Private Statistical Estimation via Truncation

Manolis Zampetakis, Felix Zhou

Differential Privacy

Differential privacy is the gold standard for privacy-preserving statistical analysis.

Definition ((ε, δ)-Differential Privacy) [Dwork, McSherry, Nissim, Smith 2006]
A randomized algorithm $\mathcal{A} : X^n \to Y$ is (ε, δ)-*DP* if for all neighboring datasets $D \sim D' \in X^n$ that differ by one entry, and all events $S \in \mathrm{range}(\mathcal{A})$:

$$\Pr[\mathcal{A}(D) \in S] \le e^{\varepsilon} \cdot \Pr[\mathcal{A}(D') \in S] + \delta.$$

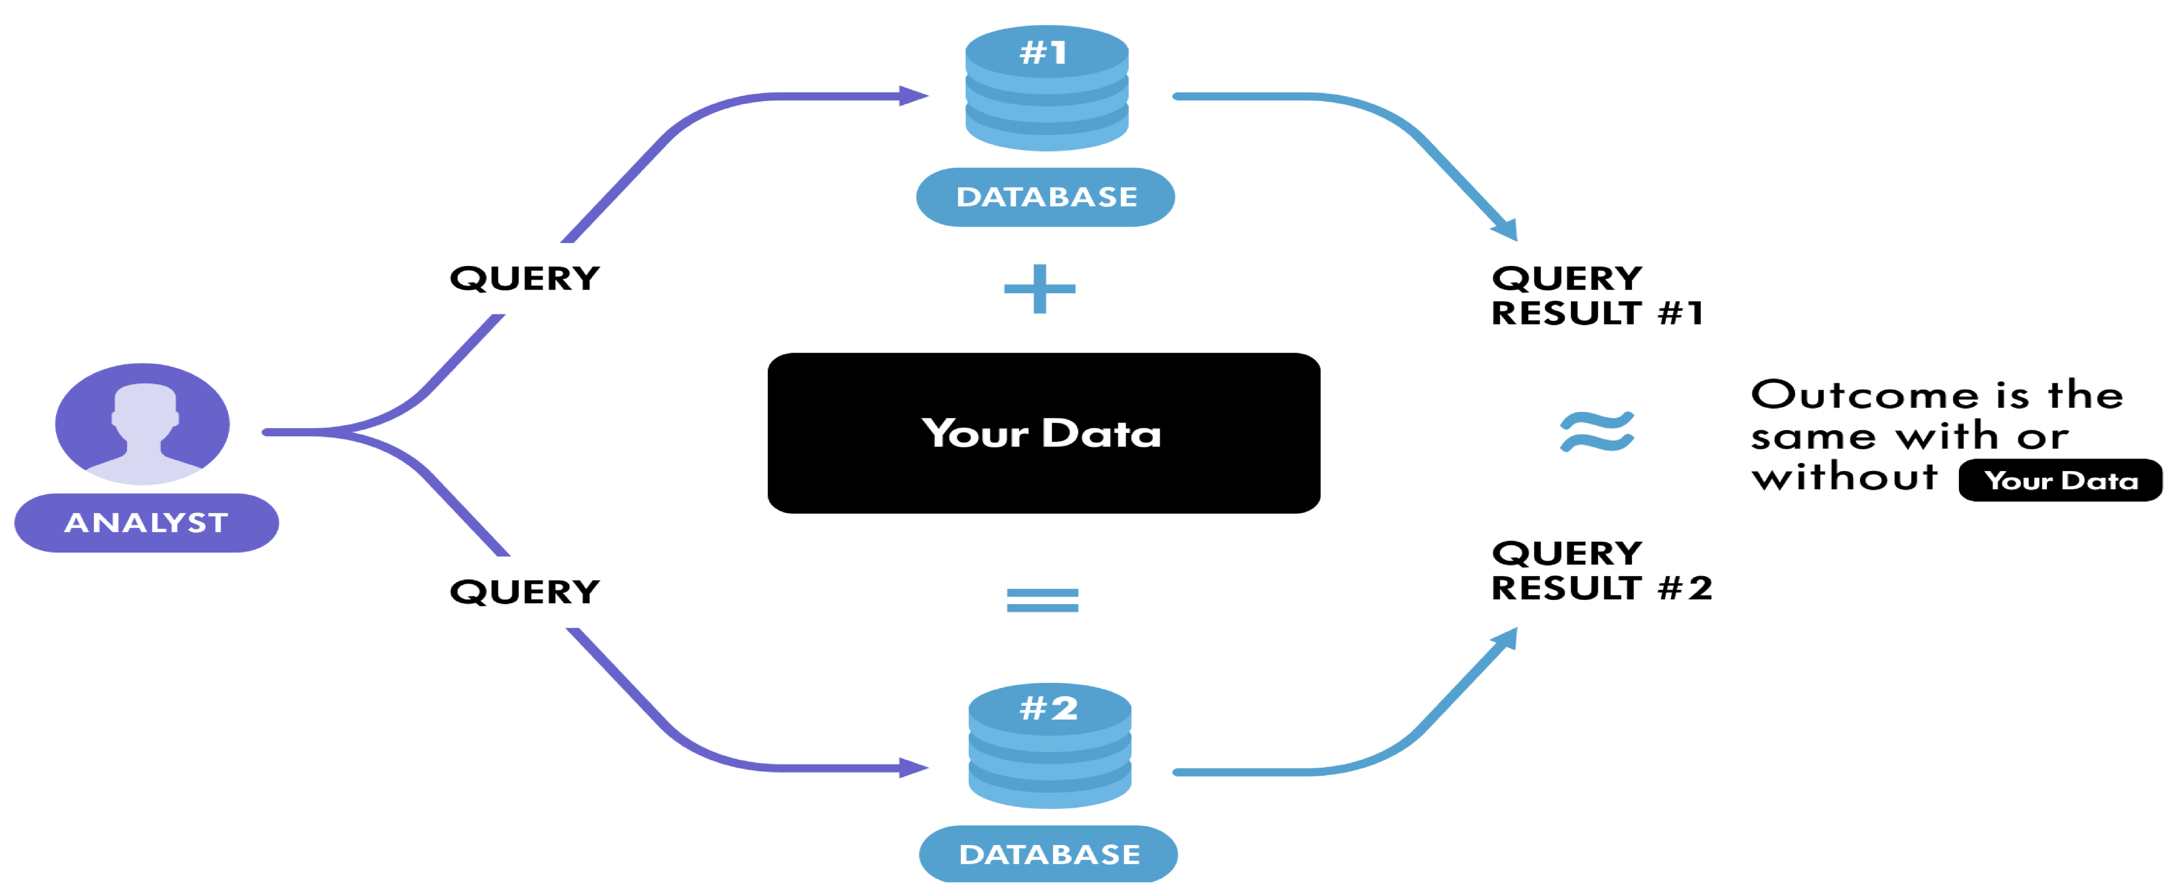

Challenges of Unbounded Data & Connection to Truncated Statistics

Data Sensitivity DP algorithms inject noise into intermediary steps of the algorithm, calibrated to the privacy parameters ε, δ and the *sensitivity* of the data.

Clipping A common folklore technique to reduce the sensitivity of datasets drawn from an unbounded distribution requires *clipping (projecting)* the data to a bounded region containing $1 - o(1)$ of the mass.

Truncation Truncated statistics is concerned with learning from *truncated* samples, where samples outside a known survival set are discarded, as long as the survival set has mass at least $\rho = 0.5 = \Omega(1)$.

Can we formally establish a relationship between differential privacy and truncated statistics?
No prior work studies this intuitive connection.

