right)\end{aligned}$$

In addition, for $j \neq i$:

$$\begin{aligned}\mathbb{E}[(\tilde{p}_i - p_i)(\tilde{p}_j - p_j)] &= \mathcal{O}\left(\frac{1}{n}\right) \\ \mathbb{E}[(\tilde{p}_i - p_i)^2(\tilde{p}_j - p_j)] &= \mathcal{O}\left(\frac{1}{n^2}\right)\end{aligned}$$

We are also interested in higher order cross moments, so we apply Lemma 3.4:

$$\begin{aligned}\mathbb{E}[(\tilde{p}_i - p_i)^2(\tilde{p}_j - p_j)^2] &\leq \sqrt{\mathbb{E}[(\tilde{p}_i - p_i)^4] \mathbb{E}[(\tilde{p}_j - p_j)^4]} \\ &= \sqrt{\mathcal{O}\left(\frac{1}{n^2} + \frac{\theta}{n} + \theta^2\right)^2} \\ &= \mathcal{O}\left(\frac{1}{n^2} + \frac{\theta}{n} + \theta^2\right) \\ \mathbb{E}[(\tilde{p}_i - p_i)(\tilde{p}_j - p_j)^3] &\leq \sqrt{\mathbb{E}[(\tilde{p}_i - p_i)^2] \mathbb{E}[(\tilde{p}_j - p_j)^6]} \\ &= \sqrt{\mathcal{O}\left(\frac{1}{n} + \theta\right) \mathcal{O}\left(\frac{1}{n^3} + \frac{\theta}{n^2} + \frac{\theta^2}{n} + \theta^3\right)} \\ &= \mathcal{O}\left(\frac{1}{n^2} + \frac{\theta^{\frac{1}{2}}}{n^{\frac{3}{2}}} + \frac{\theta}{n} + \frac{\theta^{\frac{3}{2}}}{n^{\frac{1}{2}}} + \theta^2\right) \\ \mathbb{E}[(\tilde{p}_i - p_i)^2(\tilde{p}_j - p_j)^3] &\leq \sqrt{\mathbb{E}[(\tilde{p}_i - p_i)^4] \mathbb{E}[(\tilde{p}_j - p_j)^6]} \\ &= \sqrt{\mathcal{O}\left(\frac{1}{n^2} + \frac{\theta}{n} + \theta^2\right) \mathcal{O}\left(\frac{1}{n^3} + \frac{\theta}{n^2} + \frac{\theta^2}{n} + \theta^3\right)} \\ &= \mathcal{O}\left(\frac{1}{n^{\frac{5}{2}}} + \frac{\theta^{\frac{1}{2}}}{n^2} + \frac{\theta}{n^{\frac{3}{2}}} + \frac{\theta^{\frac{3}{2}}}{n} + \frac{\theta^2}{n^{\frac{1}{2}}} + \theta^{\frac{5}{2}}\right) \\ \mathbb{E}[(\tilde{p}_i - p_i)^3(\tilde{p}_j - p_j)^3] &\leq \sqrt{\mathbb{E}[(\tilde{p}_i - p_i)^6] \mathbb{E}[(\tilde{p}_j - p_j)^6]} \\ &= \sqrt{\mathcal{O}\left(\frac{1}{n^3} + \frac{\theta}{n^2} + \frac{\theta^2}{n} + \theta^3\right)^2}\end{aligned}$$

$$= \mathcal{O}\left(\frac{1}{n^3} + \frac{\theta}{n^2} + \frac{\theta^2}{n} + \theta^3\right)$$

We take a second order Taylor expansion (with remainder) of \tilde{H} around the true entropy H , and examine the quantity $(\tilde{H} - H)^2$:

$$\begin{aligned} (\tilde{H} - H)^2 &= \left[H - \sum_{i=1}^d (1 + \log p_i)(\tilde{p}_i - p_i) - \frac{1}{2} \sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^2}{p_i} \right. \\ &\quad \left. + \frac{1}{6} \sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{[(1 - \theta)p_i + \theta\tilde{p}_i]^2} - H \right]^2 \\ &= \left[\sum_{i=1}^d (1 + \log p_i)(\tilde{p}_i - p_i) \right]^2 \\ &\quad + \frac{1}{4} \left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^2}{p_i} \right]^2 \\ &\quad + \frac{1}{36} \left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{[(1 - c)p_i + c\tilde{p}_i]^2} \right]^2 \\ &\quad + \left[\sum_{i=1}^d (1 + \log p_i)(\tilde{p}_i - p_i) \right] \left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^2}{p_i} \right] \\ &\quad - \frac{1}{3} \left[\sum_{i=1}^d (1 + \log p_i)(\tilde{p}_i - p_i) \right] \left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{[(1 - c)p_i + c\tilde{p}_i]^2} \right] \\ &\quad - \frac{1}{6} \left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^2}{p_i} \right] \left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{[(1 - c)p_i + c\tilde{p}_i]^2} \right] \end{aligned}$$

where $c \in [0, 1]$. We next take $\mathbb{E}[(\tilde{H} - H)^2]$, and by linearity of expectations, we examine the expectation of each term separately.

The first term, $\mathbb{E}\left[\left[\sum_{i=1}^d (1 + \log p_i)(\tilde{p}_i - p_i)\right]^2\right]$, involves the $\mathcal{O}\left(\frac{1}{n}\right)$ terms $\mathbb{E}[(\tilde{p}_i - p_i)^2]$ and $\mathbb{E}[(\tilde{p}_i - p_i)(\tilde{p}_j - p_j)]$, so we compute it analytically:

$$\begin{aligned} \mathbb{E}\left[\left[\sum_{i=1}^d (1 + \log p_i)(\tilde{p}_i - p_i)\right]^2\right] &= \sum_{i=1}^d (1 + \log p_i)^2 \mathbb{E}[(\tilde{p}_i - p_i)^2] \\ &\quad + 2 \sum_{i=1}^{d-1} \sum_{j=i+1}^d (1 + \log p_i)(1 + \log p_j) \\ &\quad \quad \mathbb{E}[(\tilde{p}_i - p_i)(\tilde{p}_j - p_j)] \\ &= \frac{1}{n} \left[\left(\sum_{i=1}^d p_i \log^2 p_i \right) - H^2 \right] \\ &\quad + \frac{1}{n^2} 2Mb^2 \sum_{i=1}^d (1 + \log p_i)^2 \\ &= \frac{1}{n} \left[\left(\sum_{i=1}^d p_i \log^2 p_i \right) - H^2 \right] \\ &\quad + \mathcal{O}(d \log^2(p_{\min})\theta) \\ &= \mathcal{O}\left(\frac{\sigma^2}{n}\right) + \mathcal{O}(d \log^2(p_{\min})\theta) \end{aligned}$$

To ensure that the mean squared error of \tilde{H} converges to that of \hat{H} , we want the second term to be dominated by the first, that is:

$$d \log^2(p_{\min})\theta = o\left(\frac{\sigma^2}{n}\right)$$

Given that $\sigma^2 = \mathcal{O}(d^2)$, the aforementioned condition is satisfied when:

$$\log^2(p_{\min})\theta = o\left(\frac{1}{n}\right) \Leftrightarrow \varepsilon = \omega\left(\sqrt{\frac{M}{n}} N^{k+1} |\log(p_{\min})|\right)$$

Note that the condition $\frac{\theta}{p_{\min}} = o\left(\frac{1}{n}\right)$, which appeared in Theorem 3.5, implies the condition we want to satisfy here.

We proceed with the remaining terms, based on Lemma 3.5, and on the assumption that $\log^2(p_{\min})\theta = o\left(\frac{1}{n}\right)$, which implies that $\theta = o\left(\frac{1}{n}\right)$.

$$\begin{aligned} \mathbb{E}\left[\left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^2}{p_i}\right]^2\right] &= \sum_{i=1}^d \frac{\mathbb{E}[(\tilde{p}_i - p_i)^4]}{p_i^2} \\ &\quad + 2 \sum_{i=1}^{d-1} \sum_{j=i+1}^d \frac{\mathbb{E}[(\tilde{p}_i - p_i)^2(\tilde{p}_j - p_j)^2]}{p_i p_j} \\ &= \mathcal{O}\left(\frac{d^2}{(p_{\min})^2}\right) \mathcal{O}\left(\frac{1}{n^2} + \frac{\theta}{n} + \theta^2\right) \\ &= \mathcal{O}\left(\frac{d^2}{(p_{\min})^2}\right) \left[\mathcal{O}\left(\frac{1}{n^2}\right) + o\left(\frac{1}{n^2}\right)\right] \\ \mathbb{E}\left[\left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{[(1-c)p_i + c\tilde{p}_i]^2}\right]^2\right] &\leq \mathbb{E}\left[\left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{(1-c)^2 p_i^2}\right]^2\right] \\ &= \sum_{i=1}^d \frac{\mathbb{E}[(\tilde{p}_i - p_i)^6]}{(1-c)^4 p_i^4} \\ &\quad + 2 \sum_{i=1}^{d-1} \sum_{j=i+1}^d \frac{\mathbb{E}[(\tilde{p}_i - p_i)^3(\tilde{p}_j - p_j)^3]}{(1-c)^4 p_i^2 p_j^2} \\ &= \mathcal{O}\left(\frac{d^2}{(p_{\min})^4}\right) \mathcal{O}\left(\frac{1}{n^3} + \frac{\theta}{n^2} + \frac{\theta^2}{n} + \theta^3\right) \\ &= \mathcal{O}\left(\frac{d^2}{(p_{\min})^4}\right) \left[\mathcal{O}\left(\frac{1}{n^3}\right) + o\left(\frac{1}{n^3}\right)\right] \\ \mathbb{E}\left[\sum_{i=1}^d (1 + \log p_i)(\tilde{p}_i - p_i) \sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^2}{p_i}\right] \\ &= \sum_{i=1}^d \frac{(1 + \log p_i)}{p_i} \mathbb{E}[(\tilde{p}_i - p_i)^3] + \sum_{(i,j):i \neq j} \frac{1 + \log p_i}{p_j} \mathbb{E}[(\tilde{p}_i - p_i)(\tilde{p}_j - p_j)^2] \\ &= \mathcal{O}\left(\frac{d^2 \log(p_{\min})}{p_{\min}}\right) \mathcal{O}\left(\frac{1}{n^2}\right) \\ \mathbb{E}\left[\sum_{i=1}^d (1 + \log p_i)(\tilde{p}_i - p_i) \sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{[(1-c)p_i + c\tilde{p}_i]^2}\right] \end{aligned}$$

$$\begin{aligned}
&\leq \mathbb{E} \left[\sum_{i=1}^d (1 + \log p_i) (\tilde{p}_i - p_i) \sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{(1-c)^2 p_i^2} \right] \\
&= \sum_{i=1}^d \frac{1 + \log p_i}{(1-c)^2 p_i^2} \mathbb{E}[(\tilde{p}_i - p_i)^4] + \sum_{(i,j):i \neq j} \frac{1 + \log p_i}{(1-c)^2 p_j^2} \mathbb{E}[(\tilde{p}_i - p_i)(\tilde{p}_j - p_j)^3] \\
&= \mathcal{O}\left(\frac{d^2 \log(p_{\min})}{(p_{\min})^2}\right) \mathcal{O}\left(\frac{1}{n^2} + \frac{\theta^{\frac{1}{2}}}{n^{\frac{3}{2}}} + \frac{\theta}{n} + \frac{\theta^{\frac{3}{2}}}{n^{\frac{1}{2}}} + \theta^2\right) \\
&= \mathcal{O}\left(\frac{d^2 \log(p_{\min})}{(p_{\min})^2}\right) \left[\mathcal{O}\left(\frac{1}{n^2}\right) + o\left(\frac{1}{n^2}\right)\right] \\
\mathbb{E} &\left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^2}{p_i} \sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{[(1-c)p_i + c\tilde{p}_i]^2} \right] \\
&\leq \mathbb{E} \left[\sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^2}{p_i} \sum_{i=1}^d \frac{(\tilde{p}_i - p_i)^3}{(1-c)^2 p_i^2} \right] \\
&= \sum_{i=1}^d \frac{\mathbb{E}[(\tilde{p}_i - p_i)^5]}{(1-c)^2 p_i^3} + \sum_{(i,j):i \neq j} \frac{\mathbb{E}[(\tilde{p}_i - p_i)^2 (\tilde{p}_j - p_j)^3]}{p_i (1-c)^2 p_j^2} \\
&= \mathcal{O}\left(\frac{d^2}{(p_{\min})^3}\right) \mathcal{O}\left(\frac{1}{n^{\frac{5}{2}}} + \frac{\theta^{\frac{1}{2}}}{n^2} + \frac{\theta}{n^{\frac{3}{2}}} + \frac{\theta}{n} + \frac{\theta^2}{n^{\frac{1}{2}}} + \theta^{\frac{5}{2}}\right) \\
&= \mathcal{O}\left(\frac{d^2}{(p_{\min})^3}\right) \left[\mathcal{O}\left(\frac{1}{n^{\frac{5}{2}}}\right) + o\left(\frac{1}{n^{\frac{5}{2}}}\right)\right]
\end{aligned}$$

We want all the resulting terms to go to zero faster than the dominant $\mathcal{O}\left(\frac{\sigma^2}{n}\right)$ term. Thus, we need to ensure that all five terms are $o\left(\frac{\sigma^2}{n}\right)$, which gives us five conditions to satisfy. It is easy to see that the stronger condition, which implies the remaining four ones is:

$$\frac{1}{(p_{\min})^2 n} = o(1) \Leftrightarrow p_{\min} = \omega\left(\frac{1}{\sqrt{n}}\right)$$

which completes the proof. \square

3.5 Noisy Majority Voting

As we argued, a major limitation of the Sharing the Sufficient Statistics approach is the high communication cost it incurs; recall that the joint frequency distribution of the attributes in a set \mathcal{A}' is a vector of dimension $|d_{\mathcal{A}'}|$, which can be quite large. To tackle this problem, we may be tempted to ask each data holder to instead share the local (noisy) value of the scoring function for each candidate attribute-parent pair. However, this approach is also problematic:

- Each data holder has to share an even larger number of scores, and specifically $\sum_{i=1}^{k+1} \binom{N}{i}$ scores.
- The analyst has to re-access the data to learn the parameters.
- The empirical mutual information is non-linear, so the analyst cannot compose the global score from local scores. A naive solution (which does not work well in practice) would be to take the mean or median of the local scores. An alternative solution, which we leave as an open question, would be to use a linear scoring function.

Motivated by these observations, we propose a second answer (Algorithm 3.6) to the *What to share?* question, which is based on the notion of majority voting from the distributed machine learning literature. Specifically, each data holder incrementally reports (a noisy version of) the highest mutual information attribute-parent pair that it would add to the Bayesian Network. The analyst collects all votes, and adds the most-voted pair to the structure.

Algorithm 3.6: Noisy Majority Voting	
Input:	Datasets D_1, \dots, D_M , BN degree k
Output:	BN structure \mathcal{G}
1	Initialize $\mathcal{G} = \emptyset$ and $V = \emptyset$
2	Arbitrarily select an attribute X from \mathcal{A} ; add (X, \emptyset) to \mathcal{G} and X to V
3	for $i=1$ to $N-1$ do
4	Initialize multi-set $votes = \emptyset$
5	for $j = 1$ to M do
6	QUERY (D_j): get $(X, \mathbf{Pa}(X))$ with the highest $\tilde{I}(X, \mathbf{Pa}(X))$, subject to: $X \in A \setminus V$ and $\mathbf{Pa}(X) \subseteq V$
7	Add $(X, \mathbf{Pa}(X))$ to $votes$
8	Find most-voted $(X, \mathbf{Pa}(X))$ in $votes$ (break ties arbitrarily)
9	Add $(X, \mathbf{Pa}(X))$ to \mathcal{G} and X to V
10	Return \mathcal{G}

In order to satisfy differential privacy, each data holder responds to the queries on its dataset using the exponential mechanism with the empirical mutual information as scoring function and with privacy budget $\frac{\epsilon_1}{N-1}$. This is, in fact, exactly what we describe in Algorithm 3.2 (line 12).

Once the Bayesian Network structure is known, Algorithm 3.4 is used to learn the parameters; the analyst now has to retrieve the N required frequency distributions to estimate the parameters (*retrieved* = False), and the data holders respond to its queries using the Laplace mechanism, with privacy budget $\frac{\epsilon_2}{N}$.

Theorem 3.7. *Let ϵ_1 and ϵ_2 be the total privacy budget that each data holder uses in responding to the analyst's queries, during the structure learning phase (Algorithm 3.6) and the parameter learning phase (Algorithm 3.4 with input *retrieved* = False), respectively. Then, the overall algorithm satisfies $(\epsilon_1 + \epsilon_2)$ -differential privacy for any dataset D_j ($j \in \{1, \dots, M\}$).*

3.6 Sharing the Noisy Model

In the final approach we examine, the data holders first learn a local (noisy) model, and then share it with the analyst. Therefore, the answer to the *What to share?* question is *the model*, and this is, in fact, equivalent with having each data holder locally run PrivBayes and publish the learned model. The analyst collects the local models and aggregates them using a technique inspired by knowledge probing; it generates a synthetic dataset using each local model (of size proportional to that of the local dataset that was used to learn the local model), and then it learns a global model based on the synthetic data.

Each data holder constructs its local model using Algorithms 3.2 and 3.3, which were shown to jointly satisfy ϵ -differential privacy (Theorem 3.1). Since differential privacy is immune to post-processing, we have the following theorem.

Algorithm 3.7: Sharing the Noisy Model**Input:** Datasets D_1, \dots, D_M , BN degree k **Output:** Global BN structure \mathcal{G}

- 1 Initialize $D_{synth} = \emptyset$
- 2 **for** $k = 1$ **to** M **do**
- 3 **QUERY**(D_j): get local model structure $\mathcal{G}^{(j)}$, parameters $\Theta^{(j)}$ and dataset size n_j
- 4 Generate $D_{j,synth}$ using Prior Sampling such that $|D_{j,synth}| = n_j$
- 5 Add $D_{j,synth}$ to D_{synth}
- 6 Run Algorithm 3.2 with input D_{synth} , k and without using the exponential mechanism in line 12, and return resulting \mathcal{G}

Theorem 3.8. *Assume that each data holder shares an ε -differentially private local model. Then, the global model that results by aggregating the local models according to Algorithm 3.7 also satisfies ε -differential privacy for any dataset D_j ($j \in \{1, \dots, M\}$).*

Notice that, once the analyst collects the local models, it does not need to access the data again. In particular, the analyst learns both the structure and the parameters of the global model using the synthetic data it generates. Thus, Algorithm 3.7 has the exact same two advantages as Algorithm 3.5; any algorithm can be used to address the structure learning optimization problem, and the parameter learning Algorithm 3.4 utilizes frequency distributions originating from the synthetic data.

3.7 Experimental Evaluation

In this section, we experimentally evaluate our algorithms, both on synthetic and on real-world data. To conduct our experiments, we build a system in Python programming language, utilizing popular libraries (namely: NumPy, pandas and NetworkX), as well as novel libraries we develop (prob_utils.py, pgm_utils.py).

We use two *baselines* for comparison.

- As a lower bound for performance, we use the maximum likelihood estimate of the distribution that assumes all attributes to be independent (naive-Bayes), with no privacy constraints.
- As an upper bound, we use a probabilistic graphical model learned with the greedy extension of the Chow-Liu algorithm (Algorithm 3.2), by centrally collecting all data and with no privacy constraints. For the case of synthetic data, we also evaluate for comparison the actual underlying Bayesian Network, which is known.

3.7.1 Evaluation Metric

The metric we use in our experiments is the *cross entropy* between the actual data distribution, \mathbf{p} , and the distribution encoded by the learned Bayesian Network, $\mathbf{p}_{\mathcal{B}}$. The cross entropy between the aforementioned two distributions, over the set of attributes \mathcal{A} (with $|d_{\mathcal{A}}| = d$) is defined as:

$$\begin{aligned}
 H(\mathbf{p}, \mathbf{p}_{\mathcal{B}}) &= \mathbb{E}_{\mathbf{p}}[-\log_2 \mathbf{p}_{\mathcal{B}}] \\
 &= -\sum_{i=1}^d p_i \log_2 p_{\mathcal{B}i}
 \end{aligned}$$

$$\begin{aligned}
&= - \sum_{i=1}^d p_i \log_2 p_i + p_i \log_2 \frac{p_{\mathcal{B}i}}{p_i} \\
&= H(\mathbf{p}) + D_{KL}(\mathbf{p} \parallel \mathbf{p}_{\mathcal{B}})
\end{aligned}$$

where:

- $H(\mathbf{p})$ is the entropy of \mathbf{p} , which measures the expected number of bits needed to encode an element from $d_{\mathcal{A}}$, using a coding scheme that is optimized for \mathbf{p} ,
- $D_{KL}(\mathbf{p} \parallel \mathbf{p}_{\mathcal{B}})$ is the Kullback-Leibler divergence (relative entropy) of $\mathbf{p}_{\mathcal{B}}$ from \mathbf{p} , which measures the expected number of extra bits needed to encode an element from $d_{\mathcal{A}}$, using a coding scheme that is optimized for $\mathbf{p}_{\mathcal{B}}$ rather than \mathbf{p} ,

and, thus, $H(\mathbf{p}, \mathbf{p}_{\mathcal{B}})$ measures the expected number of bits needed to encode an element from $d_{\mathcal{A}}$, using a coding scheme that is optimized for the learned distribution $\mathbf{p}_{\mathcal{B}}$, rather than the actual distribution \mathbf{p} .

