A Simulation-Based Approach to Differential Privacy

Younes Boulaguiem, PhD
Postdoctoral Researcher in Statistics

University of Geneva – Faculty of Medicine
Clinical Research Center, HUG

A simulation-based fiducial matching approach for DP inference.
Delivered May 2023

View Preprint

Motivating Problem

- Suppose we are implementing a survey to gather some **sensitive information** from a population of n individuals.
- For example we are interested to know if the proportion of individuals that have committed a traffic light violation within a time frame. We denote the information of each individual i, $X_i \in \{0, 1\}$.
- The analyst then gathers each individual response Y_i which may or may not correspond to the true value X_i:

$$Y_i = \begin{cases} X_i & \text{with probability } 1/2 + \delta \\ \\ 1 - X_i & \text{with probability } 1/2 - \delta \end{cases}$$

- The analyst is interested in getting an estimate $\tilde{p} = \frac{1}{n}\sum_{i=1}^{n} Y_i$ of $p = \frac{1}{n}\sum_{i=1}^{n} X_i$.

The Non-private and the Perfectly Private frameworks

$$Y_i = \begin{cases} X_i & \text{with probability } 1/2 + \delta \\ 1 - X_i & \text{with probability } 1/2 - \delta \end{cases}$$

- In the **non-private** framework $\delta = 1/2$ and p can be **retrieved precisely** with $\tilde{p}$ as $\tilde{p} = \frac{1}{n}\sum_{i=1}^{n} X_i = p$. However, the analyst has **complete access** to the private information X. There is no privacy.

- In the **perfectly private** framework, $\delta = 0$ and the distribution of Y_i is a uniform and is **completely independent** of X_i. We have complete privacy, but on the flip side, $\tilde{p}$ is **completely uninformative** about p.

The Randomized Response Algorithm

$$Y_i = \begin{cases} X_i & \text{with probability } 1/2 + \delta \\ 1 - X_i & \text{with probability } 1/2 - \delta \end{cases}$$

- Let's consider an intermediate strategy, and put $\delta = 1/4$. By observing $Y_i = X_i$ with probability $3/4$, meaning that the individual's true information could still be $1 - Y_i$ with probability $1/4$, it is said that the individual has been granted **"plausible deniability"**.
- How "deniable" the response is corresponds to **the level of privacy** they have been afforded, i.e., **the smaller the δ**, the greater the deniability, or in other words, **the greater the privacy**.
- Moreover, defining $\tilde{p} = \frac{1}{n}\sum_{i=1}^{n} \frac{1}{2\delta}\left(Y_i + \delta - 1/2\right)$, it follows that $|\tilde{p} - p| = O_p(\frac{1}{\delta\sqrt{n}})$.

Central Differential Privacy

- The previous example falls under the more general framework of Differential Privacy (DP), a broad formalization of the notion of "plausible deniability".
- We consider in the following a version of DP called central DP.
- It involves a trusted curator who receives raw data from individuals, runs a randomized algorithm and publicly outputs its results.
- DP is a property of this algorithm stating that no individual data has a significant impact on the output of the algorithm.

ε-Differential Privacy

ε-Differential Privacy

Consider any two neighbouring datasets $x, y \in \chi^n$ that differ in exactly one entry, which we denote $x \sim y$.
A randomized algorithm $\mathcal{M} : \chi^n \to \mathcal{Y}$ is ε-differentially private if for all neighbouring datasets x, y, and any event $S \subseteq \mathcal{Y}$

$$\Pr\left(\mathcal{M}(x) \in S\right) \leq e^{\varepsilon} \Pr\left(\mathcal{M}(y) \in S\right)$$

where ε is called the ***privacy budget*** *and the randomness is over the choices of $\mathcal{M}$.*

When ε is small, $e^{\varepsilon} \approx 1 + \varepsilon$ and the inequality states that the amount of information one would learn about a given individual would not change (much) whether or not they are in the dataset.