Truncation vs Clipping

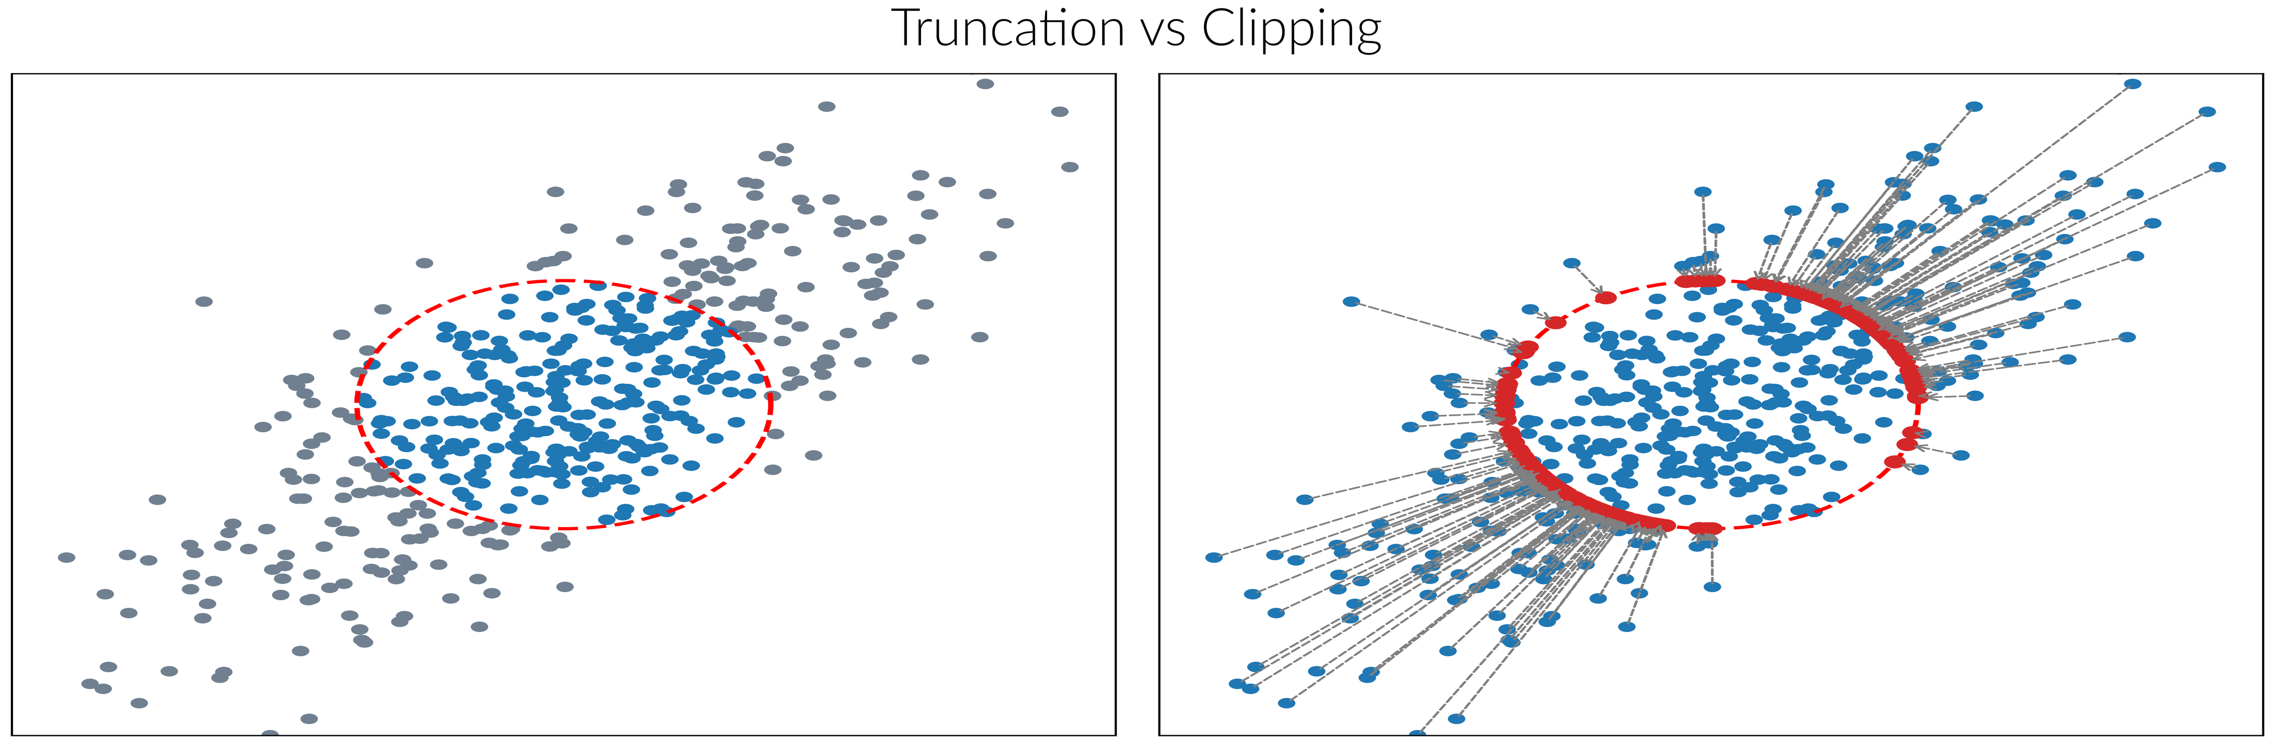

Task: Private Parameter Estimation of Exponential Families

Given samples $x_1, \dots, x_n \sim p_\theta$ from an exponential family distribution with unknown parameter $\theta \in \mathbb{R}^m$, estimate θ up to error $\alpha > 0$ under (ε, δ)-DP, where

$$p_\theta(x) = h(x) \exp\left(\theta^\top T(x) - Z(\theta)\right) .$$

- $h(x) : \mathbb{R}^d \to \mathbb{R}$ is the base measure
- $\theta \in \mathbb{R}^m$ is the natural parameter
- $T(x) : \mathbb{R}^d \to \mathbb{R}^m$ is a sufficient statistic
- $Z(\theta) : \mathbb{R}^m \to \mathbb{R}$ is the log-partition function

Example (Gaussian Mean Estimation)
Given samples $x_1, \dots, x_n \sim \mathcal{N}(\mu, I_d)$, estimate μ up to error $\alpha > 0$ under (ε, δ)-DP.

Main Results

General Exponential Families
There is an (ε, δ)-DP algorithm that outputs an estimate $\hat{\theta}$ such that $\left\|\hat{\theta} - \theta\right\| \le \alpha$ given

- $n = \tilde{O}\left(\frac{m}{\varepsilon} + \frac{m}{\alpha^2} + \frac{m}{\alpha\varepsilon}\right)$ samples
- $\mathrm{poly}(d, n)$ running time

Gaussian Mean Estimation (Special Case)
