DP-SPRT: Differentially Private Sequential Hypothesis Testing

Thomas Michel, Debabrota Basu, Emilie Kaufmann

Scool Team, Inria Center of the University of Lille

Sequential Hypothesis Testing

Setup: We observe samples $X_1, X_2, \dots$ sequentially from distribution ν_θ

Goal: Test $H_0 : \theta = \theta_0$ vs $H_1 : \theta = \theta_1$ with as few samples as possible

Constraints: False positive probability $\leq \alpha$, false negative probability $\leq \beta$

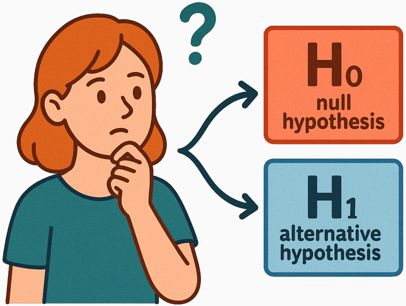

Sequential Hypothesis Testing

Setup: We observe samples $X_1, X_2, \ldots$ sequentially from distribution ν_θ
Goal: Test $H_0 : \theta = \theta_0$ vs $H_1 : \theta = \theta_1$ with as few samples as possible
Constraints: False positive probability $\leq \alpha$, false negative probability $\leq \beta$

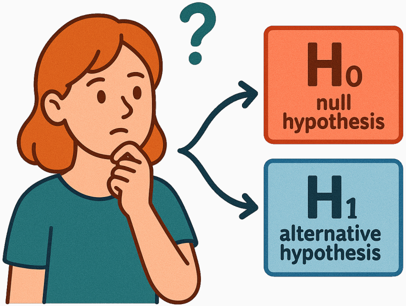

When do we need to test ?

- Clinical Trials
- A/B Testing
- Fraud Detection

Setup: We observe samples $X_1, X_2, \ldots$ sequentially from distribution ν_θ
Goal: Test $H_0 : \theta = \theta_0$ vs $H_1 : \theta = \theta_1$ with as few samples as possible
Constraints: False positive probability $\leq \alpha$, false negative probability $\leq \beta$

Sequential Probability Ratio Test (SPRT) (Wald, 1945):

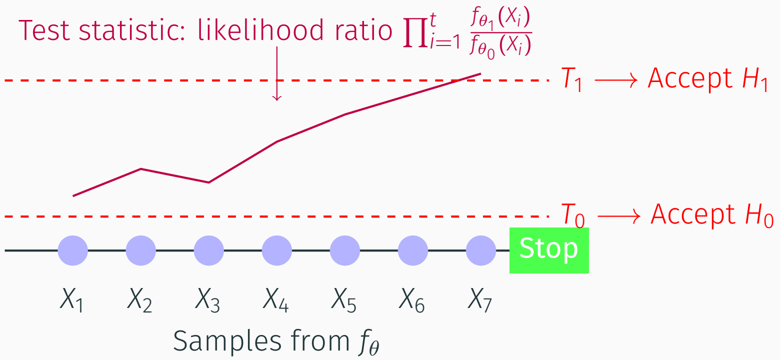

Sequential Probability Ratio Test

Setup: We observe samples $X_1, X_2, \ldots$ sequentially from distribution ν_θ
Goal: Test $H_0 : \theta = \theta_0$ vs $H_1 : \theta = \theta_1$ with as few samples as possible
Constraints: False positive probability $\leq \alpha$, false negative probability $\leq \beta$

SPRT for one dimensionnal exponential families: Stop at $\tau = \min(\tau_0, \tau_1)$ with decision $\hat{d} = i$ if $\tau = \tau_i$ where

$$\tau_0 = \inf\left\{n : \bar{X}_n \leq \mu_0 + \frac{\mathrm{KL}(\nu_{\theta_0}, \nu_{\theta_1}) - \log(1/\beta)/n}{\theta_1 - \theta_0}\right\}$$
$$\tau_1 = \inf\left\{n : \bar{X}_n \geq \mu_1 - \frac{\mathrm{KL}(\nu_{\theta_1}, \nu_{\theta_0}) - \log(1/\alpha)/n}{\theta_1 - \theta_0}\right\}.$$

$\mathrm{KL}(\nu_{\theta_0}, \nu_{\theta_1})$ is the KL divergence between the two distributions and $\bar{X}_n$ is the empirical mean of the samples up to time n.

