DISTRIBUTED GAUSSIAN DIFFERENTIAL PRIVACY VIA SHUFFLING

Kan Chen
Graduate Group of Applied Math and Computational Science
University of Pennsylvania
Philadelphia, PA 19104, USA
kanchen@sas.upenn.edu

Qi Long
Department of Biostatistics and Epidemiology
University of Pennsylvania
Philadelphia, PA 19104, USA
qlong@pennmedicine.upenn.edu

ABSTRACT

Traditionally, there are two models for implementing differential privacy: local model and centralized model. *Shuffled model* is a relatively new model that aims to provide greater accuracy while preserving privacy by shuffling batches of similar data. In this paper, we consider the analytic privacy study of a *shuffled model* for "f-differential privacy"(f-DP), a new relaxation of traditional (ϵ, δ)-differential privacy. We provide a powerful technique to import the existing *shuffled model* results proven for the (ϵ, δ)-DP to f-DP, with which we derive a simple and easy-to-interpret theorem of privacy amplification by shuffling for f-DP. Furthermore, we prove that compared with the original *shuffled model* from Cheu et al. (2019), f-DP provides a tighter upper bound in terms of the privacy analysis of sum queries. The approach of f-DP can be applied to broader classes of models to achieve more accurate privacy analysis

1 INTRODUCTION

Differential privacy Dwork et al. (2006b) Dwork et al. (2006a) Barbaro et al. (2006) Narayanan & Shmatikov (2008) Homer et al. (2008) has been developed as a valuable standard for measuring and guaranteeing data privacy. However, it was difficult to be implemented in practice until recently Abowd (2018) Abadi et al. (2016). Practitioners usually need to face the trade-off between privacy and accuracy for implementing differential privacy, with two different models, local model and centralized model.

In local model Erlingsson et al. (2014) Kairouz et al. (2014) Qin et al. (2016), users apply a privacy-preserving procedure, e.g. by adding noise, to their data before sending them to a server for analysis, which guarantees privacy without fully trusting the server at the cost of data accuracy with extra noise. Hence, the local model requires a huge amount of data in order to obtain meaningful results.

As for the centralized model Erlingsson et al. (2019) however, the users' data are directly sent to a trusted server, at which private computations are performed among these data using privacy-preserving techniques. The accuracy can be guaranteed for centralized model, requiring a far less amount of data to achieve reliable results, while the privacy is at stake when user data are centralized at a server, which could be lost, attacked or abused if not secured enough.

Shuffling is a relatively new approach for conducting richer and more reliable data analysis while preserving privacy. *Shuffled model*, originally sparked by Andrea Bittau et al.Bittau et al. (2017) entails three components, namely, Encode, Shuffle, and Analyze (ESA) as follows: user data are firstly encoded locally with two layers of encryption in the encode step. Then in shuffle step, a shuffler undoes the first layer of encryption and shuffles the data after removing the metadata and explicitly identifying features that could associate information with a specific user. Afterwards, the shuffler passes the data to the analyzer, in which the second layer of encryption will be decoded in order to access and analyze the data as the final step. The key idea behind shuffling is to take a middle step between local and centralized model so that privacy can be maintained while achieving a higher level of accuracy. Amin et al. (2020) Dwork et al. (2010a) Balle et al. (2019b) Balle et al. (2019a) Ghazi et al. (2020) Balle et al. (2020) Ghazi et al. (2019) have built on this concept and proposed different algorithms with the same basic structure.

Back to differential privacy framework, despite traditional differential privacy definition ((ϵ, δ)-DP) achieves apparent success, it does not tightly handle composition and other properties Murtagh & Vadhan (2016). As a consequence, the privacy bound is sometimes not tight enough under (ϵ, δ)-DP. Recent efforts have been devoted to developing relaxations of differential privacy in order to overcome such issue. These works include "Rényi Differential Privacy" Mironov (2017) Wang et al. (2019), "concentrated differential privacy" Dwork & Rothblum (2016) Bun & Steinke (2016), "truncated concentrated differential privacy" Bun et al. (2018), and most remarkably, "f-differential privacy" (Gaussian differential privacy) Dong et al. (2019).

f-differential privacy is a new relaxation of differential privacy that can handle this issue Bu et al. (2019) and has some nice properties Kasiviswanathan et al. (2011). Rather than providing a "divergence" based relaxation of differential privacy, f-DP gives a new definition by allowing the full trade-off between type I and type II errors in the simple hypothesis testing problem Wasserman & Zhou (2010) Kairouz et al. (2015) Dwork et al. (2010b):

$$H_0 : \text{the underlying dataset is } X$$
$$H_1 : \text{the underlying dataset is } X'.$$

In this work, we revisit the privacy analysis of shuffled model for distributed differential privacy algorithm proposed by Cheu et al. (2019), and analyze the privacy bound of shuffled model under the framework of f-differential privacy. We prove that compared with the original shuffled model results from Cheu et al. (2019), f-DP provides a tighter upper bound in terms of the privacy analysis of sum queries.

In this work, we revisit the privacy analysis of shuffled model for distributed differential privacy algorithm proposed by Cheu et al. (2019), and analyze the privacy bound of shuffled model under the framework of f-differential privacy. We prove that compared with the original shuffled model results from Cheu et al. (2019), f-DP provides a tighter upper bound in terms of the privacy analysis of sum queries. The remainder of the paper is organized as follows. In section 2, we give some preliminary background of (ϵ, δ)-differential privacy and f-differential privacy. In section 3, we provide the main results of f-DP in shuffled model and compare our results with existing results in shuffled model. In section 4, we discuss some potential extensions and conclude our paper.

2 PRELIMINARIES

We say two datasets X, X' are *neighboring* if they differ on at most one user's data, and denote $X \sim X'$.

