Online control of the false discovery rate with decaying memory

Aaditya Ramdas[⋆,†], Fanny Yang[†], Martin J. Wainwright[†,⋆], Michael I. Jordan[†,⋆]

Department of EECS[†], Department of Statistics[⋆], UC Berkeley

What we solve

Simple motivational problem:
Multiple A/B tests arrive **sequentially** over time, we make rejections using p-values

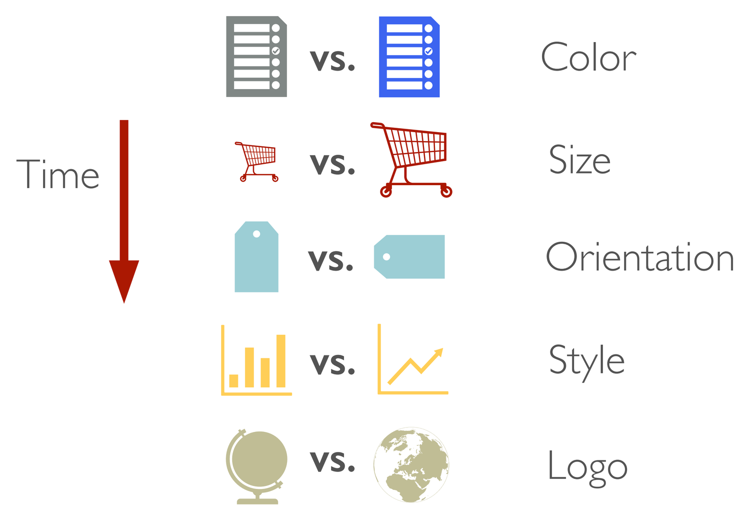

Some null hypotheses are true, others false:

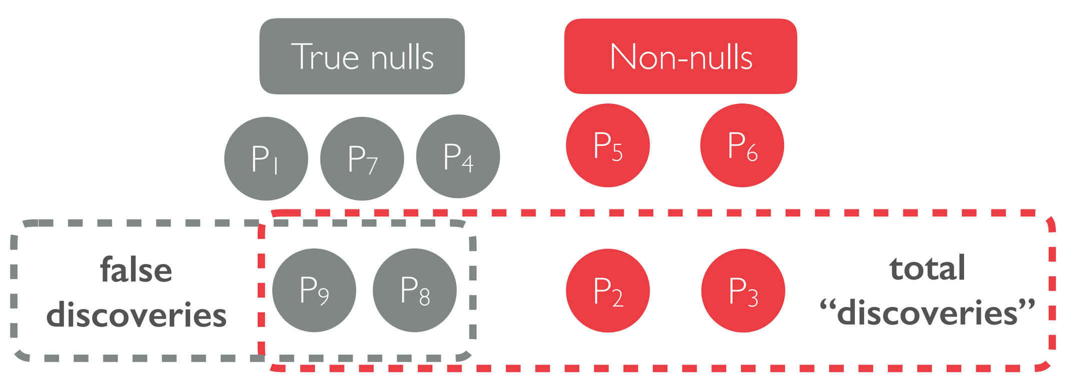

Our goal: Control **false discovery rate** (FDR) at every time T

$$\mathrm{FDR}(T) = \mathbb{E}\,\frac{\#\text{ false discoveries}}{\#\text{ total discoveries}} = \mathbb{E}\,\frac{V(T)}{R(T)}$$

Our framework, **more flexible** than existing ones, incorporates

- finite or decaying memory
- importance of each test
- prior knowledge about possible non-null locations,

is **more powerful** and guarantees **anytime FDR control**.

Generalized Alpha-investing(++)

Given α, and valid independent p-values P_t, online FDR procedure outputs significance levels α_t, and makes the decision $R_t := \mathbf{1}\{P_t \le \alpha_t\}$, where α_t is designed to be a monotone function of past rejections $R_{t-1}, \ldots, R_1$.

Generalized alpha-investing rules (GAI) update

- the wealth (with $W_0 \le \alpha$) $W(t) := W(t-1) - \phi_t + R_t\psi_t$
- the test penalty $\phi_t \le W(t-1)$
- the test reward $\psi_t \le \min\{\phi_t + b_t, \frac{\phi_t}{\alpha_t} + b_t - 1\}$ (1)
  with $b_t = B_0 = \alpha - W_0$.