There is an (ε, δ)-DP algorithm that outputs an estimate $\hat{\mu}$ such that $\|\hat{\mu} - \mu\| \le \alpha$ given

- $n = \tilde{O}\left(\frac{d}{\varepsilon} + \frac{d}{\alpha^2} + \frac{d}{\alpha\varepsilon}\right)$ samples
- $\mathrm{poly}(d, n)$ running time

Proof Sketch

We impose a *truncation* of the dataset by discarding outlier samples, and then use stochastic gradient descent to optimize the empirical negative log-likelihood function $L(\theta)$ of the *truncated* samples.

- $L(\theta)$ is $\Omega(1)$-strongly convex assuming the survival mass $\rho = \Omega(1)$.
- We can obtain unbiased stochastic gradients $\mathbb{E}[g(\theta)] = \nabla L(\theta)$ via rejection sampling.
- We prove a new uniform convergence result: the empirical minimizer is close to θ after $\tilde{O}\left(\frac{m}{\alpha^2}\right)$ samples, improving on the previous rate of $O\left(\frac{m^2}{\alpha^4}\right)$ due to Shah, Shah, and Wornell [2021].

Future Work

- Can we develop further applications of truncated statistics in private algorithm design, such as regression and linear dynamics?
- Current algorithm assumes log-concavity and polynomial sufficient statistics to ensure anti-concentration (strong convexity). Can we relax those assumptions?

SCAN ME

felix.zhou@yale.edu