The Privacy Problem: A Medical Trial

Scenario: Testing if new drug works better than placebo

H_0: Drug success rate = 30% (Same as placebo) vs H_1: Drug success rate = 70%

Each patient outcome: $X_i \in \{0, 1\}$ (failure/success)

The Privacy Problem: A Medical Trial

Scenario: Testing if new drug works better than placebo

H_0: Drug success rate = 30% (Same as placebo) vs H_1: Drug success rate = 70%

Each patient outcome: $X_i \in \{0, 1\}$ (failure/success)

Question: What was patient 7's outcome? **Success (1) or Failure (0)?**

The Privacy Problem: A Medical Trial

Scenario: Testing if new drug works better than placebo

H_0: Drug success rate = 30% (Same as placebo) vs H_1: Drug success rate = 70%

Each patient outcome: $X_i \in \{0, 1\}$ (failure/success)

Question: What was patient 7's outcome? **Success (1) or Failure (0)?**

The stopping decision reveals patient 7 had a SUCCESS!

Privacy violation: Patient 7's medical outcome is leaked by our decision to stop

Privacy in Sequential Decisions

Clinical Trial

- Testing new drug vs placebo
- Stopping pattern reveals:
  - Treatment effectiveness
  - Patient responses

A/B Testing

- Users see different versions
- When we conclude reveals:
  - User behavior
  - Conversion rates

Fraud Detection

- Monitor transactions
- Alert timing reveals:
  - Transaction patterns
  - Detection methods

Core Problem: When we stop reveals what we observed

Neighboring Datasets: Two datasets D, D' are neighboring if they differ by exactly one record.

Differential Privacy (Dwork, Roth, et al., 2014): A randomized mechanism $\mathcal{M}$ is DP if for any neigboring datasets D and D' and for any events S:

$$\varepsilon\text{-DP:} \quad \log\left(\frac{\mathbb{P}[\mathcal{M}(D) \in S]}{\mathbb{P}[\mathcal{M}(D') \in S]}\right) \leq \varepsilon$$

$$(\alpha, \varepsilon)\text{-Rényi DP:} \quad D_\alpha(\mathcal{M}(D) \| \mathcal{M}(D')) \leq \varepsilon$$

where D_α is the Rényi divergence of order $\alpha > 1$.

Stopping Time Distributions

Deterministic Algorithm

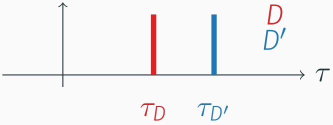

Private Algorithm

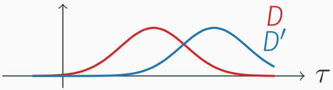

Privacy adds randomness to mask the stopping decision pattern

Our Method: DP-SPRT

DP-SPRT Algorithm (blue = privacy additions to SPRT):

Input: Hypotheses θ_0, θ_1, error probabilities α, β, noise distributions $\mathcal{D}_Z, \mathcal{D}_Y$, correction function $C(n, x)$, error allocation γ

1. Sample threshold noise $Z \sim \mathcal{D}_Z$
2. **For** $n = 1, 2, 3, \ldots$ **do**
3. Sample query noise $Y_n \sim \mathcal{D}_Y$
4. Compute noisy average $\bar{X}_n = \frac{1}{n}\sum_{i=1}^{n} X_i + \frac{Y_n}{n}$
5. Compute noisy threshold $\hat{T}_0^n = \mu_0 + \frac{\mathrm{KL}(\nu_{\theta_0}, \nu_{\theta_1}) - \log(1/(\gamma\beta))/n}{\theta_1 - \theta_0} - C(n, (1-\gamma)\beta) - \frac{Z}{n}$
6. Compute noisy threshold $\hat{T}_1^n = \mu_1 - \frac{\mathrm{KL}(\nu_{\theta_1}, \nu_{\theta_0}) - \log(1/(\gamma\alpha))/n}{\theta_1 - \theta_0} + C(n, (1-\gamma)\alpha) + \frac{Z}{n}$
7. **If** noisy average $\bar{X}_n$ is below noisy threshold $\hat{T}_0^n$ **then** Halt and accept $\mathcal{H}_0$
8. **Else if** noisy average $\bar{X}_n$ is above noisy threshold $\hat{T}_1^n$ **then** Halt and accept $\mathcal{H}_1$

