Published in final edited form as:

J Am Stat Assoc. 2018 ; 113(523): 1172–1183. doi:10.1080/01621459.2017.1336443.

Weighted False Discovery Rate Control in Large-Scale Multiple Testing

Pallavi Basu¹, T. Tony Cai², Kiranmoy Das³, and Wenguang Sun⁴

¹Department of Statistics and Operations Research, Tel Aviv University. Research supported in part by funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no[294519]-PSARPS.

²Department of Statistics, The Wharton School, University of Pennsylvania. The research of Tony Cai was supported in part by NSF Grant DMS-1403708 and NIH Grant R01 CA127334.

³Interdisciplinary Statistical Research Unit, Indian Statistical Institute, Kolkata.

⁴Department of Data Sciences and Operations, University of Southern California. The research of Wen-guang Sun was supported in part by NSF grant DMS-CAREER 1255406.

Abstract

The use of weights provides an effective strategy to incorporate prior domain knowledge in large-scale inference. This paper studies weighted multiple testing in a decision-theoretic framework. We develop oracle and data-driven procedures that aim to maximize the expected number of true positives subject to a constraint on the weighted false discovery rate. The asymptotic validity and optimality of the proposed methods are established. The results demonstrate that incorporating informative domain knowledge enhances the interpretability of results and precision of inference. Simulation studies show that the proposed method controls the error rate at the nominal level, and the gain in power over existing methods is substantial in many settings. An application to a genome-wide association study is discussed.

Keywords

Class weights; Decision weights; Multiple testing with groups; Prioritized subsets; Value to cost ratio; Weighted p -value.

1 Introduction

In large-scale studies, relevant domain knowledge, such as external covariates, scientific insights and prior data, is often available alongside the primary data set. Exploiting such information in an efficient manner promises to enhance both the interpretability of scientific results and precision of statistical inference. In multiple testing, the hypotheses being investigated often become “unequal” in light of external information, which may be reflected by differential attitudes towards the relative importance of testing units or the severity of decision errors. The use of weights provides an effective strategy to incorporate informative domain knowledge in large-scale testing problems.

In the literature, various weighting methods have been advocated for a range of multiple comparison problems. A popular scheme, referred to as the *decision weights* or loss weights approach, involves modifying the error criteria or power functions in the decision process (Benjamini and Hochberg, 1997). The idea is to employ two sets of positive constants $a = \{a_i: i = 1, \dots, m\}$ and $b = \{b_i: i = 1, \dots, m\}$ to take into account the costs and gains of multiple decisions. Typically, the choice of the weights a and b reflects the degree of confidence one has toward prior beliefs and external information. It may also be pertinent to the degree of preference that one has toward the consequence of one class of erroneous/correct decisions over another class based on various economical and ethical considerations. For example, in the spatial cluster analysis considered by Benjamini and Heller (2007), the weighted false discovery rate was used to reflect that a false positive cluster with larger size would account for a larger error. Another example arises from genome-wide association studies (GWAS), where prior data or genomic knowledge, such as prioritized subsets (Lin and Lee, 2012), allele frequencies (Lin et al., 2014) and expression quantitative trait loci information (Li et al., 2013), can often help to assess the scientific plausibility of significant associations. To incorporate such information in the analysis, a useful strategy is to up-weight the gains for the discoveries in preselected genomic regions by modifying the power functions in respective testing units (Pena et al., 2011; Sun et al., 2015; He et al., 2015). We assume in this paper that the weights have been pre-specified by the investigator. This is a reasonable assumption in many practical settings. For example, weights may be assigned according to economical considerations (Westfall and Young, 1993), external covariates (Benjamini and Heller, 2007; Sun et al., 2015) and biological insights from prior studies (Xing et al., 2010).

We mention two alternative formulations for weighted multiple testing. One popular method, referred to as the *procedural weights* approach, involves the adjustment of the p -values from individual tests. In GWAS, Roeder et al. (2006) and Roeder and Wasserman (2009) proposed to utilize linkage signals to up-weight the p -values in preselected regions and down-weight the p -values in other regions. It was shown that the power to detect association can be greatly enhanced if the linkage signals are informative, yet the loss in power is small when the linkage signals are uninformative. Another useful weighting scheme, referred to as the *class weights* approach, involves allocating varied test levels to different classes of hypotheses. For example, in analysis of the growth curve data (Box, 1950), Westfall and Young (1993, page 186) proposed to allocate a higher family-wise error rate (FWER) to the class of hypotheses related to the primary variable “gain” and a lower FWER to the secondary variable “shape”.

We focus on the decision weights approach in the present paper. This weighting scheme is not only practically useful for a wide range of applications, but also provides a powerful framework that enables a unified investigation of various weighting methods. Specifically, the proposal in Benjamini and Hochberg (1997) involves the modification of both the error rate and power function. The formulation is closely connected to the classical ideas in compound decision theory that aim to optimize the tradeoffs between the gains and losses when many simultaneous decisions are combined as a whole. Our theory reveals that if the goal is to maximize the power subject to a given error rate, then the modifications via decision weights would lead to improved multiple testing methods with sensible procedural

weights or class weights, or both. For example, in GWAS, the investigators can up-weight the power functions for discoveries in genomic regions that are considered to be more scientifically plausible or biologically meaningful; this would naturally up-weight the p -values in these regions and thus yield weighting strategies similar to those suggested by Roeder and Wasserman (2009). In large clinical trials, modifying the power functions for respective rejections at the primary and secondary end points would correspond to the allocation of varied test levels across different classes of hypotheses, leading to weighting strategies previously suggested by Westfall and Young (1993).

The false discovery rate (FDR; Benjamini and Hochberg, 1995) has been widely used in large-scale multiple testing as a practical and powerful error criterion. Following Benjamini and Hochberg (1997), we generalize the FDR to weighted false discovery rate (wFDR), and develop optimal procedures for wFDR control under the decision weights framework. We first construct an oracle procedure that maximizes the weighted power function subject to a constraint on the wFDR, and then develop a data-driven procedure to mimic the oracle and establish its asymptotic optimality. The numerical results show that the proposed method controls the wFDR at the nominal level, and the gain in power over existing methods is substantial in many settings. Our optimality result marks a clear departure from existing works in the literature on covariate-assisted inference, which aims to incorporate the external information by deriving optimal procedural weights for p -values (Roeder and Wasserman, 2009; Roquain and van de Wiel, 2009) or constructing covariate-adjusted test statistics (Feringstad et al. 2008; Zablocki et al. 2014; Cai et al. 2016) subject to the unweighted FDR and power function.

Our research also makes a novel contribution to the theory of optimal ranking in multiple testing. Conventionally, a multiple testing procedure operates in two steps: ranking the hypotheses according to their significance levels and then choosing a cutoff along the rankings. It is commonly believed that the rankings remain the same universally at all FDR levels. For example, the ranking based on p -values or adjusted p -values in common practice is invariant to the choice of the FDR threshold. The implication of our theory is interesting, for it claims that there does not exist a ranking that is universally optimal at all test levels. Instead, the optimal ranking of hypotheses depends on the pre-specified wFDR level. That is, the hypotheses may be ordered differently when different wFDR levels are chosen. This point is elaborated in Section A of the Supplementary Material. See also Roeder and Wasserman (2009) and Roquain and van de Wiel (2009), where the FDR level α is included in constructing the procedural weights.

The rest of the article is organized as follows. Section 2 discusses a general framework for weighted multiple testing. Sections 3 and 4 develop oracle and data-driven wFDR procedures and establish their optimality properties. Practice guidelines and numerical results are presented in Section 5. Section 6 concludes the article with a discussion of related and future works. Proof of Theorem 1 is given in Section 7. Proofs of other theoretical results and additional discussions are provided in the Supplementary Material.

2 Problem Formulation

This section discusses a decision weights framework for weighted multiple testing. We first introduce the model and notation and then discuss modified error criteria and power functions.

2.1 Model and notation

Suppose that m hypotheses $H_1, \dots, H_m$ are tested simultaneously based on observations $X_1, \dots, X_m$. Let $\theta = (\theta_1, \dots, \theta_m) \in \{0, 1\}^m$ denote the true state of nature, where 0 and 1 stand for null and non-null cases, respectively. Assume that observations X_i are independent and distributed according to the following model

$$X_i | \theta_i \sim (1 - \theta_i)F_{0i} + \theta_i F_{1i}, \quad (2.1)$$

where F_{0i} and F_{1i} are the null and non-null distributions for X_i , respectively. Denote by f_{0i} and f_{1i} the corresponding density functions. Suppose that the unknown states θ_i are Bernoulli (p_i) variables, where $p_i = P(\theta_i = 1)$. The mixture density is denoted by $f_i = (1 - p_i)f_{0i} + p_i f_{1i}$.

