Optimal Best Arm Identification under Differential Privacy

Marc Jourdan[⋆†] and Achraf Azize[⋆‡]

[†] EPFL, Lausanne, Switzerland [‡] FairPlay Joint Team, CREST, ENSAE Paris

BAI with Differential Privacy

Setting: Clinical trials with K candidate medicines

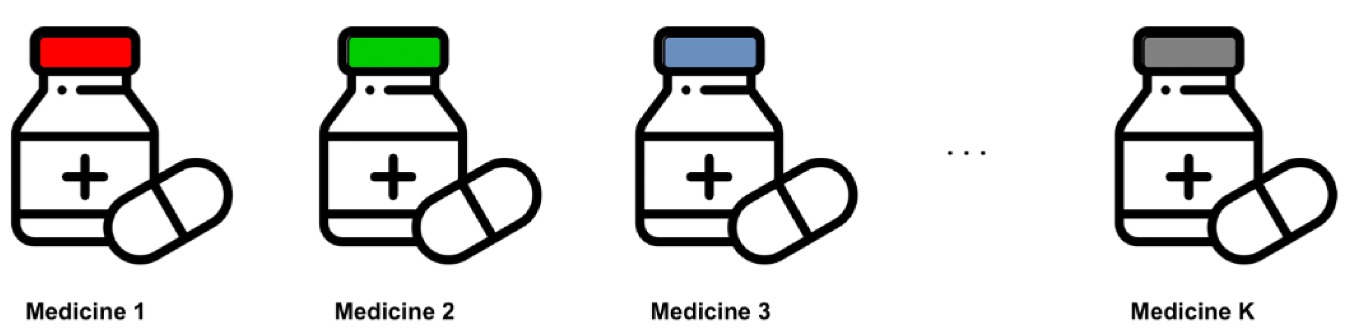

Goal: Find the medicine with the highest mean $a^\star \triangleq \arg\max_{a\in[K]} \mu_a$.

Constraint: Protect the privacy of the patients at level $\epsilon > 0$.
A patient's reaction to a medicine can reveal sensitive information about their health conditions.

Interaction Protocol: For the n-th patient in the study:

1. The doctor π chooses a Medicine $a_n \in \{1, \dots, K\}$,
2. The doctor observes a binary response $X_n \sim \nu_{a_n} \triangleq$ Ber(μ_{a_n}) .

Stop the interaction at time $\tau_{\epsilon,\delta}$ and recommend a final answer $\widetilde{a} \in [K]$.

Correctness: Let $\delta \in (0,1)$. A BAI strategy π is **δ-correct** for a class $\mathcal{M}$, if for every instance $\boldsymbol{\nu} \in \mathcal{M}$, $\mathbb{P}_{\boldsymbol{\nu}\pi}(\tau_{\epsilon,\delta} < \infty, \widetilde{a} \neq a^\star(\boldsymbol{\nu})) \le \delta$.

Definition: π satisfies **ϵ-global DP**, if $\forall R \sim R', \forall (T+1, \widetilde{a}, (a_1, \dots, a_T))$,

$$\Pr[\pi(R) = (T+1, \widetilde{a}, (a_1,\dots,a_T))] \le e^\epsilon \Pr[\pi(R') = (T+1, \widetilde{a}, (a_1,\dots,a_T))] .$$