Post-processing Property

Let $\mathcal{M} : \chi^n \to \mathcal{Y}$ be ε-differentially private, and let $f : \mathcal{Y} \to \mathcal{Y}'$ be an arbitrary randomized mapping. Then $f \circ \mathcal{M}$ is ε-differentially private.

The Laplace Mechanism I

ℓ_1-sensitivity

Let $f : \chi^n \to \mathbb{R}^k$. The ℓ_1-sensitivity of f is

$$\Delta f = \max_{X \sim X'} ||f(X) - f(X')||_1$$

The Laplace Distribution

The Laplace distribution with location and scale parameters 0 *and* b*, respectively, has the following density:*

$$p(x) = \frac{1}{2b} \exp\left(-\frac{|x|}{b}\right).$$

We note $Y \sim Lap(b)$*, with* $\mathbb{E}[Y] = 0$ *and* $\mathrm{var}(Y) = 2b^2$*. Furthermore, the following holds:*

$$\Pr\left(|Y| \geq tb\right) = \exp(-t).$$

The Laplace Mechanism II

The Laplace Mechanism

Let $f : \chi^n \to \mathbb{R}^k$. The Laplace mechanism is defined as

$$\mathcal{M}(X) = f(X) + (Y_1, \ldots, Y_k),$$

where $Y_i \overset{i.i.d}{\sim} Lap(\frac{\Delta f}{\varepsilon})$.

- Let us apply it in our initial problem where $f(X) = \frac{1}{n}\sum_{i=1}^{n} X_i$ and $X_i \in \{0, 1\}$.
- We have $\Delta f = 1/n$, therefore $\tilde{p} = f(X) + Y$ where $Y \sim Lap(\frac{1}{\varepsilon n})$.
- Recalling that $p = f(X)$, $\mathbb{E}[\tilde{p}] = p$ and $\text{var}(\tilde{p}) = \text{var}(Y) = \frac{2}{\varepsilon^2 n^2}$, we get:

  $$|\tilde{p} - p| = O_p\left(\frac{1}{\varepsilon n}\right),$$

  which is quadratically smaller than the accuracy of the Randomized Response algorithm.
- Notice that while $p \in [0, 1]$, $\tilde{p} \in \mathbb{R}$.

Accuracy and ε-DP of the Laplace Mechanism

The Laplace Mechanism is ε-DP

Let X and X' be two neighbouring datasets, and $p_X(z)$ and $p_{X'}(z)$ the pdfs of $\mathcal{M}(X)$ and $\mathcal{M}(X')$ evaluated at an arbitrary point $z \in \mathbb{R}^k$. Then

$$\frac{p_X(z)}{p_{X'}(z)} \leq e^{\varepsilon}.$$

Accuracy of the Laplace Mechanism

Let $f : \chi^n \to \mathbb{R}^k$, and $\mathcal{M}(X) = f(X) + (Y_1, \dots, Y_k)$, where $Y_i \overset{i.i.d}{\sim} Lap(\frac{\Delta f}{\varepsilon})$. Then $\forall \beta \in (0, 1]$

$$\Pr\left(||f(X)_i - \mathcal{M}(X)_i||_\infty \geq \ln\left(\frac{k}{\beta}\right)\left(\frac{\Delta f}{\varepsilon}\right)\right) = \beta$$

which yields

$$|f(X) - \mathcal{M}(x)| = O_p\left(\frac{\Delta f}{\varepsilon}\right).$$

Inference in the DP Framework

- Given a suitable (sufficient) statistic that we privatize, we have three ways to make inference:
  1. Determine the **exact distribution** of the privatized statistic under the null.
  2. Use a **suitable approximation** of its distribution under the null.
  3. Use the **non-private distribution** under the null.
- To determine an exact distribution is not always trivial: the convolutions of the added noise and the sampling error is usually intractable.
- To use an approximate distribution or the non-private one is more straightforward but is often far from accurate in finite samples.