Procedure	α_t	ϕ_t	ψ_t	Condition
[FS'08] AI	$\frac{\phi_t}{1+\phi_t}$	$\le W(t-1)$	$\phi_t + B_0$	—
[AR'14] ASR	$\kappa\phi_t$	$cW(t-1)$	satisfy (1)	$c \le 1$
[JM'17] LORD	ϕ_t	$\gamma_t W_0 + B_0 \sum_{j:\tau_j<t} \gamma_{t-\tau_j}$	B_0	$\sum_{i=1}^{\infty}\gamma_i = 1$

We propose an enhanced procedure (GAI++) with

$$b_t = \begin{cases} \alpha - W_0 & \text{when } R(t-1) = 0 \\ \alpha & \text{otherwise} \end{cases} \in \mathcal{F}^{t-1}.$$

Theorem: Any monotone GAI++ rule has $\mathrm{FDR}(T) \le \alpha$ at all times T, and is **more powerful** than corresponding GAI rule.

Simulations: more power with GAI++

With π_1 proportion of non-nulls

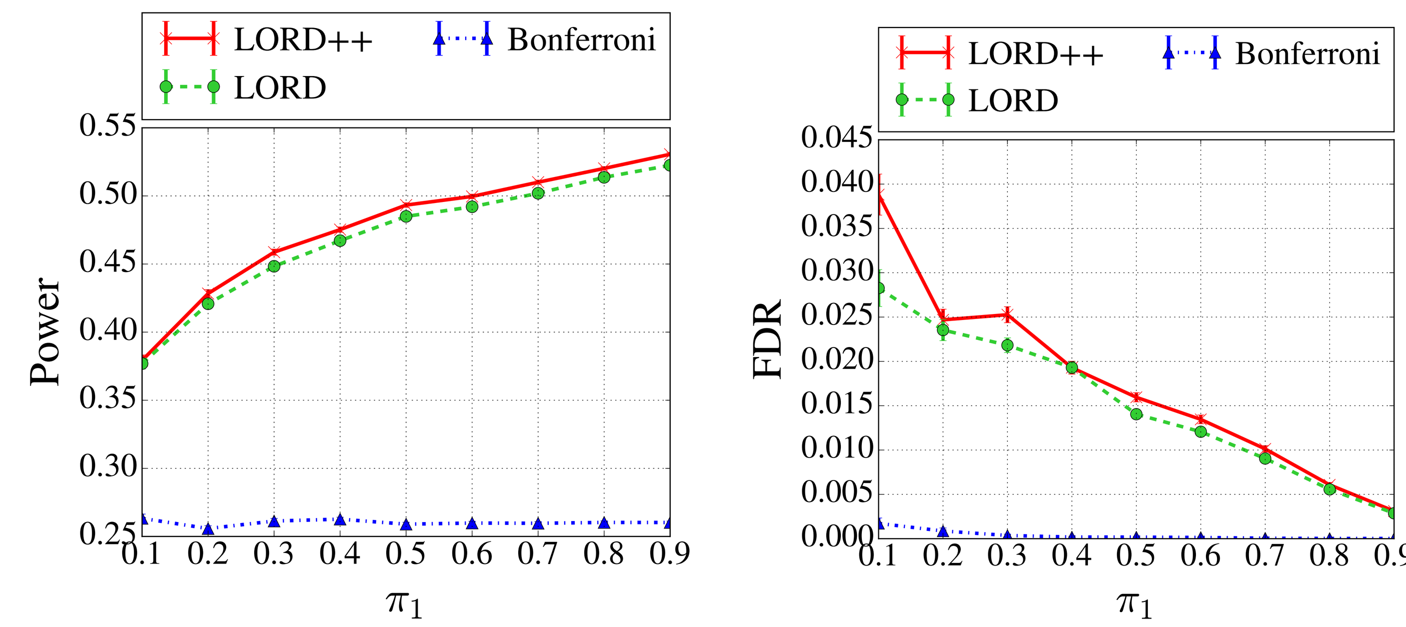

$\Longrightarrow$ GAI++ consistently more powerful!

