

Structure-Adaptive Sequential Testing for Online False Discovery Rate Control

Bowen Gang¹, Wenguang Sun², and Weinan Wang³

Abstract

Consider the online testing of a stream of hypotheses where a real-time decision must be made before the next data point arrives. The error rate is required to be controlled at all decision points. Conventional simultaneous testing rules are no longer applicable due to the more stringent error constraints and absence of future data. Moreover, the online decision-making process may come to a halt when the total error budget, or alpha-wealth, is exhausted. This work develops a new class of structure-adaptive sequential testing (SAST) rules for online false discover rate (FDR) control. A key element in our proposal is a new alpha-investment algorithm that precisely characterizes the gains and losses in sequential decision making. SAST captures time varying structures of the data stream, learns the optimal threshold adaptively in an ongoing manner and optimizes the alpha-wealth allocation across different time periods. We present theory and numerical results to show that the proposed method is valid for online FDR control and achieves substantial power gain over existing online testing rules.

Keywords: Alpha-investing; Conditional local false discovery rate; Covariate-assisted inference; Structured multiple testing; Time series anomaly detection

¹Department of Mathematics, University of Southern California.

²Department of Data Sciences and Operations, University of Southern California. Corresponding Email: wenguans@marshall.usc.edu. The research of Wenguang Sun was supported in part by NSF grant DMS-1712983.

³Snap Inc.

1 Introduction

The online testing problem is concerned with the investigation of a possibly infinite stream of null hypotheses $\{H_1, H_2, \dots\}$ in an ongoing manner based on sequentially collected data $\{X_1, X_2, \dots\}$. At each time point, the investigator must make a real-time decision after X_t arrives, without knowing future data $\{X_{t+1}, X_{t+2}, \dots\}$. The control of multiplicity in sequential testing typically involves imposing serial constraints on error rates over time, which requires that, for example, the family wise error rate (FWER) or false discovery rate (FDR; Benjamini and Hochberg, 1995) must fall below a pre-specified level α at all decision points.

The online testing problem may arise from a range of applications. For example, the quality preserving database (QPD) framework (Aharoni et al., 2010) has been widely employed by many research teams from diverse backgrounds. Some notable databases include Stanford's HIVdb that serves the community of anti-HIV treatment groups, WTCCC's large-scale database that is distributed to assist various whole-genome association studies, and the National Health Institute (NIH) influenza virus resource (IVR) that has been intensively queried by numerous researchers for designing new vaccines and treatments. The proper and efficient management of these large databases calls for new analytical tools for handling thousands of hypothesis tests with real-time decisions made in a sequential fashion. For instance, the NIH IVR has been used to investigate thousands of biomedical hypotheses and, per the record in PubMed, has lead to more than 1,000 scientific publications as of January 2020. It has become increasingly important to develop a powerful and effective monitoring system to control the false positive findings over time. Another important application scenario, which is frequently encountered in finance, social media and mobile computing, is the real-time detection of anomalies based on high-frequency and large-scale time series data. For example, large travel service providers closely monitor the number of changes or cancellation requests of existing itineraries. An abnormal spike usually signifies an unexpected event. It is important for the company to detect such events early and make necessary adjustments. The development of online detection system plays a key role for providing novel and timely marketing insights and avoiding adverse financial losses.

Large-scale testing under the online setup poses several new issues that are not present in conventional “online” setup. First, a real-time decision must be made before the next data point arrives. This makes conventional step-wise testing methods no longer applicable. For instance, the well-known Holm’s procedure (Holm, 1979) for FWER control and Benjamini–Hochberg’s procedure for FDR control both involve first ordering all observed p-values and then choosing a threshold along the ranking. However, the ranking step becomes impossible due to the absence of future data. Second, in contrast with conventional FWER and FDR criteria that only require an overall assessment of the multiplicity in simultaneous testing, the online methods must proceed with more stringent error constraints that are imposed sequentially at every decision point. This not only leads to decreased power in detecting signals but also calls for more carefully designed online testing rules. Third, the data stream often encodes useful local structures, including signal magnitudes, sparsity levels and grouping patterns, that may vary over time. It is crucial to develop flexible and adaptive online rules to exploit the underlying domain knowledge and informative structures. Fourth, the online decision-making process, which proceeds sequentially without the knowledge of future, may come to a halt when the total error budget, or alpha–wealth, is exhausted. As a result, the investigator may miss all potential discoveries in the future. This concern must be carefully addressed because in many applications the hypothesis tests are conducted in an ongoing manner with unpredictable patterns – even the total number of hypotheses to be investigated can be unknown. Finally, how to wisely allocate and invest the alpha–wealth to ensure the validity in error control while maintaining high statistical power of online testing rules in the long run has remained as a key issue that requires much research.

The online FDR control problem has received much recent attention and great progresses have been made. The alpha-investing (AI) idea (Foster and Stine, 2008) and its various generalizations (Aharoni and Rosset, 2014; Ramdas et al., 2017; Javanmard et al., 2018) have served as the basic framework and proved to be effective. Carefully designed AI rules are capable of handling an infinite stream of hypotheses and incorporating informative domain knowledge into the dynamic decision-making process. Beginning with a pre-specified alpha–wealth, the key idea in AI algorithms is that each rejection gains extra alpha–wealth, which

may be subsequently used to make more discoveries at later time points. The generalized AI (GAI) algorithms (Aharoni and Rosset, 2014; Robertson and Wason, 2018; Lynch et al., 2017) are developed for a wider class of pay-out functions, enabling the construction of new online rules with increased power. The GAI++ framework (Ramdas et al., 2017) improves the power of GAI methods uniformly and is capable of dealing with more general settings. The new class of weighted GAI++ methods are flexibly designed to allow “indecisions” and are capable of integrating prior domain knowledge. To alleviate the “piggybacking” and “alpha-death” issues of AI rules, Ramdas et al. (2017) discussed the concept of decaying memory FDR. To effectively incorporate structural information into online inference, the SAFFRON procedure (Ramdas et al., 2018) derived a sequence of thresholds that are adaptive to estimated sparsity levels and showed that the power can be much improved.

This article develops a new class of structure-adaptive sequential testing (SAST) rules for online FDR control with several new features. First, in contrast with existing AI and GAI rules whose building blocks are p-values, the class of SAST rules are built upon the conditional local false discovery rate (Clfdr), which optimally adapts to important local structures in the data stream. Second, the sequential rejection rule based on Clfdr leads to a novel alpha-investing framework that is fundamentally different from that in Foster and Stine (2008). The new framework precisely characterizes the tradeoffs between different actions in online decision making, which provides key insights for designing more powerful online FDR rules. The new AI framework also reveals that SAST automatically avoids the “alpha-death” issue in the sense that its operation always reserves budget to reject new hypotheses, and can proceed in an ongoing manner to any time point in the future. Finally, by adaptively learning from past experiences and dynamically allocating the alphaweight, SAST can effectively avoid the “piggybacking” issue and improve its performance as more data are acquired. Our theoretical and numerical results demonstrate that SAST is effective for online FDR control, and achieves substantial power gain over existing methods in many settings.

The article is organized as follows. Section 2 first introduces the model and problem formulation, and then develops the oracle SAST procedure for online FDR control by assuming that model parameters are known. Section 3 discusses computational algorithms, proposes the

data-driven SAST rule and establishes its theoretical properties. Simulation is conducted in Section 4 to investigate the finite sample performance of SAST and compare it with existing methods. SAST is illustrated in Section 5 through applications for identifying differentially expressed genes and detecting anomalies in time series data. The proofs are provided in the online supplementary material.

2 Oracle and Adaptive Rules for Online FDR Control

We first describe the model and problem formulation in Section 2.1, then discuss three key elements in the proposed SAST rule in turn: a new test statistic to capture the structural information in the data stream (Sections 2.2 and 2.3); a new alpha-investing framework to characterize the gains and losses in sequential decision making (Section 2.4); and a new adaptive learning algorithm to optimize the alpha-wealth allocation (Sections 2.5).

2.1 Model and Problem Formulation

Denote T a continuous temporal domain and $t \in T$ a time point. Let $T \rightarrow T$ be a discrete, ordered and evenly spaced index set for time labels¹. Suppose we are interested in testing a sequence of null hypotheses $\{H_t : t \in T\}$ based on data stream $\mathbf{X} = (X_t : t \in T)$. To describe the true states of nature, define Bernoulli variables \sqrt{t} , where $\sqrt{t} = 0/1$ if H_t is true/false. Let $\{\hat{\pi}_t \in P(\sqrt{t} = 1) : t \in T\}$ denote the local sparsity levels that may vary over time. The observations can be described using a hierarchical model:

$$\sqrt{t} \sim \text{Bernoulli}(\hat{\pi}_t), \quad X_t | \sqrt{t} \sim F_t = (1 - \sqrt{t})F_0 + \sqrt{t}F_{1t}, \quad (2.1)$$

where F_0 and F_{1t} are the null and non-null distributions, respectively. Denote f_0 and f_{1t} the corresponding density functions. We assume that F_0 is known and identical for all $t \in T$. By contrast, $\hat{\pi}_t$ and f_{1t} can vary smoothly in $t \in T$.

¹ T may be taken either as $\{1, 2, \dots, t\}$ on a growing domain or a set of points that lie on a fixed-domain regular grid: $\{\frac{1}{t}, \frac{2}{t}, \dots, \frac{t-1}{t}, 1\}$ with $t \geq 1$.

Remark 1. The inhomogeneity assumption reflects that signals may either vary in strengths or arrive at different rates over time. This structural information can be highly informative. The smoothness assumption makes it possible for pooling information from the observations in the neighborhood of t . We do not impose further assumptions on $\hat{\pi}_t$ and F_{1t} , both of which will be estimated non-parametrically.

Let $\mathbf{X}^t = (X_i : i \leq T; i \leq t)$ be the collection of summary statistics (e.g. p-values or z-values) up to time t . Consider a class of online decision rules $\delta = \{\delta_t(\mathbf{X}^t) : t \leq T\} \subset \{0, 1\}^T$, where $\delta_t(\mathbf{X}^t)$ represents a real-time decision in the sense that δ_t only depends on information available at time t , with $\delta_t = 1$ indicating that H_t is rejected and $\delta_t = 0$ otherwise. Denote $\delta^t = \{\delta_i(\mathbf{X}^i) : i \leq T; i \leq t\}$ the collection of decisions up to t . The online FDR problem is concerned with the performance of a stream of real-time decisions. For decisions up to t , let

$$FDR^t(\delta^t) = E \left(\frac{P_{i \leq t; i \leq T} (1 - \sqrt{\pi_i})_i}{\left(\sum_{i \leq t; i \leq T} \sqrt{\pi_i} \right) - 1} \right), \quad (2.2)$$

where the superscript “ t ” denotes that the FDR is evaluated at a specific time point. The goal is to construct a real-time decision rule $\delta = \{\delta_t(\mathbf{X}^t) : t \leq T\}$ that controls the FDR^t at level α for all $t \leq T$. To compare the power of different testing rules, define the average power (AP) and missed discovery rate (MDR) as

$$AP^t(\delta^t) = \frac{E \left(\sum_{i \leq t; i \leq T} \sqrt{\pi_i} \right)}{E \left(\sum_{i \leq t; i \leq T} \sqrt{\pi_i} \right)}; \quad MDR^t(\delta^t) = 1 - AP^t(\delta^t). \quad (2.3)$$

To simplify the discussion, throughout this section we assume that the distributional information such as the non-null proportion $\hat{\pi}_t$ and density function f_t in Model 2.1 are known. Section 3 considers the case where model parameters are unknown and discusses in detail related estimation and implementation issues.

2.2 The oracle rule for simultaneous testing

The goal of this section is to justify the fundamental role of Clfdr as the building block of the proposed online FDR rule.

The online decision-making process is complicated due to the serial constraints on FDR and absence of future data. To focus on the essential issue, we first consider an ideal setup where a hypothetical oracle observes all data in a local neighborhood at once and makes a batch of simultaneous decisions. Let d denote the size of a neighborhood. Consider the collection of hypotheses in a neighborhood prior to $t^{k^-} - d$: $\{H_i : t^{k^-} - d + 1 \leq i \leq t^{k^-}\}$. Denote the neighborhood $N_d(t^{k^-}) = \{t^{k^-} - d + 1, \dots, t^{k^-}\}$ and the simultaneous decisions $\kappa^- = \{\kappa_i^- ; i \in N_d(t^{k^-})\}$, where κ_i^- is allowed to depend on the entire d -vector $\mathbf{X}^{k^-} = \{X_i : i \in N_d(t^{k^-})\}$. Unlike (2.2), we only require that the FDR is controlled for the d simultaneous decisions:

$$\text{FDR}^s(\kappa^-) = E \left(\frac{P(\kappa_i^- \mid \mathbf{X}^{k^-}) (1 - \kappa_i^-)}{\sum_{i \in N_d(t^{k^-})} (1 - \kappa_i^-)} \right), \quad (2.4)$$

where the superscript “s” indicates a simultaneous-type FDR concept.

The simultaneous testing of multiple hypotheses can be conceptualized as a two-stage inferential process: firstly ranking all hypotheses according to a significance index and secondly choosing a cutoff along the ordered sequence. This process can be described by a thresholding rule of the form

$$= \{I(\kappa_i^- \geq c) : i \in N_d(t^{k^-})\},$$

where $I(\cdot)$ is an indicator function, κ_i^- is the significance index of H_i and c is the cutoff of κ_i^- . For example, the BH procedure uses the p-value as the significance index to order the hypotheses, and implements a step-up algorithm to determine a data-driven cutoff c .

However, the p-value is inefficient for online FDR analysis as it fails to capture the important structural information in the data stream. We propose to use the conditional local false discovery rate (Clfdr) as the significance index to order the hypotheses:

$$\text{Clfdr}_t(x_t) = P(\kappa_t = 0 \mid X_t = x_t) = \frac{(1 - \hat{f}_t(x_t)) f_0(x_t)}{f_t(x_t)}, \text{ for } t \in T. \quad (2.5)$$

Denote $\text{Clfdr}_{(1)}, \dots, \text{Clfdr}_{(d)}$ the ordered Clfdr values in $N_d(t^{k^-})$ and $H_{(1)}, \dots, H_{(d)}$ the corresponding hypotheses. To determine the cutoff for simultaneous testing, we apply a step-

wise algorithm

$$k = \max \left(j : \frac{1}{j} \sum_{i=1}^j \text{Clfdr}_{(i)} \leq c \right). \quad (2.6)$$

Then the threshold is $c = \text{Clfdr}_{(k)}$ and we reject $H_{(1)}, \dots, H_{(k)}$. The Clfdr rule (2.6) may be viewed as an oracle rule that sees all data in a local neighborhood at once and then makes simultaneous decisions. In Appendix C, we establish the optimality property of the Clfdr rule for simultaneous testing under the “online” setup. An infinite data stream can be approximated by sequential data points arrived in batches. Intuitively, the Clfdr statistic provides a good building block for developing new online sequential testing rules as it is optimal for simultaneous inference in each batch of data points.

