**DS 102 Data, Inference, and Decisions** **Fall 2019**

Optional Material: Online False Discovery Rate Control

*Lecturer: Michael I. Jordan*

# 1 Online False Discovery Rate Control

Classical statistics, including the Benjamini-Hochberg (BH) algorithm, focused on a batch setting in which we make decisions after all necessary data has already been collected. For instance, the BH procedure needs the p-values for all hypotheses one wishes to test *before* making any discoveries.

In turns out that it is possible to construct procedures that make decisions sequentially in time, while controlling the FDR at any given moment. In particular, these procedures allow performing tests one after the other, and at any given time they only need information about past tests, and do not require any knowledge about future tests, such as their number of the hypotheses they wish to test. Even more remarkably, such methods can provide FDR control over a lifetime, for a possibly infinite number of tests.

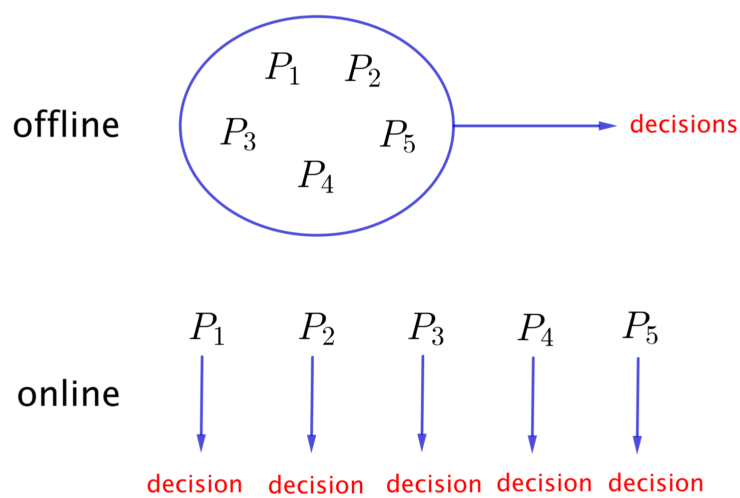


Figure 1.1: Illustration of offline versus online FDR control.

More formally, the setup of online FDR control is as follows. At every time $t \in \mathbb{N}$, a new p-value $P_t$ arrives. As soon as this p-value arrives, a decision has to be made of whether the corresponding hypothesis should be rejected. Importantly, this decision does not depend on any future tests. More precisely, online FDR algorithms determine an appropriate test level $\alpha_t$, dependent on the tests from time 1 up to time $t-1$, and reject the $t$-th hypothesis if $P_t \le \alpha_t$.

It is not clear how a procedure like BH would be adjusted to run in an online fashion. The Bonferroni correction, on the other hand, could be adjusted. The main idea behind the Bonferroni correction was to have individual test levels $\alpha_i$ sum up to $\alpha$. We can do this for infinitely many p-values by picking a sequence $\{\gamma_i\}_{i=1}^{\infty}$ such that $\sum_{i=1}^{\infty} \gamma_i = 1$. If we set $\alpha_i = \gamma_i \alpha$, we get FWER for infinitely many tests. This procedure is sometimes referred to as alpha-spending. However, notice that its construction necessarily implies that $\alpha_i \to 0$ as $i \to \infty$. In other words, after a large enough number of tests, we are very unlikely to make any discoveries.

Here is a glimmer of hope though: FDP is a ratio of two counts: the number of false discoveries, and the total number of discoveries. We can make the ratio small in one of two ways - either by making the numerator small, or by making the denominator big. We can make the denominator big by making lots of discoveries.

The first online FDR algorithm was due to Foster and Stine and is called alpha-investing, to contrast the difference with alpha-spending. The authors gave an economic interpretation to the algorithm, which also transferred to more recent online FDR algorithms. The basic idea is as follows. When initially setting the target FDR level, we give some "initial wealth" to the algorithm. For every test we start, we "invest" some of the wealth into that test. If we don't make a discovery, we end up losing our investment. If, on the other hand, we do make a discovery (regardless of whether it's a true or false one), we get rewarded for that discovery and earn back some wealth. Online FDR algorithms are constructed in a way that ensures we never run out of wealth, by never investing all the wealth we have at a given moment.