Y_n and Z are used for both conditions unlike when composing `AboveThreshold`

DP-SPRT: Privacy

Table: Comparison of DP-SPRT instantiations

	Laplace Noise	Gaussian Noise
Noise distributions	$Y_n \sim \text{Lap}(4/\varepsilon)$ $Z \sim \text{Lap}(2/\varepsilon)$	$Y_n \sim \mathcal{N}(0, \sigma_Y^2)$ $Z \sim \mathcal{N}(0, \sigma_Z^2)$
Privacy guarantee	ε-Differential Privacy	$(\alpha, \varepsilon(\alpha))$-RDP
Correction function $C(n, x)$	$\frac{6 \log(n^s \zeta(s)/x)}{n\varepsilon}$	$\frac{\sqrt{2(\sigma_Y^2+\sigma_Z^2) \log(n^s \zeta(s)/2)}}{n}$

Correctness: DP-SPRT satisfies $\mathbb{P}_{\theta_0}(\hat{d} = 1) \leq \alpha$ and $\mathbb{P}_{\theta_1}(\hat{d} = 0) \leq \beta$ when the distribution $\mathcal{D}_Y$ is symmetric and the correction function $C(n, x)$ satisfies:

$$\sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > C(n, x)\right) \leq x$$

Near-Optimal Sample Complexity (DP-SPRT with Laplace noise)

Lower Bound (any ε-DP test):

$$\mathbb{E}[\tau] \geq \frac{\log(1/\beta)}{\min(\mathrm{KL}(\nu_{\theta_0}, \nu_{\theta_1}), \varepsilon \cdot \mathrm{TV}(\nu_{\theta_0}, \nu_{\theta_1}))}$$

where TV is the total variation distance between two distributions.

Sample Complexity Upper Bound (Laplace Noise):

$$\mathbb{E}[\tau] \lesssim \max\left(\frac{\log(1/\beta)}{\mathrm{KL}(\nu_{\theta_0}, \nu_{\theta_1})}, \frac{(\theta_1 - \theta_0)\log(1/\beta)}{\varepsilon \cdot \mathrm{KL}(\nu_{\theta_0}, \nu_{\theta_1})}\right)$$

For Bernoulli distributions, we have

$$\frac{\mathrm{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0} \underset{\theta_1 \to \theta_0}{\longrightarrow} \mathrm{TV}(\nu_{\theta_0}, \nu_{\theta_1})$$

DP-SPRT with Laplace noise is near-optimal

Setup: Bernoulli($p_0 = 0.3$) vs Bernoulli($p_1 = 0.7$), $\alpha = \beta = 0.05$, 1000 trials

Sample Complexity vs Privacy Level

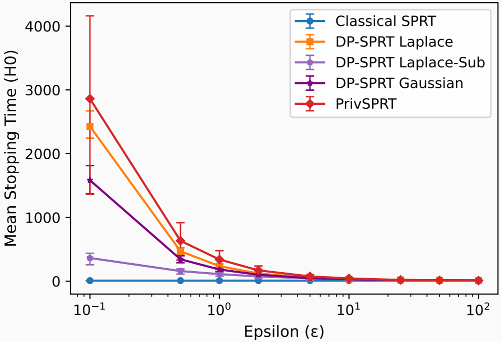

On average, **DP-SPRT variants outperform** PrivSPRT (Zhang, Mei, and Cummings, 2022) across privacy levels

Error Control Comparison

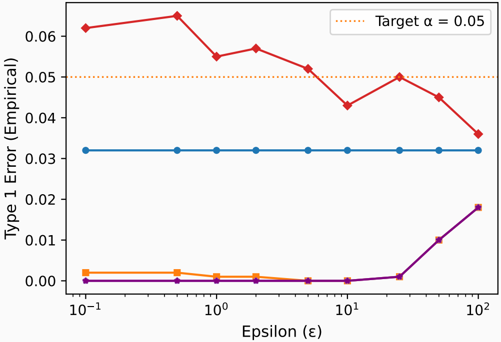

All DP-SPRT variants guarantee error control, while PrivSPRT can violate error targets due to empirical tuning