Definition 2.1 (Dwork et al. (2006b), Dwork et al. (2006a))**.** *A randomized algorithm M that takes as input a dataset consisting of individuals is (ϵ, δ)-differentially private (DP) if for any pair of datasets X, X' that differ in the record of a single individual, and any event E,*

$$\mathbb{P}[M(X) \in E] \leq e^{\epsilon}\mathbb{P}[M(X') \in E] + \delta \tag{1}$$

When $\delta = 0$, the guarantee is simply called ϵ-DP.

Consider a rejection rule $0 \leq \phi \leq 1$, with type I and type II error rates defined as

$$\alpha_\phi = \mathbb{E}_P[\phi], \quad \beta_\phi = 1 - \mathbb{E}_Q[\phi].$$

It motivates the following definition of trade-off function.

Definition 2.2 (trade-off function)**.** *For any two probability distributions P and Q on the same space, define the trade-off function $T(P, Q) : [0, 1] \to [0, 1]$ as*

$$T(P, Q)(\alpha) = \inf\{\beta_\phi : \alpha_\phi \leq \alpha\}$$

where the infimum is taken over all (measurable) rejection rule.

Trade-off function motivates the definition of f-differential privacy as follows.

Definition 2.3 (f-differential privacy Dong et al. (2019))**.** *Let f be a trade-off function. A mechanism M is said to be f-differentially private if*

$$T(M(S), M(S')) \geq f$$

for all neighboring datasets S and S'.

G_μ-differential privacy is a special type of f-differential privacy by setting $f(\alpha) = T(P, Q)(\alpha)$ where $P = \mathcal{N}(0, 1), Q = \mathcal{N}(\mu, 1)$.

Definition 2.4 (Gaussian Differential Privacy, Dong et al. (2019)). *Denote*

$$G_\mu := T(\mathcal{N}(0, 1), \mathcal{N}(\mu, 1))$$

for $\mu \geq 0$. And explicit expression for the trade-off function G_μ reads

$$G_\mu(\alpha) = \Phi(\Phi^{-1}(1 - \alpha) - \mu)$$

where Φ denotes the standard normal CDF. Then a mechanism M is said to satisfy μ-Gaussian Differential Privacy (μ-GDP) if it is G_μ-DP. That is,

$$T(M(X), M(X')) \geq G_\mu$$

for all neighboring datasets X and X'.

Corollary 2.5 (Balle & Wang (2018)). *A mechanism is μ-GDP if and only if it is $(\epsilon, \delta(\epsilon))$-DP for all $\epsilon \geq 0$, where $\delta(\epsilon) = \Phi(-\frac{\epsilon}{\mu} + \frac{\mu}{2}) - e^\epsilon \Phi(-\frac{\epsilon}{\mu} - \frac{\mu}{2})$.*

Definition 2.6. *The tensor product of two trade-off functions $f = T(P, Q)$ and $g = T(P', Q')$ is defined as*

$$f \otimes g := T(P \times P', Q \times Q')$$

Theorem 1 (Dong et al. (2019)). *Let $M_i(\cdot, y_1, \cdots, y_{i-1})$ be f_i-DP for all $y_1 \in Y_1, \cdots, y_{i-1} \in Y_{i-1}$. Then the n-fold composed mechanism $M : X \to Y_1 \times \cdots \times Y_n$ is $f_1 \otimes \cdots \otimes f_n$-DP.*

Given preliminaries of differential privacy framework and shuffled model, we are now turning into the privacy analysis.

3 Main Result

3.1 A Protocol for Boolean Sums: Privacy Analysis

The first protocol we are trying to analyze is Boolean sums. The protocol $P^{0/1}_{n,\lambda}$ takes $\mathcal{X} = \{0, 1\}$ as input domain and aims to reveal the function $f(x_1, \cdots, x_n) = \sum_{i=1}^n x_i$. And the parameters n and $\lambda \in [0, n]$ of the protocol control the trade-off between the level of privacy and accuracy as the following: Through the protocol, a random set of λ users will choose y_i randomly while the remaining $n - \lambda$ will choose y_i to be their input bit x_i as their single message/output.

Algorithm 1: Shuffled protocol $P^{0/1}_{n,\lambda}$ for computing the sum of bits

Input: $\mathbf{x} = (x_1, \cdots, x_n) \in \{0, 1\}^n, \lambda \in (0, n)$.
Output: $z \in \mathbb{N}$.
for $i \in \{1, 2, \cdots, n\}$ **do**
  $b \leftarrow \text{Ber}(\lambda/n)$
  if $b = 0$ **then**
    $y_i \leftarrow x_i$
  else
    $y_i \leftarrow Ber(1/2)$
  end if
end for
Return $z \leftarrow \frac{n}{n-\lambda}(\sum_{i=1}^n y_i - \lambda/2)$

Definition 3.1. *Define the compound binomial random variable $\mathbf{C}$ as the following*

$$\mathbf{C} \sim \textit{Bin}(s, \frac{1}{2}), \quad \textit{where } s \sim \textit{Bin}(n, \frac{\lambda}{n})$$

then the probability mass function is

$$\mathbb{P}(\mathbf{C} = k) = \sum_{l \geq k}^{n} \binom{l}{k} \binom{n}{l} (\frac{\lambda}{2n})^l (1 - \frac{\lambda}{n})^{n-l},$$

and cumulative distribution function is

$$F_\lambda(k) = \mathbb{P}(\mathbf{C} \le k) = \sum_{k'=0}^{k} \sum_{l \ge k'}^{n} \binom{l}{k'} \binom{n}{l} (\frac{\lambda}{2n})^l (1 - \frac{\lambda}{n})^{n-l}.$$

Theorem 2. *For any $\lambda \in (0, n)$, $P^{0/1}_{n,\lambda}$ is f-differentially private where*

$$f(\alpha) = 1 - F_\lambda(F_\lambda^{-1}(\alpha) + 1),$$

and F_λ is defined in Definition 3.1 with parameter λ.