Consider the widely used random mixture model (Efron et al., 2001; Storey, 2002; Genovese and Wasserman, 2002)

$$X_i \sim F = (1 - p)F_0 + pF_1. \quad (2.2)$$

This model, which assumes that all observations are identically distributed according to a common distribution F , can sometimes be unrealistic in applications. In light of domain knowledge, the observations are likely to have different distributions. For example, in the context of a brain imaging study, Efron (2008) showed that the proportions of activated voxels are different for the front and back halves of a brain. In GWAS, certain genomic regions contain higher proportions of significant signals than other regions. In the adequate yearly progress study of California high schools (Rogosa, 2003), the densities of z -scores vary significantly from small to large schools. We shall develop theory and methodology for model (2.1), which allows possibly different non-null proportions and densities and is applicable to a wide range of settings.

The multi-group model considered in Efron (2008) and Cai and Sun (2009), which has been widely used in applications, is an important case of the general model (2.1). The multi-group model assumes that the observations can be divided into K groups. Let $\mathcal{G}_k$ denote the index set of the observations in group k , $k = 1, \dots, K$. For each $i \in \mathcal{G}_k$, θ_i is distributed as Bernoulli(p_k), and X_i follows a mixture distribution:

$$(X_i | i \in \mathcal{G}_k) \sim f_{\cdot k} = (1 - p_k)f_{0k} + p_k f_{1k}, \quad (2.3)$$

where f_{0k} and f_{1k} are the null and non-null densities for observations in group k . This model will be revisited in later sections. See also Ferkingstad et al. (2008), Hu et al. (2010) and Liu et al. (2016) for related works on multiple testing with groups.

2.2 Weighted error criterion and power function

This section discusses a generalization of the FDR criterion in the context of weighted multiple testing. Denote the decisions for the m tests by $\boldsymbol{\delta} = (\delta_1, \dots, \delta_m) \in \{0, 1\}^m$, where $\delta_i = 1$ indicates that H_i is rejected and $\delta_i = 0$ otherwise. Let a_i be the weight indicating the severity of a false positive decision. For example, a_i is taken as the cluster size in the spatial cluster analyses conducted in Benjamini and Heller (2007) and Sun et al. (2015). As a result, rejecting a larger cluster erroneously corresponds to a more severe decision error. To incorporate domain knowledge in multiple testing, Benjamini and Hochberg (1997) defined the weighted FDR $\text{wFDR}_{\text{BH}} = E\{Q(\mathbf{a})\}$, where

$$Q(\mathbf{a}) = \begin{cases} \frac{\sum_{i=1}^m a_i(1 - \theta_i)\delta_i}{\sum_{i=1}^m a_i\delta_i}, & \text{if } \sum_{i=1}^m a_i\delta_i > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (2.4)$$

We consider a slightly different version of the wFDR:

$$\text{wFDR} = \frac{E \left\{ \sum_{i=1}^m a_i(1 - \theta_i)\delta_i \right\}}{E \left(\sum_{i=1}^m a_i\delta_i \right)}. \quad (2.5)$$

In Section D of the Supplementary Material, we show that the definitions in (2.4) and (2.5) are asymptotically equivalent. The main consideration of using (2.5) is to facilitate our theoretical derivations and obtain exact optimality results.

To assess the effectiveness of different wFDR procedures, we define the expected number of true positives

$$\text{ETP} = E \left(\sum_{i=1}^m b_i\theta_i\delta_i \right), \quad (2.6)$$

where b_i is the weight indicating the power gain when H_i is rejected correctly. The use of b_i provides an effective scheme to incorporate domain knowledge. In GWAS, larger b_i can be assigned to pre-selected genomic regions to reflect that the discoveries in these regions are more biologically meaningful. In spatial data analysis, correctly identifying a larger cluster that contains signal may correspond to a larger b_i indicating a greater gain.

By combining the concerns on both the error criterion and power function, the goal in weighted multiple testing is to

$$\text{maximize the ETP subject to the constraint } \text{wFDR} \leq \alpha. \quad (2.7)$$

The optimal solution to (2.7) is studied in the next section.

3 Oracle Procedure for wFDR Control

The basic framework of our theoretical and methodological developments is outlined as follows. In Section 3.1, we assume that p_i , f_{0i} and f_{1i} in the mixture model (2.1) are known by an oracle and derive an oracle procedure that maximizes the ETP subject to a constraint on the wFDR. Connections to the literature is given in Section 3.2. In Section 4, we develop a data-driven procedure to mimic the oracle and establish its asymptotic validity and optimality.

3.1 Oracle procedure

The derivation of the oracle procedure involves two key steps: the first is to derive the optimal ranking of hypotheses and the second is to determine the optimal threshold along the ranking that exhausts the pre-specified wFDR level. We discuss the two issues in turn. Consider model (2.1). Define the local false discovery rate (Lfdr, Efron et al. 2001) as

$$\text{Lfdr}_i = \frac{(1 - p_i)f_{0i}(x_i)}{f_{\cdot i}(x_i)}. \quad (3.1)$$

The wFDR problem (2.7) is equivalent to the following constrained optimization problem

$$\text{maximize } E \left\{ \sum_{i=1}^m b_i \delta_i (1 - \text{Lfdr}_i) \right\} \text{ subject to } E \left\{ \sum_{i=1}^m a_i \delta_i (\text{Lfdr}_i - \alpha) \right\} \leq 0. \quad (3.2)$$

Let $S^- = \{i: \text{Lfdr}_i \leq \alpha\}$ and $S^+ = \{i: \text{Lfdr}_i > \alpha\}$. Then the constraint in (3.2) can be equivalently expressed as

$$E \left\{ \sum_{S^+} a_i \delta_i (\text{Lfdr}_i - \alpha) \right\} \leq E \left\{ \sum_{S^-} a_i \delta_i (\alpha - \text{Lfdr}_i) \right\}. \quad (3.3)$$

Consider an optimization problem which involves packing a knapsack with a capacity given by the right hand side of equation (3.3). Every available object has a known value and a known cost (of space). Clearly rejecting a hypothesis in S^- is always beneficial as it allows the capacity to expand, and thus promotes more discoveries. Hence the key issue boils down to how to efficiently utilize the capacity (after all hypotheses in S^- are rejected) to make as

many discoveries as possible in $\mathcal{S}^+$. Observe that each rejection in $\mathcal{S}^+$ would simultaneously increase the power and decrease the capacity. Intuitively, we should sort all hypotheses in $\mathcal{S}^+$ in an decreasing order of the value to cost ratio (VCR). Equations (3.2) and (3.3) suggest that

$$\text{VCR}_i = \frac{b_i(1 - \text{Lfdr}_i)}{a_i(\text{Lfdr}_i - \alpha)}. \quad (3.4)$$

To maximize the power, the ordered hypotheses are rejected sequentially until maximum capacity is reached.

The above considerations motivate us to consider the following class of decision rules $\delta^*(t) = \{\delta_i^*(t) : i = 1, \dots, m\}$, where

$$\delta_i^*(t) = \begin{cases} 1, & \text{if } b_i(1 - \text{Lfdr}_i) > ta_i(\text{Lfdr}_i - \alpha), \\ 0, & \text{if } b_i(1 - \text{Lfdr}_i) \leq ta_i(\text{Lfdr}_i - \alpha). \end{cases} \quad (3.5)$$

We briefly explain some important operational characteristics of testing rule (3.5). First, if we let $t > 0$, then the equation implies that $\delta_i^*(t) = 1$ for all $i \in \mathcal{S}^-$; hence all hypotheses in $\mathcal{S}^-$ are rejected as desired. (This explains why the VCR is not used directly in (3.5), given that the VCR is not meaningful in $\mathcal{S}^-$.) Second, a solution path can be generated as we vary t continuously from large to small. Along the path $\delta^*(t)$ sequentially rejects the hypotheses in $\mathcal{S}^+$ according to their ordered VCRs. Denote by $H_{(1)}, \dots, H_{(m)}$ the hypotheses sequentially rejected by δ^* . (The actual ordering of the hypotheses within $\mathcal{S}^-$ does not matter in the decision process since all are always rejected.)

