Online Rules for Control of False Discovery Rate and False Discovery Exceedance

(A. Javanmard and A. Montanari, Annals of Stats 2017)

Journal club of Machine Learning, CMAP
Alain Virouleau

23 Novembre 2017

Outline

Test/Multi-Test framework (Batch)
- Test and Multiple testing
- Controlling the amount of false discoveries
- Some basic procedures

Online Multi-test framework
- What changes ?
- LORD and other Investing Rules

Outline

Test/Multi-Test framework (Batch)
- Test and Multiple testing
- Controlling the amount of false discoveries
- Some basic procedures

Online Multi-test framework
- What changes ?
- LORD and other Investing Rules

Test and multiple testing

Let $X : (\Omega, \mathcal{F}, \mathbb{P}) \to (\mathfrak{X}, \mathcal{X}, P)$ observations.
Suppose you have a set of null hypotheses and alternatives

$$H_{0,j} : P \in \mathcal{P}_{0,j} \qquad \text{against} \qquad H_{1,j}, \quad j = 1, \ldots, n$$

Test and multiple testing

Let $X : (\Omega, \mathcal{F}, \mathbb{P}) \to (\mathfrak{X}, \mathcal{X}, P)$ observations.
Suppose you have a set of null hypotheses and alternatives

$$H_{0,j} : P \in \mathcal{P}_{0,j} \qquad \text{against} \qquad H_{1,j}, \quad j = 1, \ldots, n$$

Simple example : standard linear model

$Z \sim \mathcal{N}(\mu, \sigma^2 I_n)$ and testing :

$$H_{0,j} : \mu_j = 0, \quad j = 1, \ldots, n$$

Test and multiple testing

Let $X : (\Omega, \mathcal{F}, \mathbb{P}) \to (\mathfrak{X}, \mathcal{X}, P)$ observations.
Suppose you have a set of null hypotheses and alternatives

$$H_{0,j} : P \in \mathcal{P}_{0,j} \qquad \text{against} \qquad H_{1,j}, \quad j = 1, \ldots, n$$

Simple example : standard linear model

$Z \sim \mathcal{N}(\mu, \sigma^2 I_n)$ and testing :

$$H_{0,j} : \mu_j = 0, \quad j = 1, \ldots, n$$

Basically, for each test you compute a statistic $T_j(X)$ and a p-value $p_j(X)$.

Quantifying False Discoveries

For one test,

Quantifying False Discoveries

For one test,making a false discovery is just wrongly rejected H_0. For some fixed level α we impose :

$$\mathbb{P}_{H_0}(H_0 \text{ rejected}) \leq \alpha$$

Quantifying False Discoveries

For one test,making a false discovery is just wrongly rejected H_0. For some fixed level α we impose :

$$\mathbb{P}_{H_0}(H_0 \text{ rejected}) \leq \alpha$$

leading to reject if the p-value p verify : $p \leq \alpha$.
(the p-value of a single test is a random variable designed to verify $(p \leq \alpha) \subset (H_0 \text{ rejected})$, and equality if the cdf of the test statistic is increasing)

Quantifying False Discoveries

For one test,making a false discovery is just wrongly rejected H_0. For some fixed level α we impose :

$$\mathbb{P}_{H_0}(H_0 \text{ rejected}) \leq \alpha$$

leading to reject if the p-value p verify : $p \leq \alpha$.
(the p-value of a single test is a random variable designed to verify $(p \leq \alpha) \subset (H_0 \text{ rejected})$, and equality if the cdf of the test statistic is increasing)For multiple testing :

let $R = \{i \in \{1, \dots, n\} : H_{0,i} \text{ rejected}\}$ and $\mathcal{H}_0$ the set of true null hypotheses. Then the false discoveries are :

$$FD = R \cap \mathcal{H}_0$$

Curse of multiplicity

Suppose that all null hypotheses are true.

If $R = \{i \in \{1, \ldots, n\} : p_i(X) \leq \alpha\}$ and under independency, the probability to make (at least) a false discovery is :

$$P(|FD| \neq 0) = P(\exists i \in \{1, \ldots, n\}, p_i \leq \alpha) = 1 - (1 - \alpha)^n$$

and in expectation we roughly make $n\alpha$ (false) discoveries !