By the theorem above, we have the following corollary.

Corollary 3.2. *For any $\alpha \in [0, 1]$, $P^{0/1}_{n,\lambda}$ is μ-Gaussian -differentially private where*

$$\begin{aligned} \mu &\ge \max_{0 \le \alpha \le 1} \left\{ \Phi^{-1}(1-\alpha) - \Phi^{-1}(1 - F_\lambda(F_\lambda^{-1}(\alpha) + 1)) \right\} \\ &:= \mu^*_\lambda, \end{aligned}$$

and F_λ is defined in Definition 3.1 with parameter λ.

We defer the proof of Theorem 2 and Corollary 3.2 to Appendix.

3.2 A Protocol for Real Sums: Privacy Analysis

We will then show how to extend our result from computing Boolean sum to sum of bounded real numbers. In this case the data domain is real $\mathcal{X} = [0, 1]$, while the function $f(\boldsymbol{x}) = \sum_{i=1}^n x_i$ remains the same. The main idea of the protocol is to encode each input x_i to a vector of Boolean values $(b_{i1}, b_{i2}, \cdots, b_{ir})$ in $\{0, 1\}^r$ with expected value x_i in average and then apply the similar algorithm as $P^{0/1}_{n,\lambda}$ at each bit.

Theorem 3. *For any $\lambda \in (0, n)$, $P^{\mathbb{R}}_{n,\lambda,r}$ is $f^{\circ r}$-differentially private where*

$$f(\alpha) = 1 - F_\lambda(F_\lambda^{-1}(\alpha) + 1),$$

$$f^{\circ r}(\alpha) = \left(\underbrace{f \otimes f \cdots \otimes f}_{r \text{ terms}} \right)(\alpha),$$

and F_λ is defined in Definition 3.1 with parameter λ.

Proof. The proof follows by the combination of Theorem 1 and Theorem 2. □

Algorithm 2: An encoder $E_r(\mathbf{x})$

Input: $\mathbf{x} = (x_1, \cdots, x_n) \in [0, 1]^n, r \in \mathbb{N}$.
Output: $((b_{11}, b_{12}, \cdots, b_{1r}), \cdots, (b_{n1}, b_{n2}, \cdots, b_{nr}))$ with $b_{ij} \in \{0, 1\}, 1 \le i \le n, 1 \le j \le r$.
for $i \in \{1, 2, \cdots, n\}$ **do**
  Let $\mu \leftarrow \lceil x_i r \rceil$ and $p \leftarrow x_i r - \mu + 1$
  for $j \in \{1, 2, \cdots, r\}$ **do**

$$b_{ij} = \begin{cases} 1 & j < \mu \\ Ber(p) & j = \mu \\ 0 & j > \mu \end{cases}$$

  end for
end for
Return $((b_{11}, b_{12}, \cdots, b_{1r}), \cdots, (b_{n1}, b_{n2}, \cdots, b_{nr}))$

Corollary 3.3. *For any $\alpha \in [0, 1]$, $P^{\mathbb{R}}_{n,\lambda,r}$ is μ-Gaussian differentially private where $\mu \ge \sqrt{r}\mu^*_\lambda$,*

$$\mu^*_\lambda := \max_{0 \le \alpha \le 1} \left\{ \Phi^{-1}(1-\alpha) - \Phi^{-1}(1 - F_\lambda(F_\lambda^{-1}(\alpha) + 1)) \right\},$$

and F_λ is defined in Definition 3.1 with parameter λ.

The proof of Corollary 3.3 is the same as Corollary 3.2. We compare the privacy results from our paper and Cheu et al. (2019) to demonstrate the power of f-DP in Figure 1.

Algorithm 3: Protocol $P^{\mathbb{R}}_{n,\lambda,r} = (R^{\mathbb{R}}_{\lambda,r}, S, A^{\mathbb{R}}_{\lambda,r})$

Input: $\mathbf{x} = (x_1, \cdots, x_n) \in [0,1]^n, r \in \mathbb{N}, \lambda \in (0, n)$.
Output: $z \in [0, n]$.
Let $((b_{11}, b_{12}, \cdots, b_{1r}), \cdots, (b_{n1}, b_{n2}, \cdots, b_{nr})) \leftarrow E_r(\mathbf{x})$
for $i \in \{1, 2, \cdots, n\}$ **do**
  for $j \in \{1, 2, \cdots, r\}$ **do**
    $l \leftarrow \text{Ber}(\lambda/n)$
    if $l = 0$ **then**
      $y_{ij} \leftarrow b_{ij}$
    else
      $y_{ij} \leftarrow \text{Ber}(1/2)$
    end if
  end for
end for
Return $z \leftarrow \frac{n}{r(n-\lambda)}(\sum_{j=1}^{r}\sum_{i=1}^{n} y_{ij} - \frac{\lambda r}{2})$

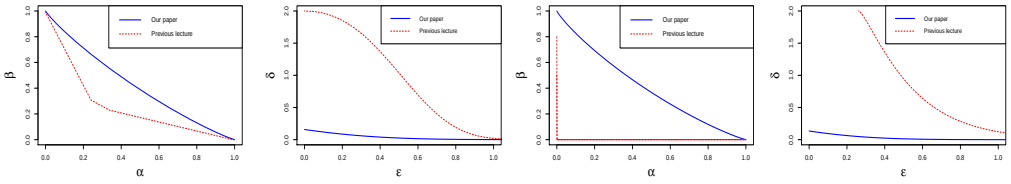

(a) (b) (c) (d)