In our experiments with synthetic data, the actual data distribution \mathbf{p} is known, so we are able to use the aforementioned formula. However, when experimenting with real-world data, \mathbf{p} is unknown, so we instead use a Monte Carlo estimate of the cross entropy, namely the *empirical cross entropy* $\hat{H}(D_{test}, \mathbf{p}_{\mathcal{B}})$. To compute $\hat{H}(D_{test}, \mathbf{p}_{\mathcal{B}})$, we reserve a fraction (20%) of the total dataset D which we do not use when learning the Bayesian Network model $p_{\mathcal{B}}$. Therefore, the reserved, test dataset D_{test} consists of independent data points that were randomly drawn from D , and hence can be assumed to follow the same distribution \mathbf{p} . We estimate the empirical cross entropy as:

$$\hat{H}(D_{test}, \mathbf{p}_{\mathcal{B}}) = - \frac{1}{|D_{test}|} \sum_{\mathbf{x} \in D_{test}} \log_2 \mathbf{p}_{\mathcal{B}}(\mathbf{x})$$

Before proceeding, we make an additional, important remark. In theory, we have assumed that there exist no *zero-probability events* in \mathbf{p} , so that all elements in $d_{\mathcal{A}}$ have non-zero probability and, hence, $p_i > 0, \forall i \in \{1, \dots, d\}$. In practice, however, the learned distribution $\mathbf{p}_{\mathcal{B}}$ may contain zero-probability events, since it is the maximum likelihood estimate of \mathbf{p} , computed using a noisy version $\tilde{\mathbf{c}}$ of the observed frequency distribution \mathbf{c} . Even if all elements in $d_{\mathcal{A}}$ are observed (which is unlikely when d is large), so that $c_i > 0, \forall i \in \{1, \dots, d\}$, the noise added by the Laplace mechanism may cause $\tilde{c}_i = 0$ for some i (it actually may cause $\tilde{c}_i \leq 0$, but we replace negative frequencies with 0), and, thus, $p_{\mathcal{B}i} = 0$. If this is the case, it is not hard to see that both $H(\mathbf{p}, \mathbf{p}_{\mathcal{B}}) = \infty$ and $\hat{H}(D_{test}, \mathbf{p}_{\mathcal{B}}) = \infty$. We tackle this problem by using *additive (add-one) smoothing*, so $\forall i$, we estimate the smoothed probability $\tilde{p}'_{\mathcal{B}i}$ based on the smoothed frequency distribution:

$$\tilde{c}'_i = \tilde{c}_i + 1 \Rightarrow \tilde{p}'_{\mathcal{B}i} = \frac{\tilde{c}'_i}{n + d}$$

3.7.2 Synthetic Data

In our first set of experiments, we use a *synthetic dataset*, where each data point consists of a (relatively small) set of attributes $\mathcal{A} = \{X_0, \dots, X_6\}$. The first six attributes ($X_0 - X_5$) are binary, whereas the seventh attribute (X_6) is ternary, so the joint domain of the 7 attributes has size $d = |d_{\mathcal{A}}| = 2^6 \times 3 = 192$. The actual data distribution $\mathbf{p}_{\mathcal{A}} = \mathbf{p}$ can be encoded by the Bayesian Network illustrated in Figure 3.3.

A tedious (but straightforward) computation gives that $H(\mathbf{p}) = 5.12$ bits, whereas for a uniform distribution \mathbf{p}' , $H(\mathbf{p}') = -\log_2(\frac{1}{192}) = 7.58$ bits. This illustrates that \mathbf{p} is far from perfectly random, and, by identifying the dependencies that are present in \mathbf{p} , we

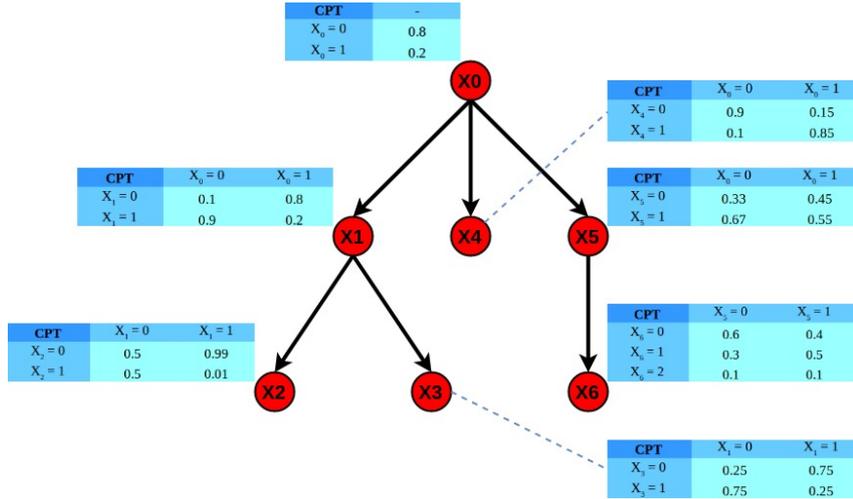


Figure 3.3: Actual data distribution

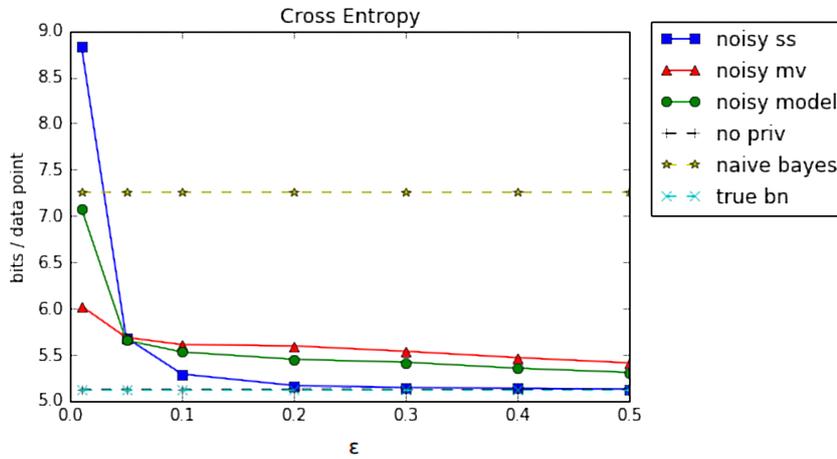


Figure 3.4: Cross entropy for varying privacy budget ϵ .
 Number of data holders: $M = 3$. Bayesian Network degree: $k = 1$.

may be able to save a significant amount of bits! If we manage to perfectly approximate \mathbf{p} , then we expect $H(\mathbf{p}, \mathbf{p}_B) = H(\mathbf{p}, \mathbf{p}) = H(\mathbf{p}) + D_{KL}(\mathbf{p}||\mathbf{p}) = 5.12 + 0 = 5.12$ bits.

In all our experiments, we average each measurement over 100 independent repetitions. In each repetition, we use *prior sampling* to generate a synthetic dataset of $n \approx 40,000$ data points, according to the distribution we described. We partition the n data points among the M data holders using a *round-robin partitioning* scheme, so that each data holder has $\lfloor \frac{n}{M} \rfloor$ data points.

In Figures 3.4 and 3.5 we observe that all algorithms perform significantly better than the naive-Bayes, and approach the optimum value of bits per data point of the true Bayesian Network as ϵ increases; Algorithm 3.5 (noisy ss) asymptotically outperforms the other algorithms, but is more sensitive to noise. We also notice that the non-private algorithm accurately learns the true Bayesian Network.

The key thing to notice in Figure 3.6 is the robustness of Algorithm 3.6 to the number of data holders. On the contrary, the performance of Algorithms 3.5 and 3.7 deteriorates as the number of data holders increases.

Finally, comparing Figures 3.4 and 3.5, and as we also observe in Figure 3.7, as the Bayesian Network degree increases, the performance of Algorithm 3.5 (noisy ss) worsens. This is expected, since a higher degree Bayesian Network implies the use and perturbation

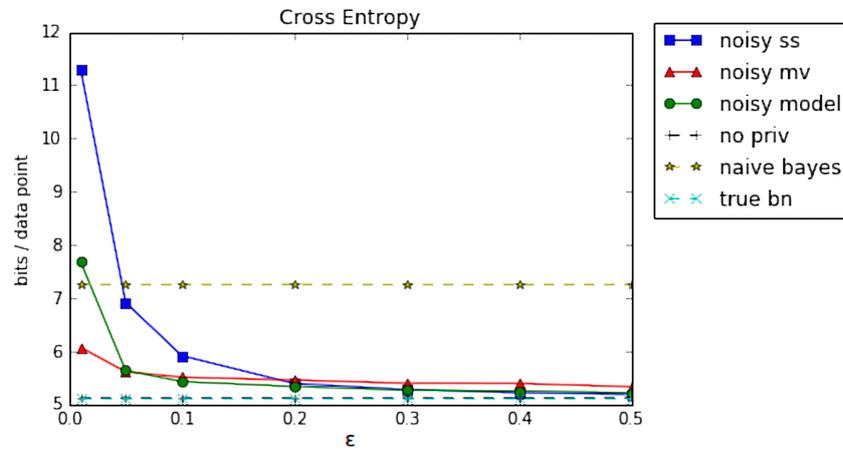


Figure 3.5: Cross entropy for varying privacy budget ϵ .
 Number of data holders: $M = 3$. Bayesian Network degree: $k = 2$.

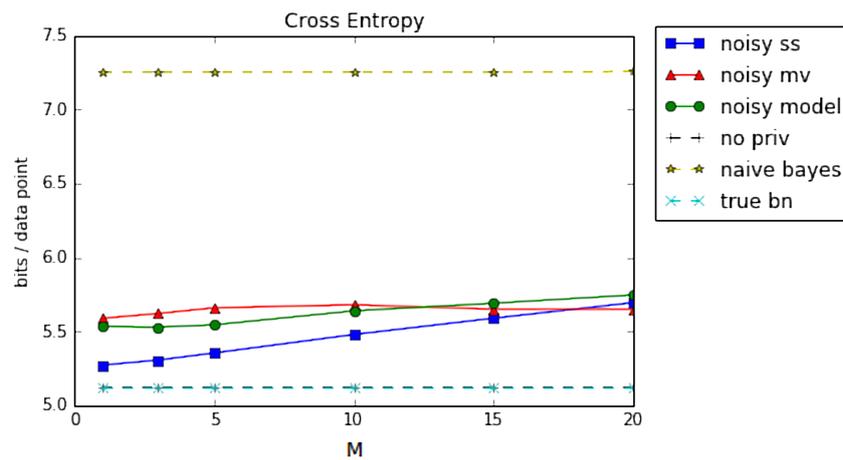


Figure 3.6: Cross entropy for varying number of data holders M .
 Privacy budget: $\epsilon = 0.1$. Bayesian Network degree: $k = 1$.

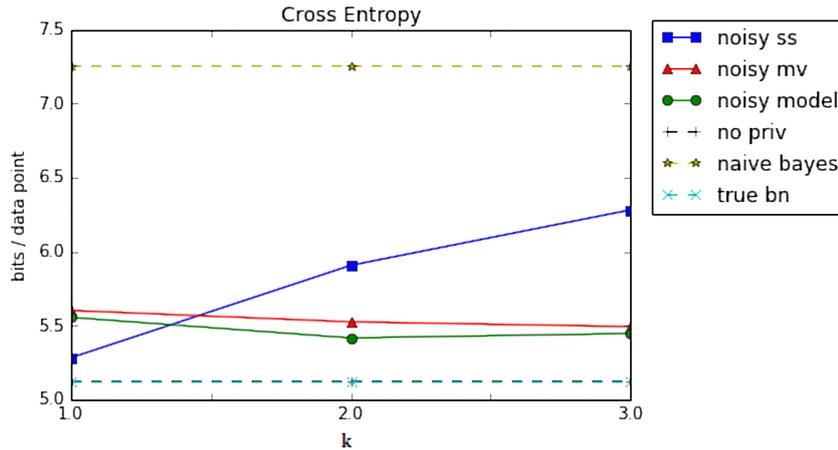


Figure 3.7: Cross entropy for varying Bayesian Network degree. Privacy budget: $\epsilon = 0.1$. Number of data holders: $M = 3$.

of higher dimensional distributions. In contrast, Algorithms 3.6 (noisy mv) and 3.7 (noisy model) slightly improve as the degree increases.

3.7.3 Real-World Data

In our second set of experiments, we utilize *real-world data* from the Diabetes dataset ([55]) of the UCI Machine Learning Repository.

The dataset was extracted from a large database that represents 10 years of clinical care at US hospitals and includes over 50 attributes representing patient and hospital outcomes. In particular, a total of 101,766 records (data points) were selected, based on the following inclusion criteria:

- It is an inpatient encounter (a hospital admission).
- It is a diabetic encounter, that is, one during which any kind of diabetes was entered to the system as a diagnosis.
- The length of stay was at least 1 day and at most 14 days.
- Laboratory tests were performed during the encounter.
- Medications were administered during the encounter.

Next, attribute (feature) selection was performed (by clinical experts) and only attributes that were potentially associated with the diabetic condition or management were retained (55 attributes, describing the diabetic encounters, including demographics, diagnoses, diabetic medications, number of visits in the year preceding the encounter, and payer information). The final attributes in each record include (among others): patient id, race, gender, age, admission type, time in hospital, medical specialty of admitting physician, number of lab test performed, HbA1c test result, diagnosis, number of medication, diabetic medications, number of outpatient, inpatient, emergency visits in the year before the hospitalization, etc.

We perform additional pre-processing to the dataset, as we describe in the next three steps.

- First, we remove identification-related attributes (e.g. patient id) and attributes with a high percentage of missing values (e.g. weight - 97 percent of values are missing). We are left with 43 attributes.

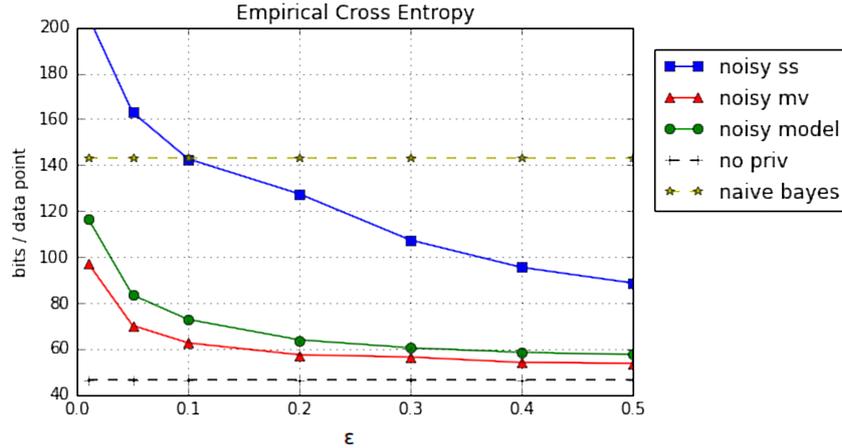


Figure 3.8: Cross entropy for varying privacy budget ϵ .
 Number of data holders: $M = 3$. Bayesian Network degree: $k = 1$.

- Second, we identify 3 diagnosis-related attributes with large domains (almost 1,000 values per attribute) and compress their domains into 9 general groups of diagnoses, as described in [55]. The final domain sizes vary from 2 to 118; we have a total of almost 10^{32} possible tuples, so (inevitably) we will not observe a huge number of events.
- Third, we remove records that refer to the same patient. Such records cannot be considered statistically independent, and, also, differential privacy is guaranteed at a record level, so the privacy of a patient that has multiple records is violated. We note that one solution that would allow each patient to have multiple records is the group differential privacy guarantee. We are left with approximately 70,000 records.

In all our experiments, we average each measurement over 20 independent repetitions. We again split the data among the data holders using a round-robin partitioning scheme. We first experiment for varying privacy budget ϵ (Figure 3.8). We observe that Algorithm 3.5 (noisy ss) is significantly worse; it appears that it is highly sensitive to the data dimension and the domain sizes of the attributes, which are significantly increased compared to the synthetic data we experimented with. Algorithm 3.6 (noisy mv) outperforms the other two algorithms and approaches the performance of the non-private algorithm.

We next experiment for varying number of data holders. Figure 3.9 confirms the robustness of Algorithm 3.6 (noisy mv) to the number of data holders, and validates that -for complex datasets- Algorithm 3.6 (noisy mv) is the clear winner. The performance of Algorithm 3.7 (noisy model) significantly worsens as number of holders increases. When the dataset is split among a larger number of data holders, each holder possesses a smaller dataset, and therefore the quality of the local models deteriorates. Finally, the behavior of Algorithm 3.5 (noisy ss) may at first seem strange, as it differs from the algorithm's behavior in the corresponding experiment with synthetic data (Figure 3.6). This behavior illustrates that when the data dimension and the domain sizes of the attributes are both large, the impact of the number of data holders is not critical. A possible explanation for the improvement in performance is that, as the number of data holders increases, more zero-mean noise terms are added to the frequency distributions, and hence the noise cancels out.

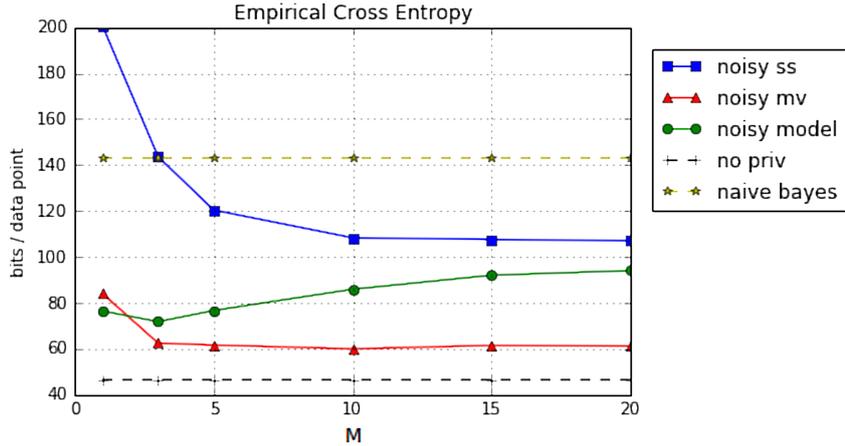


Figure 3.9: Cross entropy for varying number of data holders M . Privacy budget: $\varepsilon = 0.1$. Bayesian Network degree: $k = 1$.

3.8 Discussion

To summarize, in this chapter we addressed the problem of distributed learning of Bayesian Networks with differential privacy. After formally describing our model and identifying the challenges that arise when moving from a centralized to a distributed environment, we examined three solutions, namely Sharing the Noisy Sufficient Statistics, Noisy Majority Voting and Sharing the Noisy Model.

Our detailed experimental evaluation indicates that Noisy Majority Voting is robust to all the parameters we examined (privacy budget, number of data holders, Bayesian Network degree), and significantly outperforms the other algorithms when the data are high-dimensional. We also note that the behavior of Sharing the Noisy Sufficient Statistics is consistent with our theoretical analysis, which predicted that:

- When the privacy budget ε exceeds a certain threshold, our differentially private entropy estimates converge to the non-private ones, and thus the learned distribution accurately approximates the non-private one.
- When either the number of data holders M or the Bayesian Network degree k increases, our differentially private entropy estimates worsen, and so does the performance of the algorithm.

Chapter 4

Pan-Private Stream Density Estimation

We work in the *data stream model* ([30], [17]). The input is represented as a finite sequence of tuples (from some finite domain) which is not available for random access, but instead arrives one at a time in a stream. We thus deal with dynamic data, and we assume that each input tuple corresponds to some user from a universe \mathcal{U} , who is identified by a unique key (user id).

More specifically, in what follows, we consider the *cashier-register* streaming model: each update is of the form $(user\ id, update\ value)$, where the update values are strictly positive integers. For further simplicity, we permit only unit updates, so the stream can be viewed as a sequence of user id's. Each user is mapped to a state that accumulates the user's updates. We note that, although our work does not apply to the (most general) turnstile streaming model, where negative updates are also allowed, it can be extended to the special case where both unit insertions and unit deletions are allowed, and, at any time step, each user's state can be either 1 (user is present) or 0 (user is absent).

In general, *streaming algorithms* should be able to operate in a single pass (each tuple is examined at most once in fixed arrival order) and in real-time (each tuple's processing time must be low), and should require small space. For most problems, it is not possible to offer exact answers while satisfying these requirements, so a common approach is to offer (α, β) -approximate answers, that is, with probability at least $1 - \beta$ the computed answer is within an additive / multiplicative factor α of the actual answer.

The particular problem we examine is the following: given an input data stream, we want to estimate its *density*, that is, the fraction of \mathcal{U} that actually appears in the stream. This statistic is closely related to the well-studied distinct count (a.k.a. 0^{th} frequency moment and Hamming norm), which expresses the number of distinct users that appear in the stream. The density is more convenient in that it is normalized and takes values in the $[0, 1]$ interval.

Additionally, given that the focus of this thesis is the design of privacy-preserving algorithms, we also want to protect the *privacy* of the users that appear in the stream. We would like to offer a strong privacy guarantee, such as differential privacy, adjusted to the streaming model.

Before proceeding with the technical part of this chapter, we briefly present two potential applications, that motivate the study of this problem.

- We borrow this example from Dwork [20]. Consider a website for an *epidemic disease self-assessment*. Users interact with the website to figure out whether the symptoms they are experiencing may be indicative of the disease. An epidemiological study

would be interested in answering queries, like the following:

```
SELECT COUNT( DISTINCT ip_addr )      SELECT COUNT( DISTINCT ip_addr )
FROM data_stream                      FROM data_stream
WHERE symptom1 = "YES", ...
```

An individual’s medical record is undoubtedly an extremely sensitive piece of information, so knowing even the presence of a user in such a website would constitute a major privacy breach.