A Simple Proportion

Framework

Let $U_{0,j} \overset{i.i.d}{\sim} Unif(0,1)$ *be a given seed with with* $j = 1, \dots, n$.

From the vector of seeds $\boldsymbol{U}_0 = (U_{0,1}, \dots, U_{0,n})$, *a sample is generated as follows:*

$$Z(\theta_0, \boldsymbol{U}_0, n) = (\underbrace{\mathbb{1}_{\{U_{0,1} < \theta_0\}}}_{Z_{0,1}}, \dots, \underbrace{\mathbb{1}_{\{U_{0,n} < \theta_0\}}}_{Z_{0,n}}), \quad Z_{0,j} \overset{i.i.d}{\sim} Ber(\theta_0).$$

Let

$$X_0 := X(\theta_0, \boldsymbol{U}_0, n) = \sum_{j=1}^{n} Z_{0,j},$$

from which we compute the sample proportion

$$\hat{\theta}_0 := \hat{\theta}\left(\theta_0, \boldsymbol{U}_0, n\right) = \frac{X(\theta_0, \boldsymbol{U}_0, n)}{n}.$$

Private Inference approaches

Let us consider the following three approaches:

DP-UMP Test (*Awan and Slavković, 2018*)

Consider $Y_0 = X_0 + N_0$ *where* $N_0 \sim Tulap(0, e^{-\varepsilon}, 0)$, *then* $Y_0|X_0 \sim Tulap(X_0, e^{-\varepsilon}, 0)$.

- *Releasing* $Y_0|X_0$ *satisfies* ε*-DP.*
- *A UMP size-*α *test can be derived and p-values can be computed easily.*

DP Normal Approximation Test (*Vu and Slavković, 2009*)

Consider $\hat{\theta}_0^{(p)} = \hat{\theta}_0 + L_0$ *where* $L_0 \sim Lap\left(\frac{1}{\varepsilon n}\right)$. *The distribution of* $\hat{\theta}_0$ *can be approximated with* $N\left(\theta_0, \frac{\sigma^2}{n}\right)$ *and if* L_0 *is approximated with a* $N\left(0, \frac{2}{\varepsilon^2 n^2}\right)$, *then under the null*

$$\hat{\theta}_0^{(p)} \sim N\left(\theta_0, \frac{\sigma^2}{n} + \frac{2}{\varepsilon^2 n^2}\right).$$

Non-private JIMI

Non-private JIMI Estimator

*The **J**ust **I**dentified **M**inimal d**I**stance estimator (JIMI) in the non-private framework is defined as follows:*

Let $U_{i,j} \overset{i.i.d}{\sim} Unif(0,1)$ *be a given seed with* $i = 1, \ldots, B$ *and* $j = 1, \ldots, n$.

$$\tilde{\theta}_i \in \underset{\theta}{\operatorname{argzero}} \, \hat{\theta}\left(\theta_0, \boldsymbol{U}_0, \varepsilon, n\right) - \hat{\theta}\left(\theta, \boldsymbol{U}_i, \varepsilon, n\right),$$

which implies in the case of a proportion

$$\tilde{\theta}_i \in \left\{\theta : \frac{\sum_{j=1}^{n} \mathbb{1}\{U_{i,j} < \theta\}}{n} = \hat{\theta}\left(\theta_0, \mathbf{U}_0, n\right)\right\} \Leftrightarrow \tilde{\theta}_i \in \left\{\theta : \hat{\mathsf{F}}_n\left(\theta; \mathbf{U}_i\right) = \hat{\theta}\left(\theta_0, \mathbf{U}_0, n\right)\right\}.$$

This in turn implies that, when non-empty, the solution is an interval

$$\tilde{\theta}_i \in \hat{\mathsf{F}}_n^{-}\left(\hat{\theta}_0; \mathbf{U}_i\right) = \left[u_{i,[n\hat{\theta}_0]},\ u_{i,[n\hat{\theta}_0+1]}\right),$$

for $\hat{\theta}_0 \in (0,1)$, otherwise take $[0, u_{i,[1]}]$ if $\hat{\theta}_0 = 0$ and $[u_{i,[n]}, 1]$ if $\hat{\theta}_0 = 1$, where $q_{[x]}$ corresponds to the order statistic of rank x.