Figure 1: (a) Privacy analysis of $P^{0/1}_{n,\lambda}$. The Type I error v.s. Type II error plot of the results are that derived from Cheu et al. (2019) and our paper, setting $n = 100, \lambda = 58$. The perfect privacy curve is $\beta = 1 - \alpha$, which implies that $M(X)$ and $M(X')$ are indistinguishable. For the details, refer to Dong et al. (2019). Our curve is much closer to $\beta = 1 - \alpha$ than previous result. (b) Same analysis of $P^{\mathbb{R}}_{n,\lambda,r}$. By converting our result to (ϵ, δ)-DP using Corollary 2.5, we observe that our result is tighter than previous result since our curve is way below the curve derived from Cheu et al. (2019). (c) (d) Privacy analysis of $P^{\mathbb{R}}_{n,\lambda,r}$, setting $r = 10, n = 100, \lambda = 58$. Regular (ϵ, δ)-DP composition results the loss of privacy as shown while GDP composition is tight.

4 Conclusion and Discussion

In this work, we introduce the shuffled model as a composition of Encode, Shuffled and Analyzer protocols under the framework of f-DP, which can tightly handle composition compared to (ϵ, δ)-DP. Thanks to the tightness of composition in f-DP, we can take the privacy analysis of shuffled model to the next level. As two simple applications, we provide privacy analysis of protocols for boolean sums and real sums in shuffled model, in which we derive analytical bounds for both type I and type II errors. Compared with the applications under traditional DP, the privacy bound is much tighter in terms of (ϵ, δ) and error bound is nearly perfect under f-DP. As we can see, the improvement can be tremendous so that we'd suggest f-DP can be implemented to broader classes of shuffle models for a more accurate privacy analysis.

Acknowledgement

We thanks Dr. Weijie Su from Statistics department in Wharton School for tremendous help regarding project ideas. This work is partly supported by NIH under Award Number RF1AG063481 and R01GM124111. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

REFERENCES

Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 308–318, 2016.

John M Abowd. The us census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2867–2867, 2018.

Kareem Amin, Matthew Joseph, and Jieming Mao. Pan-private uniformity testing. In *Conference on Learning Theory*, pp. 183–218. PMLR, 2020.

Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. *arXiv preprint arXiv:1805.06530*, 2018.

Borja Balle, James Bell, Adria Gascon, and Kobbi Nissim. Differentially private summation with multi-message shuffling. *arXiv preprint arXiv:1906.09116*, 2019a.

Borja Balle, James Bell, Adria Gascón, and Kobbi Nissim. The privacy blanket of the shuffle model. In *Annual International Cryptology Conference*, pp. 638–667. Springer, 2019b.

Borja Balle, James Bell, Adria Gascon, and Kobbi Nissim. Private summation in the multi-message shuffle model. *arXiv preprint arXiv:2002.00817*, 2020.

Michael Barbaro, Tom Zeller, and Saul Hansell. A face is exposed for aol searcher no. 4417749. *New York Times*, 9(2008):8, 2006.

Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pp. 441–459, 2017.

Zhiqi Bu, Jinshuo Dong, Qi Long, and Weijie J Su. Deep learning with gaussian differential privacy. *arXiv preprint arXiv:1911.11607*, 2019.

Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pp. 635–658. Springer, 2016.

Mark Bun, Cynthia Dwork, Guy N Rothblum, and Thomas Steinke. Composable and versatile privacy via truncated cdp. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 74–86, 2018.

Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pp. 375–403. Springer, 2019.

Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *arXiv preprint arXiv:1905.02383*, 2019.

Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.

Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pp. 486–503. Springer, 2006a.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pp. 265–284. Springer, 2006b.

Cynthia Dwork, Moni Naor, Toniann Pitassi, Guy N Rothblum, and Sergey Yekhanin. Pan-private streaming algorithms. In *ICS*, pp. 66–80, 2010a.

Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pp. 51–60. IEEE, 2010b.

Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pp. 1054–1067, 2014.

Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2468–2479. SIAM, 2019.

Badih Ghazi, Noah Golowich, Ravi Kumar, Rasmus Pagh, and Ameya Velingker. On the power of multiple anonymous messages. *IACR Cryptol. ePrint Arch.*, 2019:1382, 2019.

Badih Ghazi, Noah Golowich, Ravi Kumar, Pasin Manurangsi, Rasmus Pagh, and Ameya Velingker. Pure differentially private summation from anonymous messages. *arXiv preprint arXiv:2002.01919*, 2020.

Nils Homer, Szabolcs Szelinger, Margot Redman, David Duggan, Waibhav Tembe, Jill Muehling, John V Pearson, Dietrich A Stephan, Stanley F Nelson, and David W Craig. Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays. *PLoS Genet*, 4(8):e1000167, 2008.

Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal mechanisms for local differential privacy. In *Advances in neural information processing systems*, pp. 2879–2887, 2014.

Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In *International conference on machine learning*, pp. 1376–1385. PMLR, 2015.

Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pp. 263–275. IEEE, 2017.

Jack Murtagh and Salil Vadhan. The complexity of computing the optimal composition of differential privacy. In *Theory of Cryptography Conference*, pp. 157–175. Springer, 2016.

Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pp. 111–125. IEEE, 2008.

Zhan Qin, Yin Yang, Ting Yu, Issa Khalil, Xiaokui Xiao, and Kui Ren. Heavy hitter estimation over set-valued data with local differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 192–203, 2016.

Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. Subsampled rényi differential privacy and analytical moments accountant. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1226–1235. PMLR, 2019.

Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010.