The next task is to choose a cutoff along the ranking to achieve exact wFDR control. The difficulty is that the maximum capacity may not be attained by a sequential rejection procedure. To exhaust the wFDR level, we permit a randomized decision rule. Denote the Lfdr values and the weights corresponding to $H_{(j)}$ by $\text{Lfdr}_{(j)}$, $a_{(j)}$, and $b_{(j)}$. Let

$$C(j) = \sum_{i=1}^j a_{(i)}(\text{Lfdr}_{(i)} - \alpha) \quad (3.6)$$

denote the capacity up to j th rejection. According to the constraint in equation (3.2), we choose $k = \max\{j : C(j) \leq 0\}$ so that the capacity is not yet reached when $H_{(k)}$ is rejected but would just be exceeded if $H_{(k+1)}$ is rejected. The idea is to split the decision point at $H_{(k+1)}$ by randomization.

Let U be a Uniform (0,1) variable that is independent of the truth, the observations, and the weights. Define

$$t^* = \frac{b_{(k+1)}(1 - \text{Lfdr}_{(k+1)})}{a_{(k+1)}(\text{Lfdr}_{(k+1)} - \alpha)} \quad \text{and} \quad p^* = -\frac{C(k)}{C(k+1) - C(k)}.$$

Let $\mathcal{I}_A$ be an indicator, which takes value 1 if event A occurs and 0 otherwise. We propose the *oracle decision rule* $\delta_{OR} = \{\delta_{OR}^i : i = 1, \dots, m\}$, where

$$\delta_{OR}^i = \begin{cases} 1 & \text{if } b_i(1 - \text{Lfdr}_i) > t^*a_i(\text{Lfdr}_i - \alpha), \\ 0 & \text{if } b_i(1 - \text{Lfdr}_i) < t^*a_i(\text{Lfdr}_i - \alpha), \\ \mathcal{I}_{U < p^*} & \text{if } b_i(1 - \text{Lfdr}_i) = t^*a_i(\text{Lfdr}_i - \alpha). \end{cases} \quad (3.7)$$

Remark 1 The randomization step is only employed for theoretical considerations to enforce the wFDR to be exactly α . Thus the optimal power can be effectively characterized. Moreover, only a single decision point at $H_{(k+1)}$ is randomized, which has a negligible effect in large-scale testing problems. We do not pursue randomized rules for the data-driven procedures developed in later sections.

Let $\text{wFDR}(\delta)$ and $\text{ETP}(\delta)$ denote the wFDR and ETP of a decision rule δ , respectively. Theorem 1 shows that the oracle procedure (3.7) is valid and optimal for wFDR control.

Theorem 1 Consider model (2.1) and oracle procedure δ_{OR} defined in (3.7). Let $\mathcal{D}_\alpha$ be the collection of decision rules such that for any $\delta \in \mathcal{D}_\alpha$, $\text{wFDR}(\delta) = \alpha$. Then we have

- (i). $\text{wFDR}(\delta_{OR}) = \alpha$.
- (ii). $\text{ETP}(\delta_{OR}) \geq \text{ETP}(\delta)$ for all $\delta \in \mathcal{D}_\alpha$.

3.2 Comparison with the optimality results in Spjøtvoll (1972), Benjamini and Hochberg (1997) and Storey (2007)

Spjøtvoll (1972) showed that the likelihood ratio (LR) statistic

$$T_{LR}^i = \frac{f_{0i}(x_i)}{f_{1i}(x_i)} \quad (3.8)$$

is optimal for the following multiple testing problem

$$\text{maximize } E_{\cap H_{1i}} \left(\sum_{i=1}^m \delta_i \right) \text{ subject to } E_{\cap H_{0i}} \left\{ \sum_{i=1}^m \delta_i \right\} \leq \alpha, \quad (3.9)$$

where $\cap H_{0j}$ and $\cap H_{1j}$ denote the intersections of the nulls and non-nulls, respectively. The error criterion $E_{\cap H_{0i}} \left\{ \sum_i a_i \delta_i \right\}$ is referred to as the intersection tests error rate (ITER).

A weighted version of problem (3.9) was considered by Benjamini and Hochberg (1997), where the goal is to

$$\text{maximize } E_{\cap H_{1i}} \left(\sum_{i=1}^m b_i \delta_i \right) \text{ subject to } E_{\cap H_{0i}} \left\{ \sum_{i=1}^m a_i \delta_i \right\} \leq \alpha. \quad (3.10)$$

The optimal solution to (3.10) is given by the next proposition.

Proposition 1 (Benjamini and Hochberg, 1997). Define the weighted likelihood ratio (WLR)

$$T_{IT}^i = \frac{a_i f_{0i}(x_i)}{b_i f_{1i}(x_i)}. \quad (3.11)$$

Then the optimal solution to (3.10) is a thresholding rule of the form $\delta_{IT}^i = (T_{IT}^i < t_\alpha)$, where t_α is the largest threshold that controls the weighted ITER at level α .

The ITER is very restrictive in the sense that the expectation is taken under the conjunction of the null hypotheses. The ITER is inappropriate for mixture model (2.1) where a mixture of null and non-null hypotheses are tested simultaneously. To extend intersection tests to multiple tests, define the per family error rate (PFER) as

$$\text{PFER}(\delta) = E \left\{ \sum_{i=1}^m a_i (1 - \theta_i) \delta_i \right\}. \quad (3.12)$$

The power function should be modified correspondingly. Therefore the goal is to

$$\text{maximize } E \left(\sum_{i=1}^m b_i \theta_i \delta_i \right) \text{ subject to } E \left\{ \sum_{i=1}^m a_i (1 - \theta_i) \delta_i \right\} \leq \alpha. \quad (3.13)$$

The key difference between the ITER and PFER is that the expectation in (3.12) is now taken over all possible combinations of the null and non-null hypotheses. The optimal PFER procedure is given by the next proposition.

Proposition 2 Consider model (2.1) and assume continuity of the LR statistic. Let $\mathcal{D}_{PF}^\alpha$ be the collection of decision rules such that for every $\delta \in \mathcal{D}_{PF}^\alpha$, $\text{PFER}(\delta) \leq \alpha$. Define the weighted posterior odds (WPO)

$$T_{PF}^i = \frac{a_i(1-p_i)f_{0i}(x_i)}{b_i p_i f_{1i}(x_i)}. \quad (3.14)$$

Denote by $Q_{PF}(t)$ the PFER of $\delta_{PF}^i = I(T_{PF}^i < t)$. Then the oracle PFER procedure is $\delta_{PF} = (\delta_{PF}^i : i = 1, \dots, m)$, where $\delta_{PF}^i = I(T_{PF}^i < t_{PF})$ and $t_{PF} = \sup\{t : Q_{PF}(t) = \alpha\}$.

This oracle rule satisfies:

- (i). $ETP(\delta_{PF}) = \alpha$.
- (ii). $ETP(\delta_{PF}) \leq ETP(\delta)$ for all $\delta \in \mathcal{D}_{PF}^\alpha$.

Storey (2007) proposed the optimal discovery procedure (ODP) that aims to maximize the ETP subject to a constraint on the expected number of false positives (EFP). The ODP extends the optimality result in Spjøtvoll (1972) from the intersection tests to multiple tests. However, the formulation of ODP does not incorporate weights. As a result, the ODP procedure is a symmetric rule where all hypotheses are exchangeable. Symmetric rules are in general not suitable for weighted multiple testing where it is desirable to incorporate external information and treat the hypotheses differently.

Our formulation (2.7) modifies the conventional formulations in (3.10) and (3.13) to the multiple testing situation with an FDR type criterion. These modifications lead to methods that are more suitable for large-scale scientific studies. The oracle procedure (3.7) uses the VCR (3.4) to rank the hypotheses and is an asymmetric rule. The VCR, which optimally combines the decision weights, significance measure (L_{fdr}) and test level α , produces a more powerful ranking than the WPO (3.14) in the wFDR problem. Section A in the supplementary material provides a detailed discussion on the ranking issue.

4 Data-Driven Procedures and Asymptotics

The oracle procedure (3.7) cannot be implemented in practice since it relies on unknown quantities such as L_{fdr} and t^* . This section develops a data-driven procedure to mimic the oracle. We first propose a test statistic to rank the hypotheses and discuss related estimation issues. A step-wise procedure is then derived to determine the best cutoff along the ranking. Finally, asymptotic results on the validity and optimality of the proposed procedure are presented.

4.1 Proposed test statistic and its estimation

The oracle procedure utilizes the ranking based on the VCR (3.4). However, the VCR is only meaningful for the tests in S^+ and becomes problematic when both S^- and S^+ are considered. Moreover, the VCR could be unbounded, which would lead to difficulties in both numerical implementations and technical derivations. We propose to rank the hypotheses using the following statistic (in increasing values)

$$R_i = \frac{a_i(\text{Lfdr}_i - \alpha)}{b_i(1 - \text{Lfdr}_i) + a_i|\text{Lfdr}_i - \alpha|}. \quad (4.1)$$

As shown in the next proposition, R_i always ranks hypotheses in S^- higher than hypotheses in S^+ (as desired), and yields the same ranking as that by the VCR (3.4) for hypotheses in S^+ . The other drawbacks of VCR can also be overcome by R_i : R_i is always bounded in the interval $[-1, 1]$ and is a continuous function of the Lfdr_i .