- We mentioned that our work can be extended to support both unit insertions and deletions. This extension is especially useful in the context of *dynamic graphs*; user id’s are replaced by edge id’s, and each update either inserts or deletes an edge. An example sequence of updates is the following:

$$((edge_i, insert), (edge_j, insert), (edge_i, delete), \dots)$$

The density of an undirected graph $\mathcal{G} = (V, E)$ is defined as: $d(\mathcal{G}) = \frac{2|E|}{|V|(|V|-1)}$ and expresses the fraction of edges that are present. For a dynamic graph, created through a stream \mathcal{S} , it is not hard to see that $d(\mathcal{G}) = d(\mathcal{S})$.

If our focus is, for example, on social graphs, which depict personal relations of Internet users, then any analysis performed on the graph (such as the density estimation) must not leak information about the existence of individual edges.

Table 4.1: Table of notations

Notation	Description
\mathcal{U}	Universe of users, w.l.o.g.: $\mathcal{U} = \{1, \dots, \mathcal{U} \}$
T	Stream length
s_t	Stream update at time step t : $t \in \{1, \dots, T\}$, $s_t \in \mathcal{U}$
\mathcal{S}	Input stream: $\mathcal{S} = (s_1, \dots, s_T)$
$a_u(t)$	State of user $u \in \mathcal{U}$ after t updates: $a_u(t) = \sum_{i=1}^t \mathbf{1}(s_i = u)$
$\mathbf{a}(t)$	State vector after t updates: $\mathbf{a}(t) = [a_1(t) \dots a_{ \mathcal{U} }(t)]$
$F_0(\mathcal{S})$	Number of distinct users in \mathcal{S} : $F_0(\mathcal{S}) = \ \mathbf{a}(T)\ _0$
$d(\mathcal{S})$	Density of \mathcal{S} : $d(\mathcal{S}) = \frac{F_0(\mathcal{S})}{ \mathcal{U} }$

4.1 Differential Privacy in the Streaming Model

So far, we have considered differential privacy on static datasets, and have assumed that an adversary only has access to the output of the privacy-preserving algorithm. In the model we examine in this chapter, where the data arrive dynamically in a stream, *two privacy models* (definitions) have been developed ([23],[22],[20]).

- **Pan-Privacy.** In pan-privacy, the internal state of the algorithm is also differentially private, as is the joint distribution of the internal state and the output. This protects the privacy of the data in case the data holder is subject to compulsory, non-private data release, because of a subpoena, or faces an intrusion from a hacker.

- **Continual Observation.** In continual observation, the goal is to continually monitor and report statistics about events that occur dynamically in discrete time steps, and to ensure that the output sequence is differentially private. This is the case, for example, when monitoring incidence of influenza, traffic conditions and search trends.

The two notions provide orthogonal guarantees, and some work addresses pan-privacy under continual observation. Let us present a simple example:

Example. Consider an algorithm that takes as input a data stream $\mathcal{S} = (s_1, \dots, s_T)$, and upon arrival of each element s_i , outputs some estimate $f_i = f(s_1, \dots, s_i)$. Pan-privacy ensures that the state of the algorithm is subject to a differential privacy constraint, so privacy is preserved against an adversary that observes the internal state and a (single) output of the algorithm, say f_j for some $j \in \{1, \dots, T\}$. We can think of this single output as the final output, so that $j = T$. Differential privacy under continual observation ensures that the entire sequence of outputs f_1, \dots, f_T satisfies differential privacy. The two definitions can be combined, by ensuring that the internal state, the output sequence and their joint distribution, all satisfy differential privacy.

A very important notion in differential privacy is that of *adjacency*; as we already argued, two (static) datasets are adjacent if they differ on a single record. But how should we define adjacency in the streaming model? When are two data streams considered to be adjacent? We (informally) introduce the following two notions:

- **Event-level privacy.** Two data streams are considered event-level adjacent if they differ on a single update s_t for some $t \in \{1, \dots, T\}$. Event-level privacy protects the privacy of s_t and an adversary cannot distinguish whether s_t did or did not appear. In the epidemic disease self-assessment website, event-level privacy protects the privacy of single visits to the website. In a dynamic graph, event-level privacy protects the privacy of a single insertion / deletion of an edge between any two nodes.
- **User-level privacy.** Two data streams are considered user-level adjacent if they differ on all (or some) updates that refer to a user $u \in \mathcal{U}$. User-level privacy protects the privacy of u and an adversary cannot distinguish whether u did or did not ever appear, independently of the actual number of appearances of u . In the epidemic disease self-assessment website, user-level privacy protects the privacy of all the visits of a user (IP-address) to the website. In a dynamic graph, user-level privacy protects the privacy of any edge, regardless of how many times this edge has been inserted deleted.

As in the case of static datasets, the word “differ” in the phrases “differ on a single update” and “differ on all (or some) updates that refer to a user” can be interpreted in two ways. According to the first interpretation, the different updates are present in the one data stream and absent in the other; in the traditional differential privacy model, this interpretation would lead to unbounded differential privacy. According to the second interpretation, the different updates are present in both data streams, but refer to two different users; in the traditional differential privacy model, this interpretation would lead to bounded differential privacy.

The privacy level affects the amount of perturbation used, so a much stronger guarantee like user-level privacy requires excessively more perturbation. Kellaris et al. [42] attempt to bridge the gap between event-level and user-level privacy and develop a framework in-between; they introduce w -event privacy, which protects any event sequence occurring

within a window of w time steps.

For completeness, and before proceeding with our specific problem (density estimation), we present the basic definitions for both pan-privacy and continual observation, along with some related work.

4.1.1 Pan-Privacy

As we already argued, the definition of pan-privacy aims to protect against an adversary that can, on rare occasions, observe the internal state of the algorithm. This is, of course, in addition to the standard -for differential privacy- assumptions about the adversary having arbitrary control over the input, arbitrary prior knowledge and arbitrary computational power.

Before formally defining the model, we remark that ordinary streaming algorithms based on sampling and sketching techniques do not provide the pan-privacy guarantee. Sampling techniques maintain information about a subset of the users, so an intruder with access to the sample (the internal state of the algorithm) would violate the privacy of the sampled users. Sketching techniques, like the FM Sketch [26] and the Count-Min Sketch [14], which are based on hashing also cannot protect the privacy of the users against an adversary who has access to the hash functions used; the hash functions are part of the internal state.

Definition 4.1 (Adjacency). *Data streams \mathcal{S} and \mathcal{S}' are (user-level) adjacent, denoted as $\text{adj}(\mathcal{S}, \mathcal{S}')$, if they differ only in the presence or absence of all (or some) occurrences of a single user $u \in \mathcal{U}$.*

Notice that in Definition 4.1, we use the first interpretation of the word “differ” in defining adjacent data streams, following Dwork et al. [23]. Therefore, two adjacent data streams cannot have the same length. Mir et al. [46], for instance, adopt the second interpretation, as one of the pan-private algorithms they propose requires that the total sum of updates over all users stays the same; they replace all (or some of) the updates that refer to user u by updates that refer to another user u' (with the same total sum).

Definition 4.2 (Pan-privacy). *Let \mathbf{Alg} be a randomized algorithm. Let \mathcal{I} denote its set of internal states, and let \mathcal{O} denote its set of possible outputs. Then \mathbf{Alg} , mapping data streams of finite length T to the range $\mathcal{I} \times \mathcal{O}$, is ε -pan-private against a single intrusion, if for all sets $I \subseteq \mathcal{I}$ and $O \subseteq \mathcal{O}$, and for all pairs of adjacent data streams $\mathcal{S}, \mathcal{S}'$*

$$\mathbb{P}[\mathbf{Alg}(\mathcal{S}) \in (I, O)] \leq e^\varepsilon \mathbb{P}[\mathbf{Alg}(\mathcal{S}') \in (I, O)]$$

where the probability space is over the coin flips of \mathbf{Alg} .

We remark that Definition 4.2 speaks only of a single intrusion. To handle multiple intrusions, we must consider interleavings of observations of internal states and output sequences. We also have to differentiate between announced and unannounced intrusions. In the former case (e.g. subpoena), the algorithm knows that an intrusion occurred, so it can re-randomize its state and handle multiple announced intrusions. In the latter case (e.g. hacking), the algorithm can only tolerate a single unannounced intrusion and strong negative results have been proved for even two unannounced intrusions.

So far, two works have addressed the challenge of developing pan-private streaming algorithms. They both examine variants of the same problems, applying different techniques. Dwork et al. [23] work in the cashier-register streaming model and develop algorithms based on sampling and randomized response. Mir et al. [46] work in the turnstile

streaming model and rely on both existing and novel sketches; they also develop a general noise-calibrating technique for sketches. The problems examined are the following:

- Density estimation / Distinct count: the fraction of \mathcal{U} that appears / the number of users with nonzero state.
- t -cropped mean / t -cropped first moment: the average / sum, over all users, of the minimum of t and the number of appearances of the user.
- Fraction of k -heavy hitters / k -heavy hitters count: the fraction / number of users that appear at least k times.
- t -incidence estimation: the fraction of users that appear exactly t times.

In our work, we aim to offer the user-level pan-privacy guarantee (against a single intrusion). Our objective is to estimate the density of the given input stream, and our algorithms are based on the density estimator of Dwork et al. [23]. The algorithm proposed by Mir et al. [46] for the (similar) distinct count problem relies on the ℓ_0 Sketch, which is due to Cormode et al. [15] and utilizes a family of distributions called p -stable (Indyk [39]). Despite being particularly interesting theoretically, the pan-private algorithm of Mir et al. [46] is extremely impractical as it involves an application of the exponential mechanism, which requires sampling twice from a space of 2^S sketches (where S is the bit size of the sketch). To make matters worse, in order to evaluate the exponential mechanism's scoring function, a Hamming norm minimization problem is solved for every single sketch, which translates to solving 2^S problems.

4.1.2 Continual Observation

As we already stated, differential privacy under continual observation addresses the need for privacy in applications that involve repeated computations over dynamic data, and require continually producing updated outputs. We stick to Definition 4.1 for adjacent data streams, which means that the following definition (4.3) refers to user-level differential privacy under continual observation.

Definition 4.3 (Differential privacy under continual observation). *Let \mathbf{Alg} be a randomized algorithm, and let \mathcal{O} denote its set of possible output sequences. Then \mathbf{Alg} , mapping data streams of finite length T to the range \mathcal{O} , is ϵ -differentially private under continual observation if for all sets $O \subseteq \mathcal{O}$, for all pairs of adjacent data streams $\mathcal{S}, \mathcal{S}'$, and for all $t \in \{1, \dots, T\}$,*

$$\mathbb{P}[\mathbf{Alg}(\mathcal{S}(t)) \in O] \leq e^\epsilon \mathbb{P}[\mathbf{Alg}(\mathcal{S}'(t)) \in O]$$

where the probability space is over the coin flips of \mathbf{Alg} .

Although our focus is on algorithms that produce a single output, we remark that Dwork ?? shows how to modify the pan-private density estimator developed by Dwork et al. [23] (which, we repeat, is the baseline for our work), to produce output continually. The resulting continual observation density estimator guarantees *user-level pan-privacy under continual observation*. A similar technique can be applied to all the algorithms we develop to allow them to produce continual output while still preserving privacy.

Dwork [20] and Dwork et al. [22] initiate the study of differential privacy under continual observation; they examine the problem of continually releasing differentially private counts (both event-level and user-level), and they also optionally guarantee pan-privacy. Chan et al. [8] address the exact same counter problem (independently), but they only focus on event-level privacy, and Bolot et al. [6] extend their work by proposing decayed

sums, which emphasize on recent data more than the data of the past (in addition to proposing a new notion of decayed privacy). Chan et al. [9] report the heavy hitters instead of the counter values in various models (e.g. untrusted aggregator), also ensuring event-level privacy.

A few works examine problems beyond counting. Fan et al. [25] propose a framework to release real-time aggregate statistics under differential privacy, based on filtering and adaptive sampling; they focus on user-level privacy. The next authors examine event-level privacy. Friedman et al. [28] develop a framework for privacy-preserving distributed stream monitoring based on the geometric method [53], which enables monitoring arbitrary functions over statistics derived from distributed data streams. Bonomi et al. [7] compute the longest increasing subsequence in a differentially private manner, and Upadhyay [58] examines the connections between (ϵ, δ) -differential privacy and linear algebra in the streaming model, for low-rank approximation, linear regression and matrix multiplication.

4.2 Dwork’s Density Estimator

Dwork et al. [23] proposed the first user-level pan-private algorithm for the density estimation problem (Algorithm 4.1). In this section, we provide a detailed presentation and analysis of their algorithm, which we call Dwork’s Density Estimator.

Their (randomized) algorithm takes as input the data stream \mathcal{S} whose density we wish to estimate, as well as the desired privacy budget and accuracy parameters. For simplicity, we assume that the length of the stream is known in advance (that is, we know that the output will be requested after T updates), but this assumption does not affect the analysis of the algorithm, and can easily be dropped, so that the algorithm outputs the stream density when it receives a special signal.

The first random choice is a sampling step - the algorithm maintains information only about a random subset of the users, in order to keep its state small. The state of the algorithm is a bitarray with one entry per sampled user. The entries are random bits, generated as described below; this is the second random choice of the algorithm. For users that have not appeared in \mathcal{S} , their entry is drawn from a uniform Bernoulli distribution, while for users that have appeared, their entry is drawn from a slightly biased (towards 1) Bernoulli distribution, no matter how many times they have appeared. The two distributions should be close enough to guarantee that the state satisfies differential privacy, but far enough to allow collection of aggregate statistics about the fraction of users that appear at least once. The last random choice is performed via the use of the Laplace mechanism, when outputting the final density estimate, which guarantees that the output also satisfies differential privacy.

We now make two important remarks concerning potential extensions of Algorithm 4.1. The techniques described in the remarks apply (slightly modified) to all the algorithms we present, so we do not revisit them in our work.

- Algorithm 4.1 can tolerate a single (announced or unannounced) intrusion. Dwork et al. [23] show how to handle multiple announced intrusions, by re-randomizing the state after each intrusion has occurred.
- Algorithm 4.1 works in the cashier-register streaming model, that is, once a user appears in the stream, it cannot be deleted. However, as we mentioned in the introduction of this chapter, our work also applies to the case where a user u may be both inserted and deleted (later on) from the stream. In particular, if an update of the form $(u, insert)$ arrives, u ’s bit is drawn Bernoulli($\frac{1}{2} + \frac{\epsilon}{4}$), whereas if an update of the form $(u, delete)$ arrives, u ’s bit is re-drawn Bernoulli($\frac{1}{2}$). This allows Algorithm 4.1 to perform pan-private graph density estimation as well.

Algorithm 4.1: Dwork's Density Estimator**Input:** Data stream \mathcal{S} , Privacy budget ε , Accuracy parameters (α, β) **Output:** Density $\tilde{d}(\mathcal{S})$

- 1 Pick $m = \mathcal{O}(\frac{1}{\varepsilon^2 \alpha^2} \log \frac{1}{\beta})$ (or compute m^* as described in subsection 4.2.3 and set $m = m^*$)
- 2 Sample a random subset $M \subseteq \mathcal{U}$ of m users (without replacement) and define an arbitrary ordering over M
- 3 Create a bitarray $\mathbf{b} = [b_1 \dots b_m]$ and map $M[i] \rightarrow b_i, \forall i \in \{1, \dots, m\}$
- 4 Initialize \mathbf{b} randomly: $b_i \sim \text{Bernoulli}(\frac{1}{2}), \forall i \in \{1, \dots, m\}$
- 5 **for** $t = 1$ **to** T **do**
- 6 **if** $s_t \in M$ **then**
- 7 Find $i : M[i] = s_t$
- 8 Re-sample: $b_i \sim \text{Bernoulli}(\frac{1}{2} + \frac{\varepsilon}{4})$
- 9 Return $\tilde{d}(\mathcal{S}) = \frac{4}{\varepsilon} (\frac{1}{m} \sum_{i=1}^m b_i - \frac{1}{2}) + \text{Laplace}(0, \frac{1}{\varepsilon m})$

4.2.1 Privacy Analysis

We first examine the privacy guarantees of Algorithm 4.1. We have the following theorem.

Theorem 4.1. *Assume $\varepsilon \leq \frac{1}{2}$. Then Algorithm 4.1 satisfies 2ε -pan-privacy.*

Proof. Let $\mathcal{S}, \mathcal{S}'$ be two adjacent streams that differ on all occurrences of user $u \in \mathcal{U}$. Assume w.l.o.g. that $u \in \mathcal{S}$ and $u \notin \mathcal{S}'$.

Firstly, the *state* of Algorithm 4.1 satisfies ε -differential privacy. All the information that Algorithm 4.1 stores as its state is the bitarray \mathbf{b} . We distinguish the following two cases:

- If $u \notin M$: perfect privacy is guaranteed, as no information is stored on user u .
- If $u \in M$: let (w.l.o.g.) b_1 be the entry that corresponds to u in the bitarray. Then, assuming that u has already appeared in the stream when the adversary views \mathbf{b} , $b_1(\mathcal{S})$ is drawn from $\text{Bernoulli}(\frac{1}{2} + \frac{\varepsilon}{4})$ and $b_1(\mathcal{S}')$ is drawn from $\text{Bernoulli}(\frac{1}{2})$. We thus have to bound the following probability ratios according to the differential privacy definition:

$$\begin{aligned} \frac{\mathbb{P}(\mathbf{b}(\mathcal{S}) = [1 \ b_2 \ \dots \ b_m])}{\mathbb{P}(\mathbf{b}(\mathcal{S}') = [1 \ b_2 \ \dots \ b_m])} &= \frac{\mathbb{P}(b_1(\mathcal{S}) = 1)}{\mathbb{P}(b_1(\mathcal{S}') = 1)} \\ &= \frac{\frac{1}{2} + \frac{\varepsilon}{4}}{\frac{1}{2}} = 1 + \frac{\varepsilon}{2} \\ \frac{\mathbb{P}(\mathbf{b}(\mathcal{S}) = [0 \ b_2 \ \dots \ b_m])}{\mathbb{P}(\mathbf{b}(\mathcal{S}') = [0 \ b_2 \ \dots \ b_m])} &= \frac{\mathbb{P}(b_1(\mathcal{S}) = 0)}{\mathbb{P}(b_1(\mathcal{S}') = 0)} \\ &= \frac{\frac{1}{2} - \frac{\varepsilon}{4}}{\frac{1}{2}} = 1 - \frac{\varepsilon}{2} \end{aligned}$$

Noting that $e^{-\varepsilon} \leq 1 + \frac{\varepsilon}{2} \leq \sum_{k=0}^{\infty} \varepsilon^k k! = e^\varepsilon$ and $e^{-\varepsilon} \leq 1 - \frac{\varepsilon}{2} \leq e^\varepsilon, \forall \varepsilon \in [0, \frac{1}{2}]$, it becomes clear that user u is guaranteed ε -differential privacy against an adversary that observes u 's entry in \mathbf{b} .

The *output* of Algorithm 4.1 (conditioned on the state) also satisfies ε -differential privacy, as it is computed by independently applying the Laplace mechanism. Specifically, the sensitivity of \tilde{d} is

$$\Delta \tilde{d} = \max_{\mathcal{S}, \mathcal{S}': \text{adj}(\mathcal{S}, \mathcal{S}')} |\tilde{d}(\mathcal{S}) - \tilde{d}(\mathcal{S}')| = \frac{1}{m}$$

so we add noise $\sim \text{Laplace}(0, \frac{1}{\varepsilon m})$.

The *overall* Algorithm 4.1 satisfies 2ε -pan-privacy. Specifically, for all possible states (bitarrays) $\underline{\mathbf{b}}$ and outputs (estimated densities) $\underline{\tilde{d}}$, it holds that

$$\begin{aligned} \mathbb{P}(\mathbf{b}(\mathcal{S}) = \underline{\mathbf{b}}, \tilde{d}(\mathcal{S}) = \underline{\tilde{d}}) &= \mathbb{P}(\mathbf{b}(\mathcal{S}) = \underline{\mathbf{b}}) \mathbb{P}(\tilde{d}(\mathcal{S}) = \underline{\tilde{d}} \mid \mathbf{b}(\mathcal{S}) = \underline{\mathbf{b}}) \\ &\leq e^\varepsilon \mathbb{P}(\mathbf{b}(\mathcal{S}') = \underline{\mathbf{b}}) \\ &\quad e^\varepsilon \mathbb{P}(\tilde{d}(\mathcal{S}') = \underline{\tilde{d}} \mid \mathbf{b}(\mathcal{S}') = \underline{\mathbf{b}}) \\ &= e^{2\varepsilon} \mathbb{P}(\mathbf{b}(\mathcal{S}') = \underline{\mathbf{b}}, \tilde{d}(\mathcal{S}') = \underline{\tilde{d}}) \end{aligned}$$

so the definition of pan-privacy is satisfied. \square

4.2.2 Accuracy Analysis

As far as the accuracy of Algorithm 4.1 is concerned, we have two theorems. The first quantifies the bias and mean squared error of the estimator. Although Dwork et al. [23] demonstrate that their estimator is unbiased, they do not discuss its mean squared error.