Appendix

A NEYMAN-PEARSON LEMMA

Suppose we are performing a hypothesis test between two simple hypothesis $H_0 : \theta = \theta_0, H_1 : \theta = \theta_1$ using the likelihood ratio test with likelihood-ratio threshold η, which rejects H_0 in favour of H_1 at a significant level of

$$\alpha = \mathbb{P}(\Lambda(x) < \eta | H_0)$$

where

$$\Lambda(x) = \frac{\mathcal{L}(\theta_0|x)}{\mathcal{L}(\theta_1|x)}$$

and $\mathcal{L}$ is the likelihood function. Then, the Neyman-Pearson lemma states that the likelihood ratio $\Lambda(x)$ is the most powerful test at significant level α.

B PROOF OF THEOREM 1

To analyze $P^{0/1}_{n,\lambda}$, we need C_λ and C_H as intermediate steps to do the analysis. And we will prove Theorem 2 by using the following claims.

Algorithm 4: $C_\lambda(x_1, \cdots, x_n)$

Input: $\mathbf{x} = (x_1, \cdots, x_n) \in \{0,1\}^n, \lambda \in (0, n)$.
Output: $y \in \mathbb{N}$.
Sample $s \leftarrow Bin(n, \lambda/n)$
Define $\mathcal{H}_s = \{H \subset [n] : |H| = s\}$ and choose $\mathbf{H} \leftarrow \mathcal{H}_s$ uniformly at random
Return $y \leftarrow \sum_{i \notin \mathbf{H}} x_i + Bin(s, 1/2)$

Algorithm 5: $C_H(x_1, \cdots, x_n)$

Input: $\mathbf{x} = (x_1, \cdots, x_n) \in \{0,1\}^n, H \subset \{1, 2, \cdots, n\}$.
Output: $y_H \in \mathbb{N}$.
$\mathbf{B} \leftarrow Bin(|H|, 1/2)$
Return $y_H \leftarrow \sum_{i \notin H} x_i + \mathbf{B}$

Claim B.1 (Cheu et al. (2019))**.** *For every $n \in \mathbb{N}, x \in \{0,1\}^n$, and every $r \in \{0, 1, 2, \cdots, n\}$,*

$$\mathbb{P}[C_\lambda(X) = r] = \mathbb{P}[\sum_{i=1}^{n} R_{n,\lambda}(x_i) = r]$$

where $R_{n,\lambda}(x)$ is the local randomizer which takes input x, output $y = x$ if $b = 0, b \sim Ber(\lambda/n)$ or output $y = k, k \sim Ber(1/2)$ if $b = 1$.

Proof. Fix any $r \in \{0, 1, 2, \cdots, n\}$.

$$\mathbb{P}\left(C_\lambda(X) = r\right) = \sum_{H \subset [n]} \mathbb{P}\left(C_\lambda(X) = r \cap H\right)$$

$$= \sum_{H \subset [n]} \mathbb{P}\left(\sum_{i \notin H} x_i + Bin(|H|, 1/2) = r\right) (\frac{\lambda}{n})^{|H|} (1 - \frac{\lambda}{n})^{n-|H|}$$

$$= \sum_{H \subset [n]} \mathbb{P}\left(\sum_{i \notin H} x_i + \sum_{i \in H} Ber(1/2) = r\right) (\frac{\lambda}{n})^{|H|} (1 - \frac{\lambda}{n})^{n-|H|}$$

Denote G be the random set of people for whom $b_i = 1$ in $P^{0/1}_{n,\lambda}$. Then

$$\mathbb{P}\left(\sum_{i=1}^{n} R_{n,\lambda}(x_i) = r\right) = \sum_{K\subset[n]} \left(\sum_i R_{n,\lambda}(x_i) = r \cap G = K\right)$$
$$= \sum_{K\subset[n]} \left(\sum_{i\notin K} x_i + \sum_{i\notin K} Ber(1/2) = r\right) (\frac{\lambda}{n})^{|K|}(1-\frac{\lambda}{n})^{n-|G|}$$
$$= \mathbb{P}\left(C_\lambda(X) = r\right)$$

which completes the proof. □

Claim B.2 (Cheu et al. (2019))**.** *If C_λ is f-differentially private, then $P^{0/1}_{n,\lambda}$ is f-differentially private.*

Before we start our proof, we need to state the lemma below which we will use later.

Lemma B.3 (Post-Processing (f-DP))**.** *If M is f-differentially private, then for every A, $A \circ M$ is f-differentially private.*

Proof of Claim B.2. Fix any number of users n. Consider the randomized algorithm $A : \{0, 1, 2, \cdots, n\} \to \{0, 1\}^n$ that takes a number r and outputs a uniformly random string z that has r ones. If C_λ is f-DP, then $A \circ C_\lambda$ is f-DP by the post-processing. Now we can complete our proof by showing $(A \circ C_\lambda)(X)$ has the same distribution as $S(R_\lambda(x_1), \cdots, R_\lambda(x_n))$. Fix some vector $Z \in \{0, 1\}^n$ with sum r

$$\mathbb{P}\left(A(C_\lambda(X)) = Z\right) = \mathbb{P}\left(A(r) = Z\right)\mathbb{P}\left(C_\lambda(X) = r\right)$$
$$= \binom{n}{r}^{-1} \mathbb{P}\left(C_\lambda(X) = r\right)$$
$$= \binom{n}{r}^{-1} \sum_{K\in\{0,1\}^n:|K|=r} \mathbb{P}\left(R_{n,\lambda}(X) = K\right)$$
$$= \mathbb{P}\left(S(R_{n,\lambda}(X)) = Z\right)$$

Therefore, if C_λ is f-DP, so does $P^{0/1}_{n,\lambda}$. □

Claim B.4. *For any $\alpha \in [0, 1]$, C_H is f-differential privacy where*

$$f(\alpha) = 1 - F(F^{-1}(\alpha) + 1),$$

and F is the cumulative density function of binomial distribution with parameters $|H|$ and $1/2$.