Proposition 3 (i) *The rankings generated by the decreasing values of VCR (3.4) and increasing values of R_i (4.1) are the same in both S^- and S^+ .* (ii) *The ranking based on increasing values of R_i always puts hypotheses in S^- ahead of hypotheses in S^+ .*

Next we discuss how to estimate R_i ; this involves the estimation of the Lfdr statistic (3.1), which has been studied extensively in the multiple testing literature. We give a review of related methodologies. If all observations follow a common mixture distribution (2.2), then we can first estimate the non-null proportion p and the null density f_0 using the methods in Jin and Cai (2007), and then estimate the mixture density f using a standard kernel density estimator (e.g. Silverman, 1986). If all observations follow a multi-group model (2.3), then we can apply the above estimation methods to separate groups to obtain corresponding estimates $\hat{p}_k$, $\hat{f}_{0k}$, and $\hat{f}_{\cdot k}$, $k = 1, \dots, K$. The theoretical properties of these estimators have been established in Sun and Cai (2007) and Cai and Sun (2009). In practice, estimation problems may arise from more complicated models. Related theories and methodologies have been studied in Storey (2007), Ferkingstad et al. (2008), and Efron (2008, 2010); theoretical supports for these estimators are yet to be developed.

The estimated Lfdr value for H_i is denoted by $\widehat{\text{Lfdr}}_i$. By convention, we take $\widehat{\text{Lfdr}}_i = 1$ if $\widehat{\text{Lfdr}}_i > 1$. This modification only facilitates the development of theory and has no practical effect on the testing results (since rejections are essentially only made for small $\widehat{\text{Lfdr}}_i$'s). The ranking statistic R_i can therefore be estimated as

$$\hat{R}_i = \frac{a_i(\widehat{\text{Lfdr}}_i - \alpha)}{b_i(1 - \widehat{\text{Lfdr}}_i) + a_i|\widehat{\text{Lfdr}}_i - \alpha|}. \quad (4.2)$$

The performance of the data driven procedure relies on the accuracy of the estimate $\widehat{\text{Lfdr}}_i$; some technical conditions are discussed in the next subsection. The finite sample performance of different Lfdr estimates are investigated in Section E.4 of the Supplementary Material.

4.2 Proposed testing procedure and its asymptotic properties

Consider $\hat{R}_i$ defined in (4.2). Denote by $\hat{N}_i = a_i(\widehat{\text{Lfdr}}_i - \alpha)$ the estimate of excessive error rate when H_i is rejected. Let $\hat{R}_{(1)}, \dots, \hat{R}_{(m)}$ be the ordered test statistics (in increasing values). The hypothesis and estimated excessive error rate corresponding to $\hat{R}_{(i)}$ are denoted by $H_{(i)}$ and $\hat{N}_{(i)}$. The idea is to choose the largest cutoff along the ranking based on $\hat{R}_i$ so that the maximum capacity is reached. Motivated by the constraint in (3.2), we propose the following step-wise procedure.

Procedure 1 (*wFDR control with general weights*). Rank hypotheses according to $\hat{R}_i$ in

increasing values. Let $k = \max \left\{ j : \sum_{i=1}^j \hat{N}_{(i)} \leq 0 \right\}$. Reject $H_{(i)}$, for $i = 1, \dots, k$.

It is important to note that in Procedure 1, $\hat{R}_i$ is used in the ranking step whereas $\hat{N}_i$ (or a weighted transformation of $\widehat{\text{Lfdr}}_i$) is used in the thresholding step. The ranking by $\widehat{\text{Lfdr}}_i$ is in general different from that by $\hat{R}_i$. In some applications where the weights are proportional, i.e. $a = c \cdot b$ for some constant $c > 0$, then the rankings by $\hat{R}_i$ and $\widehat{\text{Lfdr}}_i$ are identical. Specifically $\hat{R}_i$ is then monotone in $\widehat{\text{Lfdr}}_i$. Further, choosing the cutoff based on $\hat{N}_i$ is equivalent to that of choosing by a weighted $\widehat{\text{Lfdr}}_i$. This leads to an Lfdr based procedure (Sun et al., 2015), which can be viewed as a special case of Procedure 1.

Procedure 2 (*wFDR control with proportional weights*). Rank hypotheses according to $\widehat{\text{Lfdr}}_i$ in increasing values. Denote the hypotheses and weights corresponding to $\widehat{\text{Lfdr}}_{(i)}$ by $H(i)$ and $a(i)$. Let

$$k = \max \left\{ j : \left(\sum_{i=1}^j a(i) \right)^{-1} \sum_{i=1}^j a(i) \widehat{\text{Lfdr}}_{(i)} \leq \alpha \right\}.$$

Reject $H_{(i)}$, for $i = 1, \dots, k$.

Next we investigate the asymptotic performance of Procedure 1. We first give some regularity conditions for the weights. Our theoretical framework requires that the decision weights must be obtained from external sources such as prior data, biological insights, or economical considerations. In particular, *the observed data $\{X_i : i = 1, \dots, m\}$ cannot be used to derive the weights*. The assumption is not only crucial in theoretical developments, but also desirable in practice (to avoid using data twice). Therefore given the domain knowledge, the decision weights do not depend on observed values. Moreover, a model with random (known) weights is employed for technical convenience, as done in Genovese et al. (2006) and Roquain and van de Wiel (2009). We assume that the weights are independent with each other across testing units. Formally, denote e_i the external domain knowledge for hypothesis i , we require the following condition.

Condition 1 (i) $(a_i, b_i | X_i, \theta_i, e_i) \stackrel{d}{\sim} (a_i, b_i | e_i)$ for $1 \leq i \leq m$. (ii) (a_i, b_i) and (a_j, b_j) are independent for $i \neq j$.

In weighted multiple testing problems, the analysis is always carried out in light of the external information e_i implicitly. The notation of conditional distribution on e_i will be suppressed when there is no ambiguity. In practice, the weights a_i and b_i are usually bounded. We need a weaker condition in our theoretical analysis.

Condition 2 (Regularity conditions on the weights.) Let C and c be two positive constants. $E(\sup_i a_i) = \alpha(m)$, $E(\sup_i b_i) = \alpha(m)$, $E(a_i^4) \leq C$, and $\min\{E(a_i), E(b_i)\} \geq c$.

A consistent Lfdr estimate is needed to ensure the large-sample performance of the data-driven procedure. Formally, we need the following condition.

Condition 3 $\widehat{Lfdr}_i - Lfdr_i = o_p(1)$. Also, $\widehat{Lfdr}_i \xrightarrow{d} Lfdr$, where $Lfdr$ is an independent copy of $Lfdr_i$.

Remark 2 Condition 3 is a reasonable assumption in many applications. We give a few important scenarios where Condition 3 holds. Suppose we observe z -values from random mixture model

$$Z_i \sim N(\mu_i, \sigma_i^2), (\mu_i, \sigma_i) \sim F(\mu, \sigma),$$

where $(\mu_i, \sigma_i) = (\mu_0, \sigma_0)$ if $\theta_i = 0$, $(\mu_i, \sigma_i) \neq (\mu_0, \sigma_0)$ if $\theta_i = 1$, and $F(\mu, \sigma)$ is a general bivariate distribution. Let p denote the proportion of non-null cases. Then $\hat{p}_{jc}$, the estimator proposed in Jin and Cai (2007), satisfies $\hat{p}_{jc} \xrightarrow{P} p$ under mild regularity conditions. Moreover, it is known that the kernel density estimator satisfies $E\|\hat{f} - f\|^2 \rightarrow 0$. It follows from Sun and Cai (2007) that Condition 3 holds when the above estimators are used. If the null distribution is unknown, then we can use the method in Jin and Cai (2007) to estimate the null parameters (μ_0, σ_0) . Then under certain regularity conditions, we can show that $E\|\hat{f}_0 - f_0\|^2 \rightarrow 0$, and Condition 3 holds. For the multi-group model (2.3), let $\hat{p}_k$, $\hat{f}_{k0}$, and $\hat{f}_k$ be estimates of p_k , f_{k0} , and f_k such that $\hat{p}_k \xrightarrow{P} p_k$, $E\|\hat{f}_{k0} - f_{k0}\|^2 \rightarrow 0$, $E\|\hat{f}_k - f_k\|^2 \rightarrow 0$, $k = 1, \dots, K$. Let $\widehat{Lfdr}_i = (1 - \hat{p}_k)\hat{f}_{k0}(x_i)/\hat{f}_k(x_i)$ if $i \in \mathcal{G}_k$. It follows from Cai and Sun (2009) that Condition 3 holds when we apply Jin and Cai's estimators and standard kernel estimates to the groups separately.

