Exact Privacy Guarantees for Markov Chain Implementations of the Exponential Mechanism

Jeremy Seeman, Matthew Reimherr, and Aleksandra Slavkovic

Department of Statistics, Pennsylvania State University

Contact: jhs5496@psu.edu

NeurIPS 2021

December 9, 2021

Contributions

1. We propose an exact finite-runtime algorithm for implementing the exponential mechanism with exact ϵ-DP guarantees using atomic regeneration.
2. We show that our proposed algorithm relies on the confidential data to avoid the worst-case mixing times found in distance convergence arguments.
3. We derive two modifications of the previous algorithm that demonstrate a new three way trade-off between, privacy, utility, and runtime.

Differential privacy

Notation:

$$\begin{cases} \mathcal{X} \triangleq \text{: sample space of 1 individual's data} \\ (\mathcal{Y}, \mathcal{F}) \triangleq \text{output space} \\ \mathcal{M} \triangleq \{\mu_X \mid X \in \mathcal{X}^n\} \text{ release mechanism} \end{cases}$$

Differential privacy (Dwork et al, 2006)

A mechanism $\mathcal{M}$ satisfies (ϵ, δ)-DP if, for all $B \in \mathcal{F}$ and adjacent $X, X' \in \mathcal{X}^n$:

$$\mu_X(B) \leq e^{\epsilon} \mu_{X'}(B) + \delta.$$

When $\delta = 0$, we say a mechanism satisfies ϵ-DP.

Private selection:

- Goal: minimize loss function L_X while satisfying ϵ-DP

$$L_X : \mathcal{X}^n \times \mathcal{Y} \mapsto [0, \infty]$$

- Key ingredient: bounded sensitivity of L_X. For all adjacent $X, X' \in \mathcal{X}^n$:

$$|L_X(y) - L_{X'}(y)| \leq \Delta_L < \infty.$$

Exponential mechanism (McSherry and Talwar, 2007)

A sample from density f_X with the form:

$$f_X(y) \propto \exp\left(-\frac{\epsilon L_X(y)}{2\Delta_L}\right),$$

with respect to a common base measure $\nu(y)$ over $(\mathcal{Y}, \mathcal{F})$ satisfies ϵ-DP.

Problem: intractable distributions

Many generic ϵ-DP algorithms are not exactly implementable!

Why can't we just use MCMC?

- MCMC approximation has a privacy cost
- Heuristic MCMC convergence measures tell us nothing about said cost

Approximation δ cost (Li et al, 2016)

A sequence of mechanisms $\mathcal{M}_m \triangleq \{\mu_{m,X} \mid x \in \mathcal{X}\}$ approximating the exponential mechanism, $\mathcal{M}_m$, as $m \geq \tau(\alpha)$ is $(\epsilon, \delta_\alpha)$-DP where $\delta_\alpha \triangleq \alpha(1 + e^\epsilon)$ if

$$\tau(\alpha) \triangleq \sup_{X \in \mathcal{X}^n} \inf\{t \geq 0 \mid \|\mu_{t,X} - \mu_X\|_{\mathrm{TV}} \leq \alpha\}.$$

Existing work

Current approach: bounding distributional distances between the MCMC approximation and the target distribution (ex: Ganesh and Talwar, 2020)

Problems with existing approaches:

- Asymptotic rates **can't be used to calculate finite-chain privacy loss**
- Need to bound distances for **worst-case slowest mixing chains**
- **Methods don't exactly satisfy** ϵ-DP

Proposed alternative: perfect sampling

For a Markov chain with stationary distribution μ_X and transition kernel Π_X:

- Associate with each state a binary indicator $\rho \in \{0, 1\}$ indicating **regeneration** (i.e. return to the same state)
- Let $\{\tau\}_{t=1}^{\infty}$ be the sequence of **regeneration times** for the MC (i.e. time between states when $\rho = 1$)

Each τ_t is IID (can drop t index) $\implies$ μ_X has an infinite mixture form:

$$\mu_X(A) = \sum_{m=1}^{\infty} \frac{\mathbb{P}(\tau \geq m)}{\mathbb{E}[\tau]} \mathbb{P}(Y_m \in A \mid \tau \geq m).$$

(Lee et al, 2014) show that if the regeneration state is a **singleton atom**, then we can sample from μ_X using Bernoulli factories (Huber, 2013).

