# Differentially Private Sequential Probability Ratio Tests

**Thomas Michel**, Debabrota Basu, Emilie Kaufmann

Inria, Univ. Lille, CNRS, Centrale Lille, CRIStAL

AISTATS 2026 Spotlight



# Sequential Hypothesis Testing & SPRT (Wald, 1945)

**Setup**: We observe samples $X_1, X_2, \ldots$ sequentially from distribution $\nu_\theta$
**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 $\leq \beta$

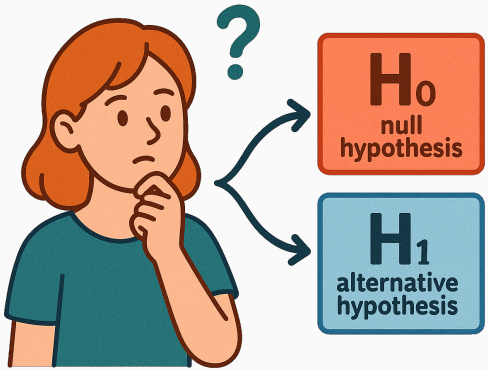

# Sequential Hypothesis Testing & SPRT (Wald, 1945)

**Setup**: We observe samples $X_1, X_2, \ldots$ sequentially from distribution $\nu_\theta$
**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 $\leq \beta$

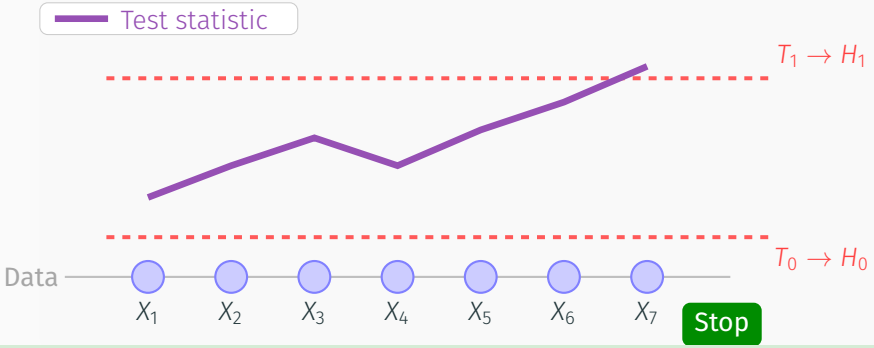


**SPRT.** Stop the moment the test statistic exits $[T_0, T_1]$.

**Optimal** (Wald, 1948): minimizes $\mathbb{E}[\tau]$ among all valid tests.

# ...but the stopping time leaks privacy

**Clinical trial:** $H_0$: drug success rate $= 30\%$ vs $H_1$: drug success rate $= 70\%$

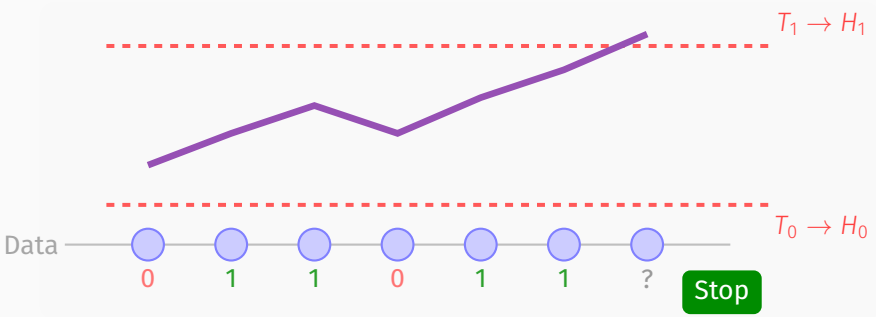

# ...but the stopping time leaks privacy

**Clinical trial:** $H_0$: drug success rate $= 30\%$ vs $H_1$: drug success rate $= 70\%$

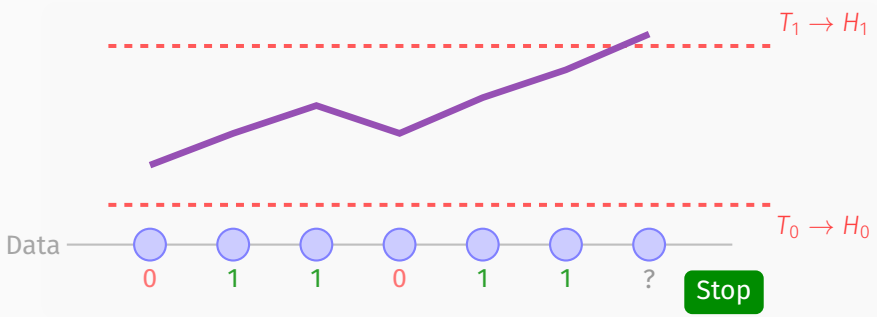


**Stopping $\Rightarrow$ patient 7 was a SUCCESS**

A failure would have moved the statistic away from $T_1$

The stopping time *itself* leaks the data.

# DP-SPRT: Noisy Statistic, Noisy Thresholds, Private Testing


Component: `OutsideInterval` — a general DP mechanism for *double-sided* private monitoring.

# DP-SPRT works and it is theoretically calibrated

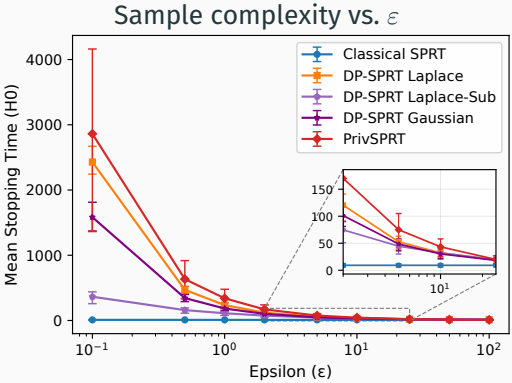


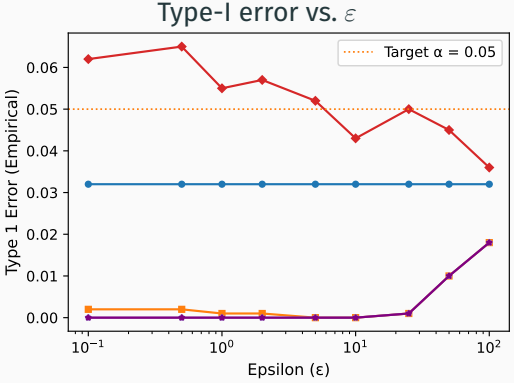


**Fewer samples** than baseline[1]

**Never violates** the error target: calibrated *theoretically*, no Monte-Carlo tuning.

**Matching lower bound** $\Rightarrow$ near-optimal for close hypotheses.



[1] W. Zhang, Y. Mei, R. Cummings. *Private Sequential Hypothesis Testing for Statisticians: Privacy, Error Rates, and Sample Size.* 2022.

# Come to the poster

On the poster you'll see:

- General `OutsideInterval` privacy mechanism
- Full `DP-SPRT` algorithm
  Laplace ($\varepsilon$-DP) and Gaussian (Rényi-DP) instantiations.
- Privacy, correctness, and sample complexity
- Privacy amplification by subsampling

Check out the full paper

See you at the poster!