The oracle procedure (3.7) provides an optimal benchmark for all wFDR procedures. The next theorem establishes the asymptotic validity and optimality of the data-driven procedure by showing that the wFDR and ETP levels of the data-driven procedure converge to the oracle levels as $m \rightarrow \infty$.

Theorem 2 Assume Conditions 1–3 hold. Denote by $wFDR_{DD}$ the $wFDR$ level of the data-driven procedure (Procedure 1). Let ETP_{OR} and ETP_{DD} be the ETP levels of the oracle procedure (3.7) and data-driven procedure, respectively. Then we have

- (i). $wFDR_{DD} = \alpha + o(1)$.
- (ii). $ETP_{DD}/ETP_{OR} = 1 + o(1)$.

Corollary 1 Suppose we choose $a_i = 1$ for all i . Then under the conditions of the theorem, our data-driven procedure controls the unweighted FDR at level $\alpha + o(1)$.

5 Practical Issues and Numerical Results

The $wFDR$ framework provides a useful approach to integrate domain knowledge in multiple testing. However, the weights must be chosen with caution to avoid improper manipulation of results. Under our formulation, we assume that the weights are pre-specified, while stressing that the task of assigning “correct” weights is a critical issue in practice. This section first states a few practical guidelines in weighted FDR analysis, then describes general strategies in choosing weights for a range of application scenarios. Finally we illustrate the proposed methodology via an application to GWAS.

5.1 Practical guidelines and examples on choosing weights

We first state a few important practical guidelines to ensure that the weighted FDR analysis is conducted properly. The first guideline is in particular important for a valid FDR control.

1. The weights must be “external” in the sense that they do not depend on the primary data such as the z -values or p -values on which the multiple tests are performed. As a practical guideline, we recommend “choosing external weights before seeing the data” as a standard rule in all weighted FDR analyses.
2. The choice of weights requires scientific motivations or economic considerations. For a statistical analysis involving any weighting schemes, the procedure for obtaining weights must be specified beforehand in the analysis plan, and needs be disclosed and justified carefully.
3. To prevent researchers from setting weights to find desired significant results, extreme weights are in general not acceptable. It is recommended to use moderate weights to incorporate domain knowledge and carry out a sensitivity analysis afterwards to avoid manipulations of results.

Next we discuss a range of scenarios that naturally give rise to weighted FDR analyses. We use these examples to illustrate how to assign weights properly in light of various economic and scientific considerations. Moreover, we emphasize how the above practical guidelines can be implemented in respective contexts to avoid manipulations of results.

Example 1. Spatial cluster analysis.—In multiple testing, relevant domain knowledge or external data can be exploited to form more meaningful testing units such as sets or clusters. For example, in a spatial setting (Pacífico et al. 2004; Benjamini and Heller 2007; Sun et al. 2015), individual locations can be aggregated into clusters to increase the signal to

noise ratio and scientific interpretability. It is natural to consider weighted error rates when clusters are heterogeneous. We state a few guidelines for the choice of weights.

- (i) The weights should be pre-specified by experts to reflect various practical and economic considerations based on relevant spatial covariates such as cluster sizes, population densities, poverty rates, and urban vs. rural regions.
- (ii) In cluster-wise analyses, weights may be used to modify both the error rate and power function. For example, symmetric weights (i.e. $a_i = b_i$ for all i) may be assigned to reflect that the gains and losses in the decision process are proportional. A careful investigation of the optimal ranking of clusters (Section A in the Supplementary Material) is helpful to improve the power of analysis.

Example 2. Multiple endpoints clinical trials.—In comparing treatments with multiple endpoints (e.g. Dmitrienko et al. 2003), the error rates may be modified by assigning different weights to primary and secondary end points; this promises to increase the power in discovering more important clinical findings (Young and Gries 1984; Westfall and Young 1993). Benjamini and Cohen (2016) discussed two weighting regimes in clinical trials: the first regime requires that any primary endpoint carries a larger weight than any secondary one; and the second regime in addition requires that the combined weight of the primary endpoints is greater than the combined weight of the secondary ones. To avoid manipulations of results, Benjamini and Cohen (2016) suggested that a weighted FDR analysis should obey the following rules:

- (i) The tradeoffs between primary and secondary end points should be evaluated prior to conducting the trial.
- (ii) The weights should be given before seeing the data.
- (iii) The details of obtaining weights should be revealed in the trial protocol.

Example 3. Prioritized subset analysis.—In GWAS data analysis, finding disease-causing SNPs in the immense data is very challenging. Weighted FDR analysis provides a powerful approach to integrate genomic knowledge in inference with prioritized SNP subsets. For example, Roeder et al. (2006) and Zablocki et al. (2014) assigned higher weights to pre-selected genomic regions to increase the power in detecting disease-associated SNPs that are more biologically plausible. In large-scale multiple testing with prioritized subset, we have the following requirements regarding the choice of weights:

- (i). It is not allowed to use the primary data (e.g. significant levels of disease association) to derive weights; such practice would distort the null distribution of p -values and lead to the failure of the wFDR procedure.
- (ii). The weights must be obtained based on either prior knowledge regarding the biological importance of an association, or external knowledge such as the linkage disequilibrium (LD) correlations between SNPs. This requirement can effectively avoid selection bias because the weights are conditionally independent from the significance levels of SNP associations (cf. our Condition 1). See also Section 5.(a) in Benjamini et al. (2009) for related discussions.

- (iii). To avoid further aggravating the multiplicity issue in large-scale testing problems, the weights should only be used to modify the power function. The unweighted FDR criterion is recommended to remain unchanged. This ensures greater power to reject more biologically relevant hypotheses without increasing the FDR (cf. our Corollary 1). See also Theorem 1 in Genovese et al. (2006) and Section 5.(a) in Benjamini et al. (2009) for related theoretical results and discussions.

5.2 Decision weights vs. procedural weights

The formulation (2.7) is different from the “procedural weights” approach, for example, in Roeder et al. (2006) and Zablocki et al. (2014), where the goal is to exploit external information by constructing weights to increase the total number of rejections (covariate-assisted inference). The decision weights framework integrates the external information in a very different way; the goal is to improve the interpretability of results instead of statistical power (although the power may be gained as a byproduct, see Section E.5 in the Supplementary Material). Under the wFDR framework, the hypotheses are of “unequal” importance in light of domain knowledge. For instance, in spatial cluster analysis, a false positive cluster with larger size would account for a larger error; hence it is not sensible to use the conventional FDR definition that treats all testing units equally. The wFDR framework is different from those in Genovese et al. (2006), Ferkingstad et al. (2008) and Cai et al. (2016), where the gains and losses in various decisions are exchangeable, and conventional FDR and power definitions are suitable.

5.3 Application to Framingham Heart Study (FHS)

In this section, we implement the proposed method for analyzing a data set from Framingham Heart Study (Fox et al., 2007; Jaquish, 2007). A brief description of the study, the implementation of our methodology, and the results are discussed in turn.

The goal of the study is to decipher the genetic architecture behind the cardiovascular disorders for the Caucasians. Started in 1948 with nearly 5,000 healthy subjects, the study is currently in its third generation of the participants. The biomarkers responsible for the cardiovascular diseases, for e.g., body mass index (BMI), weight, blood pressure, and cholesterol level, were measured longitudinally. Since the mutation SNPs (within the block of DNA) are commonly passed on to descendants, it is a standard practice to collect data over several generations. The specific goal of our study is to identify disease-associated SNPs in the second generation while paying more attention to those that remain stable over the first two generations. Identifying significant SNPs (or SNP blocks) that are preserved across generations can improve both biological plausibility and statistical replicability. In our study design, we use 399 subjects in the first generation to serve as the baseline to construct weights, and use the information from the 578 subjects in the second generation (with 310 males and 268 females) to conduct the primary analysis. It is important to note that the goal of our weighted FDR analysis is different from that in Roeder et al. (2006) and Zablocki et al. (2014), where prior likelihood of association with the phenotype is utilized for constructing “procedural weights” to improve the unweighted power. In contrast, we use decision weights to modify the power function so that discoveries in certain preselected

regions receive a higher priority. The main consideration is not to increase the total number of discoveries, but to focus on findings are more relevant to our study goal (identifying stable mutations across generations). In our study it is not sensible to use the conventional power definition where the discoveries from all regions are treated equally.