Proof. To prove this claim, we need to apply Neyman-Pearson here. We first compute the likelihood ratio.

$$\frac{\mathbb{P}(M(X) = r)}{\mathbb{P}(M(X') = r)} = \frac{\mathbb{P}(\mathbf{B} + \sum_{i\notin H} x_i = r)}{\mathbb{P}(\mathbf{B} + \sum_{i\notin H} x'_i = r)}$$

Since $x_j = 0, x'_j = 1$ and $j \notin H$, we have $\sum_{i\notin H} x_i = \sum_{i\notin H} x'_i - 1$. Hence

$$\frac{\mathbb{P}(\mathbf{B} + \sum_{i\notin H} x_i = r)}{\mathbb{P}(\mathbf{B} + \sum_{i\notin H} x'_i = r)} = \frac{\mathbb{P}(\mathbf{B} + \sum_{i\notin H} x_i = r)}{\mathbb{P}(\mathbf{B} + \sum_{i\notin H} x_i + 1 = r)}$$
$$= \frac{\mathbb{P}(\mathbf{B} = r - \sum_{i\notin H} x_i)}{\mathbb{P}(\mathbf{B} = r - \sum_{i\notin H} x_i - 1)} = \frac{\mathbb{P}(\mathbf{B} = k_r + 1)}{\mathbb{P}(\mathbf{B} = k_r)}$$

where we denote $k_r = r - \sum_{i\notin H} x'_i$. Since $\mathbf{B} \sim Bin(|H|, 1/2)$, we obtain

$$\frac{\mathbb{P}(\mathbf{B} = k_r + 1)}{\mathbb{P}(\mathbf{B} = k_r)} = \frac{|H| - k_r}{k_r + 1} = \frac{|H| - r + \sum_{i\notin H} x'_i}{r - \sum_{i\notin H} x'_i + 1}$$
$$= \frac{|H| - r + \sum_{i\notin H} x_i + 1}{r - \sum_{i\notin H} x_i}.$$

Now, we are able to compute type I error and type II error by Neyman-Pearson Lemma. For type I error,

$$\alpha(t) = \mathbb{P}_{r\sim\mathbf{B}+\sum_{i\notin H} x'_i}\left(\frac{|H| - r + \sum_{i\notin H} x'_i}{r - \sum_{i\notin H} x'_i + 1} > t\right)$$
$$= \mathbb{P}_{r\sim\mathbf{B}+\sum_{i\notin H} x'_i}\left(|H| - r + \sum_{i\notin H} x'_i > t(r - \sum_{i\notin H} x'_i + 1)\right)$$
$$= \mathbb{P}_{r\sim\mathbf{B}+\sum_{i\notin H} x'_i}\left(\frac{|H| + (t+1)\sum_{i\notin H} x'_i - t}{t+1} > r\right)$$
$$= \mathbb{P}_{r'\sim\mathbf{B}}\left(r' < \frac{|H| - t}{t+1}\right)$$

here, we denote $r' = r - \sum_{i\notin H} x'_i$. Hence,

$$\alpha(t) = F(\frac{|H| - t}{t+1})$$

where F is the cumulative density function of binomial distribution with parameters $|H|$ and $1/2$. For type II error,

$$\beta(t) = \mathbb{P}_{r\sim\mathbf{B}+\sum_{i\notin H} x_i}\left(\frac{|H| - r + \sum_{i\notin H} x_i + 1}{r - \sum_{i\notin H} x_i} < t\right)$$
$$= \mathbb{P}_{r\sim\mathbf{B}+\sum_{i\notin H} x_i}\left(|H| - r + \sum_{i\notin H} x_i + 1 < t(r - \sum_{i\notin H} x_i)\right)$$
$$= \mathbb{P}_{r\sim\mathbf{B}+\sum_{i\notin H} x_i}\left(\frac{|H| + (t+1)\sum_{i\notin H} x_i + 1}{t+1} < r\right)$$
$$= \mathbb{P}_{c\sim\mathbf{B}}\left(c > \frac{|H| + 1}{t+1}\right)$$

where we denote $c = r - \sum_{i\notin H} x_i$. Hence

$$\beta(t) = 1 - F(\frac{|H| + 1}{t+1}),$$

Now, we need to express $\beta(t)$ in terms of $\alpha(t)$, i.e. $\beta = f(\alpha)$. Because

$$F^{-1}(\alpha) = \frac{|H| - t}{t+1}$$

we have

$$t = \frac{|H| - F^{-1}(\alpha)}{F^{-1}(\alpha) + 1}.$$

Plug into $\beta(t)$, we obtain

$$\beta = 1 - F(F^{-1}(\alpha) + 1)$$

Therefore, we have

$$T(M(X), M(X'))(\alpha) \geq f(\alpha) = 1 - F(F^{-1}(\alpha) + 1)$$

□

Claim B.5. *For any $\alpha \in [0, 1]$, C_H is μ-Gaussian differential privacy where*

$$\mu \geq \max_{0\leq\alpha\leq 1} \left\{\Phi^{-1}(1-\alpha) - \Phi^{-1}(1 - F(F^{-1}(\alpha) + 1))\right\},$$

and F is the cumulative density function of binomial distribution with parameters $|H|$ and $1/2$.