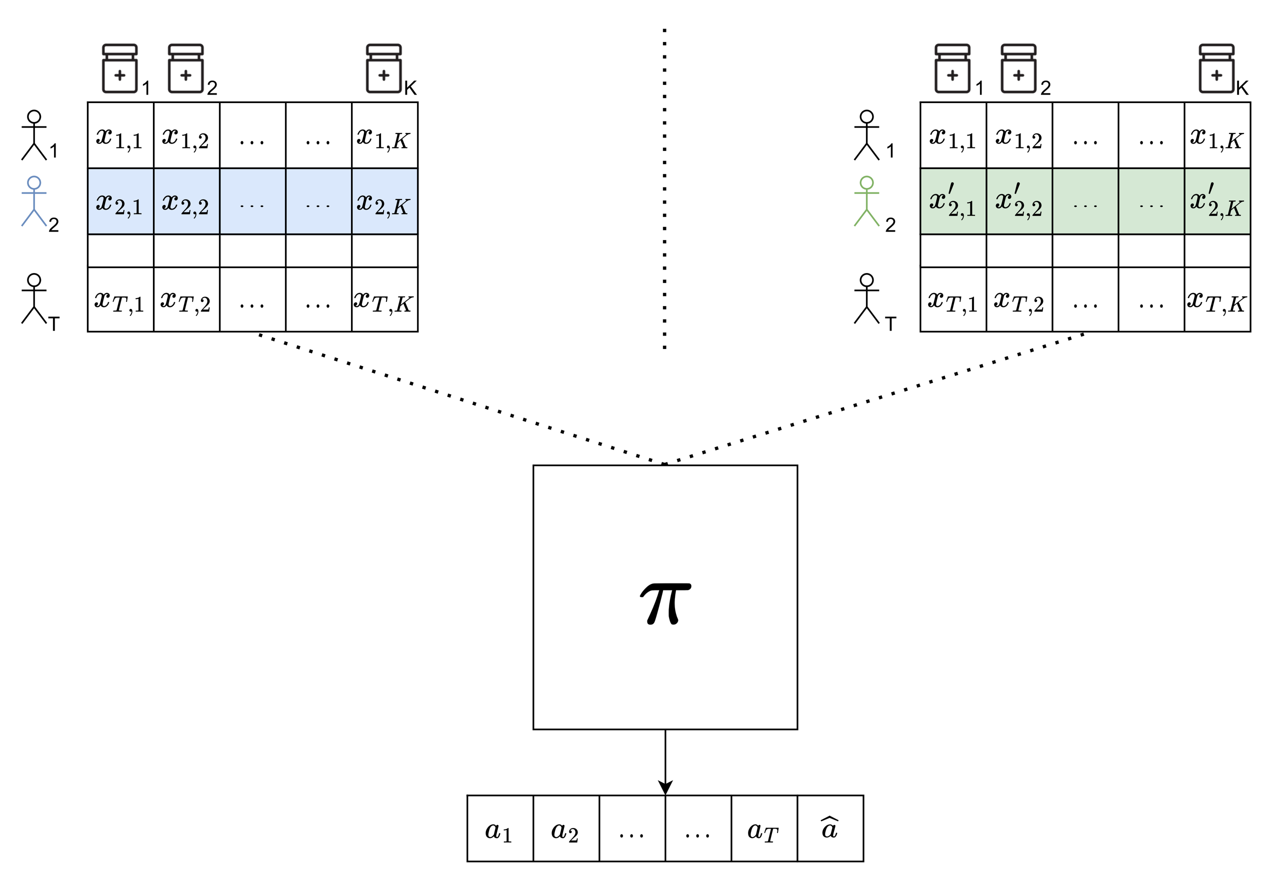

Contributions

1. Lower bound on $\mathbb{E}_{\boldsymbol{\nu}\pi}[\tau_{\epsilon,\delta}]$ under ε-global DP and δ-correctness: private characteristic time $T^\star_\epsilon(\boldsymbol{\nu})$ based on the signed divergences $\mathrm{d}^\pm_\epsilon$ interpolating between KL and TV.
2. Differentially Private Top Two algorithm: per-arm geometric batching with Laplace noise and private transportation costs.
3. Matching asymptotic upper bound $T^\star_\epsilon(\boldsymbol{\nu})$ for any privacy ϵ, when $\delta \to 0$, up to a small constant lower than 8.
4. Good empirical performance in both regimes of privacy.

Algorithm Design

Main Ingredients:

1. **Private mean estimator:** per-arm geometric grid and cumulative Laplacian noise (Lines 5 and 7). Do not use forgetting.
2. **GLR stopping rule and Top Two sampling rule:** private empirical transportation costs $C^\star_\epsilon$ based on $\mathrm{d}^\pm_\epsilon$ (Lines 10 and 13).