FWER and FDR

Control the probability of making at least one false discovery...

FWER and FDR

Control the probability of making at least one false discovery...

FWER (familywise error rate) control

$$FWER(R) = P(|FD| \neq 0) \leq \alpha$$

FWER and FDR

Control the probability of making at least one false discovery...

FWER (familywise error rate) control

$$FWER(R) = P(|FD| \neq 0) \leq \alpha$$

or the ratio of false discoveries among all the discoveries, in expectation :

FDR (false discovery rate) control

$$FDR(R) = \mathbb{E}\left[\frac{|FD|}{|R|}\right] \leq \alpha$$

FWER and FDR

Control the probability of making at least one false discovery...

FWER (familywise error rate) control

$$FWER(R) = P(|FD| \neq 0) \leq \alpha$$

or the ratio of false discoveries among all the discoveries, in expectation :

FDR (false discovery rate) control

$$FDR(R) = \mathbb{E}\left[\frac{|FD|}{|R|}\right] \leq \alpha$$

Prefer FDR control to allow more discoveries.

Procedures for FWER or FDR control

Procedures for FWER or FDR control

Bonferroni

Consider n p-values $p_1, \ldots, p_n$ computed from n tests. Let $R^{\mathrm{bonf}} = \{i = 1, \ldots, n : p_i \leq \alpha/n\}$. Then :

$$FWER(R^{\mathrm{bonf}}) \leq \alpha$$

Procedures for FWER or FDR control

Bonferroni

Consider n p-values $p_1, \ldots, p_n$ computed from n tests. Let $R^{\mathrm{bonf}} = \{i = 1, \ldots, n : p_i \leq \alpha/n\}$. Then :

$$FWER(R^{\mathrm{bonf}}) \leq \alpha$$

Benjamini and Hochberg

Consider n independent p-values $p_{(1)}, \ldots, p_{(n)}$ computed from n tests, ordered non-decreasingly.

Procedures for FWER or FDR control

Bonferroni

Consider n p-values $p_1, \dots, p_n$ computed from n tests. Let $R^{\text{bonf}} = \{i = 1, \dots, n : p_i \leq \alpha/n\}$. Then :

$$FWER(R^{\text{bonf}}) \leq \alpha$$

Benjamini and Hochberg

Consider n independent p-values $p_{(1)}, \dots, p_{(n)}$ computed from n tests, ordered non-decreasingly.
Let $\hat{i} = \max\{i = 1, \dots, n : p_{(i)} \leq \alpha i/n\}$

Procedures for FWER or FDR control

Bonferroni

Consider n p-values $p_1, \ldots, p_n$ computed from n tests. Let $R^{\text{bonf}} = \{i = 1, \ldots, n : p_i \leq \alpha/n\}$. Then :

$$FWER(R^{\text{bonf}}) \leq \alpha$$

Benjamini and Hochberg

Consider n independent p-values $p_{(1)}, \ldots, p_{(n)}$ computed from n tests, ordered non-decreasingly.
Let $\hat{i} = \max\{i = 1, \ldots, n : p_{(i)} \leq \alpha i/n\}$ and $R^{\text{BH}} = \{i = 1, \ldots, n : p_{(i)} \leq \alpha \hat{i}/n\}$,

Procedures for FWER or FDR control

Bonferroni

Consider n p-values $p_1, \ldots, p_n$ computed from n tests. Let $R^{\mathrm{bonf}} = \{i = 1, \ldots, n : p_i \leq \alpha/n\}$. Then :

$$FWER(R^{\mathrm{bonf}}) \leq \alpha$$

Benjamini and Hochberg

Consider n independent p-values $p_{(1)}, \ldots, p_{(n)}$ computed from n tests, ordered non-decreasingly.
Let $\hat{i} = \max\{i = 1, \ldots, n : p_{(i)} \leq \alpha i/n\}$ and $R^{\mathrm{BH}} = \{i = 1, \ldots, n : p_{(i)} \leq \alpha \hat{i}/n\}$, then

$$FDR(R^{\mathrm{BH}}) \leq \alpha$$

Outline

The problem

The above procedures work for data entirely available, but what if the data arrive sequentially and we do not want to / cannot wait for the whole ?

