





Contribution. Wald's SPRT decides between two hypotheses with minimal expected sample size, but its data-dependent stopping time τ leaks information about the most recent observations. We propose **DP-SPRT**, a wrapper that simultaneously meets target Type-I/II error probabilities (α, β) and an ε -DP budget on the released $(\tau, \hat{d})$. We give generic Laplace and Gaussian (Rényi-DP) calibrations with closed-form correction functions, prove **sample-complexity upper bounds matching the DP lower bound up to constants** in the small-gap regime, and obtain privacy amplification by subsampling without empirical tuning.

Sequential Hypothesis Testing

The Problem. We observe samples $X_1, X_2, \dots$ 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 $\leq \beta$.



Optimal


SPRT [1] minimizes the expected sample size among all correct tests.

Not Private

The stopping time reveals information about the recent observations.

Example: clinical trial. Patients arrive sequentially.

H_0 : drug success rate = 30% vs. H_1 : drug success rate = 70%



Privacy leak. The stopping decision reveals information about the last patient's outcome.

Differential Privacy

A randomized algorithm $\mathcal{A}$ is DP [2] if for any observation sequences $\underline{X}$ and $\underline{X}'$ differing by a single record and any event S :

$$\log \left(\frac{\mathbb{P}[\mathcal{A}(\underline{X}) \in S]}{\mathbb{P}[\mathcal{A}(\underline{X}') \in S]} \right) \leq \varepsilon$$

$$D_\alpha(\mathcal{A}(\underline{X}) \parallel \mathcal{A}(\underline{X}')) \leq \varepsilon$$

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

Stopping-time distributions



Privacy adds randomness to mask stopping patterns: neighbouring datasets must produce near-identical distributions of τ .

The DP-SPRT Algorithm

Algorithm 1 DP-SPRT

Require: Hypotheses θ_0, θ_1 ; error probabilities α, β ; noise distributions $\mathcal{D}_Z, \mathcal{D}_Y$; correction function $C(n, x)$; error allocation $\gamma \in (0, 1)$.

- 1: **Sample threshold noise** $Z \sim \mathcal{D}_Z$.
- 2: **for** $n \leftarrow 1, 2, 3, \dots$ **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 $\hat{T}_0^n = \mu_0 + \frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1}) - \log(1/(\gamma\beta))/n}{\theta_1 - \theta_0} - C(n, (1 - \gamma)\beta) - \frac{Z}{n}$.
- 6: Compute noisy threshold $\hat{T}_1^n = \mu_1 - \frac{\text{KL}(\nu_{\theta_1}, \nu_{\theta_0}) - \log(1/(\gamma\alpha))/n}{\theta_1 - \theta_0} + C(n, (1 - \gamma)\alpha) + \frac{Z}{n}$.
- 7: **if** noisy average $\bar{X}_n$ is **below** noisy threshold $\hat{T}_0^n$ **then**
- 8: **Accept** H_0 and **stop**.
- 9: **else if** noisy average $\bar{X}_n$ is **above** noisy threshold $\hat{T}_1^n$ **then**
- 10: **Accept** H_1 and **stop**.
- 11: **end if**
- 12: **end for**

Blue marks the terms added for privacy; without them DP-SPRT reduces to the classical SPRT.

Noise instantiations.

- **Laplace** (ε -DP): $Y_n \sim \text{Lap}(4/\varepsilon)$, $Z \sim \text{Lap}(2/\varepsilon)$, $C(n, x) = \frac{6 \log(n^2 \zeta(s)/x)}{n\varepsilon}$.
- **Gaussian** (RDP): $Y_n \sim \mathcal{N}(0, \sigma_Y^2)$, $Z \sim \mathcal{N}(0, \sigma_Z^2)$, $C(n, x) = \frac{\sqrt{2(\sigma_Y^2 + \sigma_Z^2)} \log(n^2 \zeta(s)/2x)}{n}$.

Privacy amplification by subsampling. Include each observation independently with probability $r \in (0, 1]$. The effective privacy guarantee improves by a factor $1/r$, significantly reducing the noise required in high-privacy regimes without empirical tuning.

Theoretical Guarantees

Privacy. DP-SPRT satisfies ε -differential privacy (Laplace) or $(\alpha, \varepsilon(\alpha))$ -RDP (Gaussian), with

$$\varepsilon(\alpha) = \frac{\alpha-1/2}{\alpha-1} \cdot \frac{\alpha}{\sigma_Z^2} + \frac{\alpha}{2\sigma_Y^2} + \frac{\log(2\mathbb{E}_{\mathcal{D}_Z}[\mathbb{E}_{\mathcal{A}(\mathcal{D}')}[\tau | Z]^2])}{2(\alpha-1)}.$$

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 correction function $C(n, x)$ satisfies

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

Lower bound (any ε -DP test):

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

Sample-complexity upper bound (Laplace noise):


$$\limsup_{\beta \rightarrow 0} \frac{\mathbb{E}_{\theta_0}[\tau]}{\log(1/\beta)} \leq 48 \max\left(\frac{1}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}, \frac{\theta_1 - \theta_0}{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})\varepsilon}\right).$$

For Bernoulli: $\frac{\text{KL}(\nu_{\theta_0}, \nu_{\theta_1})}{\theta_1 - \theta_0} \xrightarrow{\theta_1 \rightarrow \theta_0} \text{TV}(\nu_{\theta_0}, \nu_{\theta_1})$.

Near-optimal: DP-SPRT matches the lower bound up to multiplicative constants when $\theta_1 \rightarrow \theta_0$.

Experimental Results

Setting. Bernoulli(0.3) vs Bernoulli(0.7), $\alpha = \beta = 0.05$, 1000 trials. Baseline: PRIVSPRT [3].




Mean stopping time vs. ε .

Type-I error vs. ε .

DP-SPRT variants outperform PrivSPRT across privacy levels with less variance, while **respecting the $\alpha = 0.05$ target** that PRIVSPRT violates due to empirical tuning.

Stopping Time Distribution (H_0 , $\varepsilon=1.0$, $\alpha=\beta=0.05$)



Stopping-time distribution under H_0 at $\varepsilon = 1$. DP-SPRT concentrates near the lower-bound budget; PRIVSPRT spreads.

- [1] A. Wald and J. Wolfowitz. *Optimum character of the sequential probability ratio test*. Annals of Mathematical Statistics, 1948.
- [2] C. Dwork and A. Roth. *The algorithmic foundations of differential privacy*. Foundations and Trends in Theoretical Computer Science, 2014.
- [3] W. Zhang, Y. Mei, and R. Cummings. *Private Sequential Hypothesis Testing for Statisticians: Privacy, Error Rates, and Sample Size*. AISTATS, 2022.