We consider the BMI as the response variable and develop a dynamic model to detect the SNPs associated with the BMI. Let $Y_i(t_{ij})$ denote the response (BMI) from the i -th subject at time t_{ij} , $j = 1, \dots, T_i$. Consider the following model for longitudinal traits:

$$Y_i(t_{ij}) = f(t_{ij}) + \beta_k G_{ik} + \gamma_{i0} + \gamma_{i1} t_{ij} + \epsilon_i(t_{ij}), \quad (5.1)$$

where $f(\cdot)$ is the general effect of time that is modeled by a polynomial function of suitable order, selected by AIC or BIC model selection criterion, β_k is the effect of the k -th SNP on the response and G_{ik} denotes the genotype of the i -th subject for the k -th SNP. We also consider the random intercepts and random slopes, denoted γ_{i0} and γ_{i1} , respectively, for explaining the subject-specific response trajectories. A bivariate normal distribution for $\gamma_i = (\gamma_{i0}, \gamma_{i1})$ is assumed. Moreover, we assume that the residual errors are normally distributed with zero mean, and covariance matrix Σ_i with an order-one auto-regressive structure. We fit model (5.1) for each SNP and obtain the estimate of the genetic effect $\hat{\beta}_k$. If we reject the null hypothesis $H_0 : \beta_k = 0$ vs. $H_1 : \beta_k \neq 0$, then we conclude that the k -th SNP has a significant association with the BMI. Since we have nearly 5 million SNPs, the false discovery rate needs to be controlled for making scientifically meaningful inference. For each k , we take standardized $\hat{\beta}_k$ as our z -scores and obtain the estimated ranking statistic $\hat{R}_k$ as described in (4.2).

To incorporate external domain knowledge, we construct weights by utilizing the p -values obtained from first generation participants. This information is secondary in the sense that the first generation data are not directly used in multiple testing but are only employed to provide supplementary support in inference. In the FHS study, there are 399 subjects in the first generation for which we have obtained the trait values (such as BMI) and genetic information on the same set of SNPs. Using the p -values obtained from the first generation data, we partition the SNPs into three groups: less than 0.001, between 0.001 and 0.01, and greater than 0.01. The groups are denoted $\mathcal{G}_j$, $j = 1, 2, 3$. We implement the proposed wFDR method using the p -values computed from second generation participants, with group-wise weights $b_i = c_j$ if $i \in \mathcal{G}_j$. Note that b_i are the weights in respective power functions with a larger b_i indicating a higher priority or a stronger belief. Since assigning a value to the biological importance of an association can be quite subjective, we propose to use different combinations of weights to investigate the sensitivity of testing results:

Setting 1: $a_i = 1$ for all i , $b_i = c_j$ if $i \in \mathcal{G}_j$, $(c_1, c_2, c_3) = (1, 1, 1)$

Setting 2: $a_i = 1$ for all i , $b_i = c_j$ if $i \in \mathcal{G}_j$, $(c_1, c_2, c_3) = (4, 2, 1)$

Setting 3: $a_i = 1$ for all i , $b_i = c_j$ if $i \in \mathcal{G}_j$, $(c_1, c_2, c_3) = (10, 5, 1)$.

Setting 1 corresponds to an unweighted FDR analysis. Settings 2 and 3 reflect our belief that the correct rejections from groups 1 and 2 are more “powerful” or more “relevant” when compared to those from group 3. This is sensible as one of the goals in the FHS study is to discover significant SNPs that remain *stable* across generations. In our analysis, the weights $a_i = 1$ are chosen for all SNPs; this ensures that *the FDR will not be increased by weighting*. We only modify the weights b_i in the power functions to reflect that discoveries in certain pre-selected regions are of greater importance. In practice b_i may be derived from expert opinion to incorporate scientific or economic considerations. We stress that in GWAS, the details of the process in choosing weights, which requires a separate discussion or justification, should be revealed.

We implement the proposed wFDR method to select most significant SNPs. The results for the sensitivity analysis with different choices of weights are summarized in Table 1. The table shows the groups sizes, the group-wise threshold levels, and the number of SNPs selected from different “groups” for each weight combination at two different FDR levels $\alpha = 0.10$ and $\alpha = 0.05$.

From Table 1 we can see that the unweighted analysis [i.e. the weights combination (1, 1, 1)] consistently selects higher number of SNPs than the other two weights combinations. Meanwhile, it selects fewer number of SNPs from groups 1 and 2. Here we have aimed to prioritize certain pre-selected regions and discover significant SNPs that are more biologically relevant. In our analysis, groups 1 and 2 are believed to be more “informative.” At $\alpha = 0.10$, it is interesting to note that the number of rejections from groups 1 and 2 are increased by several folds but the number of rejections from group 3 is only decreased by a small proportion.

In Table 1, we also report the group-wise thresholds for the p -values. The value of the threshold is computed as the maximum of the p -values of the hypotheses rejected in the specific group and weight combination. We can see that the threshold for group 3 decreases in weights, which indicates that the rejection criteria become more stringent with higher weights. In contrast, the thresholds for groups 1 and 2 increase with weights, which allows for more discoveries in pre-selected regions. All these above observations are sensible and in agreement with the intuitions of our proposed methodology.

5.4 Simulation

In all simulation studies, we consider a two-point normal mixture model

$$X_i \sim (1 - p)N(0, 1) + pN(\mu, \sigma^2), i = 1, \dots, m. \quad (5.2)$$

The nominal wFDR is fixed at $\alpha = 0.10$. We consider the comparison of different methods under the scenario where there are two groups of hypotheses and within each group the weights are proportional.

The proposed method (Procedure 1 in Section 4.2) is denoted by 4 DD. Other methods to be compared include:

1. The wFDR method proposed by Benjamini and Hochberg (1997); denoted by 2 BH97. In simulations where $a_i = 1$ for all i , BH97 reduces to the well-known step-up procedure in Benjamini and Hochberg (1995), denoted by BH95.
2. A stepwise wFDR procedure, which rejects hypotheses along the WPO (3.14) ranking sequentially and stops at $k = \max \left\{ j : \sum_{i=1}^j \hat{N}_{(i)} \leq 0 \right\}$, with $\hat{N}_{(i)}$ defined in Section 4.2. The method is denoted by $\circ$ WPO. Following similar arguments in the proof of Theorem 2, we can show that the WPO method controls the wFDR at the nominal level asymptotically. This is also verified by our simulation results. Meanwhile, we expect that the WPO method will be outperformed by the proposed method (4 DD), which operates along the more efficient VCR ranking.
3. The adaptive z -value method in Sun and Cai (2007), denoted by + AZ. AZ is valid and optimal in the unweighted case but suboptimal in the weighted case.

To save space, this section only presents results on group-wise weights. Our setting is motivated by our application to GWAS, where the hypotheses can be divided into groups: those in preselected regions and those in other regions. It is desirable to assign varied weights to separate groups to reflect that the discoveries in preselected regions are more biologically meaningful. In Section E.2 of the Supplementary Material, we compare our methods with existing methods using general weights a_i and b_i that are generated from probability distributions. We also provide additional numerical results such as the comparison of various wFDR definitions, the finite sample performance of Lfdr, and the impacts of weights on the power of different wFDR procedures.

The first simulation study investigates the effect of weights. Consider two groups of hypotheses with group sizes $m_1 = 3000$ and $m_2 = 1500$. In both groups, the non-null proportion is $p = 0.2$. The null and non-null distributions are $N(0,1)$ and $N(1.9,1)$, respectively. We fix $a_i = 1$ for all i . Hence BH97 reduces to the method proposed in Benjamini and Hochberg (1995), denoted by BH95. The wFDR reduces to the regular FDR, and all methods being considered are valid for FDR control. For hypotheses in group 1, we let $c_1 = a_i/b_i$. For hypotheses in group 2, we let $c_2 = a_i/b_i$. We choose $c_1 = 3$ and vary c_2 . Hence the weights are proportional within respective groups and vary across groups. The second simulation study investigates the impacts of signal strength. The results are presented in Section E.1 of the Supplementary material.

In the simulation, we apply the four methods above to the simulated data set and obtain the wFDR and ETP levels by averaging the multiple testing results over 200 replications. In Figure 1, we plot the wFDR levels and ETP of different methods as functions of c_2 , which is varied over $[0.1, 0.8]$. Panel (a) shows that all methods control the wFDR under the nominal level, and the BH97 method is conservative. Panel (b) shows that the proposed method dominates all existing methods. The proposed method is followed by the WPO method, which outperforms all unweighted methods (AZ and BH95) since b_i , the weights in the power function, are incorporated in the testing procedure. The BH97 (or BH95) has the smallest ETP. As c_2 approaches 1 or the weights a_i and b_i equalizes, the relative difference of the various methods (other than BH95) becomes less.