Remark 2. In the “online” setup for simultaneous testing with a covariate sequence, which includes the Clfdr rule (2.6) as a special case, Cai et al. (2019) develops asymptotic optimality theory. We can similarly show that (2.6) is asymptotically optimal in the sense that it achieves the benchmark of a hypothetical oracle. However, the optimality issue in the online setup, which depends on many other factors such as the optimal allocation of alpha-wealth and prediction of future patterns over time, is still an open issue and requires much research.

2.3 Adapting to local structures by Clfdr: an illustration

The incorporation of structural information and domain knowledge promises to improve the power of existing FDR procedures (Genovese et al., 2006; Cai and Sun, 2009; Hu et al., 2010; Lei and Fithian, 2018; Cai et al., 2019). For example, the works by Hu et al. (2010), Li and Barber (2019) and Xia et al. (2020) showed that the weighted p-values can be constructed to capture the varying sparsity levels of ordered or grouped hypotheses. In contrast with the p-value, the Clfdr takes into account important structural information such as $\hat{\tau}_t$ and f_t , which makes Clfdr an ideal building block for multiple testing with inhomogeneous data streams. We present an example to illustrate the advantage of the Clfdr rule.

Consider the following situation where the data stream $\{X_1, X_2, \dots, X_t, \dots\}$ obeys a ran-

dom mixture model with varying sparsity levels:

$$X_t \leftarrow (1 - \hat{\pi}_t)N(0, 1) + \hat{\pi}_t N(\mu, 1). \quad (2.7)$$

Model (2.7) is a special case of Model (2.1): the null and alternative densities are fixed and the dynamic part is fully captured by the varying proportion $\hat{\pi}_t$. The key idea of Clfdr and weighted p-value (in the form of p_t/w_t , where w_t is the weight for H_t) is to upweight the hypotheses in a local neighborhood where signals appear more frequently (e.g. in clusters).

To compare the effectiveness of different weighting methods, we simulate a data stream for testing $m = 5000$ hypotheses. The top row in Figure 1 sets $\hat{\pi}_t = 0.5$ in blocks $[1001 : 1150]$, $[2001 : 2150]$, $[3001 : 3100]$ and $[4001 : 4150]$, and $\hat{\pi}_t = 0.01$ elsewhere. We vary μ from 2 to 4. The bottom row sets $\mu = 2.5$ and vary $\hat{\pi}_t$ from 0.2 to 0.9 in the above blocks. The block structure is highly informative and can be exploited by Clfdr and weighted p-values to improve the power. We apply the following methods at FDR level $\alpha = 0.05$ by assuming that the model parameters in (2.7) are known: BH (Benjamini and Hochberg, 1995), the structure-adaptive BH algorithm (SABHA; Li and Barber, 2019) using weighted p-values with $w_t = 1/(1 - \hat{\pi}_t)$, the GAP method (Xia et al., 2020) using weighted p-values with $w_t = \hat{\pi}_t/(1 - \hat{\pi}_t)$, and the Clfdr rule (2.6). We can see that all methods control the FDR at the nominal level. In terms of the power, BH can be improved by SABHA and GAP, both of which are dominated by the Clfdr rule. Clfdr captures the varying structure in the data stream more effectively: in addition to varied $\hat{\pi}_t$, it also adapts to f_t , leading to further power improvement.

2.4 A new alpha-investing framework

Existing FDR methods such as the BH and Clfdr procedures are simultaneous inference procedures that involve first ordering the significance indices (p-value or Clfdr) of all hypotheses and then applying a step-wise algorithm to the ordered sequence to determine the threshold. However, the ranking and thresholding strategy cannot be applied to the online setting where the investigator must make real-time decisions without seeing future observations. This section discusses how to avoid the overflow of the FDR at any given time t and how to efficiently

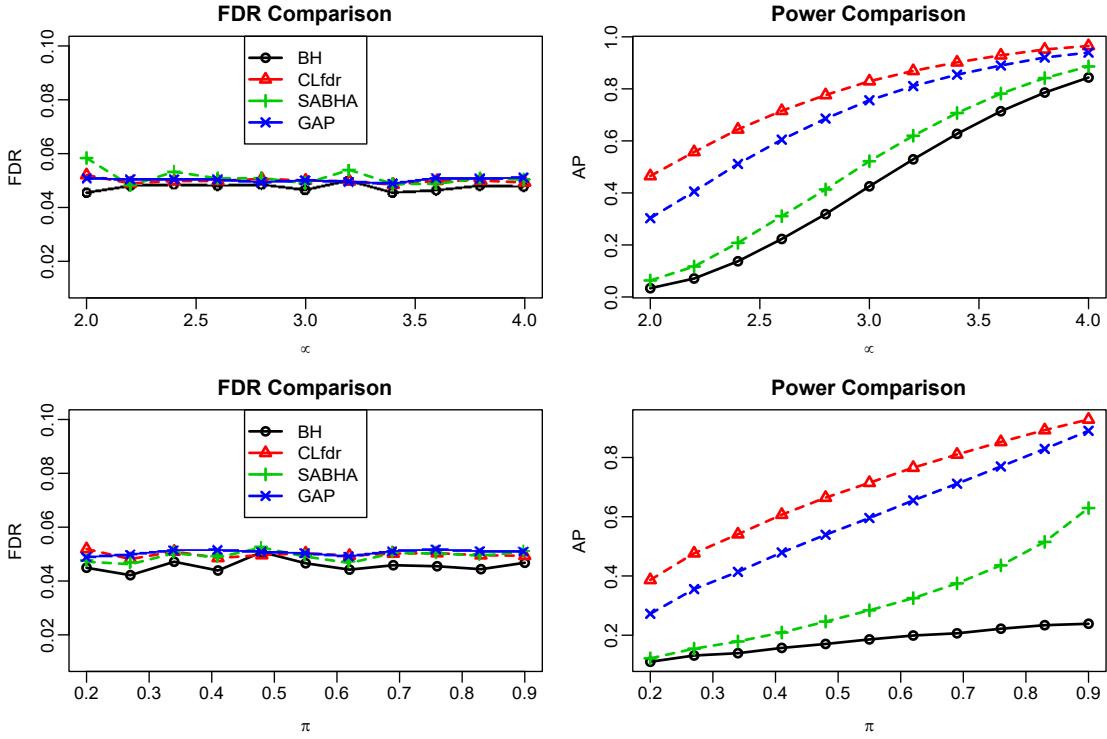


Figure 1: Structure-adaptiveness: Clfdr vs weighted p-values.

allocate the alpha-wealth to increase the power.

We start with a novel interpretation of the alpha-investing idea by recasting the Clfdr algorithm (2.6) as a varying-capacity knapsack process. Denote $R_t \rightarrow \{H_1, H_2, \dots, H_t\}$ the collection of rejected hypotheses at time t . The decision process (2.6) can be conceptualized as a sequence of comparisons of two quantities: the nominal FDR level α and the average of the rejected Clfdr values. Specifically, (2.6) motivates us to consider the constraint

$$\text{Ave}\{\text{Clfdr}_i : i \in R_t\} \leq \alpha, \quad \text{for all } t \in T, \quad (2.8)$$

where $\text{Ave}(A)$ denotes the average of the elements in set A . The simultaneous testing setup is only concerned with one constraint at the last time point when all data have been observed. By contrast, the online setup poses a series of constraints, e.g. (2.8) must be fulfilled for every t to avoid the overflow of FDR^t (2.2).

We view (2.8) as a dynamic decision process resembling a knapsack problem, where H_t can

only be rejected when the following constraint is satisfied:

$$\text{Clfdr}_t - \epsilon \geq C_t \Leftrightarrow \sum_{i=1}^t (\text{Clfdr}_i - \epsilon), \text{ for } t = 1, 2, \dots \quad (2.9)$$

where C_t is the capacity (of the knapsack) at time t with the default choice $C_1 = 0$. The capacity may either expand or shrink over time, depending on the sequential decisions along the data stream. This dynamic process can be described as follows. The initial capacity is $C_1 = 0$. Starting from $t = 1$, we reject H_t if (2.9) is fulfilled. If H_t with $\text{Clfdr}_t < \epsilon$ is rejected, then the capacity C_t increases by $\epsilon - \text{Clfdr}_t$ (gain); hence we earn bonus room. By contrast, if H_t with $\text{Clfdr}_t > \epsilon$ is rejected, then C_t decreases by $\epsilon - \text{Clfdr}_t$ (loss).

The decision process (2.9) provides a new alpha-investing framework that precisely characterizes the gains and losses in sequential testing. In contrast with the alpha-investing framework in [Foster and Stine \(2008\)](#), which views each rejection as a gain of extra alpha-wealth, the new characterization (2.9) reveals that not all rejections are created equal: rejections with small Clfdr will lead to increased alpha-wealth whereas rejections with large Clfdr will lead to decreased alpha-wealth. This view provides key insights for designing more powerful online FDR rules. Moreover, the new AI framework reveals that utilizing Clfdr rules can automatically avoid the “alpha-death” issue. Specifically, the process (2.9) can always reject new hypotheses with $\text{Clfdr} < \epsilon$ regardless of the current budget, and can proceed in an ongoing manner to any time point in the future.

2.5 Oracle-assisted adaptive learning and the SAST algorithm

To efficiently allocate the alpha-wealth, we need to further refine the online algorithm (2.9) to avoid making imprudent rejections that can potentially eat up all the budget. The specific issue is referred to as “piggybacking” ([Ramdas et al., 2017](#)), which, in a vivid way, describes the phenomenon that a string of bad decisions were made due to previously acquired budget.

To see the necessity of taking careful actions, suppose that we have accumulated some bonus room over time before observing a very large Clfdr_t satisfying (2.9). Although rejecting H_t is an action that obeys the FDR constraint, the action can be unwise since it is possible

that we can invest the extra “cost”, $\text{Clfdr}_t \leftarrow \cdot$, to make more discoveries at later time points. A practical strategy is to incorporate a “barrier” τ_t and modify (2.9) as

$$\text{Clfdr}_t < \tau_t \text{ and } \text{Clfdr}_t \leftarrow \max_{H_i \in R_{t-1}} \text{Clfdr}_i \leftarrow \cdot. \quad (2.10)$$

The barrier can effectively prevent “piggybacking” by filtering out large Clfdr_t and hence saving budget for future.

The choice of τ_t depends on the pattern of future hypotheses. However, all online methods must proceed without seeing the future. To resolve the issue, consider the oracle Clfdr rule (2.6) that sees all data in a local neighborhood at once. If we assume that the hypothesis stream is “locally stable” in its patterns, then τ_t may be informed by the oracle rule (2.6) simultaneously conducted on a local neighborhood $N_d(t) = \{t-d+1, \dots, t\}$. The rationale is to use recent past data to get some ideas about the patterns of hypotheses to arrive in the near future. Concretely, we first order $\{H_i : i \in N_d(t)\}$ according to their Clfdr values, then run the “online” algorithm (2.6) to set the barrier $\tau_t = \text{Clfdr}_{(k+1)}$. The online algorithm, by acting as if it sees the future, can effectively filter out large Clfdr values and hence avoid inefficient investments. The operation of algorithm (2.6) also implies that the barrier τ_t may be either raised or lowered according to the varied \hat{t}_t and f_t in the dynamic model, which is desirable in practice for dealing with inhomogeneous data streams. In Section 4.3, we illustrate that the incorporating of the barrier can greatly reduce the MDR (2.3).

Finally, we present the proposed structure-adaptive sequential testing (SAST) rule (oracle version with known parameters) in Algorithm 1. The SAST algorithm essentially utilizes the sequential constraints (2.10) with barriers set by the online algorithm (2.6).

We can see that Algorithm 1 runs two parallel procedures: an online procedure for making real-time decisions and an “online” procedure for determining the barrier. Thus the information of every data point has been used twice: first X_t is used for real-time decision-making at time t , then X_t is stored as past data so that we can “learn from experiences” via the online oracle. The following theorem shows that Algorithm 1 is valid for online FDR control.

Theorem 1. Consider the online FDR procedure $\tau = (\tau_t : t \in T)$, where τ_t is determined by

Algorithm 1. The oracle SAST rule.

Initialization: $A_0 = \emptyset$, $\alpha_0 = \alpha$.
 Updating the barrier: Let $N_d(t) = \{t-d+1, \dots, t\}$. Sort $\{\text{Clfdr}_i : i \in N_d(t)\}$ from the smallest to largest and denote the ordered statistics as $\{\text{Clfdr}_{(1)}^t, \text{Clfdr}_{(2)}^t, \dots\}$. If $\text{Clfdr}_{(1)}^t > \alpha$, keep the same barrier $t = t-1$. Otherwise let $k = \max\{j : Q^t(j) \leq \alpha\}$, where $Q^t(j) = \frac{1}{P} \sum_{i=1}^P \text{Clfdr}_{(i)}^t$, and update the barrier as $t = \text{Clfdr}_{(k+1)}^t$.
 Decision: Let $R_t = \{i \leq t : i = 1\}$ and denote $|R_t|$ its cardinality. If $\text{Clfdr}_t < \alpha$ and $\{R_{t-1}\} + 1 \leq \sum_{i \in R_{t-1}} \text{Clfdr}_i + \text{Clfdr}_t \leq \alpha$, then $t = 1$. Otherwise $t = 0$.

Algorithm 1. Denote $t = (\{i : i \leq t; i \in T\})$. Assume that the Clfdr values are known. Then we have $\text{FDR}^t(\alpha) \leq \alpha$, for all $t \in T$.

3 Data-Driven SAST and its Theoretical Properties

We first develop estimation methodologies and computational algorithms to implement the SAST rule in Section 3.1, then establish the theoretical properties of the data-driven procedure in Section 3.2.

3.1 Data-driven procedure and computational algorithms

We assume that the null distribution of z-values f_0 is known, which is a standard practice in the literature². The key quantities remained to be estimated are \hat{t}_t and $f_t(x)$. In our motivating applications such as queries of QPDs and anomaly detection in high-frequency time series, the databases or servers have already collected large amounts of data at the beginning of the online FDR analysis. Let $\{X_{-k_0}, \dots, X_{-1}, X_0\}$ denote the available data and suppose we start online testing at $t = 1$ with a data stream $\{X_1, X_2, \dots\}$ ³.

