Slowly Scaling Per-Record Differential Privacy

Brian Finley[1], Anthony M Caruso[1], Justin C Doty[1], Ashwin Machanavajjhala[2], Mikaela R Meyer[3], David Pujol[2], William Sexton[2], Zachary Terner[3]

Any opinions and conclusions expressed herein are those of the authors and do not reflect the views of the U.S. Census Bureau.

[1] U.S. Census Bureau
[2] Tumult Labs
[3] The MITRE Corporation

United States Census Bureau

Gist

- We develop formal privacy mechanisms for long-tailed data (*e.g.*, establishments' payroll, revenue, etc.)
- Reduce privacy loss for large records without clipping data (clipping creates bias)
- Mechanisms work by adding noise to transformations of queries or by adding fat-tailed noise
- But first, a quick overview of formal privacy

Very Quick Intro to Formal Privacy - I

- Attacker wants to determine whether your record, r, is in a dataset
- Attacker knows everything except whether r is present
  - Knows value of r, knows the rest of the dataset, D_0
  - Is only trying to decide whether dataset is D_0 or $D_0 \cup \{r\}$
- Attacker's knowledge means no inherent privacy from publishing statistics on large groups
  - Suppose we just publish the number of observations. If dataset is D_0, count is $|D_0|$; if dataset is $D_0 \cup \{r\}$, count is $|D_0| + 1$. Attacker knows $|D_0|$, so can tell which dataset it is.

United States® Census Bureau

Very Quick Intro to Formal Privacy - II

- Instead, add randomness to any statistic, q, from dataset and publish noisy statistic, $\tilde{q}$
  - *E.g.*, add zero-mean Gaussian random variable to q
- Attacker now tries to infer whether r present via Bayesian reasoning, hypothesis testing, or similar (see, *e.g.*, Kifer et al. (2022))
- We inhibit attacker's inferences by ensuring that the distributions of $\tilde{q}(D_0)$ and $\tilde{q}(D_0 \cup \{r\})$ are similar
  - Ensures either database could plausibly have generated realized $\tilde{q}$
  - Quantify "privacy loss" with some measure of distributions' dissimilarity

United States® Census Bureau

Very Quick Intro to Formal Privacy - III

- Per-Record Zero-Concentrated Differential Privacy (PRzCDP) guarantees someone with record r has privacy loss $\leq P(r)$ (Seeman et al., 2023)
- Definition (PRzCDP): Let f_D be the PDF of $\tilde{q}(D)$ and $\ominus$ denote the symmetric set difference. $\tilde{q}$ satisfies P-PRzCDP iff

$$\underbrace{\max_{\alpha \in (1,\infty);\, D, D' \text{ such that } D \ominus D' = \{r\}} \frac{1}{\alpha(\alpha-1)} \int \left(\frac{f_D(x)}{f_{D'}(x)}\right)^{\alpha} f_{D'}(x)dx\,,}_{\text{"Privacy loss"}} \leq P(r).$$

Slowly Scaling Privacy Loss

- Let dataset be a single nonnegative scalar variable and q be a sum over it (more general results in paper)
- $P(r)$ grows with r
- Traditional fix: cap r at some value. Bounds privacy loss, but creates bias
  - See, e.g., Covington et al. (2024)
- Unit splitting: Split r into subrecords capped at some value, apply traditional mechanism, aggregate subrecords' privacy losses (Seeman et al., 2023)
  - $P(r) = O(r^2)$
- Want $P(r)$ to scale more slowly with r, but without bias from capping records
  - Strong protection for small r and weaker, but still meaningful protection for large r
- We contribute two families of mechanisms with slowly scaling $P(r)$

United States® Census Bureau

Additive Mechanism - I

- Additive mechanism simply adds noise to the query from a fat-tailed distribution
- With $Z \sim f_Z(z) \propto e^{-f(|z|)}$

$$\tilde{q} \equiv q + Z$$

- $f(\cdot)$ is a user-chosen, increasing, and concave function
- Privacy guarantee is $P(r) = f(r) - f(0)$
  - Choose slowly scaling f for slowly scaling $P(r)$