Proof. This result directly comes from the definition of Gaussian differential privacy.

$$T(M(X), M(X'))(\alpha) \geq f(\alpha) \geq \Phi(\Phi^{-1}(1-\alpha) - \mu)$$

Then, we have

$$1 - F(F^{-1}(\alpha) + 1) \geq \Phi(\Phi^{-1}(1-\alpha) - \mu).$$

Therefore,

$$\mu \geq \max_{0 \leq \alpha \leq 1} \left\{ \Phi^{-1}(1-\alpha) - \Phi^{-1}(1 - F(F^{-1}(\alpha) + 1)) \right\}.$$

□

Notice that C_H is a simplified version of C_λ by directly taking batch H as input. As a result, the privacy analysis of C_H will serve as a basic framework followed by privacy analysis of C_λ.

Claim B.6. *For any $\alpha \in [0, 1]$, C_λ is f-differential privacy where*

$$f(\alpha) = 1 - F_\lambda(F_\lambda^{-1}(\alpha) + 1),$$

and F_λ is defined in Definition 3.1 with parameter λ.

Proof. Similar to Claim B.4, we need to apply Neyman-Pearson here. We first compute the likelihood ratio.

$$\frac{\mathbb{P}(M(X) = r)}{\mathbb{P}(M(X') = r)} = \frac{\mathbb{P}(\mathbf{C} + \sum_{i \notin \mathbf{H}} x_i = r)}{\mathbb{P}(\mathbf{C} + \sum_{i \notin \mathbf{H}} x'_i = r)}$$

Since $x_j = 0, x'_j = 1$ and $j \notin \mathbf{H}$, we have $\sum_{i \notin \mathbf{H}} x_i = \sum_{i \notin \mathbf{H}} x'_i - 1$. Hence

$$\begin{aligned}
\frac{\mathbb{P}(\mathbf{C} + \sum_{i \notin \mathbf{H}} x_i = r)}{\mathbb{P}(\mathbf{C} + \sum_{i \notin \mathbf{H}} x'_i = r)} &= \frac{\mathbb{P}(\mathbf{C} + \sum_{i \notin \mathbf{H}} x_i = r)}{\mathbb{P}(\mathbf{C} + \sum_{i \notin \mathbf{H}} x_i + 1 = r)} \\
= \frac{\mathbb{P}(\mathbf{C} = r - \sum_{i \notin \mathbf{H}} x_i)}{\mathbb{P}(\mathbf{C} = r - \sum_{i \notin \mathbf{H}} x_i - 1)} &= \frac{\mathbb{P}(\mathbf{C} = k_r + 1)}{\mathbb{P}(\mathbf{C} = k_r)}
\end{aligned}$$

where we denote $k_r = r - \sum_{i \notin H} x'_i$. Since $\mathbf{C}$ follows the compound binomial distribution defined in Definition 3.1 , we obtain

$$\begin{aligned}
\eta &= \frac{\mathbb{P}(\mathbf{C} = k_r + 1)}{\mathbb{P}(\mathbf{C} = k_r)} \\
&= \frac{\sum_{l \geq k_r + 1}^{n} \binom{l}{k_r + 1} \binom{n}{l} (\frac{\lambda}{2n})^l (1 - \frac{\lambda}{n})^{n-l}}{\sum_{l \geq k_r}^{n} \binom{l}{k_r} \binom{n}{l} (\frac{\lambda}{2n})^l (1 - \frac{\lambda}{n})^{n-l}} \\
&= \frac{\sum_{l \geq k_r}^{n} \binom{l}{k_r + 1} \binom{n}{l} (\frac{\lambda}{2n})^l (1 - \frac{\lambda}{n})^{n-l}}{\sum_{l \geq k_r}^{n} \binom{l}{k_r} \binom{n}{l} (\frac{\lambda}{2n})^l (1 - \frac{\lambda}{n})^{n-l}}
\end{aligned}$$

Notice that $\binom{l}{k} = 0$ for $l < k$. Since for $k_r \leq l \leq n$,

$$\frac{\binom{l}{k_r + 1} \binom{n}{l} (\frac{\lambda}{2n})^l (1 - \frac{\lambda}{n})^{n-l}}{\binom{l}{k_r} \binom{n}{l} (\frac{\lambda}{2n})^l (1 - \frac{\lambda}{n})^{n-l}} = \frac{l - k_r}{k_r + 1}$$

If $\frac{l - k_r}{k_r + 1} < t$ for every $k_r \leq l \leq n$, then $\eta < t$. If $\frac{l - k_r}{k_r + 1} > t$, then $\eta > t$. Therefore, $\{r : \frac{l - k_r}{k_r + 1} < t\} \subset \{r : \eta < t\}$ and $\{r : \frac{l - k_r}{k_r + 1} > t\} \subset \{r : \eta > t\}$. As a consequence, the probability measure of $\{r : \eta < t\}$ or $\{r : \eta > t\}$ is greater than or equal to probability measure of $\{r : \frac{l - k_r}{k_r + 1} < t\}$ or $\{r : \frac{l - k_r}{k_r + 1} > t\}$.