The conditional density f_t can be estimated using standard (one-sided) bivariate kernel

²In situations where the empirical null is more appropriate (Efron, 2004), f_0 can be first estimated using the method in Jin and Cai (2007) and then treated as known.

³In situations where the online FDR analysis must start without prior data, we suggest applying existing methods such as LOND first and then switch to SAST as more data are acquired.

methods (Silverman, 1986):

$$\hat{f}_t(x) = \frac{\sum_{j=t-d+1}^{t-1} K_{h_t}(j-t) K_{h_x}(x_j - x)}{\sum_{j=t-d+1}^{t-1} K_{h_t}(j-t)}, \quad (3.1)$$

where $d \geq K_0$ is the length of the moving window that includes a pre-specified number of observations, $K(t)$ is a kernel function, h_t and h_x are the bandwidths, with $K_h(t) = h^{-1}K(t/h)$.

Remark 3. In analysis of large-scale high-frequency time series data such as the NYC taxi data (Section 5), we can pre-specify d , say, to be 1000 to speed up the computation. This virtually has no impact on the estimator \hat{f}_t (compared to using all previous data). Otherwise we can always set $d = t$. Note that our estimator has followed the standard practice in density estimation, which does not include X_t when estimating $f_t(x)$ at time t .

Next we propose a weighted screening approach to estimate the unknown proportion $\{\hat{\pi}_t : t \leq T\}$. The key idea is to use a kernel, which weights observations by their distance to t , to pool information from nearby time points. Let h_t be the bandwidth⁴ and K a kernel function satisfying $\int K(t)dt = 1$, $\int tK(t)dt = 0$ and $\int t^2K(t)dt < 1$. Consider a screening procedure $T_t(\theta) = \{t-d+1 \leq i \leq t-1 : P_i > \theta\}$, where θ is a pre-specified threshold. We propose the following estimator based on Cai et al. (2019):

$$\hat{\pi}_t = 1 - \frac{\sum_{i \in T_t(\theta)} K_{h_t}(t-i)}{(1-\theta) \sum_{i=t-d+1}^{t-1} K_{h_t}(t-i)}. \quad (3.2)$$

Now we provide some intuitions of the estimator (3.2). First, at time t , define $v_h(t, i) = K_{h_t}(|t-i|)/K_{h_t}(0)$. We can view $m_t = \sum_{i=t-d+1}^{t-1} v_h(t, i)$ as the “total” number of observations at time t . Suppose we are interested in counting how many null p-values are greater than θ among the m_t “observations” at t . The empirical count is given by $\sum_{i \in T_t(\theta)} v_h(t, i)$, whereas the expected count is given by $\sum_{i=t-d+1}^{t-1} v_h(t, i) \{1 - \hat{\pi}_t\} (1 - \theta)$. Equation (3.2) can be derived by first setting equal the expected and empirical counts and then solving for $\hat{\pi}_t$. In Section

⁴We recommend using the same h_t in both (3.1) and (3.2) to stabilize the performance.

3.2 we show that $\hat{\pi}_t^{\beta}$ is a consistent estimator of

$$\hat{\pi}_t^{\beta} = 1 - (1 - \beta)^{-1} P(P_t > \beta), \quad (3.3)$$

which always underestimates $\hat{\pi}_t$ and guarantees (conservative) FDR control (Propositions 4).

Remark 4. There is a bias-variance tradeoff in the choice of β for the proposed estimator $\hat{\pi}_t$. We shall see that when β increases, the “purity” of the screening subset $T(\beta)$ increases, which decreases the approximation bias of $\hat{\pi}_t^{\beta}$ (desirable). At the same time, when β increases, the sample size for estimating $\hat{\pi}_t^{\beta}$ will decrease, thereby increasing the variance of the estimator $\hat{\pi}_t^{\beta}$ (undesirable). The common choice of β is 0.5. In Section 4.1, we discuss a data-driven algorithm that chooses β adaptively.

Combining (3.1) and (3.2), we propose to estimate the Clfdr as

$$\hat{\text{Clfdr}}_t = \min \left(\frac{(1 - \hat{\pi}_t^{\beta}) f_0(x_t)}{\hat{f}_t(x_t)}, 1 \right), \quad t \geq T. \quad (3.4)$$

Our proposed data-driven rule implements Algorithm 1 by substituting $\hat{\text{Clfdr}}_t$ in place of Clfdr_t . The data-driven algorithm is summarized in Algorithm 2.

Algorithm 2. The data-driven SAST.

Initialization: $R_0 = \emptyset$, $\pi_0 = \emptyset$.

Estimation: $\hat{\text{Clfdr}}_t = \min \left(\frac{(1 - \hat{\pi}_t^{\beta}) f_0(x_t)}{\hat{f}_t(x_t)}, 1 \right)$, where $\hat{\pi}_t^{\beta}$ and \hat{f}_t are defined by (3.2) and (3.1), respectively.

Updating the barrier: Let $N_d(t) = \{t-d+1, \dots, t\}$. Sort $\{\hat{\text{Clfdr}}_i : i \in N_d(t)\}$ from the smallest to largest and denote the ordered statistics as $\{\hat{\text{Clfdr}}_{(1)}^t, \hat{\text{Clfdr}}_{(2)}^t, \dots\}$. If $\hat{\text{Clfdr}}_{(1)}^t > \beta$, keep the same barrier $\pi_t = \pi_{t-1}$. Otherwise let $k = \max\{j : Q^t(j) \geq \beta\}$, where $Q^t(j) = \frac{1}{j} \sum_{i=1}^j \hat{\text{Clfdr}}_{(i)}^t$, and update the barrier as $\pi_t = \hat{\text{Clfdr}}_{(k+1)}^t$.

Decision: Let $R_t = \{i \in t : i = 1\}$ and denote $|R_t|$ its cardinality. If $\text{Clfdr}_t < \pi_t$ and $|R_{t-1}| + 1 \leq \hat{\text{Clfdr}}_t + \hat{\text{Clfdr}}_{t-1}$, then $\pi_t = 1$. Otherwise $\pi_t = 0$.

3.2 Theoretical properties of data–driven SAST

This section aims to show that the data–driven SAST procedure is asymptotically valid for online FDR control. Our theoretical analysis is divided into three steps. The first step (Proposition 1) shows that a hypothetical rule, which substitutes

$$\text{Clfdr}_t^{\square} = \frac{(1 - \hat{\gamma}_t^{\square})f_0(x_t)}{f_t(x_t)} \quad (3.5)$$

in place of Clfdr_t in Algorithm 1, is conservative for online FDR control.

Proposition 1. Consider $\hat{\gamma}_t^{\square}$ defined by (3.3), then we have $\hat{\gamma}_t^{\square} \leq \hat{\gamma}_t$ and $\text{Clfdr}_t \leq \text{Clfdr}_t^{\square}$. Hence the hypothetical rule using (3.5) is valid (and conservative) for online FDR control.

The second step (Proposition 2) shows that Clfdr_t is a consistent estimator of Clfdr_t^{\square} . We prove the result by appealing to the infill–asymptotics framework (Stein, 2012), which converts the set of time points $\{1, 2, \dots, t\}$ on a growing domain to a set of points that lie on a fixed-domain regular grid: $\{\frac{1}{t}, \frac{2}{t}, \dots, \frac{t-1}{t}, 1\}$. The discussions in Stein (2012) indicate that the infill model is equivalent to the growing domain model under mild conditions: When $t \rightarrow \infty$, the asymptotic arguments, which respectively correspond to letting the grid become denser and denser in the fixed interval $(0, 1]$ and letting the domain $\{1, 2, \dots, t\}$ to grow to infinity, can be essentially established in the same manner. We state the fixed domain theory as it naturally connects to the familiar density estimation theory, where the notations and regularity conditions are standard and easy to understand. The growing domain version of the theory is briefly discussed in Appendix B.3. □

We can similarly define the bivariate density estimator and the following conditional proportion estimator:

$$\hat{\gamma}_t^{\square} = 1 - \frac{\sum_{i=t+1}^t K_{h_t}(1 - i/t)}{\sum_{i=t}^t K_{h_t}(1 - i/t)}. \quad (3.6)$$

The two estimators (3.2) and (3.6) are essentially identical (with rescaled bandwidths).

We state the following regularity conditions. Condition (A1) requires that $f_t(x)$ is smooth in t . Conditions (A2) to (A4) are standard in density estimation theory; see, for example, (Wand and Jones, 1994).

(A1): For any $s \in (0, 1]$ and $\epsilon > 0, \eta$ such that if $|s - s^0| \geq \eta, s^0 \in (0, 1]$ then $\int_0^1 |f_s(x) - f_{s^0}(x)| dx < \epsilon$.

(A2): $h_x \geq 0, h_t \geq 0$ and $th_x h_t \leq 1$.

(A3): $f_j(x) \leq C$ and $\int_0^1 |f_j^{(0)}(x)| dx \leq C$ for all j .

(A4): $dh_t \leq 1$ and $d \leq c h_t$ for some $c > 0$.

Proposition 2. Suppose (A1)–(A4) hold, then $\text{Clfdr}_t \leq \text{Clfdr}_t^*$.

In the third step of our theoretical analysis (Theorem 2), we establish the asymptotic validity of the data-driven SAST procedure for online FDR control.

Theorem 2. Assume the conditions in Proposition 2 hold. Then for any given time t , the data-driven SAST rule (Algorithm 2) controls the FDR t at level α asymptotically.

3.3 Theory for data streams with fixed distributions

SAST learns from past decisions and improves its performance over time through the assistance from an online oracle. The barrier $\hat{\alpha}_t$ would become more informative as more tests are conducted. Specifically, the initial barrier is set to be α at time $t = 1$, which is very conservative. In the special case when the mixture model has fixed $\hat{\alpha}_t$ and f_t over time, we can show that the barrier $\hat{\alpha}_t$ would converge to α_{OR} , where α_{OR} is the optimal threshold of the “online” oracle procedure in Section 2.2. Hence, provided that the capacity allows, the operation of (2.10) implies that SAST behaves like an oracle that sees all data points (including future ones). Our numerical results show that the FDR levels of SAST are conservative at the beginning but the FDR becomes closer to α as we sequentially update the barrier with information from more time points.

Theorem 3. Assume conditions from Theorem 2 holds. Then under the model with $\hat{\alpha}_t \geq \hat{\alpha}$ and $f_t \geq f$, the data-driven barrier $\hat{\alpha}_t \geq \alpha_{\text{OR}}$ when $t \geq 1$, where α_{OR} is the optimal threshold of the oracle FDR procedure for simultaneous testing defined in Section 2.2.

4 Simulation

In this section, we first provide some details in implementation. Simulation studies are conducted in Section 4.2 to compare the oracle and data-driven SAST procedures with other existing online FDR rules. Section 4.3 presents an example to illustrate the merit of including a barrier in online sequential testing.

4.1 Implementation Details

In our simulation, the conditional density function $\hat{f}_t(x)$ is estimated using R function `density`, where the bandwidths h_x and h_t are chosen based on [Silverman \(1986\)](#). A key step in the SAST algorithm is to estimate $\hat{\pi}_t$. We propose to choose a data-driven $\hat{\pi}_{BH}$ by running BH at $\alpha = 0.5$. Roughly speaking, in the subset $T_t(\hat{\pi}_{BH}) = \{t \in \tilde{d+1} : \pi_t \leq 1 : P_i < \hat{\pi}\}$, 50% of the cases come from the null (e.g. the expected proportion of false positives made by BH). It is anticipated that in the remaining set $T_t(\hat{\pi}) = \{t \in \tilde{d+1} : \pi_t \leq 1 : P_i > \hat{\pi}_{BH}\}$, which is used to construct our estimator, majority of the cases should come from the null. This data-driven scheme ensures a small bias in approximation, while maintaining a larger sample size compared to the standard choice of $\alpha = 0.5$.

4.2 Comparisons of online FDRs and MDRs

We compare the proposed SAST procedure with its competitors for online FDR control. The following methods are included in the comparison:

- SAST with known π_t and f_t (SAST.OR, Algorithm 1)
- SAST with estimated model parameters (SAST.DD, Algorithm 2)
- LOND: the method proposed by [Javanmard, A. and Montanari, A. \(2016\)](#).
- LORD++: the GAI++ rule proposed by [Ramdas et al. \(2017\)](#).

For the general simulation setup, we choose $m = 5000$ and the pre-specified FDR level

$\alpha = 0.05$. The data are simulated from the following model:

$$X_t \leftarrow (1 - \hat{\pi}_t)N(0, 1) + \hat{\pi}_t N(\mu, 1).$$

For the data-driven method, we need an initial burn-in period. In simulation we generate 500 data points prior to $t = 1$ to form an initial density estimate. The varying density and proportion estimates are updated every 200 time points. The following simulation settings are considered:

1. Block Pattern: $\hat{\pi}_t = 0.01$, for $t \in [1, 1000] \cup [1200, 2000] \cup [2200, 3000] \cup [3200, 4000] \cup [4200, 5000]$; $\hat{\pi}_t = 0.6$, for $t \in [1000, 1200] \cup [2000, 2200]$; $\hat{\pi}_t = 0.8$, for $t \in [3000, 3200] \cup [4000, 4200]$. Vary μ from 2 to 4.2 with step size 0.2.
2. Constant Pattern: $\hat{\pi}_t = 0.05$, $t = 1, \dots, m$. Vary μ from 2 to 4.2 with step size 0.2.
3. Linear Pattern: Vary $\hat{\pi}_t$ linearly from 0 to 0.5. Vary μ from 2 to 4.2 with step size 0.5.
4. Sine Pattern: $\hat{\pi}_t = (\sin \frac{2\pi t}{m} + 1)/4$, $\hat{\pi}_t$ ranges between 0 to 0.5, vary μ from 2 to 4.2 with step size 0.5.

We apply different methods at $\alpha = 0.05$. The empirical FDR and MDR levels are evaluated using the average of the false discovery proportions and missed discovery proportions from 1000 replications. To investigate the performance of different methods in the online setting, we display the empirical FDR^t and MDR^t levels at various time points, where the intermediate evaluation points ranges from 1500 to 5000 with step size 500. The results for block and constant patterns are summarized in Figure 2, and the results for the linear and sine patterns are summarized in Figure 3.

The following observations can be made from the simulation results.

- (a) All methods control FDR^t at the nominal level at all decision points being considered. SAST.OR achieves the nominal level very precisely. SAST.DD is conservative. LOND and LOND++ are more conservative compared to SAST.DD.