Conclusion

Privacy:

- Real-world applications **NEED** privacy (regulations, competition, ethics)
- Sequential decisions leak sensitive information

Conclusion

Privacy:

- Real-world applications NEED privacy (regulations, competition, ethics)
- Sequential decisions leak sensitive information

Our Contributions:

1. **Theoretically calibrated** private sequential test with guaranteed error control
2. **Near-optimal sample complexity** matching lower bounds up to a constant in some regimes
3. **Practical implementation** with no empirical tuning required, low variance in stopping times, and subsampling amplification in high-privacy regimes

Conclusion

Privacy:

- Real-world applications **NEED** privacy (regulations, competition, ethics)
- Sequential decisions leak sensitive information

Our Contributions:

1. **Theoretically calibrated** private sequential test with **guaranteed error control**
2. **Near-optimal sample complexity** matching lower bounds up to a constant in some regimes
3. **Practical implementation** with **no empirical tuning** required, **low variance** in stopping times, and **subsampling amplification** in high-privacy regimes

Thank you! Questions?

References

Dwork, Cynthia, Aaron Roth, et al. (2014). **"The algorithmic foundations of differential privacy".** In: *Foundations and Trends® in Theoretical Computer Science* 9.3–4, pp. 211–407.

Wald, A. (1945). **"Sequential Tests of Statistical Hypotheses".** In: *Annals of Mathematical Statistics* 16(2), pp. 117–186.

Zhang, Wanrong, Yajun Mei, and Rachel Cummings (2022). **"Private Sequential Hypothesis Testing for Statisticians: Privacy, Error Rates, and Sample Size".** In: *AISTATS*.

Distribution of Stopping Time

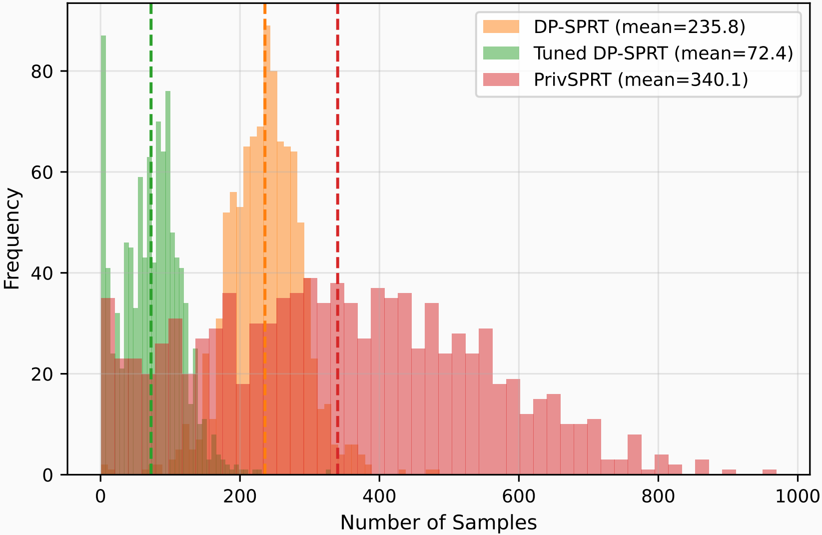

Theorem (Privacy)
Let $\mathcal{A}$ be the DP-SPRT algorithm and τ the (random) stopping time.