The new setting

Data, hypotheses $(H_{0,1}, \dots, H_{0,n}, \dots)$ and p-values $(p_1, \dots, p_n, \dots)$ arrive sequentially in a stream, and at each step we must decide whether or not reject the new hypothese, based on the previous decision.

The new setting

Data, hypotheses $(H_{0,1}, \ldots, H_{0,n}, \ldots)$ and p-values $(p_1, \ldots, p_n, \ldots)$ arrive sequentially in a stream, and at each step we must decide whether or not reject the new hypothese, based on the previous decision.

Online testing procedure

Provide a sequence of significence levels α_i with decision rule $R_i = \mathbb{1}_{p_i \leq \alpha_i}$ (1 for rejection, 0 if not)

The new setting

Data, hypotheses $(H_{0,1}, \dots, H_{0,n}, \dots)$ and p-values $(p_1, \dots, p_n, \dots)$ arrive sequentially in a stream, and at each step we must decide whether or not reject the new hypothese, based on the previous decision.

Online testing procedure

Provide a sequence of significence levels α_i with decision rule $R_i = \mathbb{1}_{p_i \leq \alpha_i}$ (1 for rejection, 0 if not)with the following constraints :

$$\forall i \in \{1, \dots, n\}, \quad \alpha_i = \alpha_i(R_1, \dots, R_{i-1})$$

Online control of the FWER

Choose $(\boldsymbol{\alpha})$ a sequence that sum to α, for example $\alpha_i = \alpha 2^{-i}$, then

$$FWER(n) \leq \alpha$$

Online control of the FWER

Choose $(\boldsymbol{\alpha})$ a sequence that sum to α, for example $\alpha_i = \alpha 2^{-i}$, then

$$FWER(n) \leq \alpha$$

proof :
let T be the subset of true null hypotheses

$$P(|FD(n)| \neq 0) = P(\exists i \in T : R_i = 1)$$

Online control of the FWER

Choose $(\boldsymbol{\alpha})$ a sequence that sum to α, for example $\alpha_i = \alpha 2^{-i}$, then

$$FWER(n) \leq \alpha$$

proof :
let T be the subset of true null hypotheses

$$\begin{aligned} P(|FD(n)| \neq 0) &= P(\exists i \in T : R_i = 1) \\ &\leq \sum_{i \in T} P(p_i \leq \alpha_i) \end{aligned}$$

Online control of the FWER

Choose $(\boldsymbol{\alpha})$ a sequence that sum to α, for example $\alpha_i = \alpha 2^{-i}$, then

$$FWER(n) \leq \alpha$$

proof :
let T be the subset of true null hypotheses

$$\begin{aligned} P(|FD(n)| \neq 0) &= P(\exists i \in T : R_i = 1) \\ &\leq \sum_{i \in T} P(p_i \leq \alpha_i) \\ &\leq \sum_{i \in T} \alpha_i \leq \alpha \end{aligned}$$

Online control of the FWER

Choose $(\boldsymbol{\alpha})$ a sequence that sum to α, for example $\alpha_i = \alpha 2^{-i}$, then

$$FWER(n) \leq \alpha$$

proof :
let T be the subset of true null hypotheses

$$\begin{aligned} P(|FD(n)| \neq 0) &= P(\exists i \in T : R_i = 1) \\ &\leq \sum_{i \in T} P(p_i \leq \alpha_i) \\ &\leq \sum_{i \in T} \alpha_i \leq \alpha \end{aligned}$$

But as explained before we are more interested in controlling the FDR.

LORD procedures

LORD stands for Levels Based On Recent Discoveries, which means what it means...

LORD procedures

LORD stands for Levels Based On Recent Discoveries, which means what it means... The (quite surprising, but same as in BH procedure) idea is to reward each discovery ! (because you really deserved it !).

LORD procedures

LORD stands for Levels Based On Recent Discoveries, which means what it means... The (quite surprising, but same as in BH procedure) idea is to reward each discovery ! (because you really deserved it !).

Let ω_0 and b_0 positive constants such that $\omega_0 + b_0 \leq \alpha$ (α still being the target FDR level), and choose $\boldsymbol{\gamma} = (\gamma_i)$ a non-negative sequence summing to 1.