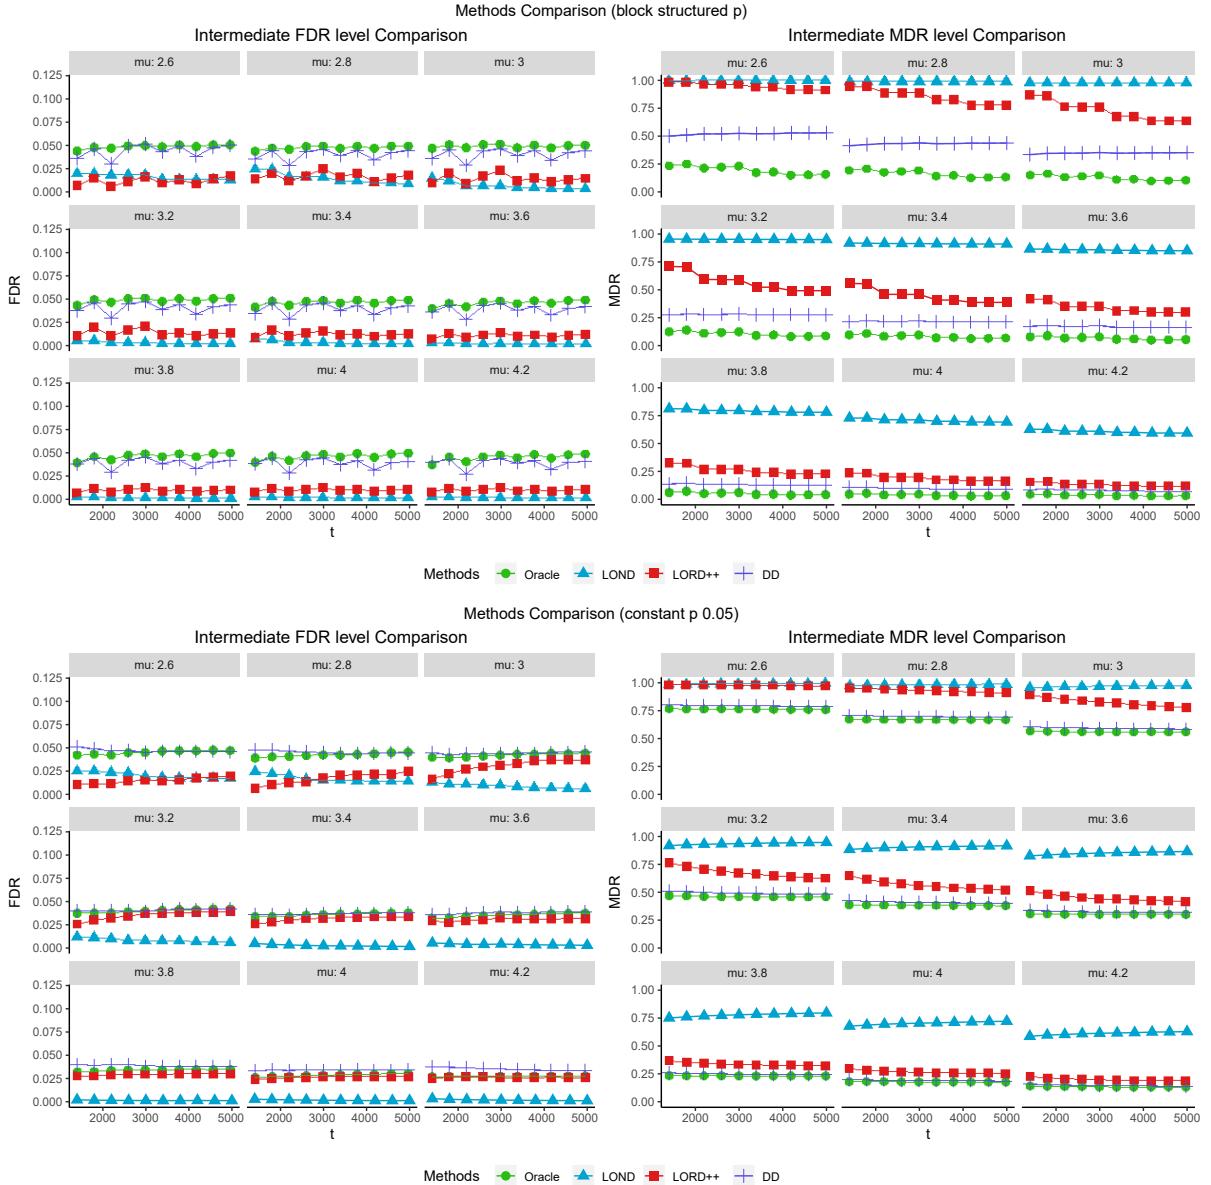


Figure 2: Simulation results for Settings 1 and 2: signal proportions are varied in a block fashion and kept constant respectively. Various signal strengths are investigated as well. Our data-driven and oracle procedures provide significantly more power while controlling FDR under the nominal level in comparison with others.

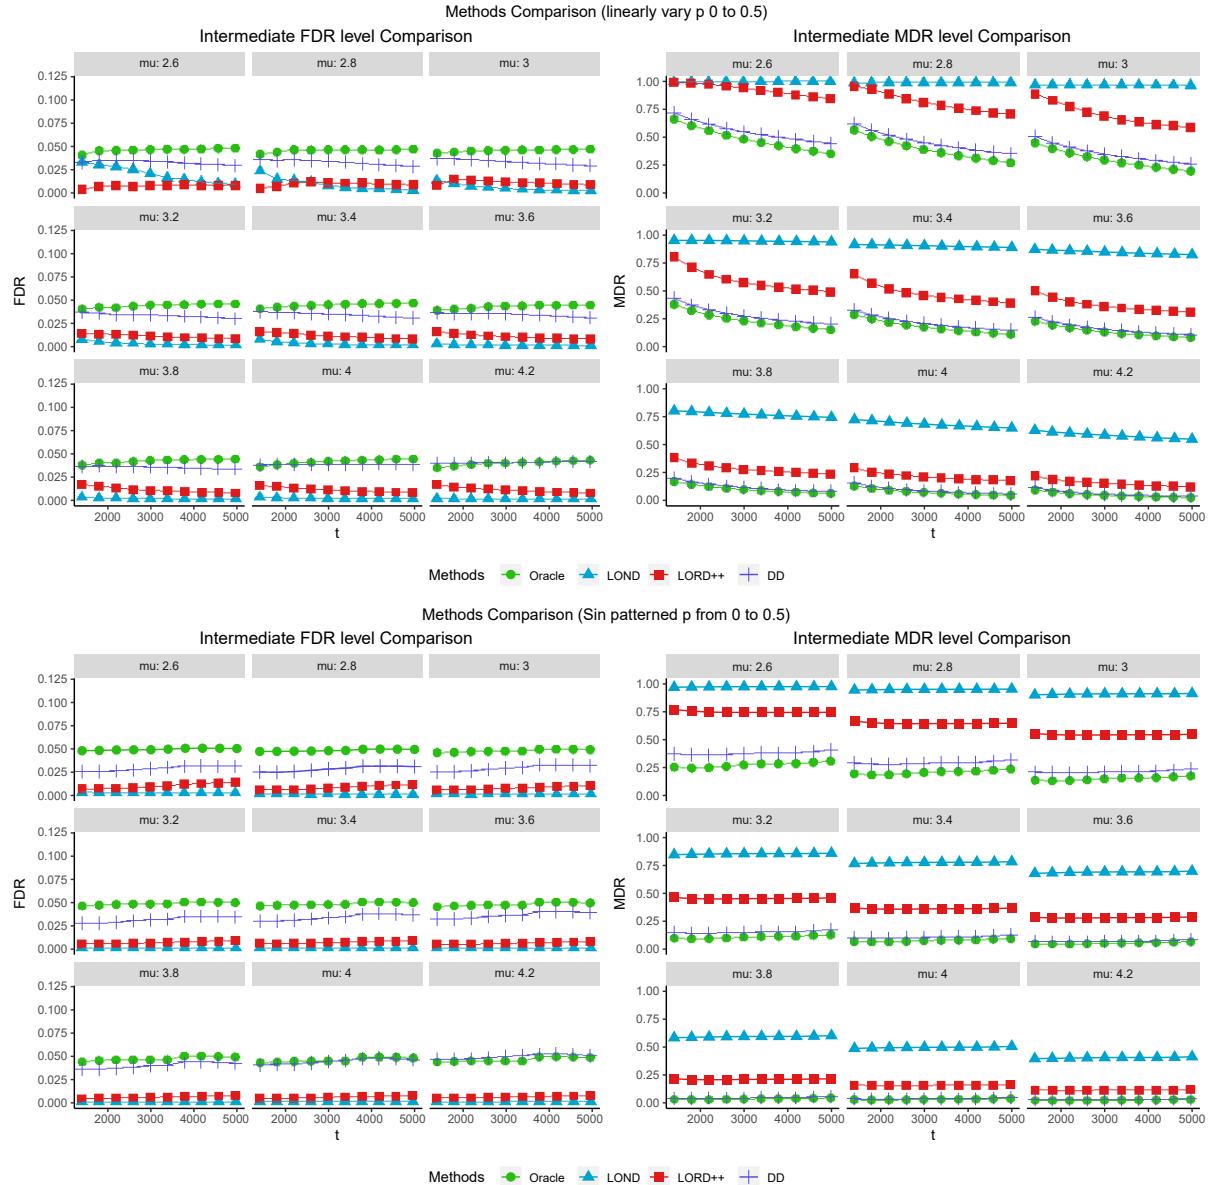


Figure 3: Simulation results for Settings 3 and 4: signal proportions are varied in linear and sine patterns, respectively. Our data-driven and oracle procedures provide significantly more power while controlling FDR under the nominal level in comparison with others.

- (b) SAST.DD is inferior compared to SAST.OR. This is largely due to the conservativeness of the estimator $\hat{\gamma}_t^{\text{DD}}$. The gap in the performances between SAST.DD and SAST.OR narrows as the signal strength becomes stronger, in which situation the estimator $\hat{\gamma}_t^{\text{DD}}$ becomes more precise.
- (c) In general LOND can be much improved by LORD++, which can be further improved by SAST.DD. The gap in power performances between SAST.DD and LORD++ narrows as the signal strength becomes stronger, in which situation it is easier to separate the signals from null cases.
- (d) When $\hat{\gamma}_t$ is fixed over time, the signals arrive at a constant rate and there is no informative structural information in the data stream (Setting 2: constant pattern). SAST.DD still outperforms LOND and LORD++ because our AI framework based on Clfdr precisely characterizes the gains and losses of different decisions; this not only leads to more precise FDR control but also optimizes the alpha-wealth allocation in the online setting.

4.3 Effects of the barrier

This section presents a toy example to illustrate that the barrier, which aims to prevent the “piggybacking” issue (Ramdas et al. 2017), can greatly reduce the MDR by allocating existing alpha-wealth in a more cost-effective way. Consider the previous block structured setting (Setting 1 in Section 4.2). Figure 4 shows the FDR and MDR comparisons for the following methods at FDR level $\alpha = 0.05$: (i) oracle SAST rule (OR); (ii) oracle SAST rule with no barrier (OR_nob); (iii) data-driven SAST rule with estimated parameters (DD, Section 3); (iv) data-driven SAST rule with no barrier (DD_nob).

We can see from the comparison that although the FDR levels between the two oracle methods are roughly the same, the MDR levels are greatly reduced by incorporating the barrier (hence the alpha-wealth is invested more efficiently). The same patterns can be observed for the two data-driven procedures.

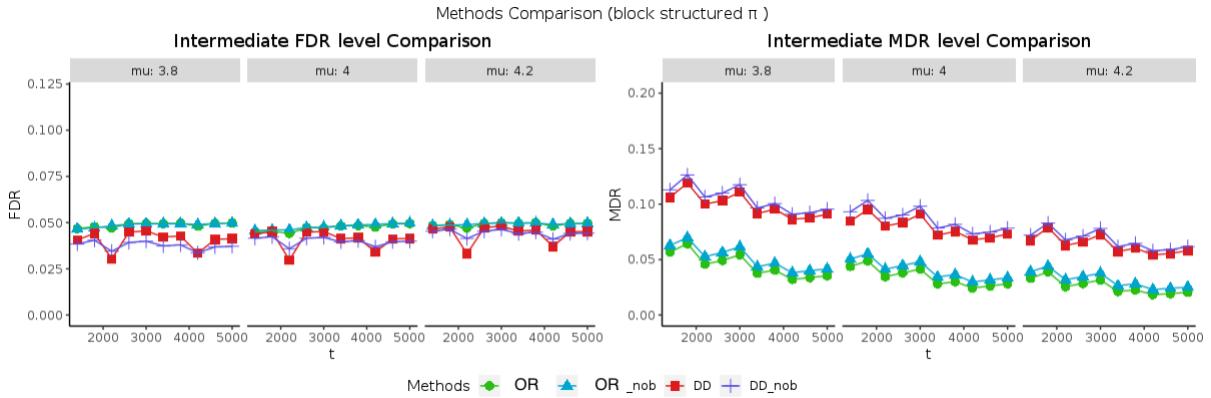


Figure 4: The incorporation of the barrier greatly reduces the MDR levels.

5 Applications

Online FDR rules are useful for a wide range of scenarios. We discuss two applications, respectively for anomaly detection in large-scale time series data and genotype discovery under the QPD framework.

5.1 Time series anomaly detection

The NYC taxi dataset can be downloaded from the Numenta Anomaly Benchmark (NAB) repository (Ahmad et al., 2017), which contains useful tools and datasets for evaluating algorithms for anomaly detection in streaming, real-time applications. The dataset records the counts of NYC taxi passengers every 30 minutes from July 1, 2014 to January 31, 2015, during which period five known anomalies had occurred (the NYC marathon, Thanksgiving, Christmas, New Years day and a snow storm). In Figure 5, we plot the time series, with the known anomalous intervals displayed in red rectangles.

We formulate the anomaly detection problem as an online sequential multiple testing problem. The basic setup can be described as follows. The null hypothesis H_t corresponds to no anomaly at time t . We claim that an anomaly occurs at t if H_t is rejected. A rejection within the red intervals is considered to be a true discovery.