(i) ε***-DP:*** *If noise mechanisms $\mathcal{D}_Z$ and $\mathcal{D}_Y$ satisfy ε_Z-DP (sensitivity 1) and ε_Y-DP (sensitivity 2), then DP-SPRT satisfies $(\varepsilon_Z + \varepsilon_Y)$-DP.*

(ii) (α, ε)***-RDP:*** *If mechanisms have RDP profiles $\varepsilon_Z(\alpha)$ and $\varepsilon_Y(\alpha)$, then:*

$$\mathbb{D}_\alpha(\mathcal{A}(D) \| \mathcal{A}(D')) \leq \frac{\alpha - 1/2}{\alpha - 1} \varepsilon_Z(2\alpha) + \varepsilon_Y(\alpha) + \frac{\log\left(2\mathbb{E}_{Z \sim \mathcal{D}_Z}[\mathbb{E}_{(\tau, \hat{d}) \sim \mathcal{A}(D')}[\tau | Z]^2]\right)}{2(\alpha - 1)}$$

Correctness

Theorem (Correctness)
For any error allocation $\gamma \in (0,1)$, DP-SPRT is (α, β)-correct if:

$$\forall \delta \in (0,1), \quad \sum_{n=1}^{\infty} \mathbb{P}\left(\frac{Y_n}{n} - \frac{Z}{n} > C(n,\delta)\right) \leq \delta$$

Proof idea: Union bound decomposes error into SPRT error + privacy noise error

Theorem (Sample Complexity Upper Bound)
Assume correction condition holds. For error allocation $\gamma \in (0,1)$ and $i \in \{0,1\}$:

$$\mathbb{E}_{\theta_0}[\tau] \leq 1 + (1-\gamma)\beta + \frac{1}{1 - \exp\left(-\frac{(TV(\nu_{\theta_0}, \nu_{\theta_1}))^4}{2(\theta_1 - \theta_0)^2}\right)} + N(\theta_0, \theta_1, \beta, \gamma)$$

$$\mathbb{E}_{\theta_1}[\tau] \leq 1 + (1-\gamma)\alpha + \frac{1}{1 - \exp\left(-\frac{(TV(\nu_{\theta_0}, \nu_{\theta_1}))^4}{2(\theta_1 - \theta_0)^2}\right)} + N(\theta_1, \theta_0, \alpha, \gamma)$$

where $N(\theta, \theta', \delta, \gamma) = \inf\left\{n : \frac{\log(1/(\delta\gamma))/n}{|\theta - \theta'|} + 2C(n, (1-\gamma)\delta) \leq \frac{1}{2}\frac{KL(\nu_\theta, \nu_{\theta'})}{|\theta - \theta'|}\right\}$

Theorem (Lower Bound)
Any ε-DP test with $\mathbb{P}_{\theta_0}(\hat{d} = 1) \leq \alpha$, $\mathbb{P}_{\theta_1}(\hat{d} = 0) \leq \beta$ satisfies:

$$\mathbb{E}_{\theta_0}[\tau] \geq \frac{kl(\alpha, 1-\beta)}{\min\left(KL(\nu_{\theta_0}, \nu_{\theta_1}), \varepsilon \cdot TV(\nu_{\theta_0}, \nu_{\theta_1})\right)}$$

$$\mathbb{E}_{\theta_1}[\tau] \geq \frac{kl(\beta, 1-\alpha)}{\min\left(KL(\nu_{\theta_1}, \nu_{\theta_0}), \varepsilon \cdot TV(\nu_{\theta_0}, \nu_{\theta_1})\right)}$$

where $kl(x, y) = x \log(x/y) + (1-x)\log((1-x)/(1-y))$ is binary relative entropy.

Laplace Implementation & Near-Optimality

Corollary (Laplace DP-SPRT Sample Complexity)
For DP-SPRT with Laplace noise ($Y_n \sim Lap(4/\varepsilon)$, $Z \sim Lap(2/\varepsilon)$):

$$N(\theta_0, \theta_1, \beta, \gamma) \leq \frac{2\log(1/(\gamma\beta))}{KL(\nu_{\theta_0}, \nu_{\theta_1})} + \frac{24(\theta_1 - \theta_0)\log(\zeta(s)/(1-\gamma)\beta)}{\varepsilon KL(\nu_{\theta_0}, \nu_{\theta_1})} + o_{\beta\to 0}(\log(1/\beta))$$

Two Regimes:

- Statistics-dominated: ε large $\Rightarrow$ complexity $\approx \frac{\log(1/\beta)}{\mathrm{KL}}$ (like non-private SPRT)
- Privacy-dominated: ε small $\Rightarrow$ complexity $\approx \frac{(\theta_1-\theta_0)\log(1/\beta)}{\varepsilon\cdot\mathrm{KL}}$