Algorithm 1 Differentially Private Top Two (DP-TT)

1: **Input:** setting parameters $(\epsilon,\delta) \in \mathbb{R}^\star_+ \times (0,1)$, hyperparameters $(\eta,\beta) \in \mathbb{R}^\star_+ \times (0,1)$, e.g., $(\eta,\beta) = (1, 1/2)$, and stopping threshold c.
2: **Output:** Stopping $\tau_{\epsilon,\delta}$, recommendation $\widetilde{a}_{\tau_{\epsilon,\delta}}$, actions $(a_n)_{n<\tau_{\epsilon,\delta}}$
3: **Initialization:** $\forall a \in [K]$, pull arm a, observe $X_a \sim \nu_a$ and draw $Y_{1,a} \sim \mathrm{Lap}(1/\epsilon)$. Set $n = K+1$. $\forall a \in [K]$, $\widetilde{S}_{n,a} = X_a + Y_{1,a}$, $k_{n,a} = 1$, $T_1(a) = n$, $N_{n,a} = \widetilde{N}_{n,a} = 1$, $\widetilde{\mu}_{n,a} = \widetilde{S}_{n,a}/\widetilde{N}_{n,a}$, $L_{n,a} = 0$ and $N^a_{n,a} = 0$.
4: **for** $n \ge K+1$ **do**
5:   **if** there exists $a \in [K]$ such that $N_{n,a} \ge (1+\eta)^{k_{n,a}}$ **then**
6:     Change phase $k_{n,a} \leftarrow k_{n,a}+1$; $(T_{k_{n,a}}(a), \widetilde{N}_{n,a}) = (n, N_{T_{k_{n,a}}(a),a})$
7:     Set $\widetilde{S}_{k_{n,a},a} = \sum_{t=T_{k_{n,a}-1}(a)}^{T_{k_{n,a}}(a)-1} X_t \mathbb{1}(a_t = a) + Y_{k_{n,a},a} + \widetilde{S}_{k_{n,a}-1,a}$
    with $Y_{k_{n,a},a} \sim \mathrm{Lap}(1/\epsilon)$, and update $\widetilde{\mu}_{n,a} = \widetilde{S}_{k_{n,a},a}/\widetilde{N}_{n,a}$
8:   **end if**
9:   Set $\widetilde{a}_n \in \arg\max_{a\in[K]} [\widetilde{\mu}_{n,a}]_0^1$
10:   **if** $C^\star_\epsilon(\widetilde{a}_n, a, \widetilde{\mu}_n, \widetilde{N}_n) > c(\widetilde{N}_{n,\widetilde{a}_n}, \epsilon, \delta) + c(\widetilde{N}_{n,a}, \epsilon, \delta)$ for all $a \neq \widetilde{a}_n$ **then**
11:     **return** $(n, \widetilde{a}_n, (a_t)_{t<n})$
12:   **end if**
13:   Set $B_n = \widetilde{a}_n$ and $C_n \in \arg\min_{a \neq B_n} \{C^\star_\epsilon(B_n, a, \widetilde{\mu}_n, N_n) + \log N_{n,a}\}$
14:   Set $a_n = B_n$ if $N^{B_n}_{n,B_n} \le \beta L_{n+1,B_n}$, and $a_n = C_n$ otherwise
15:   Pull a_n, observe and store $X_n \sim \nu_{a_n}$
16:   Update $(N_{n+1,a_n}, L_{n+1,B_n}, N^{B_n}_{n+1,B_n}) = (N_{n,a_n}, L_{n,B_n}, N^{B_n}_{n,B_n}) + (1, 1, \mathbb{1}(B_n = a_n))$
17: **end for**

Expected Sample Complexity Upper Bound

Privacy analysis: For observations in $[0,1]$, DP-TT is ϵ-global DP.