The application of online FDR rules requires summarizing the stream of counts data as a sequence of p-values or CLfdr statistics. However, directly calculating the p-values based



Figure 5: NYC Taxi passenger count time series from July 1st 2014 to Jan 31st 2015. Blue lines are Loess smoothed time series indicating the overall trend change.

on this dataset would be problematic as the data demonstrate strong trend and seasonality patterns. We first use the R package `stlplus` to carry out an STL decomposition (Seasonal Trend decomposition using Loess smoothing; [Cleveland et al., 1990](#)) to remove the seasonal and trend components. The residuals, displayed in the top 3 rows of Figure 6, are standardized and modeled using a two-component mixture (2.1). However, as can be seen from the histogram at the bottom of Figure 6, the null distribution is approximately normal but deviates from a standard normal. Following the method in [Jin and Cai \(2007\)](#), we estimate the empirical null distribution as $N(0.028, 0.618)$. We apply the BH (pretending all observations are seen at once), LOND, LORD++ and SAST.DD at FDR level 0.0001. For the SAST.DD method, the neighborhood size d and initial burn-in period are both chosen to be 500. In calculating the Clfdr, $f_0(x)$ is taken as the density of the estimated empirical null \hat{F}_0 . Moreover, the p-values are obtained by the formula $P_i = 2\hat{F}_0(|Z_i|)$, where z-scores are computed based on the residuals. Figure 7 summarizes the anomaly points detected by different methods.

We can see that for the several anomaly time periods labeled, SAST can detect more points than other methods. Table 1 summarizes the total number of rejections within the labeled time windows. It may appear counter-intuitive that SAST, being an online procedure, rejects

Method	Number of Discoveries
Off-line BH procedure	179
Online SAST.DD (Proposed)	201
Online LOND	137
Online LORD++	178

Table 1: Number of discoveries made by various online and off-line FDR procedures for the NYC taxi dataset, nominal FDR level at 0.0001.

more null hypotheses than the off-line BH procedure. The reason is that the anomalies tend to appear in clusters. This structural information is captured by the Clfdr statistic, which forms the building block of SAST and leads to improved power in detecting structured signals (Section 2.3).

5.2 IMPC dataset Genotype Discovery

In this section, we demonstrate the SAST procedure on a real dataset from the International Mouse Phenotyping Consortium (IMPC). This dataset, which has been analyzed in Karp et al. (2017), involves a large study to functionally annotate every protein coding gene by exploring the impact of gene knockouts. This dataset and resulting family of hypotheses are constantly growing as new results come in. Karp et al. (2017) tested both the roles of genotype and sex as modifiers of genotype effects, resulting in two sets of p-values: one set for testing genotype effects, and the other for sexual dimorphism. This dataset has been widely used for comparing online FDR algorithms. Currently it is available as one of the application datasets in the R-package `OnlineFDR` that implements methods such as LORD, LOND and LORD++. In order to implement our proposed SAST procedure, we need the original z-scores instead of p-values. However, the directions of effects cannot be determined based on p-value alone. Hence, we transform the p-values into z-scores by introducing a Bernoulli random variable to ensure asymptotic symmetry around 0: $z = X^{-1}(p/2) - (1 - X)^{-1}(p/2)$, where $X \sim \text{Ber}(0.5)$ ⁵.

Table 2 summarizes the total number of discoveries made by each method. We can see that SAST makes more discoveries than other alpha-investing methods. Similar to the analysis

⁵We recommend that in the future the biomedical community should report, in addition to p-values, the effect sizes. Thus we also know the direction and magnitude of an interesting signal. In fact, converting z-scores to p-values may lead to loss of information (cf. Suh and Cai, 2007).

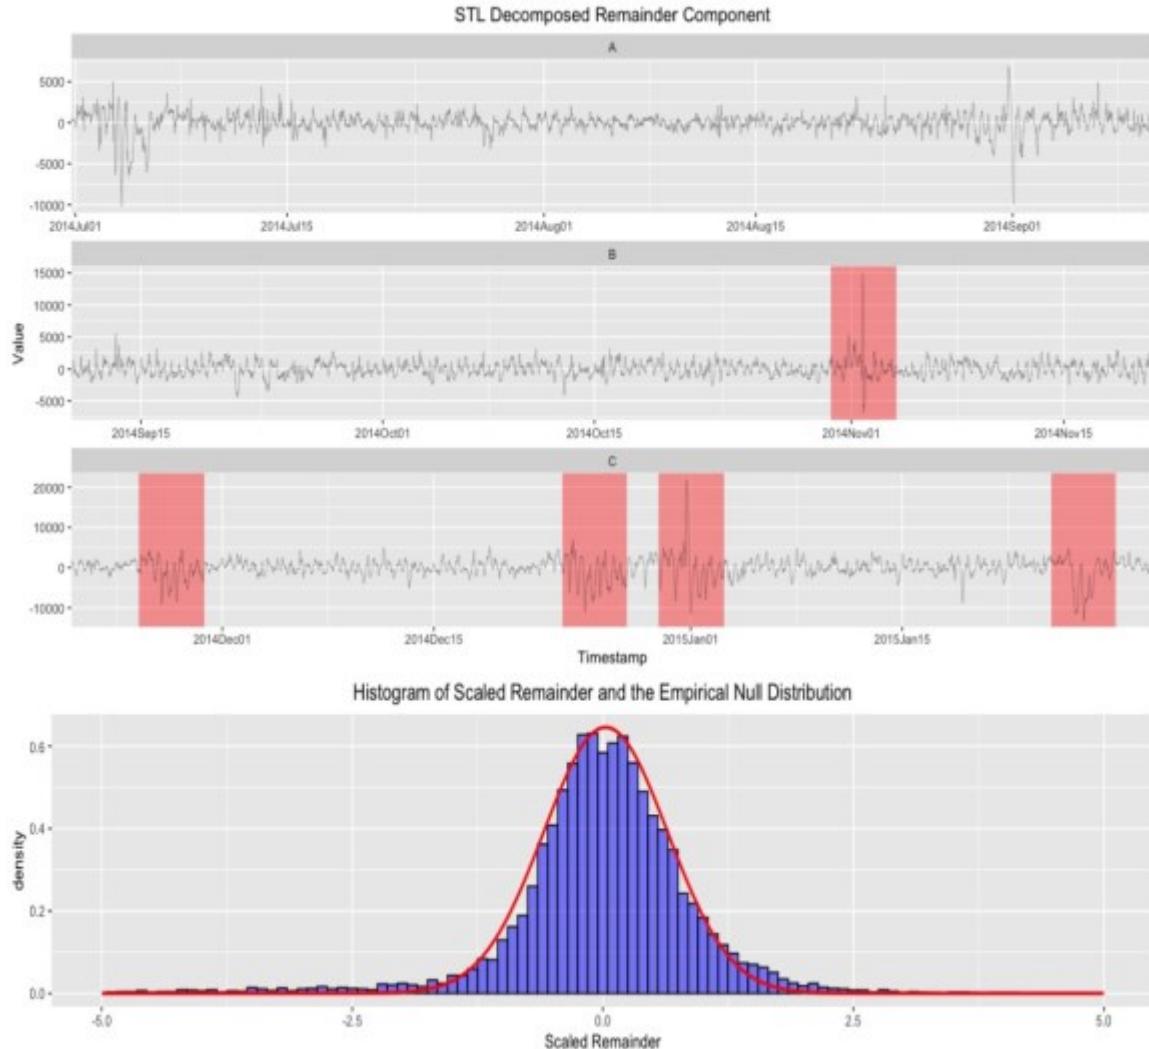


Figure 6: Top three rows: Time series of remainder component from STL decomposition with the known anomaly regions marked in red rectangles. Bottom row: Histogram of the remainder term from STL decomposition, the red curve indicates the estimated empirical null distribution $N(0.028, 0.618)$.

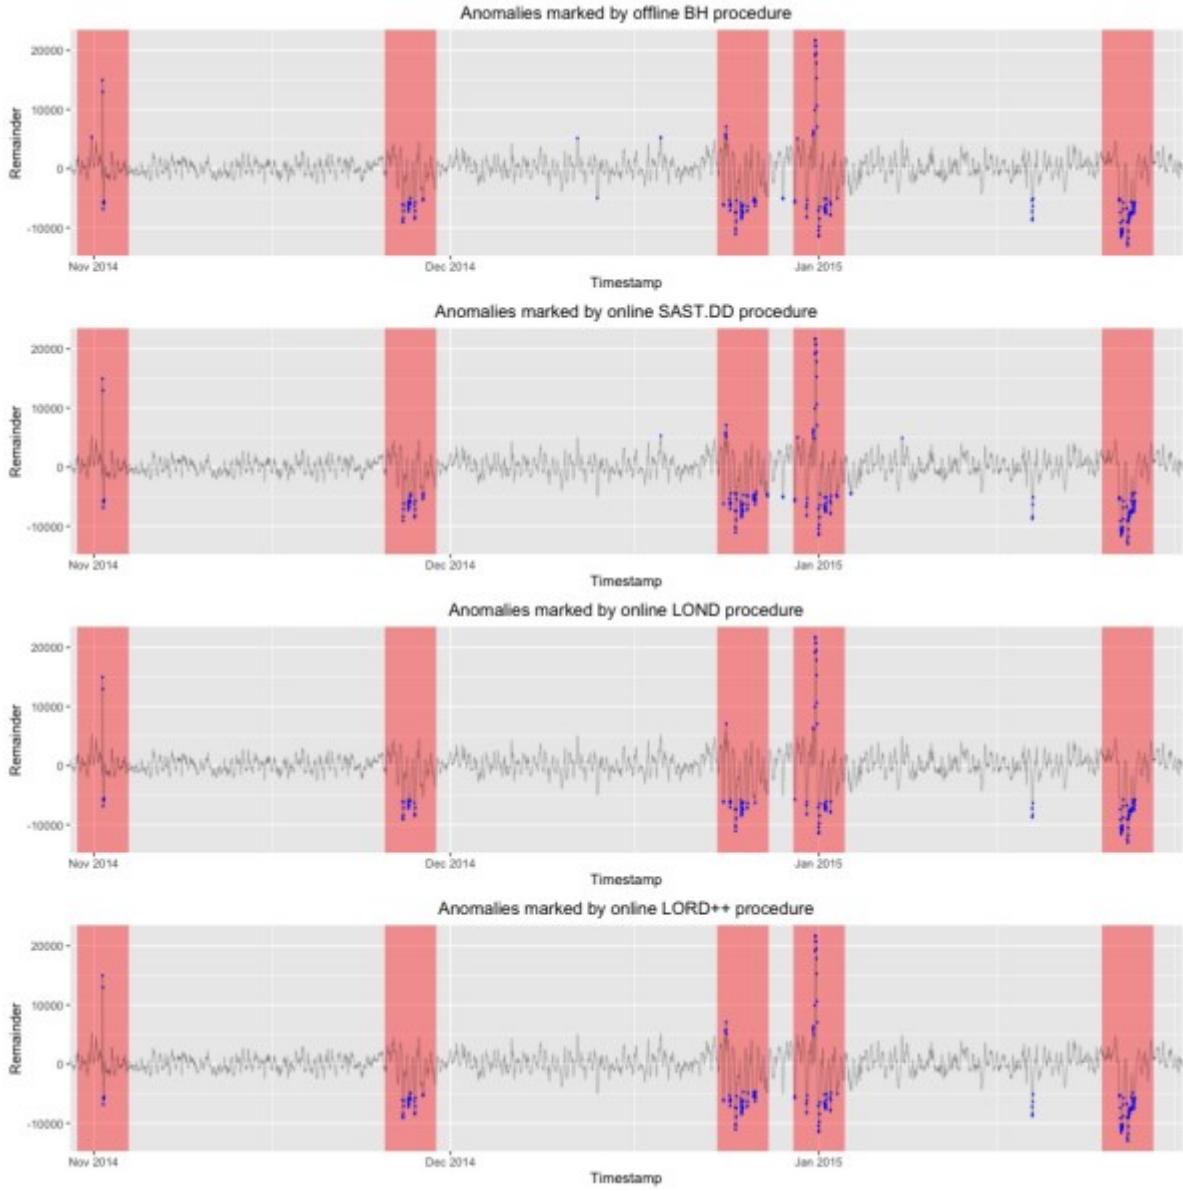


Figure 7: Anomaly points detected by various algorithms, our data-driven SAST procedure detects the most anomaly points within the labeled window marked by red rectangles. Nominal significance level chosen as 0.0001.

in Section 5, SAST rejects more hypotheses than the offline BH procedure. One possible explanation is that Clfdr is more powerful than p-values since it captures useful structural information in the data stream.

Method	Genotype	Method type
SAST	12975	online
BH	12907	online
LORD++	8517	online
LOND	2905	online
Fixed threshold 0.0001	4158	online

Table 2: Number of discoveries made by various online and offline FDR procedures for the IMPC dataset on Genotypes, nominal FDR level at 0.05.

References

Aharoni, E., Neuvirth, H., and Rosset, S. (2010). The quality preserving database: A computational framework for encouraging collaboration, enhancing power and controlling false discovery. *IEEE/ACM transactions on computational biology and bioinformatics*, 8(5):1431–1437.

Aharoni, E. and Rosset, S. (2014). Generalized ϵ -investing: definitions, optimality results and application to public databases. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(4):771–794.

Ahmad, S., Lavin, A., Purdy, S., and Agha, Z. (2017). Unsupervised real-time anomaly detection for streaming data. *Neurocomputing*, 262:134–147.

Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. B*, 57:289–300.

Cai, T. T. and Sun, W. (2009). Simultaneous testing of grouped hypotheses: Finding needles in multiple haystacks. *J. Amer. Statist. Assoc.*, 104:1467–1481.

Cai, T. T., Sun, W., and Wang, W. (2019). CARS: Covariate assisted ranking and screening for large-scale two-sample inference (with discussion). *J. Roy. Statist. Soc. B*, 81(2):187–234.

Cleveland, R. B., Cleveland, W. S., McRae, J. E., and Terpenning, I. (1990). Stl: a seasonal-trend decomposition. *Journal of official statistics*, 6(1):3–73.

Efron, B. (2004). Large-scale simultaneous hypothesis testing: The choice of a null hypothesis. *Journal of the American Statistical Association*, 99(465):96–104.

Foster, D. P. and Stine, R. A. (2008). ℓ -investing: a procedure for sequential control of expected false discoveries. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(2):429–444.

Genovese, C. R., Roeder, K., and Wasserman, L. (2006). False discovery control with p-value weighting. *Biometrika*, 93(3):509–524.

Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, pages 65–70.

Hu, J. X., Zhao, H., and Zhou, H. H. (2010). False discovery rate control with groups. *Journal of the American Statistical Association*, 105(491):1215–1227.

Javanmard, A., Montanari, A., et al. (2018). Online rules for control of false discovery rate and false discovery exceedance. *The Annals of statistics*, 46(2):526–554.