A more recent (and slightly simpler) online FDR algorithm is due to Javanmard and Montanari and is called LORD. We state the steps of this procedure below.

**Algorithm 1** The LORD Procedure

**input:** FDR level $\alpha$, non-increasing sequence $\{\gamma_t\}_{t=1}^{\infty}$ such that $\sum_{t=1}^{\infty} \gamma_t = 1$, initial wealth $W_0 \le \alpha$
Set $\alpha_1 = \gamma_1 W_0$
**for** $t = 1, 2, \ldots$ **do**
  p-value $P_t$ arrives
  if $P_t \le \alpha_t$, reject corresponding hypothesis
  $\alpha_{t+1} = \gamma_{t+1} W_0 + \gamma_{t+1-\tau_1}(\alpha - W_0)\mathbf{1}\{\tau_1 < t\} + \alpha \sum_{j=2}^{\infty} \gamma_{t+1-\tau_j} \mathbf{1}\{\tau_j < t\}$,
  where $\tau_j$ is time of $j$-th rejection $\tau_j = \min\{k : \sum_{l=1}^{k} \mathbf{1}\{P_l \le \alpha_l\} = j\}$
**end**

Immediately we see that the LORD procedure is an online algorithm, as p-values $P_t$ arrive at each time step, and decisions about whether to accept or reject the corresponding hypothesis must also be made at time step $t$. As discussed above, this is done by comparing $P_t$ to a time-dependent significance threshold $\alpha_t$.

To gain some intuition, let's take a closer look at the update step that determines $\alpha_{t+1}$ in every round. The last term is a sum over all of the previous discovery declarations, indexed by $j$. Each discovery gains $\alpha$ total wealth, and the $\gamma_{t+1-\tau_j}$ term in the sum serves to spread that wealth over all future time steps. The first term can be thought as doing the same for a "borrrowed" initial wealth $W_0$, whereas the second term pays off that borrowing over time.

We'll investigate this algorithm further in the section, but for now there are two major takeaways: (1) online control of the FDR is possible, and (2) there are in fact implementable algorithms for achieving online FDR control arbitrarily far into the future!

## 2 Analyzing the LORD Algorithm

There are several different versions of the LORD algorithm, and for simplicity, in this section we analyze a version slightly different from the one from previous section.

Let $r_t$ the time of the last rejection before time $t$, and let $\{\gamma_t\}_{t=1}^{\infty}$ be a non-negative infinite sequence that sums to 1. The LORD version that we consider assigns significance levels at each time step $t$ as:

$$\alpha_t = \begin{cases} \gamma_t \alpha, & \text{if no rejection has yet been made} \\ \gamma_{t-r_t}\alpha, & \text{otherwise} \end{cases} .$$

First, we show that this update guarantees an upper bound on estimate of the FDP:

$$\widehat{\text{FDP}} := \frac{\sum_{i=1}^{t} \alpha_i}{\sum_{i=1}^{t} \mathbf{1}\{P_i \leq \alpha_i\}} \leq \alpha.$$

Notice that the denominator is equal to the total number of discoveries. And more generally, any version of the LORD algorithm guarantees this inequality (you can check the update rule and prove this yourself!).

We will first show that the LORD update satisfies $\widehat{\text{FDP}} \leq \alpha$, and then we will show (approximately) that this is sufficient for FDR control.