Solution Types

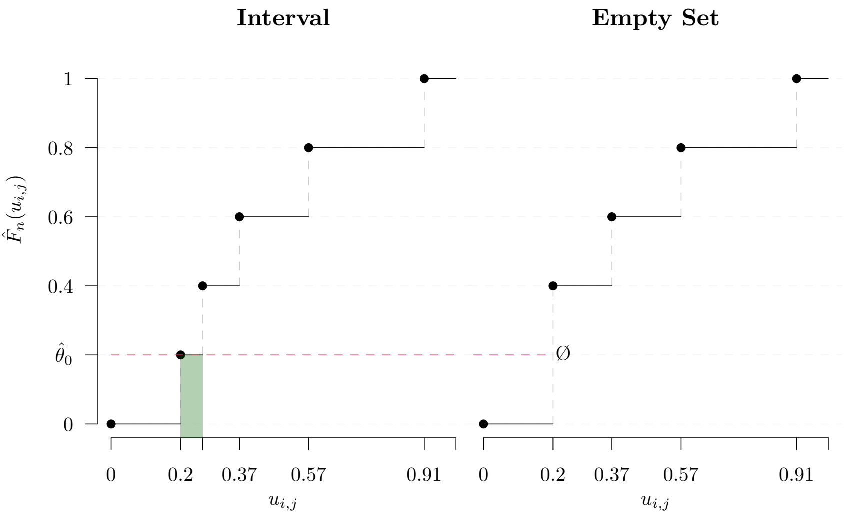

Selection Criterion

Selection Criterion

In order to have a unique solution, we need to impose a selection rule on $\tilde{\theta}_i$. Let D be a selection criterion:

$$\tilde{\theta}_i = Q\left(\hat{\theta}_0, D, \mathbf{U}_i\right) = U_{i,[n\hat{\theta}_0]} + D\left(U_{i,[n\hat{\theta}_0+1]} - U_{i,[n\hat{\theta}_0]}\right). \tag{1}$$

It turns out that the exact distribution of $\tilde{\theta}_i$ can sometimes be derived depending on the choice of D as demonstrated by Hannig (2009). Particularly,

Distribution of D	Distribution of $\tilde{\theta}_i$
Beta $(1/2, 1/2)$	Beta $\left(n\hat{\theta}_0 + 1/2, n(1-\hat{\theta}_0) + 1/2\right)$
Beta $(1, 1)$	Beta $\left(n\hat{\theta}_0 + 1, n(1-\hat{\theta}_0) + 1\right)$

Therefore, in stead of solving numerically, simply use the exact distribution when available.

Differentially Private JIMI

DP Proportion

Consider a simple DP proportion

$$\hat{\pi}\left(\theta_0, \boldsymbol{U}_0, W_0, \varepsilon, n\right) := \frac{X\left(\theta_0, \boldsymbol{U}_0, n\right)}{n} + Y\left(W_0, \varepsilon, n\right),$$

where

$$Y\left(W_0, \varepsilon, n\right) := -\frac{1}{\varepsilon n} sign\left(W_0\right) \log\left(1 - |W_0|\right)$$

with $W_i \sim Unif(-1/2, 1/2)$. *Thus*

$$Y\left(W_0, \varepsilon, n\right) \sim Laplace\left(0, (\varepsilon n)^{-1}\right)$$