Jin, J. and Cai, T. T. (2007). Estimating the null and the proportion of nonnull effects in large-scale multiple comparisons. *J. Amer. Statist. Assoc.*, 102:495–506.

Karp, N. A., Mason, J., Beaudet, A. L., Benjamini, Y., Bower, L., Braun, R. E., Brown, S. D., Chesler, E. J., Dickinson, M. E., Flenniken, A. M., et al. (2017). Prevalence of sexual dimorphism in mammalian phenotypic traits. *Nature communications*, 8:15475.

Lei, L. and Fithian, W. (2018). Adapt: an interactive procedure for multiple testing with side information. *J. Roy. Statist. Soc. B*, 80(4):649–679.

Li, A. and Barber, R. F. (2019). Multiple testing with the structure-adaptive benjamini-hochberg algorithm. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(1):45–74.

Lynch, G., Guo, W., Sarkar, S. K., Finner, H., et al. (2017). The control of the false discovery rate in fixed sequence multiple testing. *Electronic Journal of Statistics*, 11(2):4649–4673.

Ramdas, A., Yang, F., Wainwright, M. J., and Jordan, M. I. (2017). Online control of the false discovery rate with decaying memory. In *Advances In Neural Information Processing Systems*, pages 5650–5659.

Ramdas, A., Zrnic, T., Wainwright, M., and Jordan, M. (2018). Sadron: an adaptive algorithm for online control of the false discovery rate. In *International Conference on Machine Learning*, pages 4286–4294.

Robertson, D. S. and Wason, J. (2018). Online control of the false discovery rate in biomedical research. *arXiv preprint arXiv:1809.07292*.

Silverman, B. W. (1986). *Density estimation for statistics and data analysis*, volume 26. CRC press.

Stein, M. L. (2012). *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media.

Sun, W. and Cai, T. T. (2007). Oracle and adaptive compound decision rules for false discovery rate control. *J. Amer. Statist. Assoc.*, 102:901–912.

Wand, M. P. and Jones, M. C. (1994). *Kernel smoothing*. Chapman and Hall/CRC.

Xia, Y., Cai, T. T., and Sun, W. (2020+). Gap: A general framework for information pooling in two-sample sparse inference. *J. Am. Statist. Assoc.*, to appear.

Online Supplementary Material for “Structure-Adaptive Sequential Testing for Online False Discovery Rate Control”

This supplement contains the proofs of main theorems (Section A), other theoretical results (Section B), and optimality theory on simultaneous testing (Section C).

A Proof of main theorems

A.1 Proof of Theorem 1

Note that the Clfdr is defined as $\text{Clfdr}_i = P(\check{Y}_i = 0 | X_i)$. Then by the definition of FDR and double expectation theorem, we have:

$$\text{FDR}_t = \mathbb{E}_X \left(\frac{1}{|R_t| - 1} \sum_{i \in R_t} \text{Clfdr}_i \right).$$

By construction of the decision rule, $(|R_t| - 1)^{-1} \sum_{i \in R_t} \text{Clfdr}_i \leq \epsilon$ for all realization of X . It follows that $\text{FDR}_t \leq \epsilon$.

A.2 Proof of Theorem 2

We need the following lemma:

Lemma 1. Suppose $a_n \geq 0$ and $|a_n|$ is bounded for all n , then $\lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n a_i}{n} \geq 0$.

The proof of lemma 1 is elementary thus omitted. By definition of our algorithm, if $\text{Clfdr}_t \leq \epsilon$ then $t = 1$. Note that, for any $\epsilon' > 0$,

$$\begin{aligned} \mathbb{P} \left(\sum_{i=1}^{|R_t|} \mathbb{I}(\text{Clfdr}_i \leq \epsilon) < 1 \right) &= \mathbb{P} \left(\sum_{i=1}^{|R_t|} \mathbb{I}(\text{Clfdr}_i \leq \epsilon) < M \right) \\ &\leq \mathbb{P} \left(\sum_{i=1}^M \mathbb{I}(\text{Clfdr}_i \leq \epsilon) < M \right). \end{aligned}$$

Note that Clfdr_i is a random variable from random mixture model (2.1) with a non-vanishing proportion of nonzero signals, we have

$$P \left(\bigcup_{i=1}^n I(\text{Clfdr}_i^{\frac{1}{2}} < \epsilon) < M \right) = 0$$

for every M . We have $P \left\{ \bigcup_{i=1}^n I(\text{Clfdr}_i^{\frac{1}{2}} < \epsilon) < 1 \right\} = 0$. Now, $\bigcup_{i=1}^n I(\text{Clfdr}_i^{\frac{1}{2}} < \epsilon) < 1$ would imply $|\text{Clfdr}_i - \text{Clfdr}_i^{\frac{1}{2}}| > \epsilon$ infinitely many times. By Proposition 2, $P(|\text{Clfdr}_i - \text{Clfdr}_i^{\frac{1}{2}}| > \epsilon) > 0$. It follows that $P \left(\bigcup_{i=1}^n I(\text{Clfdr}_i^{\frac{1}{2}} < \epsilon) < 1 \right) = 0$, hence $|R_t| \leq 1$. By Proposition 2

and Lemma 1, we have

$$P \left(\bigcup_{i \in R_t} I(\text{Clfdr}_i - \text{Clfdr}_i^{\frac{1}{2}}) \neq \emptyset \right) = 0.$$

Finally, the operation of Algorithm 2 implies that $\frac{P \left(\bigcup_{i \in R_t} I(\text{Clfdr}_i - \text{Clfdr}_i^{\frac{1}{2}}) \neq \emptyset \right)}{|R_t|} \rightarrow 0$, It follows that

$$\text{FDR}(\cdot) = E_x \left[\frac{P \left(\bigcup_{i \in R_t} I(\text{Clfdr}_i - \text{Clfdr}_i^{\frac{1}{2}}) \neq \emptyset \right)}{|R_t|} \right] = E_x \left[\frac{P \left(\bigcup_{i \in R_t} I(\text{Clfdr}_i^{\frac{1}{2}}) \neq \emptyset \right)}{|R_t|} \right] + o(1) \rightarrow 0 + o(1).$$

A.3 Proof of theorem 3

Note when both f_t and \hat{f}_t are fixed over time, the Clfdr statistic reduces to $\text{Lfdr}_i = \frac{(1 - \hat{f}_t(x_i))}{f_t(x_i)}$.

The optimal threshold in the online simultaneous testing setup would be independent of time t and the chosen neighborhood. The oracle online rule coincides with the oracle procedure described in Section 3.2 of Sun and Cai (2007).

We now introduce some notations:

- $\hat{U}^t(\cdot) = t^{-1} \sum_{i=1}^t I(\text{Clfdr}_{(i)} < \cdot)$
- $U^t(\cdot) = t^{-1} \sum_{i=1}^t I(\text{Clfdr}_{(i)}^{\frac{1}{2}} < \cdot)$
- $U_1^t(\cdot) = E\{I(\text{Clfdr}_{(1)}^{\frac{1}{2}} < \cdot)\}$.
- $\tau_1 = \sup\{t \in (0, 1), U_1^t(\cdot) \geq 0\}$ is the “ideal” threshold.

Note that \hat{U}^t is discrete. To facilitate the theoretical analysis, we define, for $\text{Clfdr}_{(i)} < \cdot <$

$\hat{C}Ifdr_{(i+1)}$, a continuous version of \hat{U}^t :

$$\hat{U}_C^t(\cdot) = \frac{CIfdr_{(i)}}{CIfdr_{(i+1)}} \hat{U}_i^t + \frac{CIfdr_{(i+1)}}{CIfdr_{(i+1)}} \hat{U}_{i+1}^t,$$

where $\hat{U}_i^t = \hat{U}^t(CIfdr_{(i)})$. It is easy to verify that \hat{U}_C^t is continuous and monotone. Hence its inverse \hat{U}_C^{t-1} is well defined, continuous and monotone.

Next we show the following two results in turn: (i) $\hat{U}^t(\cdot) \xrightarrow{P} U_1^t(\cdot)$ and (ii) $\hat{U}_C^{t-1}(0) \xrightarrow{P} 1$.

Proof of (i). Note that $U^t(\cdot) \xrightarrow{P} U_1^t(\cdot)$ by the WLLN, so that we only need to establish that $\hat{U}^t(\cdot) \xrightarrow{P} U^t(\cdot)$. We need to following lemma:

Lemma 2. Let $V_i = \frac{\downarrow}{\downarrow} (CIfdr_i^{\frac{1}{2}} - \epsilon) I\{CIfdr_i^{\frac{1}{2}} < \cdot\}$ and $\hat{V}_i = \frac{\downarrow}{\downarrow} (CIfdr_i - \epsilon) I\{CIfdr_i < \cdot\}$. Then $E \hat{V}_i - V_i = o(1)$.

Proof of Lemma 2. Using the definitions of \hat{V}_i and V_i , we can show that

$$\begin{aligned} \hat{V}_i - V_i &= \frac{\downarrow}{\downarrow} (CIfdr_i - CIfdr_i^{\frac{1}{2}}) I\{CIfdr_i^{\frac{1}{2}} < \cdot, CIfdr_i^{\frac{1}{2}} < \cdot\} + \frac{\downarrow}{\downarrow} (CIfdr_i - \epsilon) I\{CIfdr_i < \cdot, CIfdr_i > \cdot\} \\ &\quad + (CIfdr_i^{\frac{1}{2}} - \epsilon)^2 I\{CIfdr_i > \cdot, CIfdr_i^{\frac{1}{2}} < \cdot\}. \end{aligned}$$

Let us refer to the three sums on the right hand as I, II, and III respectively. By step 2 in the proof of Theorem 2, $I = o(1)$. Then let $\epsilon > 0$, and consider that

$$\begin{aligned} P \{CIfdr_i > \cdot, CIfdr_i^{\frac{1}{2}} < \cdot\} &= P \{CIfdr_i > \cdot, CIfdr_i^{\frac{1}{2}} < \cdot, \epsilon < CIfdr_i < \cdot\} + P \{CIfdr_i > \cdot, CIfdr_i^{\frac{1}{2}} < \cdot, CIfdr_i < \epsilon\} \\ &\quad + P \{CIfdr_i^{\frac{1}{2}} < \cdot, CIfdr_i < \epsilon\} + P \{CIfdr_i > \cdot, CIfdr_i < \epsilon\} \end{aligned}$$

The first term on the right hand is vanishingly small as $\epsilon \rightarrow 0$ because $CIfdr_i$ is a continuous random variable. The second term converges to 0 by Proposition 2. Noting that $0 \leq CIfdr_i \leq 1$, we conclude $II = o(1)$. In a similar fashion, we can show that $III = o(1)$, thus proving the lemma.

Let $S_t = \sum_{i=1}^t (\hat{V}_i - V_i)$, by Lemma 2 and the Cauchy-Schwartz inequality,

$$E \left| \sum_{i=1}^n \hat{V}_i - V_i \right|^2 = o(1).$$

It follows that

$$\begin{aligned} \text{Var} \left(\sum_{i=1}^t S_i \right) &= t^{-2} \text{Var}(S_t) \sum_{i=1}^t E \left| \hat{V}_i - V_i \right|^2 \\ &+ O\left(\frac{1}{t^2} \sum_{i,j:i=j} \left| \hat{V}_i - V_i \right| \left| \hat{V}_j - V_j \right| \right) \\ &= o(1). \end{aligned}$$

By Proposition 2, $E \left| \sum_{i=1}^t S_i \right| \neq 0$, applying Chebyshev's inequality, we obtain

$$\sum_{i=1}^t S_i = \hat{U}_C^t \neq 0,$$

establishing (i).

Proof of (ii). Since \hat{U}_C^t is continuous, for any $\epsilon > 0$, we can find $\delta > 0$ such that $\hat{U}_C^{t-1}(0) - \hat{U}_C^{t-1}(\hat{U}_C^t(\epsilon)) < \epsilon$ if $\hat{U}_C^t(\epsilon) < \delta$. It follows that

$$P \left| \hat{U}_C^t(\epsilon) \right| > \delta \quad P \left| \hat{U}_C^{t-1}(0) - \hat{U}_C^{t-1}(\hat{U}_C^t(\epsilon)) \right| > \epsilon.$$

Proposition 2 and the WLLN imply that $\hat{U}_C^t(\epsilon) \xrightarrow{P} U_1^t(\epsilon)$. Note that $U_1^t(\epsilon) = 0$, then,

$$P \left| \hat{U}_C^t(\epsilon) \right| > \delta \neq 0.$$

Hence, we have

$$\hat{U}_C^{t-1}(0) \xrightarrow{P} \hat{U}_C^{t-1}(\hat{U}_C^t(\epsilon)) \xrightarrow{P} 0, \quad (A.1)$$

completing the proof of (ii).

B Proof of propositions

B.1 Proof of Proposition 1

Let $A_{\alpha} = \{x : P_0(x) > \alpha\}$, where $P_0(x)$ is the p-value of x . Then

$$\begin{aligned}
 (1 - \alpha)^{-1} P(P_t > \alpha) &= (1 - \alpha)^{-1} \int_{A_{\alpha}}^Z f_0(x)(1 - \hat{f}_t) + \hat{f}_t f_{1t}(x) dx \\
 &\quad (1 - \alpha)^{-1} \int_{A_{\alpha}}^Z f_0(x)(1 - \hat{f}_t) dx \\
 &= (1 - \hat{f}_t).
 \end{aligned}$$

Hence $\hat{f}_t = 1 - (1 - \alpha)^{-1} P(P_t > \alpha) \geq 1 - (1 - \hat{f}_t) = \hat{f}_t$. By definition of Clfdr_t , we have

$$\text{Clfdr}_t^{\alpha} \leq \text{Clfdr}_t.$$

Let FDR_{OR}^{α} be the decision rule described in Algorithm 1 with Clfdr_t^{α} used in place of Clfdr_t .