6 Discussion

The FDR provides a practical and powerful approach to large-scale multiple testing problems and has been widely used in a wide range of scientific studies. The wFDR framework extends the FDR paradigm to integrate useful domain knowledge in simultaneous inference. The weights must be chosen with caution to avoid improper manipulation of results.

In the multiple testing literature, procedural, decision, and class weights are often viewed as distinct weighting schemes and have been mostly investigated separately. Although this paper focuses on the decision weights approach, the decision-theoretic framework enables a unified investigation of other weighting schemes. For example, a comparison of the LR (3.8) and WLR (3.11) demonstrates how the LR statistic may be adjusted optimally to account for the decision gains or losses. This shows that procedural weights may be derived in the decision weights framework. Moreover, the difference between the WLR (3.11) and WPO (3.14) shows the important role that p_i plays in multiple testing. In particular the WPO (3.14) provides important insights on how prior beliefs may be incorporated in a decision weights approach to derive appropriate class weights. To see this, consider the multi-class model (2.3). Following the arguments in Cai and Sun (2009), we can conclude that in order to maximize the power, different FDR levels should be assigned to different classes. Similar suggestions for varied class weights have been made in Westfall and Young (1993, pages 169 and 186). These examples demonstrate that the decision weights approach provides a powerful framework to derive both procedural weights and class weights.

Our formulation requires that the weights must be *pre-specified* based on *external* domain knowledge. It is of interest to extend the work to the setting where the weights are unknown. Due to the variability in the quality of external information, subjectivity of investigators, and complexity in modeling and analysis, a systematic study of the issue is beyond the scope of the current paper. The optimal choice of weights depends on statistical, economic and scientific concerns jointly. Notable progresses have been made, for example, in Roeder and Wasserman (2009) and Roquain and van de Wiel (2009). However, these methods are mainly focused on the weighted p -value approach under the unweighted FDR criterion, hence do not apply to the framework in Benjamini and Hochberg (1997). Moreover, the optimal decision rule in the wFDR problem in general is not a thresholding rule based on the adjusted p -values. Much work is still needed to derive decision weights that would optimally incorporate domain knowledge in large-scale studies.

7 Proofs

This section proves Theorem 1. The proofs of other results are given in the Supplementary Material.

Proof of Part (i) of Theorem 1.

To show that $\text{wFDR}(\delta_{OR}) = \alpha$, we only need to establish that

$$E_{U, \mathbf{a}, \mathbf{b}, \mathbf{X}} \left\{ \sum_{i=1}^m a_i \delta_{OR}^i (\text{Lfd}_i - \alpha) \right\} = 0,$$

where the notation $E_{U, \mathbf{a}, \mathbf{b}, \mathbf{X}}$ denotes that the expectation is taken over $U, \mathbf{a}, \mathbf{b}$, and $\mathbf{X}$. According to the definitions of the capacity function $C(\cdot)$ and threshold t^* , we have

$$\sum_{i=1}^m a_i \delta_{OR}^i (\text{Lfd}_i - \alpha) = C(k) + I(U < p^*) \{C(k+1) - C(k)\}.$$

It follows from the definition of p^* that

$E_{U| \mathbf{a}, \mathbf{b}, \mathbf{X}} \left\{ \sum_{i=1}^m a_i \delta_{OR}^i (\text{Lfd}_i - \alpha) \right\} = C(k) + \{C(k+1) - C(k)\} p^* = 0$, where the notation $E_{U| \mathbf{a}, \mathbf{b}, \mathbf{X}}$ indicates that the expectation is taken over U while holding $(\mathbf{a}, \mathbf{b}, \mathbf{X})$ fixed. Therefore

$$E_{U, \mathbf{a}, \mathbf{b}, \mathbf{X}} \left\{ \sum_{i=1}^m a_i \delta_{OR}^i (\text{Lfd}_i - \alpha) \right\} = 0, \quad (7.1)$$

and the desired result follows.

Proof of Part (ii) of Theorem 1.

Let δ^* be an arbitrary decision rule such that $wFDR(\delta^*) = \alpha$. Then

$$E_{\mathbf{a}, \mathbf{b}, \mathbf{X}} \left\{ \sum_{i=1}^m a_i E(\delta_i^* | \mathbf{a}, \mathbf{b}, \mathbf{x}) (\text{Lfd}_i - \alpha) \right\} \leq 0. \quad (7.2)$$

The notation $E(\delta_i^* | \mathbf{x}, \mathbf{a}, \mathbf{b})$ means that the expectation is taken to average over potential randomization conditional on the observations and weights.

Let $\mathcal{J}^+ = \{i : \delta_{OR}^i - E(\delta_i^* | \mathbf{x}, \mathbf{a}, \mathbf{b}) > 0\}$ and $\mathcal{J}^- = \{i : \delta_{OR}^i - E(\delta_i^* | \mathbf{x}, \mathbf{a}, \mathbf{b}) < 0\}$. For $i \in \mathcal{J}^+$, we have $\delta_{OR}^i = 1$ and hence $b_i(1 - \text{Lfd}_i) - t^* a_i (\text{Lfd}_i - \alpha)$. Similarly for $i \in \mathcal{J}^-$, we have $\delta_{OR}^i = 0$ and so $b_i(1 - \text{Lfd}_i) - t^* a_i (\text{Lfd}_i - \alpha)$. Thus

$$\sum_{i \in \mathcal{J}^+ \cup \mathcal{J}^-} \left\{ \delta_{OR}^i - E(\delta_i^* | \mathbf{x}, \mathbf{a}, \mathbf{b}) \right\} \left\{ b_i(1 - \text{Lfd}_i) - t^* a_i (\text{Lfd}_i - \alpha) \right\} \geq 0.$$

Note that δ_{OR}^i is perfectly determined by $\mathbf{X}$ except for $(k+1)$ th decision. Meanwhile, $b_{(k=0)}(1 - \text{Lfd}_{(k+1)} - t^* a_{(k+1)} (\text{Lfd}_{(k+1)} - \alpha)) = 0$ by our choice of t^* . It follows that

$$E_{a,b,X} \left[\sum_{i=1}^m \{E(\delta_{OR}^i | \mathbf{x}, \mathbf{a}, \mathbf{b}) - E(\delta_i^* | \mathbf{x}, \mathbf{a}, \mathbf{b})\} \{b_i(1 - \text{Lfdr}_i) - t^* a_i(\text{Lfdr}_i - \alpha)\} \right] \geq 0. \quad (7.3)$$

Recall that the power function is given by $\text{ETP}(\boldsymbol{\delta}) = E\left\{\sum_{i=1}^m E(\delta_i | \mathbf{x}, \mathbf{a}, \mathbf{b}) b_i(1 - \text{Lfdr}_i)\right\}$ for any decision rule $\boldsymbol{\delta}$. Combining equations (7.1) – (7.3) and noting that $t^* > 0$, we claim that $\text{ETP}(\boldsymbol{\delta}_{OR}) \geq \text{ETP}(\boldsymbol{\delta}^*)$ and the desired result follows.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- [1]. Benjamini Y, and Cohen R (2016), “Weighted False Discovery Rate Controlling Procedures for Clinical Trials,” *Biostatistics*, Advance Access.
- [2]. Benjamini Y, and Heller R (2007), “False Discovery Rates for Spatial Signals,” *Journal of the American Statistical Association*, 102, 1272–1281.
- [3]. Benjamini Y, Heller R, and Yekutieli D (2009), “Selective Inference in Complex Research,” *Philosophical Transactions of the Royal Society A*, 367, 4255–4271.
- [4]. Benjamini Y, and Hochberg Y (1995), “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing,” *Journal of the Royal Statistical Society, Series B*, 57, 289–300.
- [5]. Benjamini Y, and Hochberg Y (1997), “Multiple Hypotheses Testing with Weights,” *Scandinavian Journal of Statistics*, 24, 407–418.
- [6]. Box GEP (1950), “Problems in the Analysis of Growth and Wear Curves,” *Biometrics*, 6, 362–389. [PubMed: 14791573]
- [7]. Cai TT, and Sun W (2009), “Simultaneous Testing of Grouped Hypotheses: Finding Needles in Multiple Haystacks,” *Journal of the American Statistical Association*, 104, 1467–1481.
- [8]. Cai TT, and Sun W (2016), “Optimal Screening and Discovery of Sparse Signals with Applications to Multistage High-Throughput Studies,” *Journal of the Royal Statistical Society, Series B*, Accepted.
- [9]. Cai TT, Sun W, and Wang W (2016), “CARS: Covariate Assisted Ranking and Screening for Large-Scale Two-Sample Inference,” Technical Report.
- [10]. Dmitrienko A, Offen WW, and Westfall PH (2003), “Gatekeeping Strategies for Clinical Trials that Do Not Require All Primary Effects to be Significant,” *Statistics in Medicine*, 22, 2387–2400. [PubMed: 12872297]
- [11]. Efron B (2008), “Simultaneous Inference: When Should Hypothesis Testing Problems be Combined?” *The Annals of Applied Statistics*, 2, 197–223.
- [12]. Efron B (2010), *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*, New York: Cambridge University Press.
- [13]. Efron B, Tibshirani R, Storey JD, and Tusher V (2001), “Empirical Bayes Analysis of a Microarray Experiment,” *Journal of the American Statistical Association*, 96, 1151–1160.
- [14]. Ferkingstad E, Frigessi A, Rue H, Thorleifsson G, and Kong A (2008), “Unsupervised Empirical Bayesian Multiple Testing with External Covariates,” *The Annals of Applied Statistics*, 2, 714–735.
- [15]. Fox CS, Heard-Costa N, Cupples LA, Dupuis J, Vasan RS, and Atwood LD (2007), “Genome-Wide Association to Body Mass Index and Waist Circumference: The Framingham Heart Study 100K Project,” *BMC Medical Genetics*, 8, S18. [PubMed: 17903300]