Now, we are able to bound type I error and type II error by Neyman-Pearson Lemma. For type I error,

$$\begin{aligned}
&\alpha(t) = \mathbb{P}_{r\sim\mathbf{C}+\sum_{i\notin\mathbf{H}} x'_i}(\eta > t) \\
&\geq \prod_{l\geq k_r}^{n} \mathbb{P}_{r\sim\mathbf{C}+\sum_{i\notin\mathbf{H}} x'_i}\left(\frac{l-r+\sum_{i\notin\mathbf{H}} x'_i}{r-\sum_{i\notin\mathbf{H}} x'_i+1} > t\right) \\
&= \prod_{l\geq k_r}^{n} \mathbb{P}_{r\sim\mathbf{C}+\sum_{i\notin\mathbf{H}} x'_i}\left(l-r+\sum_{i\notin\mathbf{H}} x'_i > t(r-\sum_{i\notin\mathbf{H}} x'_i+1)\right) \\
&= \prod_{l\geq k_r}^{n} \mathbb{P}_{r\sim\mathbf{C}+\sum_{i\notin\mathbf{H}} x'_i}\left(\frac{l+(t+1)\sum_{i\notin\mathbf{H}} x'_i - t}{t+1} > r\right) \\
&= \prod_{l\geq k_r}^{n} \mathbb{P}_{r'\sim\mathbf{C}}\left(r' < \frac{l-t}{t+1}\right)
\end{aligned}$$

here, we denote $r' = r - \sum_{i\notin\mathbf{H}} x'_i$. Hence,

$$\alpha(t) \geq \prod_{l\geq k_r}^{n} F_\lambda(\frac{l-t}{t+1})$$

For type II error,

$$\begin{aligned}
&\beta(t) = \mathbb{P}_{r\sim\mathbf{C}+\sum_{i\notin\mathbf{H}} x_i}(\eta < t) \\
&\geq \prod_{l\geq k_r}^{n} \mathbb{P}_{r\sim\mathbf{C}+\sum_{i\notin\mathbf{H}} x_i}\left(\frac{l-r+\sum_{i\notin\mathbf{H}} x_i+1}{r-\sum_{i\notin\mathbf{H}} x_i} < t\right) \\
&= \prod_{l\geq k_r}^{n} \mathbb{P}_{r\sim\mathbf{C}+\sum_{i\notin\mathbf{H}} x_i}\left(l-r+\sum_{i\notin\mathbf{H}} x_i+1 < t(r-\sum_{i\notin\mathbf{H}} x_i)\right) \\
&= \prod_{l\geq k_r}^{n} \mathbb{P}_{r\sim\mathbf{C}+\sum_{i\notin\mathbf{H}} x_i}\left(\frac{l+(t+1)\sum_{i\notin\mathbf{H}} x_i+1}{t+1} < r\right) \\
&= \prod_{l\geq k_r}^{n} \mathbb{P}_{c\sim\mathbf{C}}\left(c > \frac{l+1}{t+1}\right)
\end{aligned}$$

where we denote $c = r - \sum_{i\notin\mathbf{H}} x_i$. Hence

$$\beta(t) \geq \prod_{l\geq k_r}^{n} \left(1 - F_\lambda(\frac{l+1}{t+1})\right),$$

Since $0 \leq k_r \leq n$ and we are seeking for the "best" function $\beta = f(\alpha)$ such that it is most close to $\beta = 1-\alpha$, we can take $k_r = n$ (otherwise, the curve will close to axis). Now, we need to express $\beta(t)$ in terms of $\alpha(t)$, i.e. $\beta = f(\alpha)$. Because

$$F_\lambda^{-1}(\alpha) = \frac{n-t}{t+1}$$

we have

$$t = \frac{n - F_\lambda^{-1}(\alpha)}{F_\lambda^{-1}(\alpha)+1}.$$

Plug into $\beta(t)$, we obtain

$$\beta = 1 - F_\lambda(F_\lambda^{-1}(\alpha)+1)$$

Therefore, we have

$$T(M(X), M(X'))(\alpha) \geq f(\alpha) = 1 - F_\lambda(F_\lambda^{-1}(\alpha)+1)$$

□

Proof of Theorem 2. Because C_λ and $P^{0/1}_{n,\lambda}$ share the same privacy guarantees, we can conclude our proof by using Claim B.6 and Claim B.2. □

Proof of Corollary 3.2. The same proof as Claim B.5. □

C (ϵ, δ)-DP AND f-DP OF C_H

We first state the claim from Cheu et al. (2019).

Claim C.1 (Cheu et al. (2019))**.** *For any $\delta > 0$ and any $H \subset [n]$ such that $|H| > 8 \log \frac{4}{\delta}$, C_H is $f_{\epsilon, \frac{\delta}{2}}$ differentially private for*

$$\epsilon = \ln\left(1 + \sqrt{\frac{32 \log \frac{4}{\delta}}{|H|}}\right) < \sqrt{\frac{32 \log \frac{4}{\delta}}{|H|}}$$

By using our **Theorem 3**, we are able to do the following comparison.

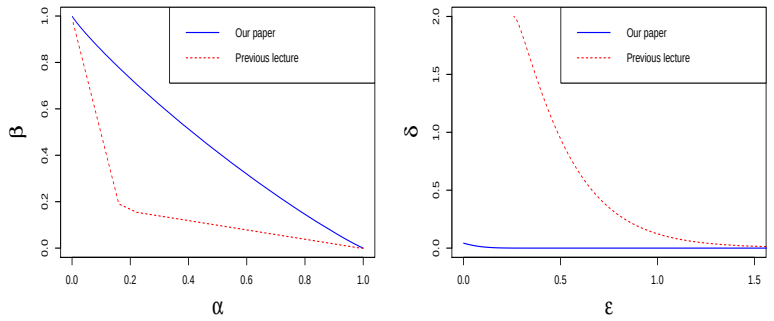

(a) (b)

Figure 2: (a) Privacy analysis of C_H. The Type I error v.s. Type II error plot of the results are that derived from Cheu et al. (2019) and our paper, setting $|H| = 50, n = 100$. The perfect privacy curve is $\beta = 1 - \alpha$, which implies that $M(X)$ and $M(X')$ are indistinguishable. For the details, refer to Dong et al. (2019). Our curve is much closer to $\beta = 1 - \alpha$ than previous result. (b) Same analysis of C_H. By converting our result to (ϵ, δ)-DP, we observe that our result is tighter than previous result since our curve is way below the curve derived from Cheu et al. (2019).