United States® Census Bureau

Additive Mechanisms - II

- Generalized Gaussian noise distribution makes $P(r) = O(\sqrt[p]{r})$ for $p \geq 1$
- Exponential polylogarithmic distribution, makes $P(r) = O(\ln(r)^p)$ for $p \geq 1$

United States® Census Bureau

Transformation Mechanism - I

- With $Z \sim N(0, \sigma^2)$, and estimator $g(.)$, transformation mechanism is:
$$\tilde{q} \equiv g(f(q) + Z)$$
- Transformation mechanism transforms q with concave function f so that $f(q)$ itself scales slowly in r
  - Privacy loss comes from differences in query with and without $r \Rightarrow$ if query scales slowly in r, privacy loss scales slowly, too
- g is an estimator of q, using the noisy, transformed q as input
  - $g = f^{-1}$ leads to bias
  - We derive mean- and median-unbiased estimators for many choices of f
- Similar idea in Webb et al. (2023) and Haney et al. (2017)

United States® Census Bureau

Transformation Mechanism - II

- Privacy guarantee is $P(r) = \frac{(f(r)-f(0))^2}{2\sigma^2} = O(f(r)^2)$
  - Slowly scaling f $\Rightarrow$ slowly scaling $P(r)$
- Possible $f(q)$ include $\sqrt[k]{q},\ \ln(q+a)$

United States® Census Bureau

Empirical Experiments

- Simulated data based on County Business Patterns (CBP)
  - CBP is annual Census series of regional establishment data
- Apply mechanisms to sums of employment, grouped by 3-digit NAICS and county
- Employment is very skewed; large values risk large privacy loss
- Transformation and additive mechanisms for 3 asymp. policy functions:

Asymptotic Policy Function ($P(r)$)	Transformation Mechanism Transformation	Additive Mechanism Noise Distribution
$O(r^2)$	Identity	Gaussian
$O(\sqrt{r})$	Fourth root	Generalized Gaussian
$O(\ln(r)^2)$	Log	Exponential Polylogarithmic

- Set variance of all mechanisms to 2 (when $q = 2$, for transformation mechanisms)

United States® Census Bureau

Privacy Loss CDFs

- Each point on CDF shows % of records with lower privacy loss
  - Faster growth is better
- Quickly scaling mechanisms better for very low privacy losses, but quickly lose out to slowly scaling mechanisms
- Transformation mechanisms have larger privacy loss (see x-axis scale)