- [16]. Genovese CR, Roeder K, and Wasserman L (2006), "False Discovery Control with P-Value Weighting," *Biometrika*, 93(3), 509–524.
- [17]. Genovese C, and Wasserman L (2002), "Operating Characteristics and Extensions of the False Discovery Rate Procedure," *Journal of the Royal Statistical Society, Series B*, 64, 499–517.
- [18]. He L, Sarkar SK, and Zhao Z (2015), "Capturing the Severity of Type II Errors in High-Dimensional Multiple Testing," *Journal of Multivariate Analysis*, 142, 106–116.
- [19]. Hu JX, Zhao H, and Zhou HH (2010), "False Discovery Rate Control With Groups," *Journal of the American Statistical Association*, 105, 1215–1227. [PubMed: 21931466]
- [20]. Jaquish CE (2007), "The Framingham Heart Study, on its Way to Becoming the Gold Standard for Cardiovascular Genetic Epidemiology?" *BMC Medical Genetics*, 8, 63. [PubMed: 17916250]
- [21]. Jin J, and Cai TT (2007), "Estimating the Null and the Proportion of Nonnull Effects in Large-Scale Multiple Comparisons," *Journal of the American Statistical Association*, 102, 495–506.
- [22]. Li L, Kabesch M, Bouzigon E, Demenais F, Farrall M, Moffatt MF, Lin X, and Liang L (2013), "Using eQTL Weights to Improve Power for Genome-Wide Association Studies: A Genetic Study of Childhood Asthma," *Frontiers in Genetics*, 4, Article 103.
- [23]. Lin W-Y, and Lee W-C (2012), "Improving Power of Genome-Wide Association Studies with Weighted False Discovery Rate Control and Prioritized Subset Analysis," *PLoS ONE*, 7, e33716. [PubMed: 22496761]
- [24]. Lin C-Y, Xing G, Ku H-C, Elston RC, and Xing C (2014), "Enhancing the Power to Detect Low-Frequency Variants in Genome-Wide Screens," *Genetics*, 196, 1293–1302. [PubMed: 24496013]
- [25]. Liu Y, Sarkar SK, and Zhao Z (2016), "A New Approach to Multiple Testing of Grouped Hypotheses," *Journal of Statistical Planning and Inference*, 179, 1–14.
- [26]. Pacifico MP, Genovese C, Verdinelli I, and Wasserman L (2004), "False Discovery Control for Random Fields," *Journal of the American Statistical Association*, 99, 1002–1014.
- [27]. P`ena EA, Habiger JD, and Wu W (2011), "Power-Enhanced Multiple Decision Functions Controlling Family-Wise Error and False Discovery Rates," *The Annals of Statistics*, 39, 556–583. [PubMed: 25018568]
- [28]. Roeder K, Bacanu S-A, Wasserman L, and Devlin B (2006), "Using Linkage Genome Scans to Improve Power on Association in Genome Scans," *The American Journal of Human Genetics*, 78, 243–252. [PubMed: 16400608]
- [29]. Roeder K, and Wasserman L (2009), "Genome-Wide Significance Levels and Weighted Hypothesis Testing," *Statistical Science*, 24, 398–413. [PubMed: 20711421]
- [30]. Rogosa D (2003), "Student Progress in California Charter Schools," 1999–2002. Available at <http://statweb.stanford.edu/rag/draft/charter9902.pdf>.
- [31]. Roquain E, and van de Wiel MA (2009), "Optimal Weighting for False Discovery Rate Control," *Electronic Journal of Statistics*, 3, 678–711.
- [32]. Silverman BW (1986), *Density Estimation for Statistics and Data Analysis*, London: Chapman & Hall.
- [33]. Spjøtvoll E (1972), "On the Optimality of Some Multiple Comparison Procedures," *The Annals of Mathematical Statistics*, 43, 398–411.
- [34]. Storey JD (2002), "A Direct Approach to False Discovery Rates," *Journal of the Royal Statistical Society, Series B*, 64, 479–498.
- [35]. Storey JD (2007), "The Optimal Discovery Procedure: A New Approach to Simultaneous Significance testing," *Journal of the Royal Statistical Society, Series B*, 69, 347–368.
- [36]. Sun W, and Cai TT (2007), "Oracle and Adaptive Compound Decision Rules for False Discovery Rate Control," *Journal of the American Statistical Association*, 102, 901–912.
- [37]. Sun W, Reich B, Cai TT, Guindani M, and Schwartzman A (2015), "False Discovery Control in Large-Scale Spatial Multiple Testing," *Journal of the Royal Statistical Society, Series B*, 77, 59–83.
- [38]. Westfall PH, and Young SS (1993), *Resampling-Based Multiple Testing*, New York: John Wiley & Sons.

- [39]. Xing C, Cohen JC, and Boerwinkle E (2010), “A Weighted False Discovery Rate Control Procedure Reveals Alleles at FOXA2 that Influence Fasting Glucose Levels,” *The American Journal of Human Genetics*, 86, 440–446. [PubMed: 20152958]
- [40]. Young SS, and Gries CL (1984), “Exploration of the Negative Correlation Between Proliferative Hepatocellular Lesions and Lymphoma in Rats and Mice – Establishment and Implications,” *Fundamental and Applied Toxicology*, 4, 632–640. [PubMed: 6479509]
- [41]. Zablocki RW, Schork AJ, Levine RA, Andreassen OA, Dale AM, and Thompson WK (2014), “Covariate-Modulated Local False Discovery Rate for Genome-Wide Association Studies,” *Bioinformatics*, 30(15), 2098–2104. [PubMed: 24711653]

Figure 1: Comparison under group-wise weights: $\square$ BH97 (or BH95), $\circ$ WPO, $\triangle$ DD (proposed), and $+$ AZ. The efficiency gain of the proposed method increases as c_1 and c_2 become more distinct.

Table 1:

Number of rejections and thresholds by varying choice of weights between groups.

Choice of Weights							
Group	Total	(1, 1, 1)	Threshold	(4, 2, 1)	Threshold	(10, 5, 1)	Threshold
$\alpha = 0.10$							
$\mathcal{G}_1$	1043	3	6.69×10^{-4}	5	1.70×10^{-3}	53	3.17×10^{-2}
$\mathcal{G}_2$	5277	23	9.01×10^{-4}	33	1.52×10^{-3}	57	4.53×10^{-3}
$\mathcal{G}_3$	220392	2019	2.92×10^{-3}	1997	2.88×10^{-3}	1663	2.31×10^{-3}
Overall	226712	2045	–	2035	–	1773	–
$\alpha = 0.05$							
$\mathcal{G}_1$	1043	1	5.22×10^{-6}	2	3.95×10^{-4}	4	7.43×10^{-4}
$\mathcal{G}_2$	5277	13	4.22×10^{-4}	19	5.20×10^{-4}	21	6.61×10^{-4}
$\mathcal{G}_3$	220392	159	1.17×10^{-4}	147	1.06×10^{-4}	131	1.01×10^{-4}
Overall	226712	173	–	168	–	156	–