Define for all i :

$$T(i) = \{l \in \{1, \dots, i-1\} : R_l = 1\}$$
$$\tau_i = \max\{l : l \in T(i)\}$$

LORD Procedures

LORD 1 : based on the last discovery

$$\alpha_i = \begin{cases} \gamma_i \omega_0 & \text{if } i \leq t_1 \\ \gamma_{i-\tau_i} b_0 & \text{if } i > t_1 \end{cases}$$

LORD 2 : based on all previous discoveries

$$\alpha_i = \gamma_i \omega_0 + \left(\sum_{l \in T(i)} \gamma_{i-l} \right) b_0$$

LORD Procedures

If the p-values are independent, it works ! !

LORD Procedures

If the p-values are independent, it works!!

proof : (for LORD 1)
$D(n)$ the number of rejections (discoveries) after n step,
$V(n)$ the number of false rejections after n steps. Then :

LORD Procedures

If the p-values are independent, it works ! !

proof : (for LORD 1)
$D(n)$ the number of rejections (discoveries) after n step, $V(n)$ the number of false rejections after n steps. Then :

$$FDR(n) = \mathbb{E}\left[\frac{V(n)}{D(n) \vee 1}\right] = \mathbb{E}\left[\sum_{\substack{i \in \mathcal{H}_0 \\ i \leq n}} \frac{R_i}{D(n) \vee 1}\right]$$

$$= \mathbb{E}\left[\sum_{\substack{i \in \mathcal{H}_0 \\ i \leq n}} \frac{R_i}{D^{-i}(n) \vee 1}\right]$$

$$= \sum_{\substack{i \in \mathcal{H}_0 \\ i \leq n}} \mathbb{E}\left[\mathbb{E}\left[\frac{R_i}{D^{-i}(n) \vee 1} | \mathcal{F}_{i-1}\right]\right]$$

where $D^{-i}(n)$ is the number of rejections when replacing p_i by 0.

Proof

Then, conditioning on $\mathcal{F}_{i-1}$, R_i and $D^{-i}(n)$ are independent so :

Proof

Then, conditioning on $\mathcal{F}_{i-1}$, R_i and $D^{-i}(n)$ are independent so :

$$\mathbb{E}\left[\frac{R_i}{D^{-i}(n) \vee 1}|\mathcal{F}_{i-1}\right] = \mathbb{E}[R_i|\mathcal{F}_{i-1}]\mathbb{E}[\frac{1}{D^{-i}(n) \vee 1}|\mathcal{F}_{i-1}]$$

recall that $i \in \mathcal{H}_0$ so

Proof

Then, conditioning on $\mathcal{F}_{i-1}$, R_i and $D^{-i}(n)$ are independent so :

$$\mathbb{E}\left[\frac{R_i}{D^{-i}(n) \vee 1}|\mathcal{F}_{i-1}\right] = \mathbb{E}[R_i|\mathcal{F}_{i-1}]\mathbb{E}[\frac{1}{D^{-i}(n) \vee 1}|\mathcal{F}_{i-1}]$$

recall that $i \in \mathcal{H}_0$ so

$$\mathbb{E}[R_i|\mathcal{F}_{i-1}] \leq \alpha_i$$

Proof

Then, conditioning on $\mathcal{F}_{i-1}$, R_i and $D^{-i}(n)$ are independent so :

$$\mathbb{E}\left[\frac{R_i}{D^{-i}(n) \vee 1}|\mathcal{F}_{i-1}\right] = \mathbb{E}[R_i|\mathcal{F}_{i-1}]\mathbb{E}[\frac{1}{D^{-i}(n) \vee 1}|\mathcal{F}_{i-1}]$$

recall that $i \in \mathcal{H}_0$ so

$$\mathbb{E}[R_i|\mathcal{F}_{i-1}] \leq \alpha_i$$

and finally :

$$FDR(n) \leq \mathbb{E}\left[\sum_{\substack{i \in \mathcal{H}_0 \\ i \leq n}} \frac{\alpha_i}{D^{-i}(n) \vee 1}\right]$$

(End of) proof

Small argument to obtain $D^{-i}(n) \geq D(n)$ and then remark that :

$$\sum_i \alpha_i \leq \max(\omega_0, b_0) D(n) \sum_i \gamma_i$$

More in the article

- More general rules ;
- How to deal with dependent pvalues ;
- Control the tail of the FDP ;
- Numerical illustrations.