Suppose that we have completed $t$ tests. At some of these $t$ time steps, we have made discoveries. Denote by $\tau_j$ the time of the $j$-th discovery. Suppose we have made $D$ total discoveries so far. For every $j \leq D$, consider the significance levels in an epoch between two discoveries, i.e. at times $\tau_j, \tau_j + 1, \ldots, \tau_{j+1}$. By definition of the significance level update, these levels are equal to $\alpha\gamma_1, \alpha\gamma_2, \alpha\gamma_3, \ldots$. If we sum up these test levels, we get at most $\alpha$, because $\sum_{t=1}^{\infty} \gamma_t = 1$. Moreover, for each such epoch we get the sum of significance levels at most $\alpha$, so:

$$\sum_{i=1}^{t} \alpha_i \leq \alpha \cdot \text{ number of "epochs"}.$$

However, notice that the number of epochs is exactly equal to the number of rejections, so

$$\sum_{i=1}^{t} \alpha_i \leq \alpha \sum_{i=1}^{t} \mathbf{1}\{P_i \leq \alpha_i\},$$

and after dividing each side by the total number of rejections, we can conclude that $\widehat{\text{FDP}} \leq \alpha$.

Now we want to show that $\widehat{\text{FDP}} \leq \alpha$ implies FDR control. A formal proof of this is slightly more contrived, so we will present a proof stating that a close approximation of the FDR is controlled. The approximation we consider is

$$\text{FDR} \approx \frac{\mathbb{E}[\sum_{i\leq t, i \text{ null}} \mathbf{1}\{P_i \leq \alpha_i\}]}{\mathbb{E}[\sum_{i\leq t} \mathbf{1}\{P_i \leq \alpha_i\}]}.$$

The difference between the approximation and exact FDR is that FDR takes an expectation of the ratio, while here we are taking a ratio of expectations. We will show that approximately FDR $\leq \alpha$ by showing

$$\mathbb{E}\left[\sum_{i\leq t, i \text{ null}} \mathbf{1}\{P_i \leq \alpha_i\}\right] \leq \alpha\, \mathbb{E}\left[\sum_{i\leq t} \mathbf{1}\{P_i \leq \alpha_i\}\right].$$

First, by the tower property, we have:

$$\mathbb{E}\left[\sum_{i\leq t, i \text{ null}} \mathbf{1}\{P_i \leq \alpha_i\}\right] = \mathbb{E}\left[\sum_{i\leq t, i \text{ null}} \mathbb{E}[\mathbf{1}\{P_i \leq \alpha_i\}|\alpha_i]\right].$$

Next, we use the fact that the expectation of an indicator of an event is the probability of that event:

$$\mathbb{E}\left[\sum_{i\leq t, i \text{ null}} \mathbb{E}[\mathbf{1}\{P_i \leq \alpha_i\}|\alpha_i]\right] = \mathbb{E}\left[\sum_{i\leq t, i \text{ null}} \mathbb{P}(P_i \leq \alpha_i|\alpha_i)\right].$$

By uniformity of null p-values, we further have:

$$\mathbb{E}\left[\sum_{i\leq t, i \text{ null}} \mathbb{P}(P_i \leq \alpha_i|\alpha_i)\right] = \sum_{i\leq t, i \text{ null}} \mathbb{E}[\alpha_i].$$

By summing up all the test levels, and not just the null ones (remember $\alpha_i \geq 0$), we get

$$\sum_{i\leq t, i \text{ null}} \mathbb{E}[\alpha_i] \leq \sum_{i\leq t} \mathbb{E}[\alpha_i].$$

Finally, we use $\widehat{\text{FDP}} \leq \alpha$ to conclude the argument:

$$\mathbb{E}\left[\sum_{i\leq t, i \text{ null}} \mathbf{1}\{P_i \leq \alpha_i\}\right] \leq \sum_{i\leq t} \mathbb{E}[\alpha_i] \leq \alpha \mathbb{E}[\sum_{i=1}^{t} \mathbf{1}\{P_i \leq \alpha_i\}].$$

And rearranging gives

$$\frac{\mathbb{E}[\sum_{i\leq t, i \text{ null}} \mathbf{1}\{P_i \leq \alpha_i\}]}{\mathbb{E}[\sum_{i\leq t} \mathbf{1}\{P_i \leq \alpha_i\}]} \leq \alpha\ .$$