Let \mathcal{S}_M be the subsequence (sub-stream) of the original stream \mathcal{S} that consists only of updates that refer to users in M . In particular, \mathcal{S}_M is constructed as:

for $i = 1$ to T
if $s_i \in M$: add s_i to \mathcal{S}_M

Lemma 4.1. *A special case of the law of total variance states that, if events A and B partition the whole outcome space of a random variable b , then:*

$$\begin{aligned} \text{var}(b) &= \text{var}(b|A) \mathbb{P}(A) + \text{var}(b|B) \mathbb{P}(B) \\ &\quad + \mathbb{E}[b|A]^2(1 - \mathbb{P}(A)) \mathbb{P}(A) + \mathbb{E}[b|B]^2(1 - \mathbb{P}(B)) \mathbb{P}(B) \\ &\quad - 2 \mathbb{E}[b|A] \mathbb{P}(A) \mathbb{E}[b|B] \mathbb{P}(B) \end{aligned}$$

Theorem 4.2. *For fixed sample M , Algorithm 4.1 provides an unbiased estimate \tilde{d} of the density of \mathcal{S}_M and has mean squared error:*

$$\mathbb{E}[(\tilde{d} - d(\mathcal{S}_M))^2] \leq \frac{2(2m+1)}{m^2 \varepsilon^2}$$

Proof. We begin with the *bias* computation. We examine the distribution of an arbitrary entry in \mathbf{b} :

$$b_i \sim \begin{cases} \text{Bernoulli}(\frac{1}{2}) & M[i] \notin \mathcal{S}_M \\ \text{Bernoulli}(\frac{1}{2} + \frac{\varepsilon}{4}) & M[i] \in \mathcal{S}_M \end{cases}$$

Note that the distribution of b_i does not depend on the number of appearances of $M[i]$; it only depends on whether it appeared or not.

Now, let $\hat{d} = \frac{1}{m} \sum_{i=1}^m b_i$. Then,

$$\begin{aligned}
\mathbb{E}[\hat{d}] &= \frac{1}{m} \sum_{i=1}^m \mathbb{E}[b_i] \\
&= \frac{1}{m} \sum_{i=1}^m \mathbb{E}[b_i | M[i] \in \mathcal{S}_M] \mathbb{P}(M[i] \in \mathcal{S}_M) \\
&\quad + \mathbb{E}[b_i | M[i] \notin \mathcal{S}_M] \mathbb{P}(M[i] \notin \mathcal{S}_M) \\
&= \frac{1}{m} \sum_{i=1}^m \left(\frac{1}{2} + \frac{\varepsilon}{4} \right) d(\mathcal{S}_M) + \frac{1}{2} (1 - d(\mathcal{S}_M)) \\
&= \frac{1}{2} + \frac{\varepsilon}{4} d(\mathcal{S}_M)
\end{aligned}$$

where we interpret the probability that a user is present in the sub-stream as the density of sub-stream. This is true if all users are considered equally likely to appear.

The final estimate (output) \tilde{d} is then computed as $\tilde{d} = \frac{4}{\varepsilon} (\hat{d} - \frac{1}{2}) + \text{Laplace}(0, \frac{1}{\varepsilon m})$, which gives that

$$\mathbb{E}[\tilde{d}] = \frac{4}{\varepsilon} (\mathbb{E}[\hat{d}] - \frac{1}{2}) + \mathbb{E}[\text{Laplace}(0, \frac{1}{\varepsilon m})] = d(\mathcal{S}_M)$$

so it is indeed an unbiased estimate.

We proceed with the *mean squared error*. Given that \tilde{d} is an unbiased estimate of $d(\mathcal{S}_M)$, its mean squared error coincides with its variance.

We make use of Lemma 4.1 and of the fact that, if $b \sim \text{Bernoulli}(p)$, then $\text{var}(b) = p(1-p)$. In our setting, we have m Bernoulli random variables b_i ($i = 1, \dots, m$), and for each user i , A is the event that $M[i] \in \mathcal{S}_M$ and B is the event that $M[i] \notin \mathcal{S}_M$. Therefore, $b_i | A \sim \text{Bernoulli}(\frac{1}{2} + \frac{\varepsilon}{4})$ and $b_i | B \sim \text{Bernoulli}(\frac{1}{2})$. By a straightforward computation:

$$\begin{aligned}
\text{var}(b_i) &= \frac{1}{4} - \frac{\varepsilon^2 d(\mathcal{S}_M)^2}{16} \\
\Rightarrow \text{var}(\hat{d}) &= \frac{1}{m^2} \sum_{i=1}^m \text{var}(b_i) = \frac{1}{4m} - \frac{\varepsilon^2 d(\mathcal{S}_M)^2}{16m} \\
\Rightarrow \text{var}(\tilde{d}) &= \left(\frac{4}{\varepsilon}\right)^2 \text{var}(\hat{d}) + \text{var}(\text{Laplace}(0, \frac{1}{\varepsilon m})) \\
&= \frac{4}{m \varepsilon^2} - \frac{d(\mathcal{S}_M)^2}{m} + \frac{2}{m^2 \varepsilon^2} \\
&\leq \frac{2(2m+1)}{m^2 \varepsilon^2}
\end{aligned}$$

which completes the proof. To derive the last inequality we used the fact that $0 \leq d(\mathcal{S}_M) \leq 1$. \square

The next theorem validates that the estimator provides the desired (α, β) -approximation of the actual stream density. In contrast to Dwork et al. [23], we parameterize the proof, so we are then able to numerically compute the tightest version of the bound we derive.

Before presenting the theorem, we prove the following useful lemma.

Lemma 4.2. *For any random variables X and Y , and for some $\alpha > 0$ and $\delta \in (0, 1)$:*

$$\mathbb{P}(|X + Y| > \alpha) \leq \mathbb{P}(|X| > \alpha\delta) + \mathbb{P}(|Y| > \alpha(1 - \delta))$$

Proof. We have that:

$$\begin{aligned}
\{(X, Y) : |X + Y| > \alpha\} &= \{(X, Y) : X + Y > \alpha \text{ or } X + Y < -\alpha\} \\
&= \{(X, Y) : X + Y > \alpha\} \\
&\quad \cup \{(X, Y) : X + Y < -\alpha\} \\
&\subseteq \{(X, Y) : X > \alpha\delta \text{ or } Y > \alpha(1 - \delta)\} \\
&\quad \cup \{(X, Y) : X < -\alpha\delta \text{ or } Y < -\alpha(1 - \delta)\} \\
&= \{(X, Y) : X > \alpha\delta \text{ or } Y > \alpha(1 - \delta) \\
&\quad \text{or } X < -\alpha\delta \text{ or } Y < -\alpha(1 - \delta)\} \\
&= \{(X, Y) : |X| > \alpha\delta \text{ or } |Y| > \alpha(1 - \delta)\}
\end{aligned}$$

so it follows that:

$$\begin{aligned}
\mathbb{P}(|X + Y| > \alpha) &\leq \mathbb{P}(|X| > \alpha\delta \text{ or } |Y| > \alpha(1 - \delta)) \\
&\leq \mathbb{P}(|X| > \alpha\delta) + \mathbb{P}(|Y| > \alpha(1 - \delta))
\end{aligned}$$

where the last inequality follows by the union bound. \square

Theorem 4.3. *If the sample maintained by Algorithm 4.1 consists of $m = \mathcal{O}(\frac{1}{\varepsilon^2 \alpha^2} \log \frac{1}{\beta})$ users from \mathcal{U} , then, for fixed input \mathcal{S} :*

$$\mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha) \leq \beta$$

where the probability space is over the random choices of the algorithm.

Proof. Let $\hat{d} = \frac{1}{m} \sum_{i=1}^m b_i$. We apply Lemma 4.2 twice, so for some $\alpha > 0$ and $\delta_1, \delta_2 \in (0, 1)$:

$$\begin{aligned}
\mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha) &= \mathbb{P}(|\tilde{d} - d(\mathcal{S}_M) + d(\mathcal{S}_M) - d(\mathcal{S})| \geq \alpha) \\
&\leq \mathbb{P}(|\tilde{d} - d(\mathcal{S}_M)| \geq \delta_1 \alpha) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1) \alpha) \\
&= \mathbb{P}(|\frac{4}{\varepsilon} \hat{d} - \frac{2}{\varepsilon} + \text{Laplace}(0, \frac{1}{\varepsilon m}) - d(\mathcal{S}_M)| \geq \delta_1 \alpha) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1) \alpha) \\
&= \mathbb{P}(|\hat{d} - \frac{1}{2} + \frac{\varepsilon}{4} \text{Laplace}(0, \frac{1}{\varepsilon m}) - \frac{\varepsilon}{4} d(\mathcal{S}_M)| \\
&\quad \geq \frac{\varepsilon}{4} \delta_1 \alpha) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1) \alpha) \\
&= \mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}] + \frac{\varepsilon}{4} \text{Laplace}(0, \frac{1}{\varepsilon m})| \geq \frac{\varepsilon}{4} \delta_1 \alpha) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1) \alpha) \\
&\leq \mathbb{P}(|\frac{\varepsilon}{4} \text{Laplace}(0, \frac{1}{\varepsilon m})| \geq \frac{\varepsilon}{4} \delta_1 \delta_2 \alpha) \\
&\quad + \mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}]| \geq \frac{\varepsilon}{4} \delta_1 (1 - \delta_2) \alpha) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1) \alpha)
\end{aligned}$$

We examine the quantities involved in the final expression. For fixed input, $d(\mathcal{S})$ is deterministic, as we do not take into account the randomness of the input. On the contrary, $d(\mathcal{S}_M)$ is random (due to sampling), \hat{d} is random (as a sum of Bernoulli

random variables), $\mathbb{E}[\hat{d}]$ is random (as a function of $d(\mathcal{S}_M)$), and \tilde{d} is random (as the sum of \hat{d} and a Laplace random variable). We therefore have 3 sources of error to control, which correspond to the following probabilities:

- $p_1 = \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1)\alpha)$
- $p_2 = \mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}]| \geq \frac{\varepsilon}{4}\delta_1(1 - \delta_2)\alpha)$
- $p_3 = \mathbb{P}(|\frac{\varepsilon}{4}\text{Laplace}(0, \frac{1}{\varepsilon m})| \geq \frac{\varepsilon}{4}\delta_1\delta_2\alpha)$

We want $p_1 + p_2 + p_3 \leq \beta$. We bound each error probability separately, so for some $\delta_3, \delta_4 \in (0, 1)$, such that $\delta_3 + \delta_4 < 1$, we want:

$$p_1 < \delta_3\beta, \quad p_2 < \delta_4\beta, \quad p_3 < (1 - \delta_3 - \delta_4)\beta$$

• First, we bound p_1 .

We define, $\forall i \in \{1, \dots, m\}$,

$$X_i = \begin{cases} 0 & M[i] \notin \mathcal{S}_M \\ 1 & M[i] \in \mathcal{S}_M \end{cases}$$

and since \mathcal{S}_M was sampled uniformly at random from \mathcal{S} , $X_i \sim \text{Bernoulli}(d(\mathcal{S}))$. It is easy to see that, since $d(\mathcal{S}_M) = \frac{1}{m} \sum_{i=1}^m X_i$, $\mathbb{E}[d(\mathcal{S}_M)] = d(\mathcal{S})$.

Since $d(\mathcal{S}_M)$ is a weighted sum of i.i.d. Bernoulli random variables, we take an additive, two-sided Chernoff bound:

$$p_1 = \mathbb{P}(|\frac{1}{m} \sum_{i=1}^m X_i - \frac{1}{m} \sum_{i=1}^m \mathbb{E}[X_i]| \geq (1 - \delta_1)\alpha) \leq 2e^{-2m\alpha^2(1-\delta_1)^2}$$

and, therefore:

$$p_1 \leq \delta_3\beta \Leftrightarrow m \geq \underbrace{\frac{1}{2\alpha^2(1-\delta_1)^2} \ln\left(\frac{2}{\beta\delta_3}\right)}_{m_1} = \mathcal{O}\left(\frac{1}{\alpha^2} \ln\left(\frac{1}{\beta}\right)\right)$$

• Next, we bound p_2 .

As we already mentioned, both \hat{d} and $\mathbb{E}[\hat{d}] = \frac{1}{2} + \frac{\varepsilon}{4}d(\mathcal{S}_M)$ are random. $d(\mathcal{S}_M)$ takes values in the set $\mathcal{D} = \{0, \frac{1}{m}, \dots, \frac{m-1}{m}, 1\}$, so using the law of total probability:

$$p_2 = \sum_{\underline{d} \in \mathcal{D}} \mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}]| \geq \frac{\varepsilon}{4}\delta_1(1 - \delta_2)\alpha \mid d(\mathcal{S}_M) = \underline{d}) \mathbb{P}(d(\mathcal{S}_M) = \underline{d})$$

We make the following two critical remarks:

- For fixed $d(\mathcal{S}_M) = \underline{d}$, $\mathbb{E}[\hat{d}]$ is no longer random.
- Let d_i , $i \in \{1, \dots, m\}$ denote the distribution from which each b_i is drawn from, so that $d_i = 0 \Rightarrow b_i \sim \text{Bernoulli}(\frac{1}{2})$ and $d_i = 1 \Rightarrow b_i \sim \text{Bernoulli}(\frac{1}{2} + \frac{\varepsilon}{4})$. Once we fix $d(\mathcal{S}_M) = \underline{d}$, the d_i 's are not independent, as $\sum_{i=1}^m d_i = m\underline{d}$.

However, the b_i 's are independent, as we impose no constraint on them and each is drawn independently from a fixed distribution. Therefore, \hat{d} is a sum of independent Poisson trials.

We again make use of an additive, two-sided Chernoff bound:

$$\mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}]| \geq \frac{\varepsilon}{4}\delta_1(1 - \delta_2)\alpha \mid d(\mathcal{S}_M) = \underline{d}) \leq 2e^{-2m\varepsilon^2\alpha^2\delta_1^2(1-\delta_2)^2\frac{1}{16}}$$

We observe that the bound is independent of \underline{d} , so:

$$p_2 \leq \sum_{\underline{d} \in \mathcal{D}} 2e^{-2m\varepsilon^2\alpha^2\delta_1^2(1-\delta_2)^2\frac{1}{16}} \mathbb{P}(d(\mathcal{S}_M) = \underline{d}) = 2e^{-\frac{1}{8}m\varepsilon^2\alpha^2\delta_1^2(1-\delta_2)^2}$$

and, therefore:

$$p_2 \leq \delta_4\beta \Leftrightarrow m \geq \underbrace{\frac{8}{\varepsilon^2\alpha^2\delta_1^2(1-\delta_2)^2} \ln\left(\frac{2}{\beta\delta_4}\right)}_{m_2} = \mathcal{O}\left(\frac{1}{\varepsilon^2\alpha^2} \ln\left(\frac{1}{\beta}\right)\right)$$

• Finally, we compute p_3 .

$$\begin{aligned} p_3 &= \mathbb{P}\left(\left|\frac{\varepsilon}{4}\text{Laplace}\left(0, \frac{1}{\varepsilon m}\right)\right| \geq \frac{\varepsilon}{4}\delta_1\delta_2\alpha\right) \\ &= \mathbb{P}\left(\left|\text{Laplace}\left(0, \frac{1}{\varepsilon m}\right)\right| \geq \delta_1\delta_2\alpha\right) \\ &= \mathbb{P}\left(\text{Exponential}(\varepsilon m) \geq \delta_1\delta_2\alpha\right) \\ &= e^{-\varepsilon\alpha m\delta_1\delta_2} \end{aligned}$$

and, therefore:

$$p_3 \leq (1 - \delta_3 - \delta_4)\beta \Leftrightarrow m \geq \underbrace{\frac{1}{\varepsilon\alpha\delta_1\delta_2} \ln\left(\frac{1}{\beta(1 - \delta_3 - \delta_4)}\right)}_{m_3} = \mathcal{O}\left(\frac{1}{\varepsilon\alpha} \ln\left(\frac{1}{\beta}\right)\right)$$

By picking $m \geq \max\{m_1, m_2, m_3\} = \mathcal{O}\left(\frac{1}{\varepsilon^2\alpha^2} \ln\left(\frac{1}{\beta}\right)\right)$, we ensure that:

$$\mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha) \leq p_1 + p_2 + p_3 \leq \delta_3\beta + \delta_4\beta + (1 - \delta_3 - \delta_4)\beta = \beta$$

which completes the proof. \square

Theorem 4.3 offers an additive error guarantee, which may not be so useful if the density of the input stream is small. Dwork et al. [23] show how to modify their algorithm to obtain a multiplicative error guarantee, which is more meaningful in such cases. We do not examine this point in our work.

As we already stated, our parameterized proof allows us to optimally tune the parameters $\delta_1, \delta_2, \delta_3, \delta_4$ and, as a result, to compute the tightest version of the bound we derive. In particular, for fixed ε, α, m , the tightest bound β on $\mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha)$ is computed by numerically solving the following optimization problem:

$$\begin{aligned} \underset{\delta_1, \delta_2}{\text{minimize}} \quad & \beta(\delta_1, \delta_2) = 2e^{-2m\alpha^2(1-\delta_1)^2} \\ & + 2e^{-\frac{1}{8}m\varepsilon^2\alpha^2\delta_1^2(1-\delta_2)^2} + e^{-\varepsilon\alpha m\delta_1\delta_2} \\ \text{subject to} \quad & 0 \leq \delta_1 \leq 1 \\ & 0 \leq \delta_2 \leq 1 \end{aligned}$$

4.2.3 Picking the Optimal Sample Size

Theorem 4.3 provides an asymptotic expression for the sample size m to achieve the desired approximation accuracy. A question that arises is how we should pick m in practice. This is again achieved by taking advantage of our parameterized proof and numerically solving a similar optimization problem, which allows us to compute the optimal (minimum) sample size m^* that achieves the desired approximation accuracy, according to the bounds we derived. Specifically, for fixed $\varepsilon, \alpha, \beta$, we define:

$$\begin{aligned} m(\delta_1, \delta_2, \delta_3, \delta_4) &= \max\{ m_1, m_2, m_3 \} \\ &= \max\left\{ \frac{1}{2\alpha^2(1-\delta_1)^2} \ln\left(\frac{2}{\beta\delta_3}\right), \right. \\ &\quad \frac{8}{\varepsilon^2 \alpha^2 \delta_1^2 (1-\delta_2)^2} \ln\left(\frac{2}{\beta\delta_4}\right), \\ &\quad \left. \frac{1}{\varepsilon \alpha \delta_1 \delta_2} \ln\left(\frac{1}{\beta(1-\delta_3-\delta_4)}\right) \right\} \end{aligned}$$

Then, we pick m^* as the solution to the following optimization problem:

$$\begin{aligned} &\underset{\delta_1, \delta_2, \delta_3, \delta_4}{\text{minimize}} && m(\delta_1, \delta_2, \delta_3, \delta_4) \\ &\text{subject to} && 0 \leq \delta_i \leq 1, \quad i \in \{1, 2, 3, 4\} \\ &&& \delta_3 + \delta_4 \leq 1 \end{aligned}$$

Figure 4.1 illustrates the proposed sample sizes for varying privacy budget. We repeat that our approach allows us to compute the tightest version of the specific bound that we derive on the probability of error and hence on the sample size. Specifically, the sample sizes we propose are smaller by more than an order of magnitude, compared to the sample sizes proposed by Roth in his lectures, who arbitrarily proposes to pick $m = \frac{128}{\varepsilon^2 \alpha^2} \ln\left(\frac{1}{\beta}\right)$. Nevertheless, we remark that the bound itself is not tight; this is an experimental observation and indicates that the same accuracy can be achieved with even smaller samples.

4.3 Optimal Bernoulli Density Estimator

In this section, we modify Algorithm 4.1 and derive a novel algorithm that significantly outperforms the original algorithm (both theoretically and experimentally). The key reason behind our algorithm's superiority is that -in contrast to Algorithm 4.1- it manages to use all the allocated privacy budget.

4.3.1 On the Use of the Allocated Privacy Budget

Recall that, in order to ensure that its state satisfies differential privacy, Algorithm 4.1 utilizes two different distributions, one for users that do not appear in the stream, and one for users that do appear. We now introduce a little extra notation; the bit that corresponds to a user from the former category is drawn from the distribution with pmf $f_{init} = \text{Bernoulli}\left(\frac{1}{2}\right)$, while the bit that corresponds to a user from the latter category is drawn from $f_{upd} = \text{Bernoulli}\left(\frac{1}{2} + \frac{\varepsilon}{4}\right)$. Although by f_{init} and f_{upd} we formally denote the probability mass functions of the two Bernoulli distributions, at some points we use the same notation to refer to the distributions themselves (abusing notation a little bit). To satisfy differential privacy, we have to ensure that, $\forall b \in \{0, 1\}$:

$$e^{-\varepsilon} \leq R(b) = \frac{f_{upd}(b)}{f_{init}(b)} \leq e^{\varepsilon}$$

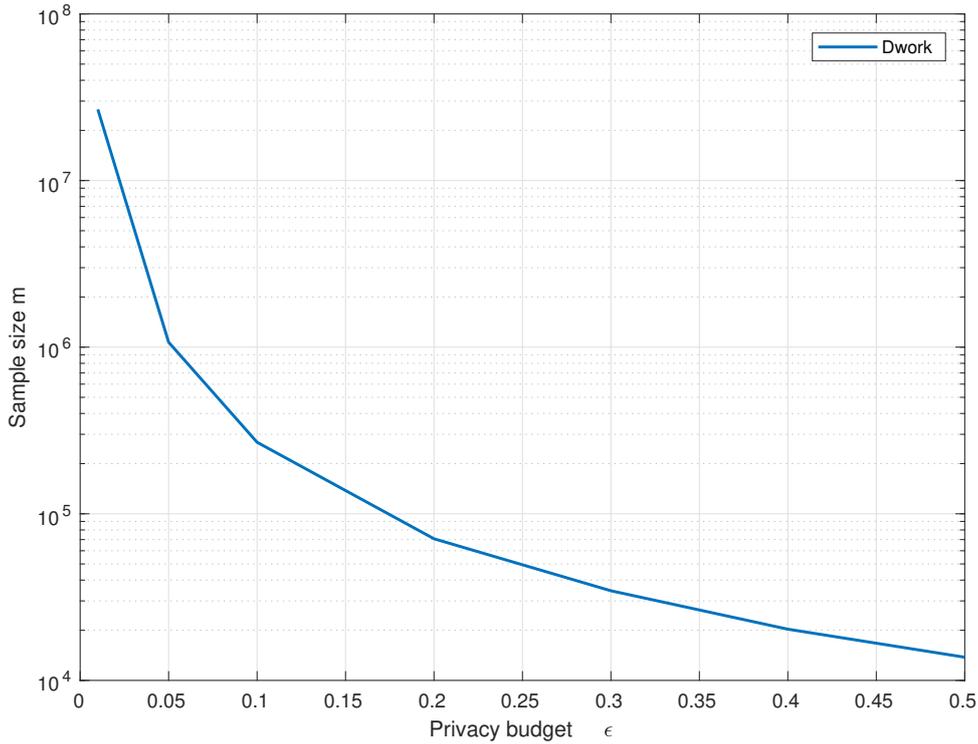


Figure 4.1: Proposed sample size

Figure 4.2 illustrates that Algorithm 4.1 does ensure that its state satisfies differential privacy (as we have already proved in Theorem 4.1).