Correctness: DP-TT is δ-correct for thresholds that satisfy

$$c(n,\epsilon,\delta) \approx \log(1/\delta) + \log(n)\log(1 + \epsilon n/\log(n)) .$$

Lemma: Let $Z_t \sim$ Bin(t,μ) and $S_t = \sum_{s \in [\log_{1+\eta} t]} Y_s$ with $Y_s \sim$ Lap$(1/\epsilon)$.

$$\forall t \in \mathbb{N}, \forall x > 0, \quad \mathbb{P}(Z_t + S_t \ge t(\mu + x)) \lesssim \exp\left(-t\mathrm{d}^-_\epsilon(\mu + x, \mu)\right) ,$$
$$\mathbb{P}(Z_t + S_t \le t(\mu - x)) \lesssim \exp\left(-t\mathrm{d}^+_\epsilon(\mu - x, \mu)\right) .$$

Novel tail concentration for convolution of probability distributions.

Upper bound on expected sample complexity: DP-TT is ϵ-global DP, δ-correct and, for any instance $\boldsymbol{\nu}$ with distinct means $\mu \in (0,1)^K$,

$$\limsup_{\delta\to 0} \frac{\mathbb{E}_{\boldsymbol{\nu}\pi}[\tau_{\epsilon,\delta}]}{\log(1/\delta)} \le 2(1+\eta) T^\star_{\epsilon,\beta}(\boldsymbol{\nu}) \underset{(\eta,\beta)=(1,1/2)}{\le} 8 T^\star_\epsilon(\boldsymbol{\nu}) .$$

Expected Sample Complexity Lower Bound

The lower bound: For any δ-correct ϵ-global DP BAI strategy,

$$\mathbb{E}_{\boldsymbol{\nu}\pi}[\tau_{\epsilon,\delta}] \ge T^\star_\epsilon(\boldsymbol{\nu}) \log\left(\frac{1}{3\delta}\right) \quad \text{with} \quad T^\star_\epsilon(\boldsymbol{\nu})^{-1} = \max_{w \in \Sigma_K} \min_{a \neq a^\star} C^\star_\epsilon(a^\star, a, \boldsymbol{\nu}, w) ,$$

where the private transportation costs are defined as

$$C^\star_\epsilon(a^\star, a, \boldsymbol{\nu}, w) = \inf_{x\in[0,1]} \{w_{a^\star} \mathrm{d}^-_\epsilon(\mu_{a^\star}, x) + w_i \mathrm{d}^+_\epsilon(\mu_a, x)\} .$$

"Distinguishability" measure: Signed divergences $\mathrm{d}^\pm_\epsilon$ based on

$$\mathrm{d}_\epsilon(\nu,\kappa) = \inf_{\phi\in\mathcal{D}} \{\epsilon \mathrm{TV}(\nu,\phi) + \mathrm{KL}(\phi,\kappa)\} .$$

Bernoullis: Let $g^-_\epsilon(\mu) = \frac{\mu e^\epsilon}{\mu(e^\epsilon - 1) + 1}$. Then, $\mathrm{d}^-_\epsilon(\mu, x) = \mathrm{d}^+_\epsilon(1-\mu, 1-x)$, and

$$\mathrm{d}^+_\epsilon(\mu,x) = \begin{cases} 0 & \text{if } x \in [0,\mu] \\ \mathrm{kl}(\mu,x) & \text{if } x \in (\mu, g^-_\epsilon(\mu)] \\ -\log(1 - x(1-e^{-\epsilon})) - \epsilon\mu & \text{if } x \in (g^-_\epsilon(\mu), 1] \end{cases} ,$$

Consequences: Low-privacy regime where privacy is for "free":

$$\epsilon \ge \max_{a\neq a^\star} \log\left(\frac{\mu_{a^\star}(1-\mu_a)}{\mu_a(1-\mu_{a^\star})}\right) \quad \Longrightarrow \quad T^\star_\epsilon(\boldsymbol{\nu}) = T^\star(\boldsymbol{\nu}) .$$

Allocation-dependent condition: $C^\star_\epsilon(a^\star, a, \boldsymbol{\nu}, w) = C^\star(a^\star, a, \boldsymbol{\nu}, w)$.