Differentially Private JIMI

DP JIMI

We define the DP JIMI estimator as

$$\tilde{\theta}_i \in \underset{\theta}{\text{argzero}}\, \hat{\pi}\left(\theta_0, \boldsymbol{U}_0, W_0, \varepsilon, n\right) - \hat{\pi}\left(\theta, \boldsymbol{U}_i, W_i, \varepsilon, n\right),$$

with $i \in 1, \ldots, B$, which implies

$$\tilde{\theta}_i \in \left\{\theta : X\left(\theta, \boldsymbol{U}_i, n\right) = \left[\hat{\pi}\left(\theta_0, \boldsymbol{U}_0, W_0, \varepsilon, n\right) - Y\left(W_i, \varepsilon, n\right)\right] n\right\}$$

However, since $X\left(\theta, \mathbf{U}_i, n\right) \in \{0, 1, \ldots, n\}$, the solution interval is rarely non-empty. Therefore, let us in stead solve

$$\tilde{\theta}_i^\star \in \underset{\theta}{\text{argmin}}\, \hat{\pi}\left(\theta_0, \mathbf{U}_0, W_0, \varepsilon, n\right) - \hat{\pi}\left(\theta, \mathbf{U}_i, W_i, \varepsilon, n\right),$$

and we can write

$$\tilde{\theta}_i^\star = Q^\star\left(\hat{\theta}_i^\star, D, \mathbf{U}_i\right) = U_{i,\lfloor n\hat{\theta}_i^\star \rceil} + D\left(U_{i,\lfloor n\hat{\theta}_i^\star + 1\rceil} - U_{i,\lfloor n\hat{\theta}_i^\star \rceil}\right), \tag{2}$$

where $\hat{\theta}_i^\star = \hat{\pi}\left(\theta_0, \mathbf{U}_0, W_0, \varepsilon, n\right) - Y\left(W_i, \varepsilon, n\right)$, and $q_{\lfloor x \rceil}$ returns the order statistic whose rank is the closest integer to x.

DP JIMI algorithm

Using the same rule as before, we get

$$\tilde{\theta}_i^\star = Q^\star\left(\hat{\theta}_i^\star, D, \mathbf{U}_i\right) \xRightarrow{D\sim\text{Beta}(1/2,1/2)} \tilde{\theta}_i^\star \sim \text{Beta}\left(n\hat{\theta}_i^\star + 1/2, n(1-\hat{\theta}_i^\star) + 1/2\right). \quad (3)$$

Alogrithm:

INPUT: $n \in \mathbb{N}$, $\varepsilon > 0$, $W_0 \sim \text{Unif}(-1/2, 1/2)$, $\mathbf{U}_0 \sim \text{Unif}(0,1)$, $\theta_0 \in (0,1)$, $\hat{\pi}\left(\theta_0, \mathbf{U}_0, W_0, \varepsilon, n\right) \in \mathbb{R}$, $i = 1, \ldots, B$

1. Generate $\hat{\theta}_i^\star := \hat{\pi}_0 - Y_i$
2. if $\hat{\theta}_i^\star > 1$ return $\tilde{\theta}_i^\star = 1$,
   if $\hat{\theta}_i^\star < 0$ return $\tilde{\theta}_i^\star = 0$,
   otherwise, generate $\tilde{\theta}_i^\star$ from $Beta\left(n\hat{\theta}_i^\star + 1/2, n(1-\hat{\theta}_i^\star) + 1/2\right)$

OUTPUT: A differentially private JIMI estimate $\tilde{\theta}_i^\star \in [0,1]$