Possible practical scenario

- Some tests are more important than others $\Longrightarrow$ weighted FDR and weighted procedure
- Know from previous experiments that for certain tests, H_0 is likely true $\Longrightarrow$ can make it easier to reject

Incorporating prior and penalty weights

[Benjamini and Hochberg '97], [Genovese et al. '06] introduced penalty and prior weights for the batch setting.
How about the online case?

We use **penalty weights** $u_t > 0$, **prior weights** $w_t > 0$ and reject according to

$$R_t := \mathbf{1}\{P_t \le \alpha_t w_t u_t\}.$$

Note: u_t, w_t may depend on $R_1, \ldots, R_{t-1}$

What if I'm interested only in the FDR in the recent past?
The global FDR at any time is controlled, how about locally?

Doubly-weighted decaying memory FDR

Adding penalty weights u_t and memory factor δ to FDR

$$\text{mem-FDR}_u(T) := \mathbb{E}\,\frac{V_u^\delta(T)}{R_u^\delta(T)}$$

with weighted and decaying discoveries

$$V_u^\delta(T) := \delta V_u^\delta(T-1) + u_T R_T \mathbf{1}\{T \in \mathcal{H}^0\}$$
$$R_u^\delta(T) := \delta R_u^\delta(T-1) + u_t R_t$$

Note: Setting $\delta = 1$, $u_t = 1$ reduces to $\mathrm{FDR}(T)$.
Let's define the (random) time of the k-th rejection as

$$\tau_k = \min_{s\in\mathbb{N}} \mathbf{1}\{\sum_{t=1}^{s} R_t = k\},$$

Doubly-weighted mem-FDR control using mem-GAI++

The doubly-weighted mem-GAI++ rules update

- the adjusted wealth
$$W(t) := \delta W(t-1) + (1-\delta)W_0\mathbf{1}\{\tau_1 > t-1\} - \phi_t + R_t\psi_t$$
$$= W_0\delta^{T-\min\{\tau_1,T\}} + \sum_{t=1}^{T}\delta^{T-t}(-\phi_t + R_t\psi_t).$$
- the adjusted penalty
$$\phi_t \le \delta W(t-1) + (1-\delta)W_0\mathbf{1}\{\tau_1 > t-1\},$$
- the adjusted reward
$$0 \le \psi_t \le \min\left\{\phi_t + u_t b_t, \frac{\phi_t}{w_t\alpha_t u_t} + u_t b_t - u_t\right\}, \quad \text{where}$$
$$b_t := \alpha - \frac{W_0}{u_t}\mathbf{1}\{\tau_1 > t-1\} \in \mathcal{F}_{t-1}.$$

Theorem: All monotone double-weighted mem-GAI++ rules have $\text{mem-FDR}_u(T) \le \alpha$ at all times T.

