CIFellows 2020-2021

Computing Innovation Fellows

Wanrong Zhang

Harvard University

Private Sequential Hypothesis Testing for Statisticians: Privacy, Error Rates, and Sample Size [AISTATS'22]

- We develop a private version of Wald's SPRT, which we call PrivSPRT for private sequential hypothesis testing
- We analyze the privacy of PrivSPRT through Renyi differential privacy
- We give bounds on the expected sample size of PrivSPRT and the Type I and Type II error
- We perform experiments to empirically validate our theoretical findings

Motivation: Clinical Trials

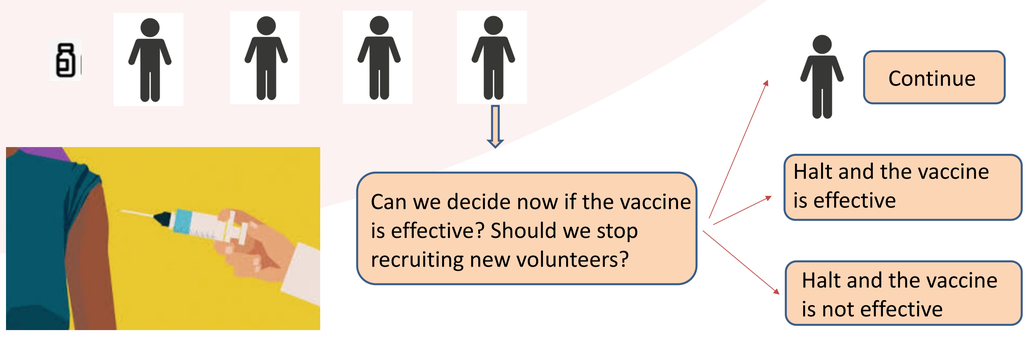

Problem Setup

Hypothesis Testing: testing a null hypothesis H_0 against an alternative hypothesis H_1

- Fixed-sample-size hypothesis testing. [Private methods [GLRV16, GR19, CKS+19,CKM+19]
- Sequential testing: Streaming data. It enables a decision to be reached earlier than with a fixed sample size test. [Private method: our work]

Model:

A sequence of data points $x_1, x_2, \ldots \sim f$

At each time, an analyst tests $H_0: f = f_0$ $H_1: f = f_1$, and makes one of the follows three decisions:

(1) Halt and accept H_0

(2) Halt and reject H_0

(3) Continue collecting observations

Criteria :

(1)Two operating characteristic functions: Type I error and Type II error (2) Two average sample number functions: $E_0[T]$ and $E_1[T]$ (3)Privacy: Renyi Differential Privacy [Mir' 17]

An algorithm $M: T^n \to R$ is $(\boldsymbol{\alpha}, \boldsymbol{\epsilon})$**-RDP** with order $\alpha \geq 1$, if $\forall$ neighboring $x, x' \in T^n$ it holds that

$$D_\alpha(M(x)||M(x')) \leq \epsilon$$

PrivSPRT Algorithm

Our approach: Run two parallel instantiations of GenAboveThresh with the truncated cumulative log-likelihood ratio as the queries.

- Input: streaming data $x_1, x_2, \ldots$, hypotheses f_0 and f_1, threshold $-a, b$, noise-adding mechanism M_1, M_2
- Let $-\hat{a} \sim M_1(-a)$, $\hat{b} \sim M_1(b)$
- Compute the cumulative log-likelihood ratio statistic for each t:

$$\ell_t(A) = \sum_{i=1}^{t} [\log \frac{f_1(x_i)}{f_0(x_i)}]_{-A}^{A}$$

- Let $\widehat{\ell_t}(A) \sim M_2(\ell_t(A))$
- If $\widehat{\ell_t}(A) \notin (-\hat{a}, \hat{b})$, then halt and reject H_0 if $\widehat{\ell_t}(A) \geq \hat{b}$, or accept H_0 if $\widehat{\ell_t}(A) \leq -\hat{a}$

Main Theorems

Privacy:

PrivSPRT satisfies $(\alpha, \frac{\alpha-(\gamma-\frac{1}{\gamma})}{\alpha-1}\frac{2\alpha A^2}{\sigma_1^2} + \frac{4\alpha A^2}{\sigma_2^2} + \frac{\log \max\{\frac{-a}{\mu_0}, \frac{b}{\mu_1}\}^{\gamma}}{\alpha-1})$-RDP.

Composition from the two Gaussian mechanisms.

Randomness of the stopping time

Sample Size:

Let $\mu_0 = -E_0[\log \frac{f_1(x)}{f_0(x)}]_{-A}^{A}$ and $\mu_1 = E_1[\log \frac{f_1(x)}{f_0(x)}]_{-A}^{A}$

The expected sample size of PrivSPRT satisfies

$$E_1[T] \leq 1 + O\left(\frac{b}{\mu_1}\right) + O\left(\frac{\sqrt{(\sigma_1^2 + \sigma_2^2)}}{\mu_1}\right),$$

Similar result holds for $E_0[T]$.

Nonprivate sample size result $O(\frac{b}{D_{KL}})$,

Error Rate:

Type I error satisfies

$$Pr_0(d = 1) \leq O(\exp(-b)) + O(\frac{\sqrt{(\sigma_1^2 + \sigma_2^2)}}{A^2})$$

Similar result holds for Type II error.

Nonprivate sample size result $O(\exp(-b))$

CRA Computing Research Association

CCC Computing Community Consortium Catalyst

NSF