Additive Mechanisms

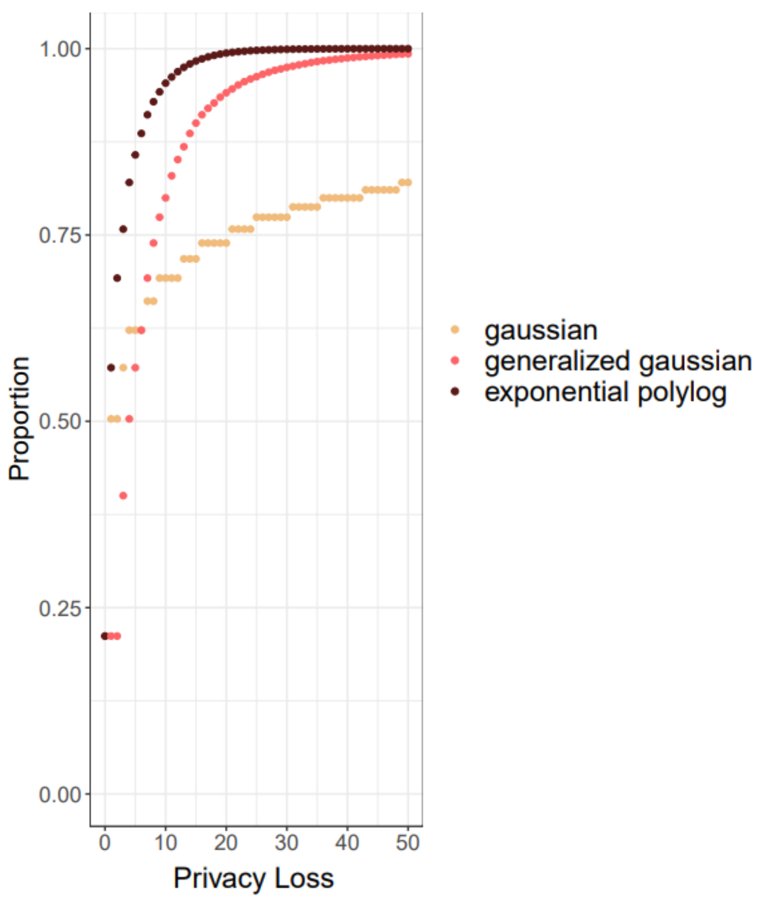

Transformation Mechanisms

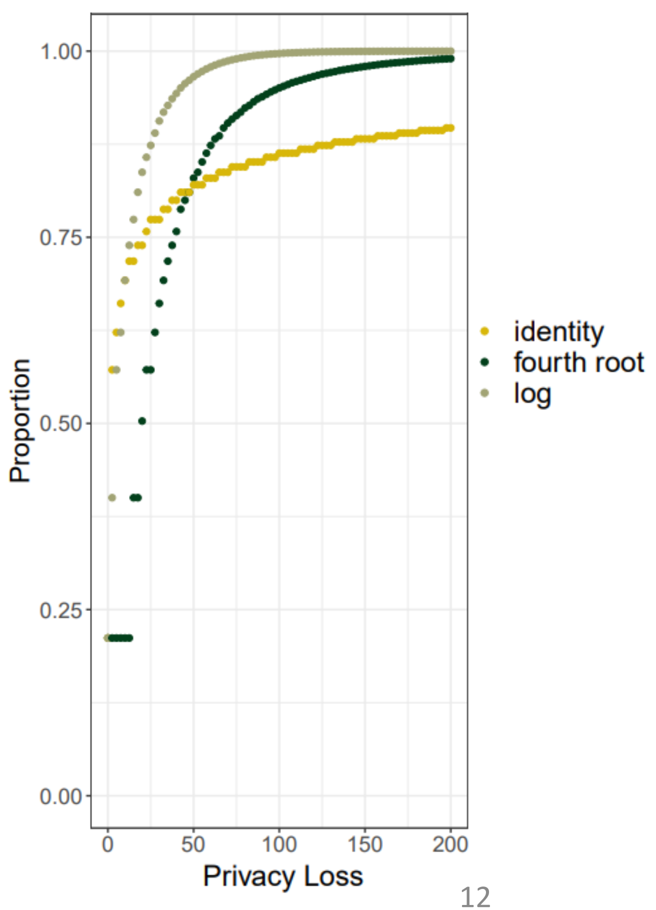

Conclusion

- Developed formally private mechanisms with slowly scaling privacy loss
  - Unbiased mechanisms with more consistent privacy loss for large and small records
- Additive Mechanisms
  - Fat-tailed distributions let privacy loss scale as slowly as log rate
- Transformation Mechanisms
  - Adding noise to transformed query lets privacy scale as slowly as log-squared rate
- Further work on how to choose a specific mechanism and policy function
- Our paper is available at arxiv.org/abs/2409.18118

United States® Census Bureau

Works Cited

- Christian Covington, Xi He, James Honaker, and Gautam Kamath. Unbiased statistical estimation and valid confidence intervals under differential privacy. Statistica Sinica, 2024.
- Samuel Haney, Ashwin Machanavajjhala, John M Abowd, Matthew Graham, Mark Kutzbach, and Lars Vilhuber. Utility cost of formal privacy for releasing national employer-employee statistics. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 1339–1354, 2017.
- Daniel Kifer, John M. Abowd, Robert Ashmead, Ryan Cumings-Menon, Philip Leclerc, Ashwin Machanavajjhala, William Sexton, and Pavel Zhuravlev. Bayesian and frequentist semantics for common variations of differential privacy: Applications to the 2020 census, 2022.
- Jeremy Seeman, William Sexton, David Pujol, and Ashwin Machanavajjhala. Privately answering queries on skewed data via per record differential privacy, 2023.
- Kaitlyn Webb, John Durrell, Daniel Kifer, Prottay Protivash, Aleksandra Slavkovic, Daniell Toth, and Danfeng Zhang. Formal privacy methodology for establishment data [poster presentation]. Theory and Practice of Differential Privacy, 2023