Simulations: with weights and decaying memory

LORD++ ($\delta = 0$) with prior weights $1 \pm a$ on non-nulls/nulls

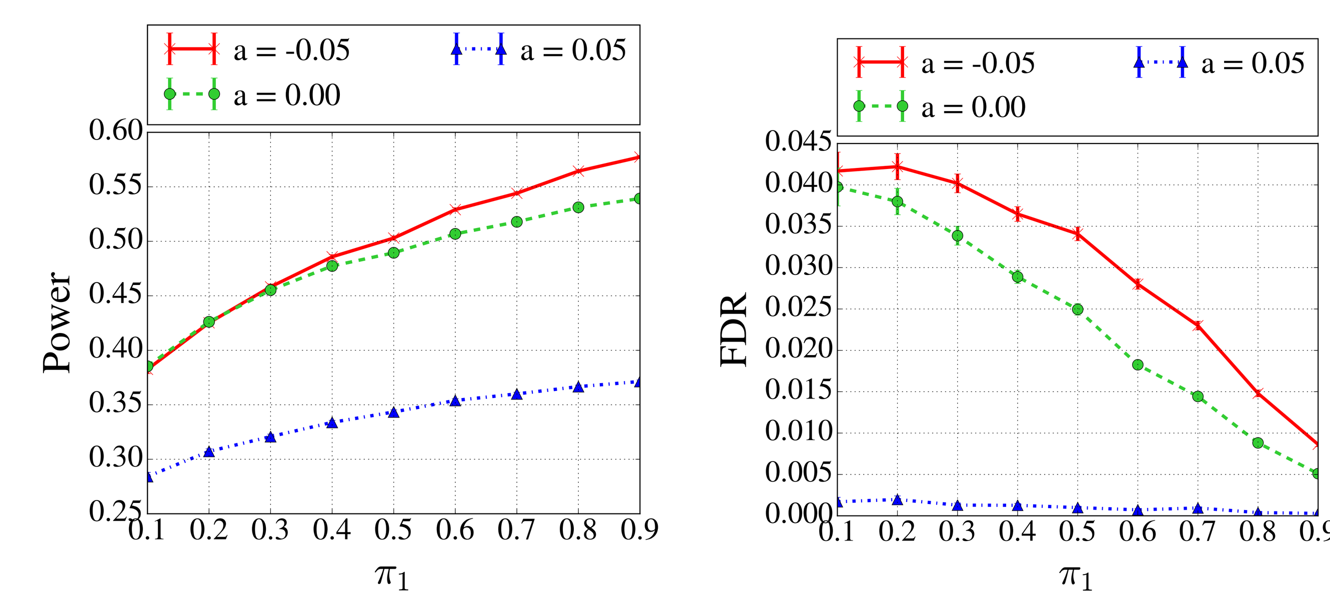

mem-LORD++ ($\delta = 0.99$) vs. LORD++ ($\delta = 0$)

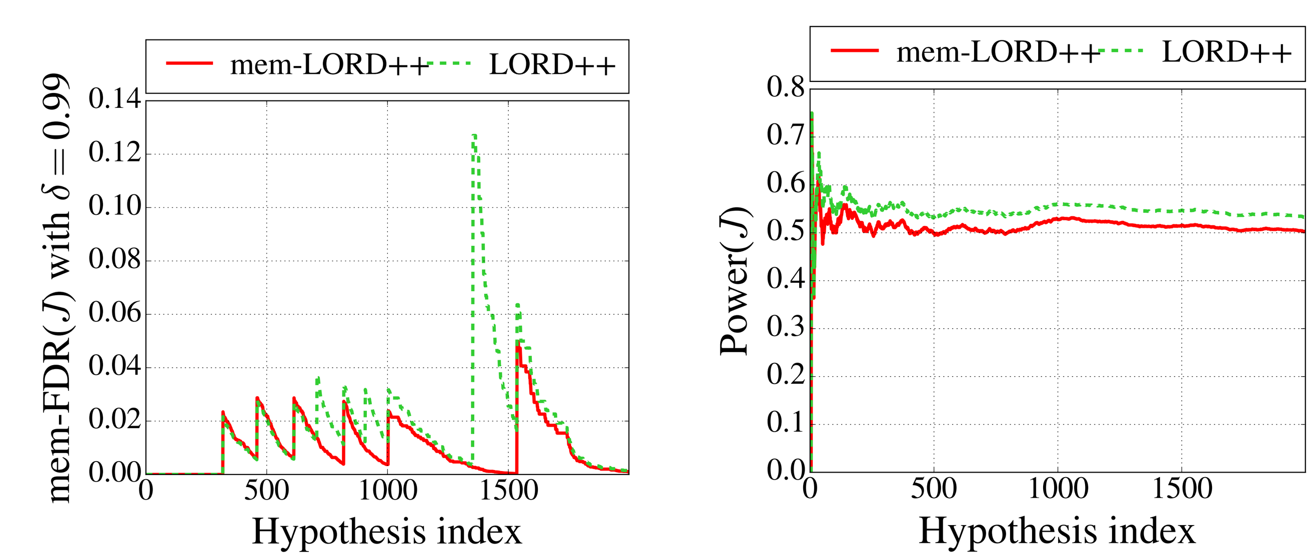

$\Longrightarrow$ mem-LORD++ prevents peaks of the local mem-FDR!

References

- D. Foster, R. Stine, "α-investing: a procedure for sequential control of expected false discoveries," *arXiv preprint arXiv:1412.8729*, 2014
- E. Aharoni, S. Rosset, "Generalized α-investing: definitions, optimality results and application to public databases," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 76, no. 4, pp. 771-794, 2014
- A. Javanmard, A. Montanari, "Online rules for control of false discovery rate and false discovery exceedance," *The Annals of Statistics*, 2017.