At the same time however, it is not hard to observe that Algorithm 4.1 fails to use all the allocated privacy budget. We have already shown (in the proof of Theorem 4.1) that $R(0) = 1 - \frac{\epsilon}{2}$ and $R(1) = 1 + \frac{\epsilon}{2}$. Let $\epsilon' > 0$ be the actual privacy budget that Algorithm 4.1 consumes. Then:

$$\left. \begin{array}{l} e^{-\epsilon'} \leq 1 - \frac{\epsilon}{2} \leq e^{\epsilon'} \\ e^{-\epsilon'} \leq 1 + \frac{\epsilon}{2} \leq e^{\epsilon'} \end{array} \right\} \Rightarrow \epsilon' \geq \max\left\{ \ln\left(1 + \frac{\epsilon}{2}\right), -\ln\left(1 - \frac{\epsilon}{2}\right) \right\}$$

Figure 4.3 illustrates the actual privacy budget used by Algorithm 4.1 for each allocated privacy budget.

4.3.2 Optimally Tuning the Bernoulli Distributions

Based on the aforementioned observation, we optimally tune the Bernoulli distributions used (with pmf's f_{init} and f_{upd}), by picking a pair of parameters that tightly satisfies ϵ -differential privacy, for any given ϵ . In particular, we maximize the distributions' distance (or, more precisely, the difference of their parameters), which allows us to more accurately distinguish between users that did not appear and users that appeared in the stream, during the density estimation computation (without sacrificing the users' privacy!).

We propose that the parameters of f_{init} and f_{upd} are picked symmetric around some value c ; although the choice of c (among the permitted values) does not seem to affect the algorithm, we explicitly show later what values c is permitted to take. For example, we could set $c = \frac{1}{2}$, or $c = \frac{1}{2} + \frac{\epsilon}{8}$, as in Algorithm 4.1. In doing so, we end up with the following distributions:

$$f_{init} = \text{Bernoulli}(c - x)$$

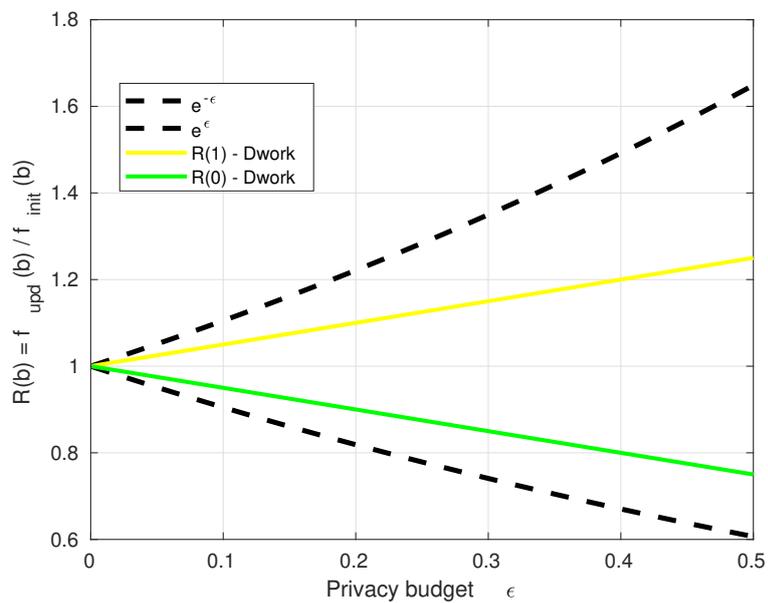


Figure 4.2: Dwork - State differential privacy

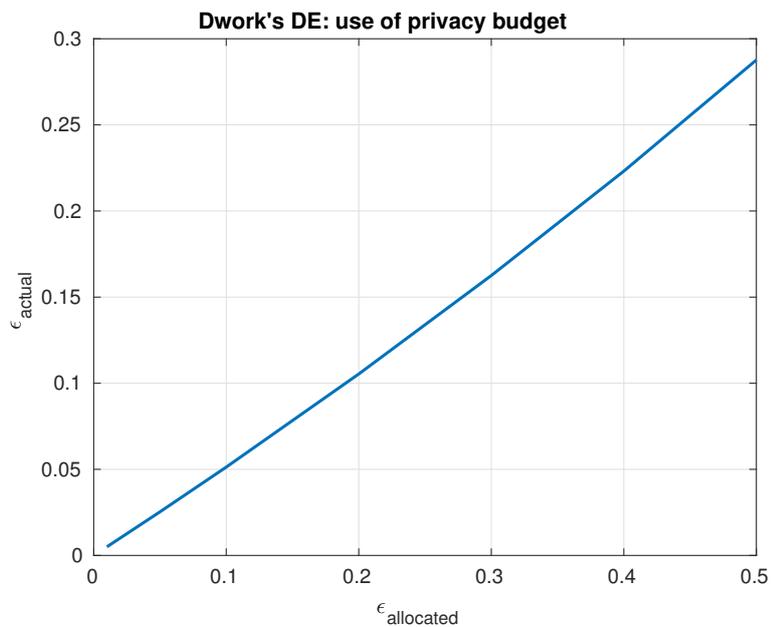


Figure 4.3: Dwork - Allocated vs actual privacy budget

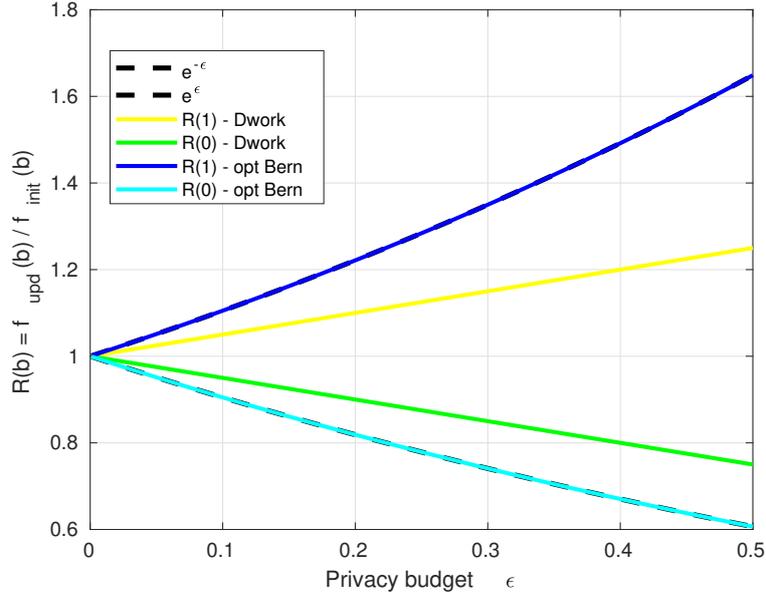


Figure 4.4: Optimal Bernoulli - State differential privacy

$$f_{upd} = \text{Bernoulli}(c + x)$$

for some $x > 0$ that corresponds to half the difference of the distributions' parameters. Clearly, we want $0 < c - x < c + x < 1$. For example, in Algorithm 4.1, $x = \frac{\epsilon}{8}$, and $0 < c - x = \frac{1}{2} < c + x = \frac{1}{2} + \frac{\epsilon}{4} < 1$, $\forall \epsilon \in (0, \frac{1}{2}]$.

The modification we propose is in the selection of x . Again, to satisfy differential privacy, we have to ensure that, $\forall b \in \{0, 1\}$:

$$\begin{aligned} e^{-\epsilon} \leq \frac{f_{upd}(b)}{f_{init}(b)} \leq e^{\epsilon} &\Leftrightarrow \begin{cases} e^{-\epsilon} \leq \frac{c+x}{c-x} \leq e^{\epsilon} \\ e^{-\epsilon} \leq \frac{c-x}{c+x} \leq e^{\epsilon} \end{cases} \\ &\Rightarrow x \leq \frac{e^{\epsilon} - 1}{e^{\epsilon} + 1} c = \tanh\left(\frac{\epsilon}{2}\right)c \end{aligned}$$

Since our goal was to maximize the difference of the distributions' parameters, we pick $x = \tanh\left(\frac{\epsilon}{2}\right)c$. The proposed distributions are:

$$\begin{aligned} f_{init} &= \text{Bernoulli}\left(c \left(1 - \tanh\left(\frac{\epsilon}{2}\right)\right)\right) \\ f_{upd} &= \text{Bernoulli}\left(c \left(1 + \tanh\left(\frac{\epsilon}{2}\right)\right)\right) \end{aligned}$$

and, since both parameters have to be greater than zero and less than one, we can determine the values that c can take for each fixed ϵ . For $0 < \epsilon \leq \frac{1}{2}$, it is easy to see that $0 < \tanh\left(\frac{\epsilon}{2}\right) < \frac{1}{4}$ (since $\tanh\left(\frac{\epsilon}{2}\right)$ is a monotonically increasing function of ϵ), so we can pick any $c \in (0, \frac{4}{5}]$, regardless of the specific value of $\epsilon \in (0, \frac{1}{2}]$.

Figure 4.4 illustrates that the proposed Bernoulli distributions indeed use all the allocated privacy budget.

4.3.3 Estimator & Analysis

We now present the modified algorithm, which we call the Optimal Bernoulli Density Estimator. For simplicity, we set $c = \frac{1}{2}$.

We proceed with the privacy and accuracy analysis of Algorithm 4.2, following the lines of our analysis of Algorithm 4.1.

Algorithm 4.2: Optimal Bernoulli Density Estimator**Input:** Data stream \mathcal{S} , Privacy budget ε , Accuracy parameters (α, β) **Output:** Density $\tilde{d}(\mathcal{S})$

- 1 Compute m^* and set $m = m^*$
- 2 Sample a random subset $M \subseteq \mathcal{U}$ of m users (without replacement) and define an arbitrary ordering over M
- 3 Create a bitarray $\mathbf{b} = [b_1 \dots b_m]$ and map $M[i] \rightarrow b_i, \forall i \in \{1, \dots, m\}$
- 4 Initialize \mathbf{b} randomly: $b_i \sim \text{Bernoulli}(\frac{1}{2}(1 - \tanh(\frac{\varepsilon}{2})))$, $\forall i \in \{1, \dots, m\}$
- 5 **for** $t = 1$ **to** T **do**
- 6 **if** $s_t \in M$ **then**
- 7 Find $i : M[i] = s_t$
- 8 Re-sample: $b_i \sim \text{Bernoulli}(\frac{1}{2}(1 + \tanh(\frac{\varepsilon}{2})))$
- 9 Return $\tilde{d}(\mathcal{S}) = \frac{1}{\tanh(\frac{\varepsilon}{2})} (\frac{1}{m} \sum_{i=1}^m b_i - \frac{1}{2} + \frac{1}{2} \tanh(\frac{\varepsilon}{2})) + \text{Laplace}(0, \frac{1}{\varepsilon m})$

Theorem 4.4. Assume $\varepsilon \leq \frac{1}{2}$. Then Algorithm 4.2 satisfies 2ε -pan-privacy and utilizes all the allocated privacy budget.

Proof. The proof is identical with that of Theorem 4.1. The only modification is on proving that the state (bitarray) satisfies ε -differential privacy, and in particular, for a user $u \in M$ that appears in stream \mathcal{S} and does not appear in stream \mathcal{S}' (again, let b_1 be the entry that corresponds to u in the bitarray), we have:

$$\begin{aligned} \frac{\mathbb{P}(\mathbf{b}(\mathcal{S}) = [1 \ \underline{b_2} \ \dots \ \underline{b_m}])}{\mathbb{P}(\mathbf{b}(\mathcal{S}') = [1 \ \underline{b_2} \ \dots \ \underline{b_m}])} &= \frac{\mathbb{P}(b_1(\mathcal{S}) = 1)}{\mathbb{P}(b_1(\mathcal{S}') = 1)} \\ &= \frac{\frac{1}{2}(1 + \frac{e^\varepsilon - 1}{e^\varepsilon + 1})}{\frac{1}{2}(1 - \frac{e^\varepsilon - 1}{e^\varepsilon + 1})} = e^\varepsilon \\ \frac{\mathbb{P}(\mathbf{b}(\mathcal{S}) = [0 \ \underline{b_2} \ \dots \ \underline{b_m}])}{\mathbb{P}(\mathbf{b}(\mathcal{S}') = [0 \ \underline{b_2} \ \dots \ \underline{b_m}])} &= \frac{\mathbb{P}(b_1(\mathcal{S}) = 0)}{\mathbb{P}(b_1(\mathcal{S}') = 0)} \\ &= \frac{\frac{1}{2}(1 - \frac{e^\varepsilon - 1}{e^\varepsilon + 1})}{\frac{1}{2}(1 + \frac{e^\varepsilon - 1}{e^\varepsilon + 1})} = e^{-\varepsilon} \end{aligned}$$

so user u is guaranteed ε -differential privacy against an adversary that observes u 's entry in \mathbf{b} . \square

Theorem 4.5. For fixed sample M , Algorithm 4.2 provides an unbiased estimate \tilde{d} of the density of \mathcal{S}_M and has mean squared error:

$$\mathbb{E}[(\tilde{d} - d(\mathcal{S}_M))^2] \leq \frac{1}{4m \tanh^2(\frac{\varepsilon}{2})} + \frac{2}{m^2 \varepsilon^2}$$

Proof. The proof is similar with that of Theorem 4.2.

We begin with the *bias* computation. In order to examine the distribution of an arbitrary entry in \mathbf{b} , we introduce the notation $p_{init} = \frac{1}{2}(1 - \tanh(\frac{\varepsilon}{2}))$ and $p_{upd} = \frac{1}{2}(1 + \tanh(\frac{\varepsilon}{2}))$. Then,

$$b_i \sim \begin{cases} \text{Bernoulli}(p_{init}) & M[i] \notin \mathcal{S}_M \\ \text{Bernoulli}(p_{upd}) & M[i] \in \mathcal{S}_M \end{cases}$$

Denoting $\hat{d} = \frac{1}{m} \sum_{i=1}^m b_i$, and applying a similar computation as in the proof of Theorem 4.2, gives:

$$\begin{aligned} \mathbb{E}[\hat{d}] &= p_{init} + (p_{upd} - p_{init})d(\mathcal{S}_M) \\ &= \frac{1}{2}(1 - \tanh(\frac{\varepsilon}{2})) + \tanh(\frac{\varepsilon}{2})d(\mathcal{S}_M) \end{aligned}$$

The final estimate (output) \tilde{d} is:

$$\begin{aligned} \tilde{d} &= \frac{\hat{d} - p_{init}}{p_{upd} - p_{init}} + \text{Laplace}(0, \frac{1}{\varepsilon m}) \\ &= \frac{1}{\tanh(\frac{\varepsilon}{2})}(\hat{d} - \frac{1}{2} + \frac{1}{2} \tanh(\frac{\varepsilon}{2})) + \text{Laplace}(0, \frac{1}{\varepsilon m}) \end{aligned}$$

so it is an unbiased estimate.

We proceed with the *mean squared error*. Again, since \tilde{d} is an unbiased estimate of $d(\mathcal{S}_M)$, its mean squared error coincides with its variance. Also, recall that we have m Bernoulli random variables b_i ($i = 1, \dots, m$), so applying the law of total variance gives us:

$$\begin{aligned} \text{var}(b_i) &= p_{init}(1 - p_{init}) + d(\mathcal{S}_M)[-p_{init}(1 - p_{init}) + p_{upd}(1 - p_{upd})] \\ &\quad + d(\mathcal{S}_M)(1 - d(\mathcal{S}_M))(p_{upd} - p_{init})^2 \\ &= \frac{1}{4}[1 - \tanh^2(\frac{\varepsilon}{2})] + d(\mathcal{S}_M)(1 - d(\mathcal{S}_M))\tanh^2(\frac{\varepsilon}{2}) \\ \Rightarrow \text{var}(\hat{d}) &= \frac{1}{m^2} \sum_{i=1}^m \text{var}(b_i) \\ &= \frac{1}{4m}[1 - \tanh^2(\frac{\varepsilon}{2})] + \frac{d(\mathcal{S}_M)(1 - d(\mathcal{S}_M))}{m} \tanh^2(\frac{\varepsilon}{2}) \\ \Rightarrow \text{var}(\tilde{d}) &= (\frac{1}{p_{upd} - p_{init}})^2 \text{var}(\hat{d}) + \text{var}(\text{Laplace}(0, \frac{1}{\varepsilon m})) \\ &= \frac{1}{4m}[\frac{1}{\tanh^2(\frac{\varepsilon}{2})} - 1] + \frac{d(\mathcal{S}_M)(1 - d(\mathcal{S}_M))}{m} + \frac{2}{m^2 \varepsilon^2} \\ &\leq \frac{1}{4m \tanh^2(\frac{\varepsilon}{2})} + \frac{2}{m^2 \varepsilon^2} \end{aligned}$$

which completes the proof. To derive the last inequality, we used the fact that, since $0 \leq d(\mathcal{S}_M) \leq 1$, it follows that $d(\mathcal{S}_M)(1 - d(\mathcal{S}_M)) \leq \frac{1}{4}$. \square

Theorem 4.6. *If the sample maintained by Algorithm 4.2 consists of $m = \mathcal{O}(\frac{1}{\varepsilon^2 \alpha^2} \log \frac{1}{\beta})$ users from \mathcal{U} , then, for fixed input \mathcal{S} :*

$$\mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha) \leq \beta$$

where the probability space is over the random choices of the algorithm.

Proof. The proof is similar with that of Theorem 4.3, so we only focus on the differences.