Our proposal: confidential artificial atoms

Implementation choices specific to DP:

- Choose an *artifical* confidential atom $a \in \mathcal{Y}$ from the set of confidential results (i.e. what we would release without privacy preservation)

$$a \in \arg\inf_{y \in \mathcal{Y}} L_X(y).$$

- First sample from $\tilde{\mu}_X$, where:

$$\tilde{\mu}_X = (1 - k)\mu_X + k\xi_a,$$

  then condition on $Y \neq a$ to sample from μ_X.
  - Assumptions about the state space (such as compact $\mathcal{X}^n$) help to satisfy our privacy AND our sampling assumptions
  - Many different possible choices for chain modification (ex: Brockwell and Kadane, 2005)

Result: modified MH for exponential mechanism

Theorem: modified Metropolis-Hastings perfect sampling for privacy

Let Π_X be the transition kernel for a Metropolis-Hastings Markov Chain with symmetric proposals q. We can construct a Markov chain on the extended space with proposals:

$$\tilde{q}(y, y') = \frac{1}{2}\left[q_X(y, y') + \mathbb{1}_{\{y'=a\}}\right],$$

And an algorithm to sample from density f_X that satisfies ϵ-DP with expected number of total proposed samples N_{prop}:

$$\mathbb{E}[N_{\mathrm{prop}}] \leq \frac{48}{k^2(1-k)^2 \inf_{y \in \mathcal{Y}} p_{\mathrm{Accept}}(y)},$$

where:

$$p_{\mathrm{Accept}}(y) \triangleq \int_{\mathcal{Y}} q_X(y, y') \min\left\{1, \frac{f_X(y')}{f_X(y)}\right\} d\nu(y').$$

Example: d-dimensional Laplace mechanism in hypercube

Key property

MCMC methods require accounting for the **slowest mixing** chain, but our method can be much faster because **the runtime depends on the realized confidential data**

Illustrative example: Laplace mechanism ($L_X(y) = \left\|\overline{X} - y\right\|_1$) with data bounded in $[0,1]^d$

- Two original Markov chains: Metropolis-Hastings (MH) with independent uniform proposals and symmetric Laplace proposals with scale α
- Closed form expressions for worst-case δ with MH MCMC (Mengersen and Tweedie, 1996)

$$\|\mu_{X,m} - \mu_X\|_{TV} \leq (1 - \beta_{\mathrm{MCMC}})^m, \tag{1}$$

$$\begin{cases} \beta_{\mathrm{MCMC,Unif}} \triangleq \left(\frac{2d}{\epsilon n}(1 - e^{-\epsilon n/2d})\right)^d \\ \beta_{\mathrm{MCMC,Lap}} \triangleq (2\alpha)^d \exp\left(-\left(\alpha d + \frac{\epsilon n}{2}\right)\right)\left(\frac{1}{\alpha}(1 - e^{-\alpha})\right)^d \end{cases}$$

(Dashed line = 64-bit double precision threshold)

a) Independent uniform proposals:

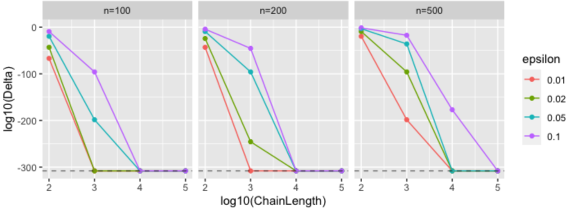

b) Symmetric Laplace proposals, scale = $\epsilon n/2$:

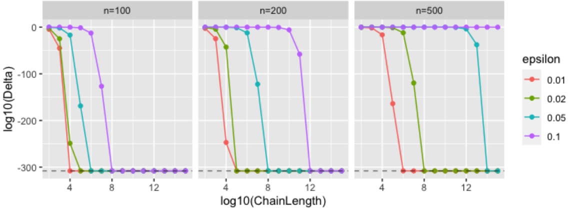

Example: effect of central concentration

Continued example, but now with $X \sim \mathrm{Beta}(\theta, \theta)$ with $\theta \in \{.1, 1, 10\}$ (Laplace proposals), demonstrating dependence on X:

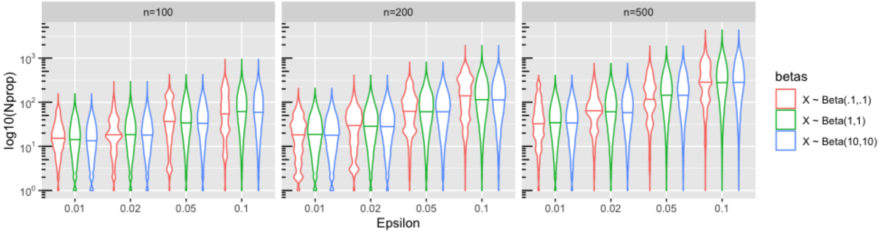

Discussion

- Benefits
  - Satisfies ϵ-DP
  - Runtime depends on realized confidential data, and *not* the confidential data for the slowest-mixing Markov Chain
  - Only requires minorizing bound, and not properties of L_X (i.e. convexity, Lipschitz, etc.)
- Limitations
  - Uniform ergodicity assumption: methods do not have finite expected runtime for unbounded state spaces, like $\mathbb{R}^d$.
  - Minorizing constant suffers from curse of dimensionality
  - Side-channel vulnerability: multiple replications of similar queries could leak information about confidential data through runtime

Extensions of main results

Traditional analysis: privacy vs. utility

Extensions of our work: privacy vs. utility vs. runtime

- Trading off utility and runtime:
  - Exponential mechanisms can be implemented exactly over enumerable discrete state spaces
  - $\implies$ corollary: if we release a sample from a discrete approximation w.p. k, then we reduce runtime at the cost of some utility
- Trading off privacy and runtime:
  - (Awan and Rao, 2021) consider rejection sampling where N_{prop} is known and can leak information
  - $\implies$ corollary: with longer artificial runtime, can release $\tilde{N}_{\text{prop}} \perp\!\!\!\perp X$ with 0-DP so that $(Y, \tilde{N}_{\text{prop}})$ is ϵ-DP

Thanks!

Contact information: jhs5496@psu.edu

This work was sponsored by NSF SES-1853209. Thanks to Alexei Novikov and Jordan Awan for helpful discussions!

Selected references:

- Anthony E Brockwell and Joseph B Kadane. Identification of regeneration times in mcmc simulation, with application to adaptive schemes. Journal of Computational and Graphical Statistics, 14(2):436–458, 2005
- Arun Ganesh and Kunal Talwar. Faster differentially private samplers via renyi divergence analysis of discretized langevin mcmc. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan,and H. Lin, editors, Advances in Neural Information Processing Systems, volume 33, pages 7222–7233. Curran Associates, Inc., 2020.
- Anthony Lee, Arnaud Doucet, and Krzysztof tatuszynski. Perfect simulation using atomic regeneration with application to sequential monte carlo. arXiv preprint arXiv:1407.5770, 2014
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07), pages 94–103. IEEE, 2007.
- K L Mengersen and R L Tweedie. Rates of Convergence of the Hastings and Metropolis Algorithms. The Annals of Statistics, 24(1):101–121, 1996