Key Lemma: ϵ-DP mechanism $\mathcal{M}$ on data distributions $(\mathbb{P}, \mathbb{Q})$,

$$\mathrm{KL}(\mathbb{M}_{\mathbb{P},\mathcal{M}}, \mathbb{M}_{\mathbb{Q},\mathcal{M}}) \le \inf_{\mathbb{L}} \{\epsilon \inf_{\mathbb{C}_{\mathbb{P},\mathbb{L}}} \{\mathbb{E}_{D,D'\sim\mathbb{C}_{\mathbb{P},\mathbb{L}}} [\mathrm{d}_{\mathsf{Ham}}(D,D')]\} + \mathrm{KL}(\mathbb{L},\mathbb{Q})\} .$$

The optimal transport on product distributions is $\sum_i \mathrm{TV}(\mathbb{P}_i, \mathbb{L}_i)$.

Experimental Analysis

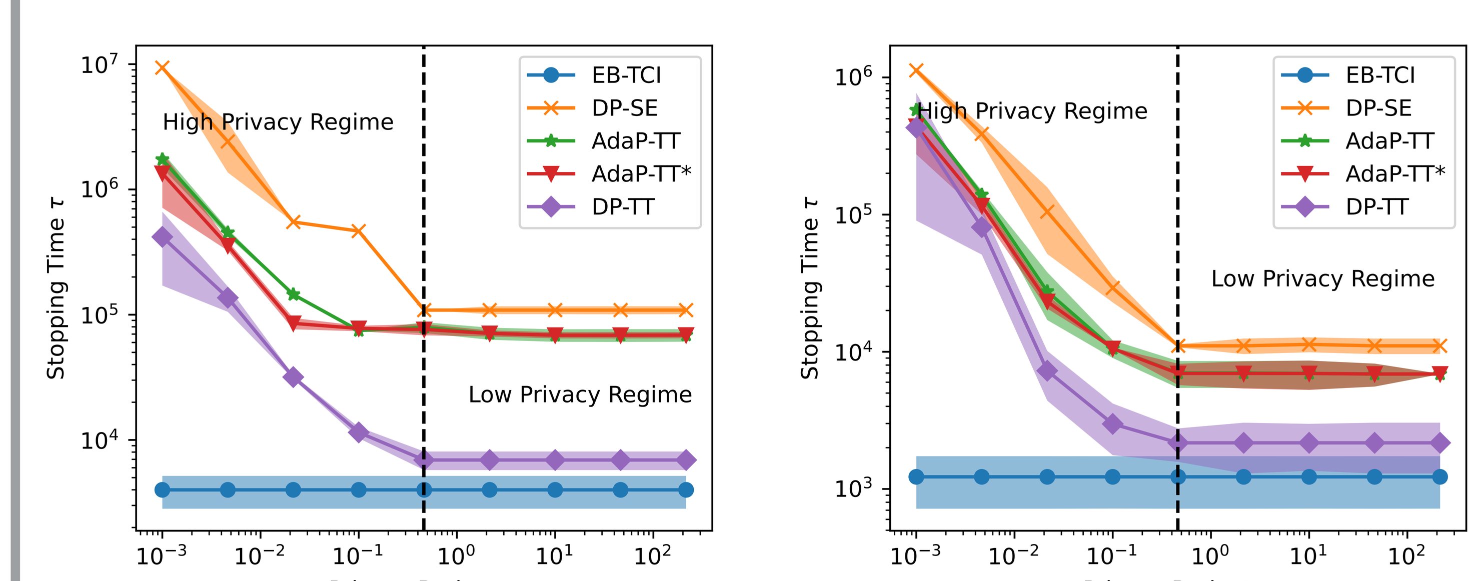

Figure 1: Empirical stopping time for $\delta = 0.01$ as function of ϵ on instances $\mu_1 = (0.95, 0.9, 0.9, 0.9, 0.5)$ and $\mu_2 = (0.75, 0.7, 0.7, 0.7, 0.7)$.

1. DP-TT outperforms DP-SE, AdaP-TT, AdaP-TT*.
2. The performance of DP-TT has two regimes: a high-privacy regime (for $\epsilon < 0.45$) and a low privacy regime (for $\epsilon > 0.45$).
3. DP-TT performs on par with EB-TCI, up to a multiplicative gap.