Let $\hat{d} = \frac{1}{m} \sum_{i=1}^m b_i$. We again apply Lemma 4.2 twice, so for some $\alpha > 0$ and $\delta_1, \delta_2 \in$

$(0, 1)$:

$$\begin{aligned}
\mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha) &= \mathbb{P}(|\tilde{d} - d(\mathcal{S}_M) + d(\mathcal{S}_M) - d(\mathcal{S})| \geq \alpha) \\
&\leq \mathbb{P}(|\tilde{d} - d(\mathcal{S}_M)| \geq \delta_1 \alpha) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1)\alpha) \\
&= \mathbb{P}\left(\left|\frac{1}{\tanh(\frac{\varepsilon}{2})}\hat{d} - \frac{1}{2}\frac{1}{\tanh(\frac{\varepsilon}{2})} + \frac{1}{2} + \text{Laplace}\left(0, \frac{1}{\varepsilon m}\right) - d(\mathcal{S}_M)\right| \geq \delta_1 \alpha\right) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1)\alpha) \\
&= \mathbb{P}\left(|\hat{d} - \mathbb{E}[\hat{d}] + \tanh\left(\frac{\varepsilon}{2}\right)\text{Laplace}\left(0, \frac{1}{\varepsilon m}\right)| \geq \tanh\left(\frac{\varepsilon}{2}\right)\delta_1 \alpha\right) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1)\alpha) \\
&\leq \mathbb{P}\left(|\tanh\left(\frac{\varepsilon}{2}\right)\text{Laplace}\left(0, \frac{1}{\varepsilon m}\right)| \geq \tanh\left(\frac{\varepsilon}{2}\right)\delta_1 \delta_2 \alpha\right) \\
&\quad + \mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}]| \geq \tanh\left(\frac{\varepsilon}{2}\right)\delta_1(1 - \delta_2)\alpha) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1)\alpha) \\
&= p_3 + p_2 + p_1 \quad (\text{respectively})
\end{aligned}$$

As in the proof of Theorem 4.3, we have 3 sources of error to control, and we want $p_1 + p_2 + p_3 \leq \beta$. We bound each error probability separately, so for some $\delta_3, \delta_4 \in (0, 1)$, such that $\delta_3 + \delta_4 < 1$, we want:

$$p_1 < \delta_3 \beta, \quad p_2 < \delta_4 \beta, \quad p_3 < (1 - \delta_3 - \delta_4)\beta$$

Bounding p_1 and p_3 is identical, as these error probabilities are unchanged. We copy the bounds we derived in Theorem 4.3:

$$\begin{aligned}
p_1 &\leq 2e^{-2m\alpha^2(1-\delta_1)^2} \leq \delta_3 \beta \\
&\Leftrightarrow m \geq \underbrace{\frac{1}{2\alpha^2(1-\delta_1)^2} \ln\left(\frac{2}{\beta\delta_3}\right)}_{m_1} = \mathcal{O}\left(\frac{1}{\alpha^2} \ln\left(\frac{1}{\beta}\right)\right) \\
p_3 &\leq e^{-\varepsilon\alpha m\delta_1\delta_2} \leq (1 - \delta_3 - \delta_4)\beta \\
&\Leftrightarrow m \geq \underbrace{\frac{1}{\varepsilon\alpha\delta_1\delta_2} \ln\left(\frac{1}{\beta(1-\delta_3-\delta_4)}\right)}_{m_3} = \mathcal{O}\left(\frac{1}{\varepsilon\alpha} \ln\left(\frac{1}{\beta}\right)\right)
\end{aligned}$$

We focus on p_2 . The difference is due to the fact that we now have $\mathbb{E}[\hat{d}] = \frac{1}{2}(1 - \tanh(\frac{\varepsilon}{2})) + \tanh(\frac{\varepsilon}{2})d(\mathcal{S}_M)$. Although $\mathbb{E}[\hat{d}]$ has changed, it still is a random mean (as $d(\mathcal{S}_M)$ is random), so we again apply the total probability theorem:

$$p_2 = \sum_{\underline{d} \in \mathcal{D}} \mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}]| \geq \tanh(\frac{\varepsilon}{2})\delta_1(1 - \delta_2)\alpha \mid d(\mathcal{S}_M) = \underline{d}) \mathbb{P}(d(\mathcal{S}_M) = \underline{d})$$

Based on observations identical to those we made when proving theorem 4.3, we again make use of an additive, two-sided Chernoff bound:

$$\mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}]| \geq \tanh(\frac{\varepsilon}{2})\delta_1(1 - \delta_2)\alpha \mid d(\mathcal{S}_M) = \underline{d}) \leq 2e^{-2m \tanh^2(\frac{\varepsilon}{2})\alpha^2\delta_1^2(1-\delta_2)^2}$$

and since the bound is independent of \underline{d} :

$$p_2 \leq \sum_{\underline{d} \in \mathcal{D}} 2e^{-2m \tanh^2(\frac{\varepsilon}{2}) \alpha^2 \delta_1^2 (1-\delta_2)^2} \mathbb{P}(d(\mathcal{S}_M) = \underline{d}) = 2e^{-2m \tanh^2(\frac{\varepsilon}{2}) \alpha^2 \delta_1^2 (1-\delta_2)^2}$$

and, therefore, noting that $\tanh(\frac{x}{2}) = \frac{x}{2} - \frac{x^3}{24} + \mathcal{O}(x^5) = \mathcal{O}(x)$:

$$p_2 \leq \delta_4 \beta \Leftrightarrow m \geq \underbrace{\frac{1}{2 \tanh^2(\frac{\varepsilon}{2}) \alpha^2 \delta_1^2 (1-\delta_2)^2} \ln\left(\frac{2}{\beta \delta_4}\right)}_{m_2} = \mathcal{O}\left(\frac{1}{\varepsilon^2 \alpha^2} \ln\left(\frac{1}{\beta}\right)\right)$$

By picking $m \geq \max\{m_1, m_2, m_3\} = \mathcal{O}\left(\frac{1}{\varepsilon^2 \alpha^2} \ln\left(\frac{1}{\beta}\right)\right)$, we ensure that:

$$\mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha) \leq p_1 + p_2 + p_3 \leq \delta_3 \beta + \delta_4 \beta + (1 - \delta_3 - \delta_4) \beta = \beta$$

which completes the proof. \square

Asymptotically, the lower bound on the sample size to achieve the desired approximation accuracy is not improved. We again compute the proposed sample size by numerically solving the optimization problem described in section 4.2.3, with the slightly modified cost function:

$$\begin{aligned} m(\delta_1, \delta_2, \delta_3, \delta_4) &= \max\{m_1, m_2, m_3\} \\ &= \max\left\{ \frac{1}{2\alpha^2(1-\delta_1)^2} \ln\left(\frac{2}{\beta\delta_3}\right), \right. \\ &\quad \left. \frac{1}{2 \tanh^2(\frac{\varepsilon}{2}) \alpha^2 \delta_1^2 (1-\delta_2)^2} \ln\left(\frac{2}{\beta\delta_4}\right), \right. \\ &\quad \left. \frac{1}{\varepsilon \alpha \delta_1 \delta_2} \ln\left(\frac{1}{\beta(1-\delta_3-\delta_4)}\right) \right\} \end{aligned}$$

The resulting proposed sample size is illustrated in Figure 4.5.

4.4 Laplace Density Estimator

In this section, we examine an alternative modification to Algorithm 4.1. Although the algorithm we derive fails to match the performance of Algorithm 4.2, it also manages to outperform the original algorithm and may provide useful theoretical insights.

4.4.1 Using Continuous Distributions

We again focus on modifying the state of Algorithm 4.1. The main idea is to replace the bitarray \mathbf{b} (which stores bits drawn from either of the two Bernoulli distributions f_{init} and f_{upd}), by an array of real numbers \mathbf{x} , drawn from two continuous distribution. Our motivation is that the algorithm's output is itself a real number, so by storing a flexible, real value per user (instead of a hard, binary value), we expect to boost accuracy. Although this approach increases the algorithm's space requirements ($\times 32$ or $\times 64$ bits per user if floating point arithmetic is used), we argue that, to achieve the same levels of accuracy, we will need to maintain a sample of less users, and we will therefore offer perfect privacy for more users.

We propose using the following two Laplace distributions:

$$f_{init} = \text{Laplace}(\mu_{init}, b)$$

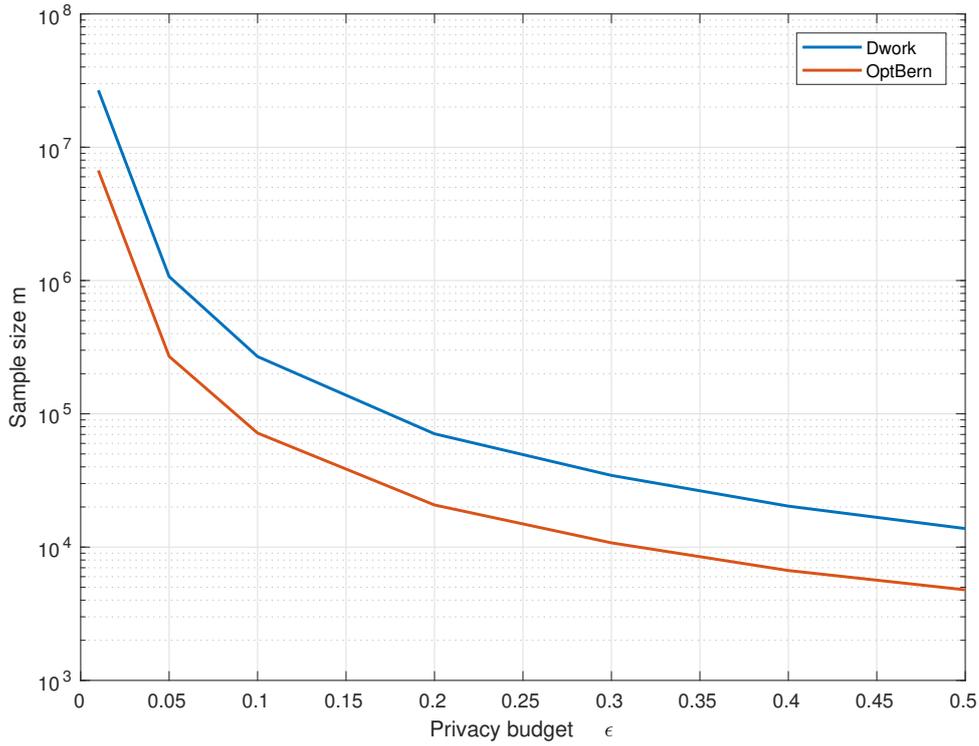


Figure 4.5: Proposed sample size

$$f_{upd} = \text{Laplace}(\mu_{upd}, b)$$

where $\mu_{init} < \mu_{upd}$ and the scale of the two distributions is the same. Therefore, if a (sampled) user does not appear in the stream, his entry in \mathbf{x} will be drawn from $\text{Laplace}(\mu_{init}, b)$, whereas if a user appears, his entry will be drawn from $\text{Laplace}(\mu_{upd}, b)$.

Before presenting the new algorithm (Algorithm 4.3), we examine the privacy guarantees of the modified state. As we show, the definition of differential privacy determines the values the parameters of the two distributions may take.

Theorem 4.7. *Assume $\epsilon \leq \frac{1}{2}$. If $\frac{\mu_{upd} - \mu_{init}}{b} = \epsilon$, then Algorithm 4.3 satisfies 2ϵ -pan-privacy.*

Proof. The proof is identical with that of Theorem 4.1. The only modification is on proving that the modified state \mathbf{x} (array of real values) satisfies ϵ -differential privacy, and in particular, for a user $u \in M$ that appears in stream \mathcal{S} and does not appear in stream \mathcal{S}' (again, let x_1 be the entry that corresponds to u in the array), we want, $\forall x \in (-\infty, \infty)$:

$$\begin{aligned} \frac{\mathbb{P}(\mathbf{x}(\mathcal{S}) = [x \ x_2 \ \dots \ x_m])}{\mathbb{P}(\mathbf{x}(\mathcal{S}') = [x \ \underline{x_2} \ \dots \ \underline{x_m}])} &= \frac{\mathbb{P}(x_1(\mathcal{S}) = x)}{\mathbb{P}(x_1(\mathcal{S}') = x)} \\ &= \frac{f_{upd}(x)}{f_{init}(x)} \\ &= \frac{\frac{1}{2b} e^{-\frac{|x - \mu_{upd}|}{b}}}{\frac{1}{2b} e^{-\frac{|x - \mu_{init}|}{b}}} \\ &= e^{\frac{|x - \mu_{init}| - |x - \mu_{upd}|}{b}} \end{aligned}$$

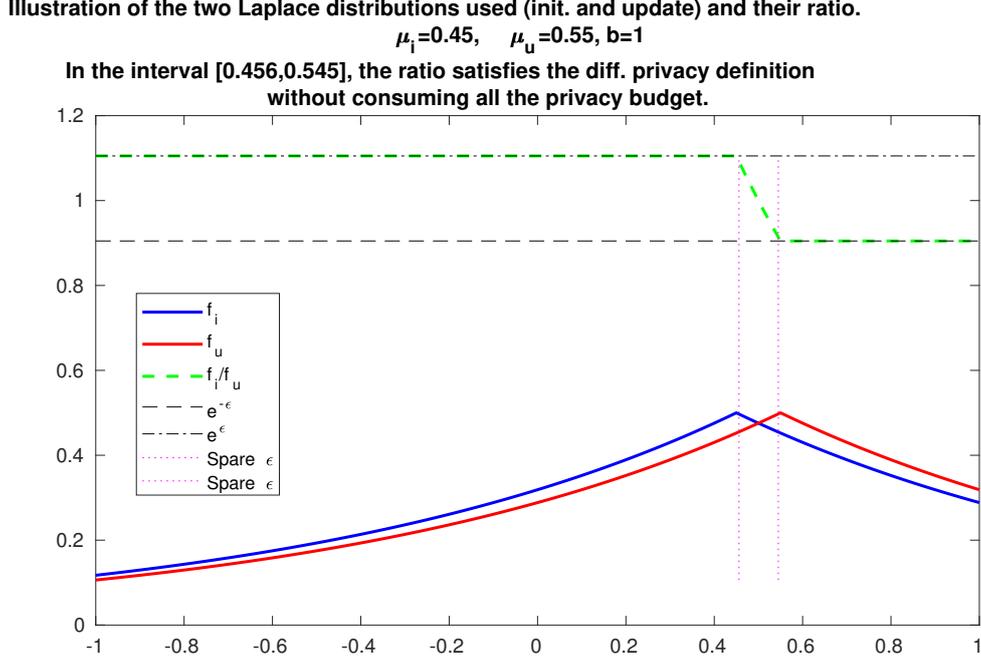


Figure 4.6: Illustration of the proposed Laplace distributions

$$\begin{aligned}
 &\leq e^{\frac{|x-\mu_{init}-x+\mu_{upd}|}{b}} \\
 &= e^{\frac{\mu_{upd}-\mu_{init}}{b}} \\
 &= e^\epsilon
 \end{aligned}$$

where the inequality holds due to the triangle inequality. Proving that $\frac{f_{upd}(x)}{f_{init}(x)} \geq e^{-\epsilon}$ is identical. Therefore, user u is guaranteed ϵ -differential privacy against an adversary that observes u 's entry in \mathbf{x} . \square

Figure 4.6 provides an example illustration of the two distributions, for $\epsilon = 0.1$. Observe that the ratio of the two distributions tightly satisfies the differential privacy definition, except for some $x \in (\mu_{init}, \mu_{upd})$. However, since $\exists x$ such that the ratio is almost equal to either e^ϵ or $e^{-\epsilon}$, we conclude that Algorithm 4.3 uses all the allocated privacy budget.

To achieve the best possible accuracy, we would want to maximize the difference $\mu_{upd} - \mu_{init}$, as well as to pick the minimum possible b . In doing so, it would have been easier to distinguish between users that did not appear and users that appeared in the stream, during the density estimation computation. However, what Theorem 4.7 tells us is that we only have one degree of freedom, so increasing the difference of the means, leads to an increase in the scale of the distributions (that is, the distributions become more uniform, and, hence, harder to distinguish).

4.4.2 Estimator & Analysis

We now present the new estimator, namely the Laplace Density Estimator.

We have already argued on the algorithm's privacy guarantees, so we proceed with the accuracy analysis.

Algorithm 4.3: Laplace Density Estimator**Input:** Data stream \mathcal{S} , Privacy budget ε , Accuracy parameters (α, β) **Output:** Density $\tilde{d}(\mathcal{S})$

- 1 Compute m^* and set $m = m^*$
- 2 Sample a random subset $M \subseteq \mathcal{U}$ of m users (without replacement) and define an arbitrary ordering over M
- 3 Create an array $\mathbf{x} = [x_1 \dots x_m]$ and map $M[i] \rightarrow x_i, \forall i \in \{1, \dots, m\}$
- 4 Arbitrarily pick μ_{init}, μ_{upd} and set $b = \frac{\mu_{upd} - \mu_{init}}{\varepsilon}$
- 5 Initialize \mathbf{x} randomly: $x_i \sim \text{Laplace}(\mu_{init}, b), \forall i \in \{1, \dots, m\}$
- 6 **for** $t = 1$ **to** T **do**
- 7 **if** $s_t \in M$ **then**
- 8 Find $i : M[i] = s_t$
- 9 Re-sample: $x_i \sim \text{Laplace}(\mu_{upd}, b)$
- 10 Return $\tilde{d}(\mathcal{S}) = \frac{1}{\mu_{upd} - \mu_{init}} \left(\frac{1}{m} \sum_{i=1}^m x_i - \mu_{init} \right) + \text{Laplace}\left(0, \frac{1}{\varepsilon m}\right)$

Theorem 4.8. For fixed sample M , Algorithm 4.3 provides an unbiased estimate \tilde{d} of the density of \mathcal{S}_M and has mean squared error:

$$\mathbb{E}[(\tilde{d} - d(\mathcal{S}_M))^2] = \frac{2(m+1)}{m^2 \varepsilon^2}$$

Proof. The proof is similar with that of Theorem 4.2.

We begin with the *bias* computation. We examine the distribution of an arbitrary entry in \mathbf{x} :

$$x_i \sim \begin{cases} \text{Laplace}(\mu_{init}, b) & M[i] \notin \mathcal{S}_M \\ \text{Laplace}(\mu_{upd}, b) & M[i] \in \mathcal{S}_M \end{cases}$$

Denoting $\hat{d} = \frac{1}{m} \sum_{i=1}^m x_i$, and applying a similar computation as in the proof of Theorem 4.2, gives:

$$\mathbb{E}[\hat{d}] = \mu_{init} + (\mu_{upd} - \mu_{init})d(\mathcal{S}_M)$$

The final estimate (output) \tilde{d} is:

$$\tilde{d} = \frac{\hat{d} - \mu_{init}}{\mu_{upd} - \mu_{init}} + \text{Laplace}\left(0, \frac{1}{\varepsilon m}\right)$$

so it is an unbiased estimate.

We proceed with the *mean squared error*. Again, since \tilde{d} is an unbiased estimate of $d(\mathcal{S}_M)$, its mean squared error coincides with its variance. We now have m Laplace

random variables x_i ($i = 1, \dots, m$) with the same scale parameter b , so:

$$\begin{aligned}
\text{var}(x_i) &= 2b^2 \\
\Rightarrow \text{var}(\hat{d}) &= \frac{1}{m^2} \sum_{i=1}^m \text{var}(x_i) = \frac{2b^2}{m} \\
\Rightarrow \text{var}(\tilde{d}) &= \left(\frac{1}{\mu_{\text{upd}} - \mu_{\text{init}}}\right)^2 \text{var}(\hat{d}) + \text{var}\left(\text{Laplace}\left(0, \frac{1}{\varepsilon m}\right)\right) \\
&= \frac{2b^2}{m(\mu_{\text{upd}} - \mu_{\text{init}})^2} + \frac{2}{m^2 \varepsilon^2} \\
&= \frac{2b^2}{m(b\varepsilon)^2} + \frac{2}{m^2 \varepsilon^2} \\
&= \frac{2(m+1)}{m^2 \varepsilon^2}
\end{aligned}$$

which completes the proof. \square

We conclude that neither the bias, nor the variance (and mean squared error) of our estimator depends on the choice of μ_{init} , μ_{upd} and b .

Before proceeding with the next theorem, we give the following definition and provide some useful lemmas.

Definition 4.4. *A random variable X with mean $\mathbb{E}[X] = \mu$ is sub-exponential if there are non-negative parameters (v_s, b_s) such that:*

$$\mathbb{E}[e^{\lambda(X-\mu)}] \leq e^{\frac{v_s^2 \lambda^2}{2}}, \quad \forall |\lambda| < \frac{1}{b_s}$$

The control on the moment generating function, when combined with the Chernoff technique allows us to prove the following concentration inequality, called the sub-exponential tail bound. We omit the proof, as it is based on a standard technique.

Lemma 4.3. *Suppose that random variable X is sub-exponential with parameters (v_s, b_s) . Then,*

$$\mathbb{P}(|X - \mathbb{E}[X]| \geq \alpha) \leq \begin{cases} e^{-\frac{\alpha^2}{2v_s^2}} & 0 \leq \alpha \leq \frac{v_s^2}{b_s} \\ e^{-\frac{\alpha}{2b_s}} & \alpha > \frac{v_s^2}{b_s} \end{cases}$$

We observe that, when α is small enough, the exponent is quadratic in α , whereas for larger α , the exponent is linear (so the bound is not as tight).

Lemma 4.4. *Let $X \sim \text{Laplace}(\mu, b)$. Then X is sub-exponential with parameters $(v_s, b_s) = (2b, \sqrt{2}b)$.*

Proof. The moment generating function of the random variable $X - \mu$ is:

$$\mathbb{E}[e^{\lambda(X-\mu)}] = \frac{1}{1 - b^2 \lambda^2}, \quad \forall |\lambda| < \frac{1}{b}$$

Taking into account that: $\frac{1}{1-x} \leq 1 + 2x \leq e^{2x}$, $\forall 0 < x < \frac{1}{2}$ and setting $x = b^2 \lambda^2$, gives:

$$\mathbb{E}[e^{\lambda(X-\mu)}] \leq e^{2b^2 \lambda^2}, \quad \forall |\lambda| < \frac{1}{\sqrt{2}b}$$

We therefore set $v_s = 2b$, which gives:

$$\mathbb{E}[e^{\lambda(X-\mu)}] \leq e^{\frac{v_s^2 \lambda^2}{2}}, \quad \forall |\lambda| < \frac{1}{\sqrt{2b}}$$

and proves the desired result, according to Definition 4.4. \square

Lemma 4.5. *Let $X = \sum_{i=1}^m X_i$, where $X_i \sim \text{Laplace}(\mu_i, b)$. Then X is sub-exponential with parameters $(v_s, b_s) = (2b\sqrt{m}, \sqrt{2b})$.*

Proof. The moment generating function of the random variable $X - \mathbb{E}[X]$ is:

$$\mathbb{E}[e^{\lambda(X-\mathbb{E}[X])}] = \prod_{i=1}^m e^{\lambda(X_i-\mu_i)} \leq \prod_{i=1}^m e^{2b^2 \lambda^2} = e^{2b^2 \lambda^2 m}, \quad \forall |\lambda| < \frac{1}{\sqrt{2b}}$$

where the inequality holds due to Lemma 4.4. We therefore set $v_s = 2b\sqrt{m}$, which gives:

$$\mathbb{E}[e^{\lambda(X-\mathbb{E}[X])}] \leq e^{\frac{v_s^2 \lambda^2}{2}}, \quad \forall |\lambda| < \frac{1}{\sqrt{2b}}$$

and proves the desired result, according to Definition 4.4. \square

Lemma 4.6. *Let $X = \sum_{i=1}^m X_i$, where $X_i \sim \text{Laplace}(\mu_i, b)$. Then,*

$$\mathbb{P}(|X - \mathbb{E}[X]| \geq \alpha) \leq \begin{cases} e^{-\frac{\alpha^2}{8b^2m}} & 0 \leq \alpha \leq 2\sqrt{2}bm \\ e^{-\frac{\alpha}{2\sqrt{2}b}} & \alpha > 2\sqrt{2}bm \end{cases}$$

Proof. The result follows by simply combining Lemma 4.3 and Lemma 4.5. \square

Theorem 4.9. *If the sample maintained by Algorithm 4.3 consists of $m = \mathcal{O}(\frac{1}{\varepsilon^2 \alpha^2} \log \frac{1}{\beta})$ users from \mathcal{U} , then, for fixed input \mathcal{S} :*

$$\mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha) \leq \beta$$

where the probability space is over the random choices of the algorithm.