Let R be the index set of hypotheses rejected by FDR_{OR}^{α} . The FDR of FDR_{OR}^{α} is

$$\begin{aligned}
 \text{FDR}(\text{FDR}_{OR}^{\alpha}) &= E \left[\frac{\sum_{i \in R} (1 - \sqrt{\alpha_i})}{|R| - 1} \right] \\
 &= E_x \left[E \left[\frac{\sum_{i \in R} (1 - \sqrt{\alpha_i})}{|R| - 1} \right] \right] \\
 &= E_x \left[\frac{1}{|R| - 1} \sum_{i \in R} \text{Clfdr}_i \right].
 \end{aligned}$$

Since $\text{Clfdr}_t^{\alpha} \leq \text{Clfdr}_t$, it follows that

$$\text{FDR}(\text{FDR}_{OR}^{\alpha}) \leq E_x \left[\frac{1}{|R| - 1} \sum_{i \in R} \text{Clfdr}_i^{\alpha} \right] \leq \text{FDR}_t^{\alpha}.$$

The last inequality is due to the definition of FDR_{OR}^{α} which guarantees that

$$\frac{1}{|R| - 1} \sum_{i \in R} \text{Clfdr}_i^{\alpha} \leq \text{FDR}_t^{\alpha}.$$

B.2 Proof of Proposition 2

Under the in-fill model, we write

$$\hat{f}_t(x) = \frac{\sum_{j=t-1}^{t-1} K_{h_t}(1 - |j/t - x|) K_{h_x}(x_j - x)}{\sum_{j=t-1}^{t-1} K_{h_t}(1 - |j/t - x|)}.$$

We first state 3 lemmas that will be proved in turn.

Lemma 3. Under the assumption of Proposition 2, $E \int_{-\infty}^{\infty} (\hat{f}_t(x) - f_t(x))^2 dx \rightarrow 0$.

Lemma 4. Under the assumptions of Proposition 2

$$E \int_{-\infty}^{\infty} (\hat{f}_t(x) - \hat{f}_t(x))^2 dx \rightarrow 0.$$

Lemma 5. Let \hat{f}_t , f_t , and f_0 be estimates such that $E \int_{-\infty}^{\infty} (\hat{f}_t(x) - f_t(x))^2 dx \rightarrow 0$, $E \int_{-\infty}^{\infty} (\hat{f}_t(x) - f_0(x))^2 dx \rightarrow 0$, and then $E \int_{-\infty}^{\infty} (\hat{f}_t(x) - f_t(x))^2 dx \rightarrow 0$.

By Lemma 3 and Lemma 4 together with the fact that f_0 is known, it follows from Lemma 5 that $E \int_{-\infty}^{\infty} (\hat{f}_t(x) - f_t(x))^2 dx \rightarrow 0$. Since convergence in second order mean implies convergence in probability, we have

$$E \int_{-\infty}^{\infty} (\hat{f}_t(x) - f_t(x))^2 dx \rightarrow 0.$$

B.3 Growing domain version of Proposition 2

In the growing domain framework, Proposition 2 takes the following form:

Proposition 3. Suppose:

(A1'): For any $\epsilon > 0$, $\exists T$ such that for all integers i, j on the interval $t \in \frac{1}{\log(t)h_t} \cup \frac{1}{h_t} \cup \frac{1}{t}$, we have $\int_{-\infty}^R |f_i(x) - f_j(x)| dx < \epsilon$.

(A2'): $h_x \rightarrow 0$, $h_t \rightarrow t$ and $h_x h_t \rightarrow 1$.

(A3'): $f_j(x) < C$ and $\int_{-\infty}^R |f_j^{(0)}(x)| dx < C$ for all j .

(A4'): $d = ch_t$ for some $c > 0$.

We have $E \int_{-\infty}^{\infty} (\hat{f}_t(x) - f_t(x))^2 dx \rightarrow 0$.

The proof follows the same line as the proof of proposition 2 thus omitted.

B.4 Proof of Lemma 3

We first compute $E\hat{f}_t(x) - f_t(x)$. Note that $E K_{h_x}(X_j - x) = \int_{-\infty}^R K(z) f_j(x - h_x z) dz$. Using Taylor expansion, we have

$$f_j(x - h_x z) = f_j(x) - h_x z f_j'(x) + \frac{1}{2} h_x^2 z^2 f_j''(x) + o(h_x^2).$$

It follows that

$$E K_{h_x}(X_j - x) - f_t(x) = f_j(x) - f_t(x) + \frac{1}{2} h_x^2 f_j''(x) - \int_{-\infty}^R z^2 K(z) dz + o(h_x^2).$$

Let $A = \sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t) \frac{1}{2} h_x^2 f_j''(x) - \int_{-\infty}^R z^2 K(z) dz + o(h_x^2)$. Then

$$E\hat{f}_t(x) - f_t(x) = \frac{\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - \frac{j}{t}) \{f_j(x) - f_t(x)\} + A}{\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - \frac{j}{t})}.$$

$$\int_{-\infty}^R |E\hat{f}_t(x) - f_t(x)| dx = O \left(\frac{\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t) \int_{-\infty}^R |f_j(x) - f_t(x)| dx}{\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t)} \right) + o(h_x^2) \rightarrow 0.$$

To see why the last expression goes to 0, note that for any $\epsilon > 0$, by Assumption (A1), we can take ϵ such that for all $i > 1$ $\int_{-\infty}^R |f_i(x) - f_t(x)| dx < \epsilon$. Hence,

$$\begin{aligned} & \frac{\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t) \int_{-\infty}^R |f_j(x) - f_t(x)| dx}{\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t)} \\ &= \frac{\sum_{j=t}^{b(1)tc} \sum_{d+1}^n K_{h_t}(1 - j/t) \int_{-\infty}^R |f_j(x) - f_t(x)| dx}{\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t)} + \frac{\sum_{j=b(1)tc+1}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t) \int_{-\infty}^R |f_j(x) - f_t(x)| dx}{\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t)} \\ &\stackrel{\text{def}}{=} \frac{\sum_{j=t}^{b(1)tc} \sum_{d+1}^n K_{h_t}(1 - j/t)}{\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t)} + \\ &\quad \frac{\sum_{j=b(1)tc+1}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t)}{\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t)} \end{aligned}$$

Note that $h_t \neq 0$, we conclude that $\sum_{j=t}^{b(1)tc} \sum_{d+1}^n K_{h_t}(1 - j/t) = O \sum_{j=t}^{b(1)tc} K_{h_t}(x) dx \rightarrow 0$. Also, since $dh_t \neq 1$ as $t \neq 1$, we have $\sum_{j=t}^{t-1} \sum_{d+1}^n K_{h_t}(1 - j/t) = c^0 h_t^{-1}$ for some c^0 .

Thus $\frac{\sum_{j=t-d+1}^{t-1} K_{h_t}(1-j/t)}{\sum_{j=t-d+1}^{t-1} K_h(1-j/t)} \neq 0$, and

$$\lim_{t \rightarrow 1} \frac{\sum_{j=t-d+1}^{t-1} K_{h_t}(1-j/t) \int_{-R}^R |f_j(x) - f_t(x)| dx}{\sum_{j=t-d+1}^{t-1} K_{h_t}(1-j/t)} = 0.$$

It follows from the boundedness of $E\hat{f}_t$ and $f_t(x)$ that

$$\int_{-R}^R |E\hat{f}_t(x) - f_t(x)|^2 dx \neq 0. \quad (B.2)$$

Next we compute $\text{Var} \hat{f}_t(x)$:

$$\begin{aligned} \text{Var} \{K_{h_x}(X_j - x)\} &= \frac{1}{h_x} \int_{-Z}^Z K(z)^2 f_j(x - h_x z) dz = \{f_j(x) + o(1)\}^2 \\ &= \frac{1}{h_x} \int_{-Z}^Z K(z)^2 (f_j(x) + o(1)) dz = \{f_j(x) + o(1)\}^2 \\ &= O \left(\frac{1}{h_x} \int_{-Z}^Z K(z)^2 dz f_j(x) + o(h_x^{-1}) \right). \end{aligned}$$

Some additional calculations give

$$\begin{aligned} \text{Var} \hat{f}_t(x) &= \frac{\sum_{j=t-d+1}^{t-1} \{K_{h_t}(1-j/t)\}^2 O \left(\frac{1}{h_x} \int_{-R}^R K(z)^2 dz f_j(x) \right)}{\sum_{j=t-d+1}^{t-1} \{K_{h_t}(1-j/t)\}^2} \\ &= O \left(\frac{1}{4} (th_x)^{-1} \int_0^R K_{h_t}^2(x) dx \right)^2 \\ &= O \left(\frac{1}{4} \left(\int_0^{R/d} K_{h_t}(x) dx \right)^2 \right)^2 \\ &= O \left(\frac{1}{4} (th_x h_t)^{-1} \int_0^R K_{h_t}(x) dx \right)^2 \end{aligned}$$

Therefore, by assumption (A3) and (A4),

$$\int_{-R}^R \text{Var} \hat{f}_t(x) dx = O((th_x h_t)^{-1}) \neq 0. \quad (B.3)$$

Since $E \int_{-R}^R \{\hat{f}_t(x) - f_t(x)\}^2 dx = \int_{-R}^R \{E\hat{f}_t(x) - f_t(x)\}^2 dx + \text{Var} \{\hat{f}_t(x)\} dx$, (B.2) and (B.3) together imply that $E \int_{-R}^R \{\hat{f}_t(x) - f_t(x)\}^2 dx \neq 0$.

B.5 Proof of lemma 4

Define $\hat{P}(P_t > \alpha) := \frac{P_{i=1}^{t-1} K_{h_t}(1 - i/t)}{P_{i=t}^{t-1} K_{h_t}(1 - i/t)}$. Let P_t be the p-value of X_t . We will show

$$E \hat{P}(P_t > \alpha) = P(P_t > \alpha) \alpha^2 \neq 0. \quad (B.4)$$

We first rewrite the term

$$\frac{P_{i=1}^{t-1} K_{h_t}(1 - i/t)}{P_{i=t}^{t-1} K_{h_t}(1 - i/t)} = \frac{P_{i=1}^{t-1} K_{h_t}(1 - i/t) I(P(x_i) > \alpha)}{P_{i=t}^{t-1} K_{h_t}(1 - i/t)}.$$

By Lemma 3 we have, $E \int_B f_t(x) - f_t(x) dx \neq 0$ for every B . In particular, take $B = \{x : P(x) > \alpha\}$ use the definition of \hat{f}_t we have

$$\frac{P_{i=1}^{t-1} K_{h_t}(1 - i/t) E \int_{P(x) > \alpha} K_{h_x}(x_j - x) dx}{P_{i=t}^{t-1} K_{h_t}(1 - i/t)} \neq P(P_t > \alpha). \quad (B.5)$$

To show the lemma, it is sufficient to show

$$E \int_{P(x) > \alpha} K_{h_x}(x_j - x) dx \neq E \{I(P(x_j) > \alpha)\}. \quad (B.6)$$

To see why (B.6) implies (B.4), note that (B.6) implies

$$\frac{P_{i=1}^{t-1} K_{h_t}(1 - i/t) E \int_{P(x) > \alpha} K_{h_x}(x_j - x) dx}{P_{i=t}^{t-1} K_{h_t}(1 - i/t)} \neq \frac{P_{i=1}^{t-1} K_{h_t}(1 - i/t) E \{I(P(x_j) > \alpha)\}}{P_{i=t}^{t-1} K_{h_t}(1 - i/t)}. \quad (B.7)$$

Next note that

$$E \left(\frac{P_{i=1}^{t-1} K_{h_t}(1 - i/t) I(P(x_i) > \alpha)}{P_{i=t}^{t-1} K_{h_t}(1 - i/t)} \right) = \frac{P_{i=1}^{t-1} K_{h_t}(1 - j/t) E \{I(P(x_j) > \alpha)\}}{P_{i=t}^{t-1} K_{h_t}(1 - j/t)}, \text{ and}$$

$$\begin{aligned}
\text{Var} \left(\frac{P_{i=t}^{t-1} \int_{d+1}^{d+1} K_{h_t} (1 - |x_i|/t) I(P(x_i) > \underline{x})}{P_{i=t}^{t-1} \int_{d+1}^{d+1} K_{h_t} (1 - |x_i|/t)} \right) &= O(4t^{-1}) \int_0^2 \frac{R_1^{d/t} K_{h_t}^2(x) dx}{\int_0^2 K_{h_t}(x) dx} \stackrel{3}{=} 0.5 \\
&= O(4(t h_t)^{-1}) \int_0^2 K_{h_t}(x) dx \stackrel{3}{=} 0.5.
\end{aligned}$$

By (A4), we have $\int_0^2 K_{h_t}(x) dx \stackrel{0.2}{=} c$ for some constant $c > 0$. Now $th_t \neq 1$, implies $\text{Var} \left(\frac{P_{i=t}^{t-1} \int_{d+1}^{d+1} K_{h_t} (1 - |x_i|/t) I(P(x_i) > \underline{x})}{P_{i=t}^{t-1} \int_{d+1}^{d+1} K_{h_t} (1 - |x_i|/t)} \right) \neq 0$.

By Chebyshev's inequality,

$$\frac{P_{i=t}^{t-1} \int_{d+1}^{d+1} K_{h_t} (1 - |x_i|/t) I(P(x_i) > \underline{x})}{P_{i=t}^{t-1} \int_{d+1}^{d+1} K_{h_t} (1 - |x_i|/t)} \stackrel{p}{\rightarrow} \frac{P_{i=t}^{t-1} \int_{d+1}^{d+1} K_{h_t} (1 - |x_j|/t) E\{I(P(x_j) > \underline{x})\}}{P_{i=t}^{t-1} \int_{d+1}^{d+1} K_{h_t} (1 - |x_j|/t)}.$$

Combining (B.7), (B.5), (A1) and (A2),

$$\hat{P}(P_t > \underline{x}) = \frac{P_{i=t}^{t-1} \int_{d+1}^{d+1} K_{h_t} (1 - |x_i|/t) I(P(x_i) > \underline{x})}{P_{i=t}^{t-1} \int_{d+1}^{d+1} K_{h_t} (1 - |x_i|/t)} \stackrel{p}{\rightarrow} P(P_t > \underline{x}).$$

Therefore (B.4) follows.

We now show (B.6). Let $\underline{h}_x = \frac{p}{h_x}$. Write

$$E \int_{P(x) > \underline{x}} K_{h_x} (x_j - x) dx = E \int_{P(x) > \underline{x}, |x_j - x| < \underline{h}_x} K_{h_x} (x_j - x) dx + E \int_{P(x) > \underline{x}, |x_j - x| > \underline{h}_x} K_{h_x} (x_j - x) dx.$$

Use the normal tail bound,

$$\begin{aligned}
&E \int_{P(x) > \underline{x}, |x_j - x| > \underline{h}_x} K_{h_x} (x_j - x) dx \stackrel{\underline{x}}{\leq} \int_{|x_j - x| > \underline{h}_x} K_{h_x} (x_j - x) dx \\
&\stackrel{\underline{x}}{\leq} 2 \exp\{-1/(2h_x)\} \stackrel{0}{=} 0.
\end{aligned}$$

Define $A_{\epsilon} = \{x_j : P(x_j) > \epsilon\}$, let f_j be the density function for X_j . Note that

$$\begin{aligned}
 E \int_{|x| \leq \epsilon} K_{h_x}(x_j - x) dx &= \int_{Z^{-1}} \int_{\substack{|x_j - x| \leq \epsilon \\ P(x) > \epsilon}} K_{h_x}(x_j - x) f_j(x_j) dx dx_j \\
 &= \int_{Z^{A_{\epsilon}^c}} \int_{\substack{|x_j - x| \leq \epsilon \\ P(x) > \epsilon}} K_{h_x}(x_j - x) dx f_j(x_j) dx_j \\
 &= \int_{Z^{A_{\epsilon}^c}} [1 - O\{2 \exp(-1/(2h_x))\}] f_j(x_j) dx_j \\
 &= \int_{A_{\epsilon}^c} f_j(x_j) dx_j = E[I(P(x_j) > \epsilon)].
 \end{aligned}$$

Hence (B.6) is proved. The lemma follows.