One-sample Proportion Tests: Power and Size

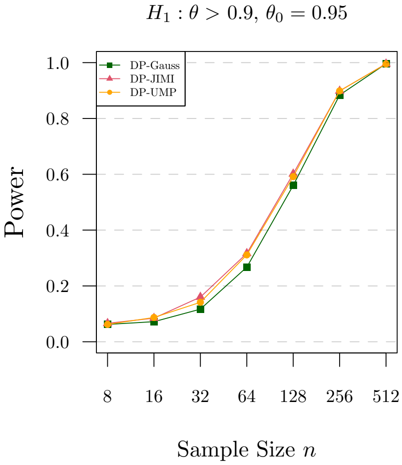

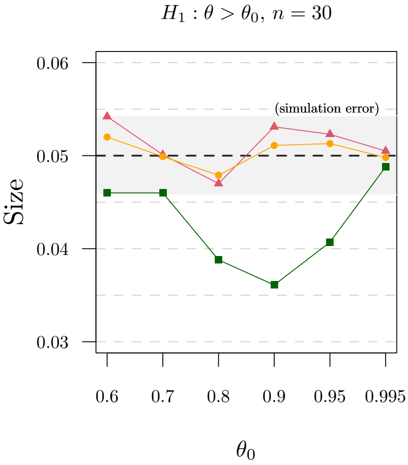

First Order Properties

- Given $n \geq 1$, $M_{i,j}, R_{i,j} \sim \Gamma\left(\frac{1}{n}, \frac{1}{\varepsilon n}\right)$, then $\sum_{j=1}^{n}\left(M_{i,j} - R_{i,j}\right) \sim Lap\left(\frac{1}{\varepsilon n}\right)$.
- Hence, the DP proportion can be written as

$$\hat{\pi}_0 = \frac{1}{n}\sum_{j=1}^{n}\left(Z_{0,j} + n(M_j - R_j)\right).$$

- Let $V_{0,j} := Z_{0,j} + n(M_{0,j} - R_{0,j})$, then the MGF of $V_{0,j}$ is

$$M_{V_{0,j}}(t) = \frac{1 - \theta_0 + \theta_0 e^t}{1 - \varepsilon^{-2}t^2},\ |t| < \varepsilon,$$

  and as $M_{V_{0,j}}(t)$ is well defined in $[-\varepsilon, \varepsilon]$ with $\varepsilon > 0$, all the moments are finite.
- Then by the Berry-Esseen theorem, it follows that

$$\sqrt{n}\,\sigma(\theta)^{-1}\left(\hat{\pi}_n(\theta) - \theta\right) \to \boldsymbol{Z},\ \text{uniformly in } \theta \in (0,1) \text{ (Assumption 2)},$$

  where $\sigma^2(\theta) = \theta(1-\theta) + 2/\varepsilon^2$. This in turn implies that

$$\sqrt{n}\,\sigma(\hat{\pi}_n)^{-1}\left(\tilde{\theta}_n^{\star} - \hat{\pi}_n\right) \to \boldsymbol{Z},\ \text{in probability (Assumption 1)},$$

  and ensures that the DP-JIMI confidence intervals are asymptotically consistent (see Lemma 23.3, in Van der Vaart (2000))

Conclusions, Limitations and Possible Extensions...

- Recall that $|f(X) - \mathcal{M}(x)| = O_p\left(\frac{\Delta f}{\varepsilon}\right)$, what if Δf is not bounded?
  - Before adding the noise, the curator will truncate/clip the data to $[a, b] \subset \mathbb{R}$ before computing $f(X)$.
  - In this case the non-private initial estimator $f(X)$ will **no longer be consistent**, and inference becomes problematic for many methods.
  - This is much less of an issue for the JIMI.

- Our method remains simple to apply (no calculations required), provided we know how to generate the data. What if the covariates are privatized?
  - We find ourselves in an **error of measurement** setting, which remains at the moment a limitation.

Selected References

Jordan Awan and Aleksandra Slavković. Differentially private uniformly most powerful tests for binomial data. *Advances in Neural Information Processing Systems*, 31, 2018.

Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4): 211–407, 2014.

Jan Hannig. On generalized fiducial inference. *Statistica Sinica*, pages 491–544, 2009.

Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.

Duy Vu and Aleksandra Slavkovic. Differential privacy for clinical trial data: Preliminary evaluations. In *2009 IEEE International Conference on Data Mining Workshops*, pages 138–143. IEEE, 2009.