Proof. The proof is similar with that of Theorem 4.3, so we again only focus on the differences.

Let $\hat{d} = \frac{1}{m} \sum_{i=1}^m x_i$ and recall that $\mu_{\text{upd}} - \mu_{\text{init}} = \varepsilon b$ (by the differential privacy requirement). We again apply Lemma 4.2 twice, so for some $\alpha > 0$ and $\delta_1, \delta_2 \in (0, 1)$:

$$\begin{aligned} \mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha) &= \mathbb{P}(|\tilde{d} - d(\mathcal{S}_M) + d(\mathcal{S}_M) - d(\mathcal{S})| \geq \alpha) \\ &\leq \mathbb{P}(|\tilde{d} - d(\mathcal{S}_M)| \geq \delta_1 \alpha) \\ &\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1) \alpha) \\ &= \mathbb{P}\left(\left| \frac{\hat{d}}{\mu_{\text{upd}} - \mu_{\text{init}}} - \frac{\mu_{\text{init}}}{\mu_{\text{upd}} - \mu_{\text{init}}} + \text{Laplace}\left(0, \frac{1}{\varepsilon m}\right) \right. \right. \\ &\quad \left. \left. - d(\mathcal{S}_M)\right| \geq \delta_1 \alpha\right) \\ &\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1) \alpha) \end{aligned}$$

$$\begin{aligned}
&= \mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}] + \varepsilon b \text{Laplace}(0, \frac{1}{\varepsilon m})| \geq \varepsilon b \delta_1 \alpha) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1) \alpha) \\
&\leq \mathbb{P}(|\varepsilon b \text{Laplace}(0, \frac{1}{\varepsilon m})| \geq \varepsilon b \delta_1 \delta_2 \alpha) \\
&\quad + \mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}]| \geq \varepsilon b \delta_1 (1 - \delta_2) \alpha) \\
&\quad + \mathbb{P}(|d(\mathcal{S}_M) - d(\mathcal{S})| \geq (1 - \delta_1) \alpha) \\
&= p_3 + p_2 + p_1 \quad (\text{respectively})
\end{aligned}$$

As in the proof of Theorem 4.3, we have 3 sources of error to control, and we want $p_1 + p_2 + p_3 \leq \beta$. We bound each error probability separately, so for some $\delta_3, \delta_4 \in (0, 1)$, such that $\delta_3 + \delta_4 < 1$, we want:

$$p_1 < \delta_3 \beta, \quad p_2 < \delta_4 \beta, \quad p_3 < (1 - \delta_3 - \delta_4) \beta$$

Bounding p_1 and p_3 is identical, as these error probabilities are unchanged. We copy the bounds we derived in Theorem 4.3:

$$\begin{aligned}
p_1 &\leq 2e^{-2m\alpha^2(1-\delta_1)^2} \leq \delta_3 \beta \\
&\Leftrightarrow m \geq \underbrace{\frac{1}{2\alpha^2(1-\delta_1)^2} \ln\left(\frac{2}{\beta\delta_3}\right)}_{m_1} = \mathcal{O}\left(\frac{1}{\alpha^2} \ln\left(\frac{1}{\beta}\right)\right) \\
p_3 &\leq e^{-\varepsilon\alpha m\delta_1\delta_2} \leq (1 - \delta_3 - \delta_4) \beta \\
&\Leftrightarrow m \geq \underbrace{\frac{1}{\varepsilon\alpha\delta_1\delta_2} \ln\left(\frac{1}{\beta(1-\delta_3-\delta_4)}\right)}_{m_3} = \mathcal{O}\left(\frac{1}{\varepsilon\alpha} \ln\left(\frac{1}{\beta}\right)\right)
\end{aligned}$$

We focus on p_2 . We again apply the total probability theorem:

$$p_2 = \sum_{\underline{d} \in \mathcal{D}} \mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}]| \geq \varepsilon b \delta_1 (1 - \delta_2) \alpha \mid d(\mathcal{S}_M) = \underline{d}) \mathbb{P}(d(\mathcal{S}_M) = \underline{d})$$

For fixed $d(\mathcal{S}_M)$, $\mathbb{E}[\hat{d}]$ is no longer random. Now \hat{d} is a weighted sum of Laplace random variables x_i , so by Lemma 4.6:

$$\begin{aligned}
&\mathbb{P}(|\hat{d} - \mathbb{E}[\hat{d}]| \geq \varepsilon b \delta_1 (1 - \delta_2) \alpha \mid d(\mathcal{S}_M) = \underline{d}) \\
&= \mathbb{P}(\left| \sum_{i=1}^m x_i - \mathbb{E}\left[\sum_{i=1}^m x_i\right] \right| \geq m \varepsilon b \delta_1 (1 - \delta_2) \alpha \mid d(\mathcal{S}_M) = \underline{d}) \\
&\leq 2e^{-\frac{[m \varepsilon b \alpha \delta_1 (1 - \delta_2)]^2}{8b^2 m}} = 2e^{-\frac{1}{8} m \varepsilon^2 \alpha^2 \delta_1^2 (1 - \delta_2)^2}
\end{aligned}$$

where we used the fact that, if $\alpha < 4$, we are in the area where the exponent in the tail bound is quadratic, since for $\varepsilon \in (0, \frac{1}{2}]$ and $\delta_1, \delta_2 \in (0, 1)$:

$$m \varepsilon b \delta_1 (1 - \delta_2) \alpha < m \varepsilon b \alpha < 2bm$$

In practice, since the density takes values in the $[0, 1]$ interval, we only care about additive deviations $\alpha \in [0, 1]$.

Again, the bound is independent of \underline{d} , so:

$$p_2 \leq \sum_{\underline{d} \in \mathcal{D}} 2e^{-\frac{1}{8} m \varepsilon^2 \alpha^2 \delta_1^2 (1 - \delta_2)^2} \mathbb{P}(d(\mathcal{S}_M) = \underline{d}) = 2e^{-\frac{1}{8} m \varepsilon^2 \alpha^2 \delta_1^2 (1 - \delta_2)^2}$$

and, therefore:

$$p_2 \leq \delta_4 \beta \Leftrightarrow m \geq \underbrace{\frac{8}{\varepsilon^2 \alpha^2 \delta_1^2 (1 - \delta_2)^2} \ln\left(\frac{2}{\beta \delta_4}\right)}_{m_2} = \mathcal{O}\left(\frac{1}{\varepsilon^2 \alpha^2} \ln\left(\frac{1}{\beta}\right)\right)$$

By picking $m \geq \max\{m_1, m_2, m_3\} = \mathcal{O}\left(\frac{1}{\varepsilon^2 \alpha^2} \ln\left(\frac{1}{\beta}\right)\right)$, we ensure that:

$$\mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha) \leq p_1 + p_2 + p_3 \leq \delta_3 \beta + \delta_4 \beta + (1 - \delta_3 - \delta_4) \beta = \beta$$

which completes the proof. \square

Observe that the bound we derive is identical with that of Algorithm 4.1 (proof of Theorem 4.3), both asymptotically and in terms of constants. Therefore, the cost function we use to compute the proposed sample size for Algorithm 4.3 is the same with that presented in section 4.2.3, so we expect the solution to the optimization problem to be the same as well.

4.4.3 Quantized Laplace Density Estimator

In order to reduce the space requirements of Algorithm 4.3, we also examine the impact of quantization on the performance of the algorithm. In particular, instead of storing floating point numbers in the array \mathbf{x} , we quantize each user's entry (which is a real value sampled either from f_{init} or f_{upd}) into L fixed quantization levels, and store the quantized value. Thus, the only modification in Algorithm 4.3 is that, after drawing a value x from either $f_{init} = \text{Laplace}(\mu_{init}, b)$, or $f_{upd} = \text{Laplace}(\mu_{upd}, b)$, we quantize it (using our pre-designed quantizer), and store $q(x)$.

In general, a quantizer consists of L quantization regions, which are separated by $L - 1$ boundaries B_1, \dots, B_{L-1} , and L representation points a_1, \dots, a_L . Assume that we wish to quantize real numbers drawn from a random variable x_i with pdf $f_{x_i}(x)$, $x \in \mathbb{R}$, and that $q(x_i)$ is the resulting random variable after the quantization. Then, the minimum mean squared error quantizer is constructed by selecting the boundaries and representation points that minimize:

$$\mathbb{E}[(x_i - q(x_i))^2] = \sum_{j=1}^L \int_{B_{j-1}}^{B_j} f_{x_i}(x) (x - a_j)^2 dx$$

where $B_0 = -\infty$ and $B_L = \infty$. If both the boundaries and the representation points are unknown, an iterative algorithm can be used to find them (e.g. Lloyd-Max). However, for fixed boundaries, it is a well known result that the representation points that minimize the mean squared error are selected as:

$$a_j = \mathbb{E}[x_i | B_{j-1} < x_i \leq B_j] = \int_{B_{j-1}}^{B_j} x f_{x_i | B_{j-1} < x_i \leq B_j}(x) dx, \quad j = 1, \dots, L$$

In what follows, we focus on the case of $L = 2$, which in fact leads us back to storing a single bit per user - this case will be proven to be particularly interesting, and it inspired us in designing the Optimal Bernoulli Density Estimator (Algorithm 4.2). We briefly present the design of our quantizer.

- The distribution of each x_i depends on the density of the stream: for a sparse stream, x_i is more likely to have been drawn from f_{init} , whereas for a dense stream, x_i is

more likely to have been drawn from f_{upd} . Since we have no prior knowledge on the density of the stream, we assume that x_i is equally likely to have been drawn from either distribution, which gives us the following mixture distribution:

$$f_{x_i}(x) = \frac{1}{2}f_{init}(x) + \frac{1}{2}f_{upd}(x), \quad x \in \mathbb{R}$$

- For $L = 2$, we set the boundary $B = \frac{\mu_{init} + \mu_{upd}}{2}$ due to symmetry. Therefore, we only have to pick the proper representation points to minimize the mean squared error.
- By simple calculus, the representation points are:

$$\begin{aligned} a_1 &= \int_{-\infty}^B x \frac{f_{x_i}(x)}{2} dx = \mu_{init} - \frac{\mu_{upd} - \mu_{init}}{\varepsilon} e^{-\frac{\varepsilon}{2}} \\ a_2 &= \int_B^{\infty} x \frac{f_{x_i}(x)}{2} dx = \mu_{upd} + \frac{\mu_{upd} - \mu_{init}}{\varepsilon} e^{-\frac{\varepsilon}{2}} \end{aligned}$$

For example, we may want to end up with $a_1 = 0$ and $a_2 = 1$. To do so, it is easy to see that we should pick: $\mu_{init} = \frac{e^{-\frac{\varepsilon}{2}}}{2e^{-\frac{\varepsilon}{2}} + \varepsilon}$ and $\mu_{upd} = \frac{e^{-\frac{\varepsilon}{2}} + \varepsilon}{2e^{-\frac{\varepsilon}{2}} + \varepsilon}$.

We now have our binary quantizer at hand. Assume that a value $x \in \mathbb{R}$ is drawn for the user whose entry in \mathbf{x} is x_i ; if $x < B = \frac{\mu_{init} + \mu_{upd}}{2}$, we set $x_i = a_1$, otherwise, we set $x_i = a_2$. We examine the distribution of $q(x_i)$, which is a binary random variable:

$$\begin{aligned} \mathbb{P}(q(x_i) = 1 | M[i] \notin \mathcal{S}) &= \mathbb{P}(\text{Laplace}(\mu_{init}, b) \geq \frac{\mu_{init} + \mu_{upd}}{2}) \\ &= \frac{1}{2} e^{-\frac{\mu_{init} + \mu_{upd} + \mu_{init}}{2b}} = \frac{1}{2} e^{-\frac{\varepsilon}{2}} \\ \mathbb{P}(q(x_i) = 1 | M[i] \in \mathcal{S}) &= \mathbb{P}(\text{Laplace}(\mu_{upd}, b) \geq \frac{\mu_{init} + \mu_{upd}}{2}) \\ &= 1 - \frac{1}{2} e^{-\frac{\mu_{init} + \mu_{upd} - \mu_{upd}}{2b}} = 1 - \frac{1}{2} e^{-\frac{\varepsilon}{2}} \end{aligned}$$

where we used the fact that $\mu_{upd} - \mu_{init} = \varepsilon b$. Therefore, by picking μ_{upd} and μ_{init} such that $a_1 = 0$ and $a_2 = 1$, the quantized Laplace Density Estimator can be viewed as a Bernoulli Density Estimator, that utilizes the following two Bernoulli distributions:

$$\begin{aligned} f_{init} &= \text{Bernoulli}\left(\frac{1}{2} e^{-\frac{\varepsilon}{2}}\right) \\ f_{upd} &= \text{Bernoulli}\left(1 - \frac{1}{2} e^{-\frac{\varepsilon}{2}}\right) \end{aligned}$$

In Figure 4.7 we examine the use of the allocated privacy budget of the derived Bernoulli distributions, and compare it with that of Algorithm 4.1 and 4.2.

4.4.4 Comments on the Use of Other Continuous Distributions

The connections of the Laplace distribution with the definition of differential privacy are well understood, so our choice of this particular continuous distribution is quite natural.

Truncated Distributions. One may argue that it would be beneficial to truncate the utilized distributions, in order to avoid large (either positive or negative) values, which may lead to a poor estimate.

For example, it would be a reasonable choice to truncate the distributions and allow them to take values in the $[0, 1]$ interval, so that each user's entry is interpreted as a

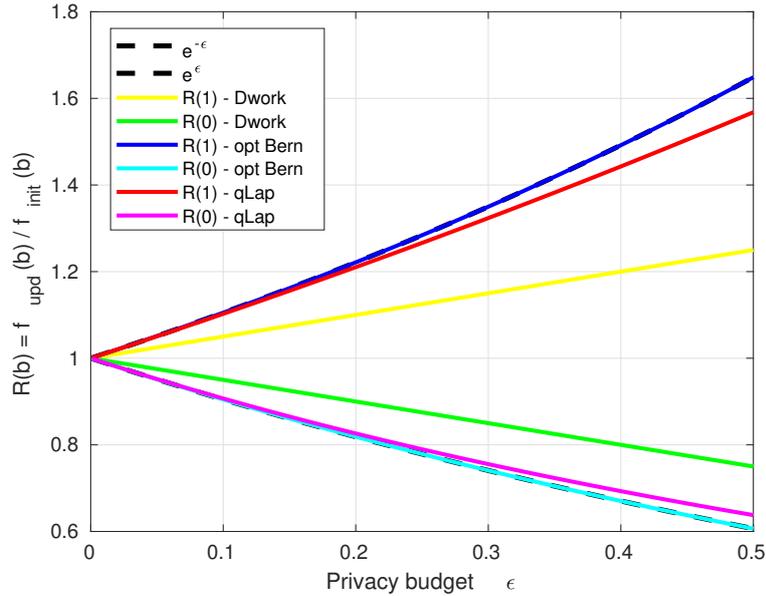


Figure 4.7: Quantized Laplace - State differential privacy

probability of the user having appeared in the stream. Taking into account that the truncated probability density that corresponds to a pdf $f(x)$, $x \in (-\infty, \infty)$, is computed as $f_{trunc}(x) = \frac{f(x)}{\int_{-\infty}^1 f(x) - \int_{-\infty}^0 f(x)}$, $\forall x \in [0, 1]$, we would pick the distributions' means symmetric around $\frac{1}{2}$, so that the normalization factors cancel out when examining the ratio $\frac{f_{upd}}{f_{init}}$ (in order to satisfy differential privacy). If user u (whose entry is x_1) has not appeared, then x_1 is more likely to be closer to 0, whereas if u has appeared, then x_1 is more likely to be closer to 1.

However, the truncation appears to be problematic. The truncated distributions' means come closer to each other (because of the truncation), while the (common) scale parameter does not change. Consequently, for fixed privacy budget, we end up with two harder-to-distinguish distributions.

Gaussian Distributions. The incompatibility of the Gaussian distribution with ϵ -differential privacy is well understood, and unfortunately appears in our setting as well. In particular, the problem arises when attempting to keep the ratio $\frac{f_{upd}}{f_{init}}$ bounded between $e^{-\epsilon}$ and e^{ϵ} , and is due to the square in the exponent of the Gaussian distribution. Although we managed to keep the ratio bounded by truncating the Gaussian distributions and properly selecting their means and (common) variance, the performance of the derived estimator is not promising.

4.5 Experimental Evaluation

In this section, we experimentally compare the algorithms we presented. We conduct all our experiments in MATLAB. In each experiment, we generate a stream of length $T = 10^5$. The universe is the set $\mathcal{U} = \{1, \dots, 10^5\}$ and the stream is either uniform or zipfian (with parameter 1). In the former case we expect to observe a stream density of 0.63, while in the latter of 0.25.

In our first set of experiments, we examine the mean squared error of the algorithms, as a function of the allocated privacy budget. We remark that in those experiments, we

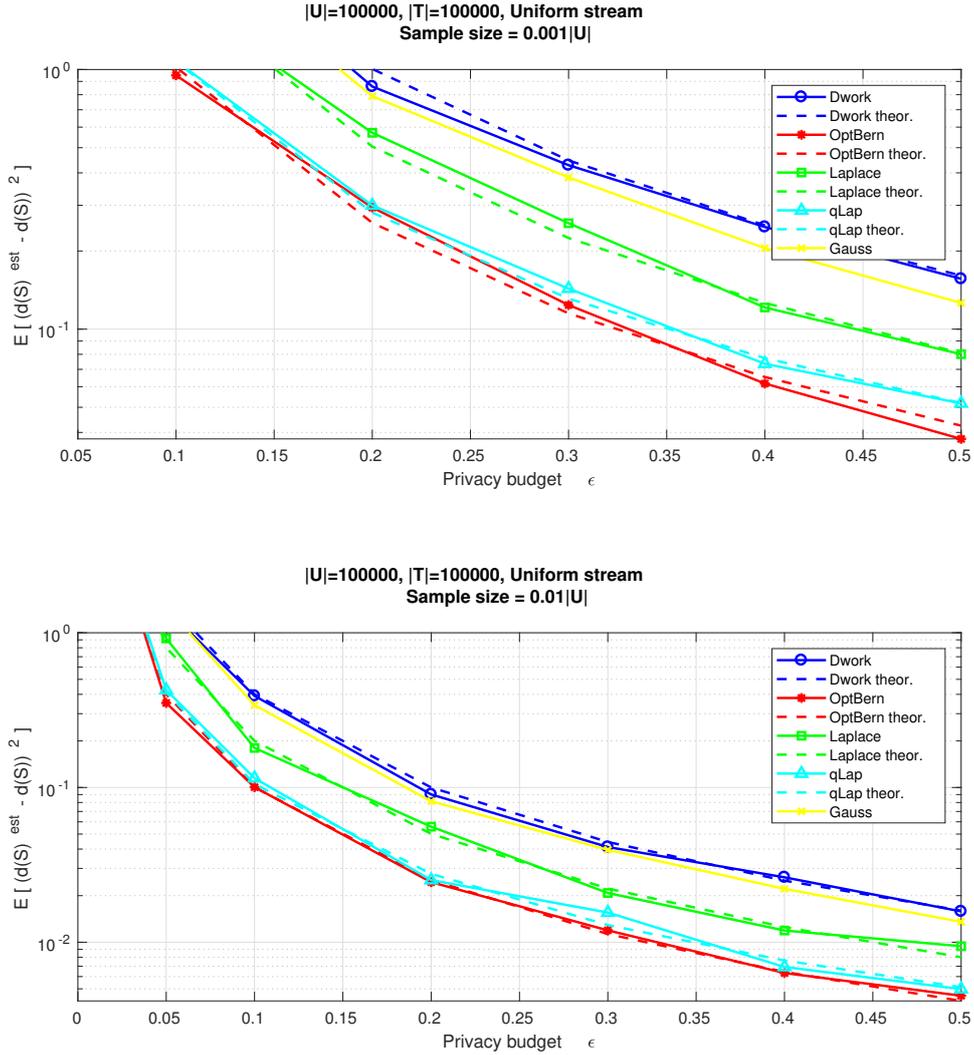


Figure 4.8: MSE for varying privacy budget, and for various sample sizes (Uniform Stream)

do not take into account the input parameters α, β in order to pick the proper sample size; instead we fix the sample size (as a fraction of the universe size) and examine how each the algorithm performs, for varying ϵ . For each ϵ , we independently repeat the experiment 300 times. Notice that in each sub-figure we also plot the theoretical mean squared error (exact or upper bound) of each algorithm.

We present the experimental results in Figures 4.8 and 4.9. The experimental and theoretical mean squared errors coincide for all algorithms. The first thing we observe is the superiority of all the algorithms we propose over Dwork's algorithm. The second is that, among our algorithms, the binary ones (Optimal Bernoulli Density Estimator denoted as optBern, Quantized Laplace Density Estimator denoted qLap) clearly beat the continuous ones (Laplace Density Estimator, Gauss Density Estimator). Finally, we mention the robustness of all algorithms to the input stream distribution; the differences between uniform and zipfian are insignificant.