B.6 Proof of lemma 5

Note that $f_t(x)$ is continuous and positive on the real line, then there exists $K_1 = [M, M]$ such that $P(x \in K_1^c) \leq 0$ as $M \geq 1$.

Let $\inf_{x \in K_1} f_t(x) = l_0$ and $A_{\epsilon}^{f_t} = \{x : |\hat{f}_t(x) - f_t(x)| \leq l_0/2\}$. Note that

$$E[\hat{f}_t(x) - f_t(x)]^2 \leq (l_0/2)^2 P(A_{\epsilon}^{f_t}), \text{ then } P(A_{\epsilon}^{f_t}) \leq 0,$$

we claim that f_t and \hat{f}_t are bounded below by a positive number for large t except for an event that has a low probability. Similar arguments can be applied to the upper bound of f_t and \hat{f}_t , as well as the cases for f_0 and \hat{f}_0 . Therefore, we conclude that f_0 , \hat{f}_0 , f_t , and \hat{f}_t are all bounded in the interval $[l_a, l_b]$, $0 < l_a < l_b < 1$ for large t except for an event A_{ϵ} that has probability tends to 0. Hence $0 < l_a < \inf_{x \in A_{\epsilon}} \min\{f_0, \hat{f}_0, f_t, \hat{f}_t\} < \sup_{x \in A_{\epsilon}^c} \max\{f_0, \hat{f}_0, f_t, \hat{f}_t\} < l_b < 1$. Next note that

$$\hat{f}_t f_t = \frac{f_0 f_t (\hat{f}_t - f_t) + (1 - \hat{f}_t) f_0 (\hat{f}_t - f_0) + (1 - \hat{f}_t) f_0 (f_t - \hat{f}_t)}{\hat{f}_t f_t},$$

we conclude that

$$\hat{f}_t f_t = C_1 (\hat{f}_t - f_t)^2 + C_2 f_0^2 + C_3 \hat{f}_t^2 \quad \text{in } A_{\epsilon}^c.$$

It is easy to see that $E\hat{C}lfdr_t^2 - Clfdr_t^2 k^2$ is bounded by some constant L , then

$$E\hat{C}lfdr_t^2 - Clfdr_t^2 k^2 \leq LP(A_{-}) + c_1 E\hat{k}_t^2 - \hat{k}_t^2 k^2 + c_2 E\hat{f}_t^2 - f_t^2 k^2 + c_3 E\hat{f}_0^2 - f_0^2 k^2.$$

According to the assumptions, we further have that for a given $\alpha > 0$, there exists $M \in \mathbb{Z}^+$ such that we can find A_{-} , $P(A_{-}) < \alpha/(4L)$, and at the same time $E\hat{k}_t^2 - \hat{k}_t^2 k^2 < \alpha/(4c_1)$, $E\hat{f}_t^2 - f_t^2 k^2 < \alpha/(4c_2)$, and $E\hat{f}_0^2 - f_0^2 k^2 < \alpha/(4c_3)$ for all $t \geq M$. Consequently, we have $E\hat{C}lfdr_t^2 - Clfdr_t^2 k^2 < \alpha$ for $t \geq M$, and the desired result follows.

C Optimality of the Clfdr rule in simultaneous testing

The optimality of the Clfdr rule in simultaneous testing is summarized in the following proposition. The idea in the proof essentially follows that in [Cai et al. \(2019\)](#). We provide it here for completeness.

Proposition 4. Consider a class of decision rules $\{\cdot\} = \{\{I(Clfdr_i < \cdot) : 1 \leq i \leq m\} : \cdot \in \mathbb{R}\}$ for simultaneous testing of hypotheses $\{H_i : i \in N_d(t)\}$ in the neighborhood of t . Denote $Q_{OR}(\cdot)$ the marginal FDR of $\{\cdot\}$. If $\alpha < Q_{OR}(1)$, then the oracle threshold $\alpha_{OR} = \sup\{\cdot : Q_{OR}(\cdot) \leq \alpha\}$ exists and is unique. Define the oracle rule $\alpha_{OR} = \{\{I(Clfdr_i \leq \alpha_{OR}) : 1 \leq i \leq m\} : \cdot \in \mathbb{R}\}$. Then α_{OR} is optimal for simultaneous testing in the sense that

$$mFDR(\alpha_{OR}) \leq \alpha, \quad ETP(\alpha_{OR}) \leq ETP(\alpha) \quad \text{for all } \alpha \text{ such that } mFDR(\alpha) \leq \alpha.$$

Proof. The proof has two parts. In (a), we establish two properties of the testing rule that thresholds the Clfdr at an arbitrary α , $\{\{I(Clfdr_i < \alpha) : 1 \leq i \leq m\} : \cdot \in \mathbb{R}\}$. We show that it produces $mFDR < \alpha$ for all α and that its $mFDR$ is monotonic in t . In (b) we show that when the threshold is α_{OR} , the testing rule, α_{OR} , exactly attains the $mFDR$ level and is optimal amongst all valid testing procedures controls $mFDR$ at level α .

Part(a). For the testing rule $\{\{I(Clfdr_i < \alpha) : 1 \leq i \leq m\} : \cdot \in \mathbb{R}\}$, let $Q_{OR}(\alpha) = \alpha$. We first show

that $\epsilon_i < \epsilon$. Since $\text{Clfdr}_i = P(\sqrt{i} = 0 | X_i = x_i)$, then

$$E \left(\sum_{i=1}^m (1 - \sqrt{i}) \right) = E \sum_{i=1}^m E_{\sqrt{i}|X} (1 - \sqrt{i}) = E \sum_{i=1}^m \text{Clfdr}_i \quad (C.8)$$

where notation E is the expected value taken over (X, \sqrt{i}) , notation $E_{\sqrt{i}|X}$ is the expectation taken over the distribution of (X) , and $E_{\sqrt{i}|X}$ is the expectation taken over \sqrt{i} , holding (X) fixed. We use (C.8) in the definition of $Q_{OR}(\epsilon)$ to get

$$E \sum_{i=1}^m (\text{Clfdr}_i - \epsilon) I(\text{Clfdr}_i < \epsilon) = 0. \quad (C.9)$$

The equality above implies that $\epsilon < \epsilon$. To see this, consider that all potentially non-zero terms arise when $\text{Clfdr}_i < \epsilon$, and when this is the case, either (i) $\epsilon \leq \text{Clfdr}_i < \epsilon$, (ii) $\text{Clfdr}_i < \epsilon$, or (iii) $\text{Clfdr}_i < \epsilon \leq \epsilon$. Notice (i) produces zero or positive terms on the LHS of (C.9), (ii) produces zero or negative terms, and (iii) produces negative terms. If $\epsilon = \epsilon$, then only (iii) is possible, which contradicts the RHS. Thus, the testing rule is valid.

Next, we show that $Q_{OR}(\epsilon)$ is nondecreasing in ϵ . That is, letting $Q(\epsilon_j) = \epsilon_j$, if $\epsilon_1 < \epsilon_2$, then $\epsilon_1 \leq \epsilon_2$. We argue by contradiction. Suppose that $\epsilon_1 < \epsilon_2$ but $\epsilon_1 > \epsilon_2$. First, it cannot be that $I(\text{Clfdr}_i < \epsilon_2) = 0$ for all i , because that implies $\epsilon_1 = \epsilon_2$ (both equal 0). Next, since $\epsilon_1 < \epsilon_2$,

$$(\text{Clfdr}_i - \epsilon_2) I(\text{Clfdr}_i < \epsilon_2) = (\text{Clfdr}_i - \epsilon_2) I(\text{Clfdr}_i < \epsilon_1) + (\text{Clfdr}_i - \epsilon_2) I(\epsilon_1 \leq \text{Clfdr}_i < \epsilon_2)$$

and rewrite $(\text{Clfdr}_i - \epsilon_2) I(\text{Clfdr}_i < \epsilon_1) = (\text{Clfdr}_i - \epsilon_1) I(\text{Clfdr}_i < \epsilon_1) + (\epsilon_1 - \epsilon_2) I(\text{Clfdr}_i < \epsilon_1)$.

If $\epsilon_2 < \epsilon_1$, then

$$\begin{aligned} & (\text{Clfdr}_i - \epsilon_2) I(\text{Clfdr}_i < \epsilon_2) - (\text{Clfdr}_i - \epsilon_1) I(\text{Clfdr}_i < \epsilon_1) + (\epsilon_1 - \epsilon_2) I(\text{Clfdr}_i < \epsilon_1) \quad (C.10) \\ & \quad + (\text{Clfdr}_i - \epsilon_1) I(\epsilon_1 \leq \text{Clfdr}_i < \epsilon_2). \end{aligned}$$

It follows that

$$E \sum_{i=1}^m (\text{Clfdr}_i - \epsilon_2) I(\text{Clfdr}_i < \epsilon_2) > 0.$$

To see this, consider the expectation of the sum over m tests for the three RHS terms of (C.10), which we reference as (i), (ii), and (iii) respectively. First, (i) is zero because of (C.9). Then for each $\text{Clfdr}_i < \alpha_2$, either (ii) is positive because $\alpha_2 < \alpha_1$, or (iii) is positive because $\alpha_1 < \alpha_2$.

However, (C.9) establishes that $E\left\{\sum_{i=1}^m (\text{Clfdr}_i - \alpha_2) I(\text{Clfdr}_i < \alpha_2)\right\} = 0$, leading to a contradiction. Hence, $\alpha_1 < \alpha_2$.

Part(b). The oracle threshold is defined as $\alpha_{\text{OR}} = \sup\{\alpha \in (0, 1) : Q_{\text{OR}}(\alpha) \geq \alpha\}$. First, let $\bar{\alpha} = Q_{\text{OR}}(1)$, which represents the largest mFDR level that the oracle testing procedure can be. By part (a), $Q_{\text{OR}}(\alpha_{\text{OR}})$ is non-decreasing. Via the squeeze theorem, for all $\alpha < \bar{\alpha}$, this implies that $Q_{\text{OR}}(\alpha_{\text{OR}}) = \alpha$.

Next, consider the power of $\alpha_{\text{OR}} = \{I(\text{Clfdr}_i < \alpha_{\text{OR}}) : 1 \leq i \leq m\}$ compared to that of an arbitrary decision rule $d_{\text{kr}} = (d_{\text{kr}}, \dots, d_{\text{kr}})$ such that $\text{mFDR}(d_{\text{kr}}) \geq \alpha$. Using the previous result

from part(a), it follows that

$$E\left(\sum_{i=1}^m (\text{Clfdr}_i - \alpha) I_{\alpha_{\text{OR}}}^i\right) = 0 \quad \text{and} \quad E\left(\sum_{i=1}^m (\text{Clfdr}_i - \alpha) d_{\text{kr}}^i\right) \geq 0.$$

Take the difference of the two expressions to obtain

$$E\left(\sum_{i=1}^m (\alpha_{\text{OR}} - d_{\text{kr}}^i) (\text{Clfdr}_i - \alpha)\right) \geq 0. \quad (\text{C.11})$$

Next apply a transformation $f(x) = (x - \alpha)/(1 - x)$ to each α_{OR} . Note that because $f'(x) = (1 - \alpha)/(1 - x)^2 > 0$, $f(x)$ is monotonically increasing. Then order is preserved: if $\text{Clfdr}_i < \alpha_{\text{OR}}$ then $f(\text{Clfdr}_i) < f(\alpha_{\text{OR}})$ and likewise for $\text{Clfdr}_i > \alpha_{\text{OR}}$. This means we can rewrite $\alpha_{\text{OR}} = I\{(\text{Clfdr}_i - \alpha)/(1 - \text{Clfdr}_i) < \alpha_{\text{OR}}\}$, where $\alpha_{\text{OR}} = (\alpha_{\text{OR}} - \alpha)/(1 - \alpha_{\text{OR}})$. It will be useful to note that, from part (a), we have $\alpha < \alpha_{\text{OR}} < 1$, which implies that $\alpha_{\text{OR}} > 0$.

Then,

$$E\left(\sum_{i=1}^m (\alpha_{\text{OR}} - d_{\text{kr}}^i) \{(\text{Clfdr}_i - \alpha) - \alpha_{\text{OR}}(1 - \text{Clfdr}_i)\}\right) \geq 0. \quad (\text{C.12})$$

To see this, consider that if $\alpha_{\text{OR}} - d_{\text{kr}}^i = 0$, then either (i) $\alpha_{\text{OR}} > d_{\text{kr}}^i$ or (ii) $\alpha_{\text{OR}} < d_{\text{kr}}^i$. If (i),

then $i_{OR} = 1$ and it follows that $\{(Clfdr_i - d^i)/(1 - Clfdr_i)\} < i_{OR}$. If (ii), then $i_{OR} = 0$ and $\{(Clfdr_i - d^i)/(1 - Clfdr_i)\} = i_{OR}$. For both cases,

$$(i_{OR} - d^i) \{(Clfdr_i - d^i)/(1 - Clfdr_i)\} \geq 0.$$

Summing over all m terms and taking the expectation yields (C.12).

Combine (C.11) and (C.12) to obtain

$$0 \leq E \left(\sum_{i=1}^m (i_{OR} - d^i) (Clfdr_i^i - d^i) \right) \leq i_{OR} E \left(\sum_{i=1}^m (i_{OR} - d^i) (Clfdr_i^i - d^i) \right).$$

Finally, since $i_{OR} > 0$, it follows that $E \left(\sum_{i=1}^m (i_{OR} - d^i) (Clfdr_i^i - d^i) \right) > 0$. After distributing the $(i_{OR} - d^i)$ term and separating the expectations for the sums of the two decision rules, we apply the definition of $ETP(\cdot) = E \sum_{i=1}^m i (Clfdr_i^i - d^i)$ to conclude that $ETP(i_{OR}) > ETP(d^i)$. \square