We next examine our second evaluation metric, that is, the probability $\mathbb{P}(|\tilde{d} - d(\mathcal{S})| \geq \alpha)$ (which we call probability of error for simplicity). Our experiments are again for fixed

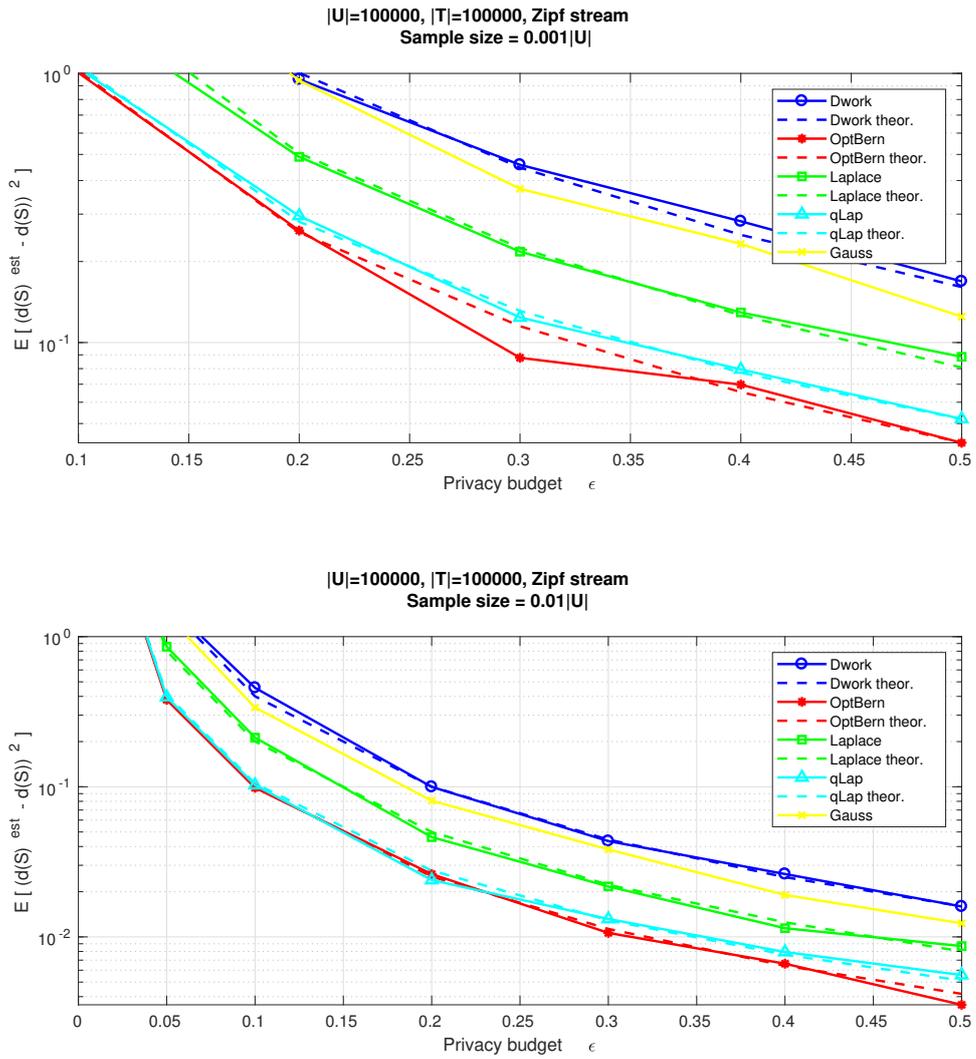


Figure 4.9: MSE for varying privacy budget, and for various sample sizes (Zipf Stream)

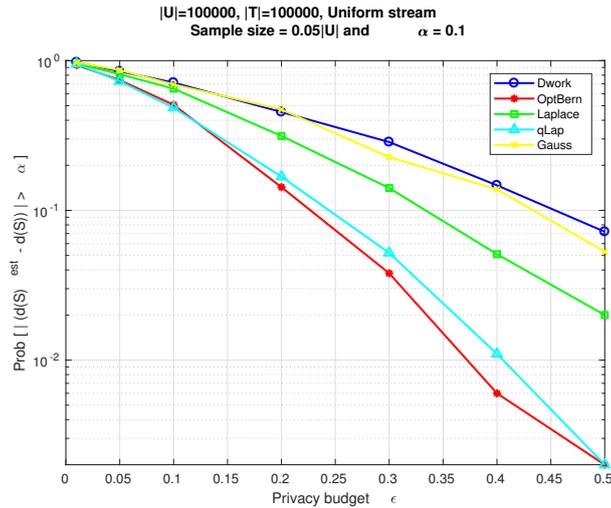


Figure 4.10: Probability of error for varying privacy budget (Uniform Stream)

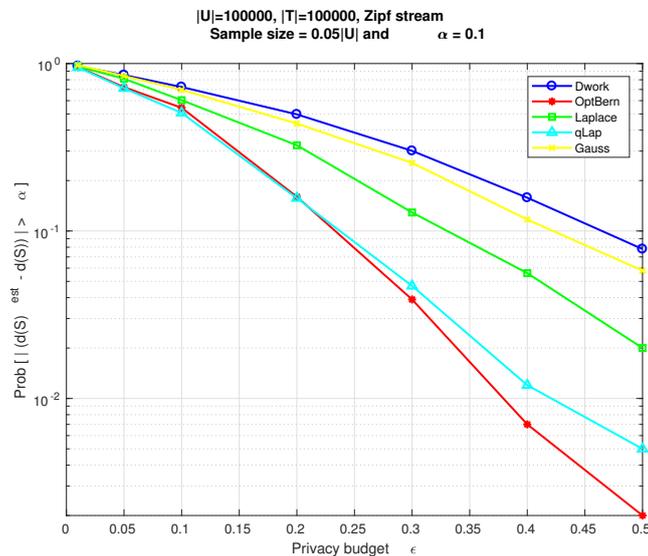


Figure 4.11: Probability of error for varying privacy budget (Zipf Stream)

sample size (as a fraction of the universe size), for fixed α , and for varying privacy budget ϵ . The illustrated probability of error is the empirical probability, computed over 1000 repetitions per ϵ .

We present the experimental results in Figures 4.10 and 4.11. Our second metric confirms the ranking of the algorithms in terms of performance, which we observed in the first set of experiments. The key thing to notice is that the bounds on the probability of error which we computed theoretically are not tight; we computed the tightest version of the bound, and the resulting probability was significantly larger than all the empirical probabilities we plot in Figures 4.10, 4.11.

4.6 Discussion

To summarize, in this chapter we examined differential privacy in the streaming model, and addressed the particular problem of pan-private stream density estimation. We analyzed the sampling-based pan-private density estimator proposed by Dwork et al. [23], and

we improved it by identifying that it does not use all the allocated privacy budget. We managed to tackle this problem by proposing modifications to Dwork's estimator that are based on optimally tuning the Bernoulli distributions it uses, as well as on using continuous distributions (Laplace, Gaussian). Based on both our theoretical and experimental results, the Optimal Bernoulli Density Estimator is the algorithm of choice.

Chapter 5

Conclusions & Future Work

In this thesis, we developed differentially private algorithms to analyze distributed and streaming data.

5.1 Distributed Model

In the distributed model, we addressed the problem of distributed learning of Bayesian Networks with differential privacy. After formally describing our model and identifying the challenges that arise when moving from a centralized to a distributed environment, we examined three solutions, namely Sharing the Noisy Sufficient Statistics, Noisy Majority Voting and Sharing the Noisy Model.

Through our theoretical analysis and our detailed experimental evaluation, we conclude that the first approach, Sharing the Noisy Sufficient Statistics, asymptotically outperforms the other approaches when the privacy budget exceeds a certain threshold. However, its performance is significantly worse below that threshold, and it is more sensitive to the data dimension, the number of data holders and the Bayesian Network degree. The second approach, Noisy Majority Voting is robust to all the parameters we examined (privacy budget, data dimensions, number of data holders, Bayesian Network degree), and significantly outperforms the other algorithms when the data are high-dimensional.

There are many directions in which the work we presented in the distributed model (Chapter 3) can be continued. Firstly, instead of having the data holders share their full, high-dimensional frequency distributions, we may ask them to send a more compact representation, such as a Sketch. Secondly, as we argued, the (global) scoring function (mutual information) used in the Bayesian Network learning problem is non-linear and hence it is not composable from its local values; however, sketch-based methods for approximating the global entropy of distributed streams have been developed (Gabel et al. [29]), and we could apply their ideas in our setting as well. Finally, although we followed a score-based approach in learning the global Bayesian Network, we remark that an approach based on Bayesian model averaging could also be examined.

5.2 Streaming Model

In the streaming model, we addressed the problem of pan-private stream density estimation. We analyzed for the first time the sampling-based pan-private density estimator proposed by Dwork et al. [23], and we improved it by identifying that it does not use all the allocated privacy budget. We managed to tackle this problem by proposing novel modifications to Dwork’s estimator that are based on optimally tuning the Bernoulli distributions it uses, as well as on using continuous distributions (Laplace, Gaussian). Based

on both our theoretical and experimental results, the Optimal Bernoulli Density Estimator is the algorithm of choice.

There are also several directions in which the work we presented in the streaming model (Chapter 3) can be extended. Firstly, we identify that the sampling step performed by our algorithm's is naive; we may apply more sophisticated sampling techniques that are optimized for our particular problem, such as Distinct Sampling [31]. Secondly, we could examine other approaches in density estimation, based on the well-known FM-Sketch and its variants [32]; although our early work on these approaches was not fruitful, there is definitely much more to try. Finally, it would be interesting to combine the model we described in Chapter 2 with that of Chapter 3, and develop differentially private algorithms for our problems in the model of distributed streams. For example, Zhang et al. [66] address the problem of learning Bayesian Network from distributed streams, although without privacy considerations.

References

- [1] Omar Abdel Wahab et al. “DARM: a privacy-preserving approach for distributed association rules mining on horizontally-partitioned data”. In: *Proceedings of the 18th International Database Engineering & Applications Symposium*. ACM. 2014, pp. 1–8.
- [2] Dima Alhadidi et al. “Secure distributed framework for achieving ϵ -differential privacy”. In: *International Symposium on Privacy Enhancing Technologies Symposium*. Springer. 2012, pp. 120–139.
- [3] Boaz Barak et al. “Privacy, accuracy, and consistency too: a holistic solution to contingency table release”. In: *Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM. 2007, pp. 273–282.
- [4] Georgij P Basharin. “On a statistical estimate for the entropy of a sequence of independent random variables”. In: *Theory of Probability & Its Applications* 4.3 (1959), pp. 333–336.
- [5] Garrett Bernstein et al. “Differentially Private Learning of Undirected Graphical Models Using Collective Graphical Models”. In: *arXiv preprint arXiv:1706.04646* (2017).
- [6] Jean Bolot et al. “Private decayed predicate sums on streams”. In: *Proceedings of the 16th International Conference on Database Theory*. ACM. 2013, pp. 284–295.
- [7] Luca Bonomi and Li Xiong. “On Differentially Private Longest Increasing Subsequence Computation in Data Stream.” In: *Transactions on Data Privacy* 9.1 (2016), pp. 73–100.
- [8] TH Hubert Chan, Elaine Shi, and Dawn Song. “Private and continual release of statistics”. In: *International Colloquium on Automata, Languages, and Programming*. Springer. 2010, pp. 405–417.
- [9] T-H Hubert Chan et al. “Differentially private continual monitoring of heavy hitters from distributed streams”. In: *International Symposium on Privacy Enhancing Technologies Symposium*. Springer. 2012, pp. 140–159.
- [10] Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. “Differentially private empirical risk minimization”. In: *Journal of Machine Learning Research* 12.Mar (2011), pp. 1069–1109.
- [11] Rui Chen et al. “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. 2015, pp. 129–138.
- [12] C Chow and Cong Liu. “Approximating discrete probability distributions with dependence trees”. In: *IEEE transactions on Information Theory* 14.3 (1968), pp. 462–467.
- [13] Thomas H Cormen et al. *Introduction to algorithms*. MIT press, 2009.

-
- [14] Graham Cormode and Shan Muthukrishnan. “An improved data stream summary: the count-min sketch and its applications”. In: *Journal of Algorithms* 55.1 (2005), pp. 58–75.
- [15] Graham Cormode et al. “Comparing data streams using hamming norms (how to zero in)”. In: *Proceedings of the 28th international conference on Very Large Data Bases*. VLDB Endowment. 2002, pp. 335–345.
- [16] Graham Cormode et al. “Differentially private spatial decompositions”. In: *2012 IEEE 28th international conference on Data engineering (ICDE)*. IEEE. 2012, pp. 20–31.
- [17] Graham Cormode et al. “Synopses for massive data: Samples, histograms, wavelets, sketches”. In: *Foundations and Trends® in Databases* 4.1–3 (2011), pp. 1–294.
- [18] Christos Dimitrakakis et al. “Robust and private Bayesian inference”. In: *International Conference on Algorithmic Learning Theory*. Springer. 2014, pp. 291–305.
- [19] Bolin Ding et al. “Differentially private data cubes: optimizing noise sources and consistency”. In: *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*. ACM. 2011, pp. 217–228.
- [20] Cynthia Dwork. “Differential privacy in new settings”. In: *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*. SIAM. 2010, pp. 174–183.
- [21] Cynthia Dwork et al. “Calibrating noise to sensitivity in private data analysis”. In: *Theory of cryptography conference*. Springer. 2006, pp. 265–284.
- [22] Cynthia Dwork et al. “Differential privacy under continual observation”. In: *Proceedings of the forty-second ACM symposium on Theory of computing*. ACM. 2010, pp. 715–724.
- [23] Cynthia Dwork et al. “Pan-Private Streaming Algorithms.” In: *Proceedings of The First Symposium on Innovations in Computer Science*. 2010, pp. 66–80.
- [24] Cynthia Dwork, Aaron Roth, et al. “The algorithmic foundations of differential privacy”. In: *Foundations and Trends® in Theoretical Computer Science* 9.3–4 (2014), pp. 211–407.
- [25] Liyue Fan and Li Xiong. “Real-time aggregate monitoring with differential privacy”. In: *Proceedings of the 21st ACM international conference on Information and knowledge management*. ACM. 2012, pp. 2169–2173.
- [26] Philippe Flajolet and G Nigel Martin. “Probabilistic counting algorithms for data base applications”. In: *Journal of computer and system sciences* 31.2 (1985), pp. 182–209.
- [27] James Foulds et al. “On the theory and practice of privacy-preserving Bayesian data analysis”. In: *arXiv preprint arXiv:1603.07294* (2016).
- [28] Arik Friedman et al. “Privacy-Preserving Distributed Stream Monitoring.” In: *Proceedings of The Network and Distributed System Security Symposium*. 2014.
- [29] Moshe Gabel, Daniel Keren, and Assaf Schuster. “Anarchists, unite: Practical entropy approximation for distributed streams”. In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM. 2017, pp. 837–846.
- [30] Minos Garofalakis, Johannes Gehrke, and Rajeev Rastogi. *Data Stream Management: Processing High-Speed Data Streams*. Springer, 2016.
- [31] Phillip B Gibbons. “Distinct sampling for highly-accurate answers to distinct values queries and event reports”. In: *VLDB*. Vol. 1. 2001, pp. 541–550.

- [32] Phillip B Gibbons. “Distinct-values estimation over data streams”. In: *Data Stream Management*. Springer, 2016, pp. 121–147.
- [33] Slawomir Goryczka, Li Xiong, and Vaidy Sunderam. “Secure multiparty aggregation with differential privacy: A comparative study”. In: *Proceedings of the Joint EDBT/ICDT 2013 Workshops*. ACM. 2013, pp. 155–163.
- [34] Jihun Hamm, Yingjun Cao, and Mikhail Belkin. “Learning privately from multiparty data”. In: *International Conference on Machine Learning*. 2016, pp. 555–563.
- [35] Bernard Harris. *The statistical estimation of entropy in the non-parametric case*. Tech. rep. WISCONSIN UNIV-MADISON MATHEMATICS RESEARCH CENTER, 1975.
- [36] Michael Hay et al. “Boosting the accuracy of differentially private histograms through consistency”. In: *Proceedings of the VLDB Endowment* 3.1-2 (2010), pp. 1021–1032.
- [37] Rodney R Howell. “On asymptotic notation with multiple variables”. In: *Dept. of Computing and Information Sciences, Tech. Rep* (2008).
- [38] Justin Hsu et al. “Differential privacy: An economic method for choosing epsilon”. In: *Computer Security Foundations Symposium (CSF), 2014 IEEE 27th*. IEEE. 2014, pp. 398–410.
- [39] Piotr Indyk. “Stable distributions, pseudorandom generators, embeddings and data stream computation”. In: *Proceedings 41st Annual Symposium on Foundations of Computer Science*. IEEE. 2000, p. 189.
- [40] Zhanglong Ji et al. “Differentially private distributed logistic regression using private and public data”. In: *BMC medical genomics* 7.1 (2014), S14.
- [41] Shiva Prasad Kasiviswanathan et al. “What can we learn privately?” In: *SIAM Journal on Computing* 40.3 (2011), pp. 793–826.
- [42] Georgios Kellaris et al. “Differentially private event sequences over infinite streams”. In: *Proceedings of the VLDB Endowment* 7.12 (2014), pp. 1155–1166.
- [43] Daniel Kifer and Ashwin Machanavajjhala. “No free lunch in data privacy”. In: *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*. ACM. 2011, pp. 193–204.
- [44] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [45] Frank McSherry and Kunal Talwar. “Mechanism design via differential privacy”. In: *Foundations of Computer Science, 2007. FOCS’07. 48th Annual IEEE Symposium on*. IEEE. 2007, pp. 94–103.
- [46] Darakhshan Mir et al. “Pan-private algorithms via statistics on sketches”. In: *Proceedings of the thirtieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM. 2011, pp. 37–48.
- [47] Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge university press, 2005.
- [48] Manas Pathak, Shantanu Rane, and Bhiksha Raj. “Multiparty differential privacy via aggregation of locally trained classifiers”. In: *Advances in Neural Information Processing Systems*. 2010, pp. 1876–1884.
- [49] Judea Pearl. “Fusion, propagation, and structuring in belief networks”. In: *Artificial intelligence* 29.3 (1986), pp. 241–288.
- [50] Haoyue Ping, Julia Stoyanovich, and Bill Howe. “DataSynthesizer: Privacy-preserving synthetic datasets”. In: *Proceedings of the 29th International Conference on Scientific and Statistical Database Management*. ACM. 2017, p. 42.

- [51] Arun Rajkumar and Shivani Agarwal. “A differentially private stochastic gradient descent algorithm for multiparty classification”. In: *Artificial Intelligence and Statistics*. 2012, pp. 933–941.
- [52] Anand D Sarwate and Kamalika Chaudhuri. “Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data”. In: *IEEE signal processing magazine* 30.5 (2013), pp. 86–94.
- [53] Izhak Sharfman, Assaf Schuster, and Daniel Keren. “A geometric approach to monitoring threshold functions over distributed data streams”. In: *ACM Transactions on Database Systems (TODS)* 32.4 (2007), p. 23.
- [54] Reza Shokri and Vitaly Shmatikov. “Privacy-preserving deep learning”. In: *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*. ACM. 2015, pp. 1310–1321.
- [55] Beata Strack et al. “Impact of HbA1c measurement on hospital readmission rates: analysis of 70,000 clinical database patient records”. In: *BioMed research international* 2014 (2014).
- [56] Sen Su et al. “Differentially private multi-party high-dimensional data publishing”. In: *2016 IEEE 32nd International Conference on Data Engineering (ICDE)*. IEEE. 2016, pp. 205–216.
- [57] Latanya Sweeney. “k-anonymity: A model for protecting privacy”. In: *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10.05 (2002), pp. 557–570.
- [58] JalaJ Upadhyay. “Differentially private linear algebra in the streaming model”. In: *arXiv preprint arXiv:1409.5414* (2014).
- [59] Stanley L Warner. “Randomized response: A survey technique for eliminating evasive answer bias”. In: *Journal of the American Statistical Association* 60.309 (1965), pp. 63–69.
- [60] Oliver Williams and Frank McSherry. “Probabilistic inference and differential privacy”. In: *Advances in Neural Information Processing Systems*. 2010, pp. 2451–2459.
- [61] Xiaokui Xiao, Guozhang Wang, and Johannes Gehrke. “Differential privacy via wavelet transforms”. In: *2010 IEEE 26th International Conference on Data Engineering (ICDE)*. 2010.
- [62] Liyang Xie, Sergey Plis, and Anand D Sarwate. “Data-weighted ensemble learning for privacy-preserving distributed learning”. In: *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2016, pp. 2309–2313.
- [63] Jia Xu et al. “Differentially private histogram publication”. In: *The VLDB Journal—The International Journal on Very Large Data Bases* 22.6 (2013), pp. 797–822.
- [64] Jun Zhang et al. “PrivBayes: private data release via bayesian networks”. In: *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*. ACM. 2014, pp. 1423–1434.
- [65] Jun Zhang et al. “Privbayes: Private data release via bayesian networks”. In: *ACM Transactions on Database Systems (TODS)* 42.4 (2017), p. 25.
- [66] Yu Zhang, Srikanta Tirthapura, and Graham Cormode. “Learning Graphical Models from a Distributed Stream”. In: *arXiv preprint arXiv:1710.02103* (2017).
- [67] Zuhe Zhang, Benjamin IP Rubinstein, Christos Dimitrakakis, et al. “On the Differential Privacy of Bayesian Inference.” In: *Proceedings AAAI’16 Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. 2016, pp. 2365–2371.

-
- [68] Shijie Zheng. “The differential privacy of Bayesian inference”. Bachelor’s Thesis. Harvard University, 2015.