



Many Phenotypes Without Many False Discoveries: Error Controlling Strategies for Multitrait Association Studies

Christine B. Peterson,¹ Marina Bogomolov,² Yoav Benjamini,³ and Chiara Sabatti^{4,5*}

¹Department of Health Research and Policy, Stanford University, Stanford, California, United States of America; ²Faculty of Industrial Engineering and Management, Technion, Haifa, Israel; ³Department of Statistics and Operations Research, Tel Aviv University, Tel Aviv, Israel; ⁴Department of Health Research and Policy, Stanford University, Stanford, California, United States of America; ⁵Department of Statistics, Stanford University, Stanford, California, United States of America

Received 21 April 2015; Revised 27 August 2015; accepted revised manuscript 19 October 2015.

Published online in Wiley Online Library (wileyonlinelibrary.com). DOI 10.1002/gepi.21942

ABSTRACT: The genetic basis of multiple phenotypes such as gene expression, metabolite levels, or imaging features is often investigated by testing a large collection of hypotheses, probing the existence of association between each of the traits and hundreds of thousands of genotyped variants. Appropriate multiplicity adjustment is crucial to guarantee replicability of findings, and the false discovery rate (FDR) is frequently adopted as a measure of global error. In the interest of interpretability, results are often summarized so that reporting focuses on variants discovered to be associated to some phenotypes. We show that applying FDR-controlling procedures on the entire collection of hypotheses fails to control the rate of false discovery of associated variants as well as the expected value of the average proportion of false discovery of phenotypes influenced by such variants. We propose a simple hierarchical testing procedure that allows control of both these error rates and provides a more reliable basis for the identification of variants with functional effects. We demonstrate the utility of this approach through simulation studies comparing various error rates and measures of power for genetic association studies of multiple traits. Finally, we apply the proposed method to identify genetic variants that impact flowering phenotypes in *Arabidopsis thaliana*, expanding the set of discoveries.

Genet Epidemiol 00:1–12, 2015. Published 2015 Wiley Periodicals, Inc.*

KEY WORDS: error control; false discovery rate; genetic association study; multiple phenotypes; multiple hypothesis testing

Introduction

Biotechnological progress has enabled the routine measurement of thousands of phenotypes that were beyond the reach of precise quantification just a couple of decades ago. Together with the reduced costs of genotyping and sequencing, this motivates research into the genetic basis of an unprecedented number of traits. Examples include expression quantitative trait loci (eQTL) studies that investigate the role of genetic variation on the expression of tens of thousands of genes [Brem et al., 2002; Cheung et al., 2005; Schadt et al., 2003]; genome-wide metabolomics studies that consider genetic influences on the levels of hundreds of metabolites [Illig et al., 2010; Keurentjes et al., 2006]; and proteomics studies investigating genetic regulation of protein abundances [Foss et al., 2007; Wu et al., 2013]. At a more macroscopic level, neuroimaging genetics aims to identify DNA variants influencing brain structures, described in thousands of voxels [Stein et al., 2010]. Looking at even higher level phenotypes, a number of large cohorts with rich phenotypic information have been or are being genotyped and will be used to map multiple

traits. Notable examples are the Kaiser Permanente Research Program on Genes, Environment, and Health (RPGEH) [<http://www.dor.kaiser.org/external/DORExternal/rpgeh/>] that has already genotyped 100,000 subjects with complete medical records, and the Million Veteran Program [<http://www.research.va.gov/MVP/>] that is aiming to genotype a million veterans with available health records.

Investigating the genetic basis of thousands of traits simultaneously offers exciting possibilities, including the hope that a comprehensive and multifaceted description of the health status of a subject can provide a strong foundation for understanding relevant genetic underpinnings. Capitalizing on these possibilities requires appropriate statistical approaches to address the challenges posed by these novel data sets. Here, we focus on one such problem: namely, the development of multiple-testing procedures to identify discoveries while controlling an appropriate measure of error. Two choices need to be made upfront: (1) what notion of error to control; and (2) what is to be considered a discovery. We discuss these at the beginning of our study. In what follows, the terms “trait” and “phenotype” are used interchangeably; similarly, and with a slight abuse, “SNP” (single nucleotide polymorphism) and “variant” are considered synonymous.

*Correspondence to: Chiara Sabatti, Department of Health Research and Policy, Stanford University, 150 Governor's Lane, Stanford, CA 94305-5405.

E-mail: sabatti@stanford.edu

The genetics community has been acutely aware of the necessity of accounting for the “look across the genome” effect. Even before genome-wide linkage (or association) studies were a possibility, sequential test procedures [Morton, 1955] and Bayesian arguments [Elston and Lange, 1975] led to the adoption of very stringent significance cutoffs. Once large marker panels became available and multiple testing became a reality, efforts focused on controlling the probability of making at least one erroneous finding, a criteria known as the familywise error rate (FWER) [Feingold et al., 1993; Lander and Kruglyak, 1995]. This is well suited to investigate the genetic basis of a single disease assumed to be substantially influenced by one or two loci, especially when following up a hit implies years of work. The nature of present-day multi-trait investigations, however, is substantially different: when one explores the genetic basis of tens of thousands of traits, as in eQTL studies, insisting on not making even one mistake is overly severe. Indeed, on the heels of the experience in analysis of gene expression data [Efron et al., 2001; Reiner et al., 2003], in eQTL and other -omics investigations, another more liberal criteria has emerged as the dominant paradigm: the false discovery rate (FDR) [Benjamini and Hochberg, 1995]. The FDR is defined as the expected proportion of findings that are erroneous, meaning that they correspond to situations where the null hypothesis is actually true. The present work adopts the point of view that such a criteria better reflects the goals of multiphenotype studies where one expects to make a sizable number of discoveries, and it is acceptable to have a few false leads as long as these represent a small proportion of the findings [Benjamini and Yekutieli, 2005; Schadt et al., 2003].

In order to control FDR one needs to define a discovery. What constitutes an interesting finding? The identification of a variant that influences a specific phenotype? The determination that there is a genetic component to the variability of a trait? The discovery that one DNA variant is not neutral? All of the above? In which order of importance? To resolve these questions it is useful to look at the typical multiphenotype genome-wide association study (GWAS): this consists in testing the hypothesis H_{it} of no association between variant i and trait t for all values of i and t . This rather simplistic approach is often preferred for its limited computational cost, its robustness to missing data, and—most importantly—the ease with which results on different phenotypes and SNPs can be compared across different studies. The collection of tested hypotheses $\{H_{it} \ i = 1, \dots, M; \ t = 1, \dots, P\}$ can be considered as a single group, but it is also quite natural to identify subgroups of hypotheses that address one specific scientific question, technically referred to as *families*. Note that—following the convention in multiple comparison literature—we here use the term “family” to indicate a collection of hypotheses rather than a group of related individuals; pedigrees do not play a role in the discussion. One can consider the families $\mathcal{P}_t = \{H_{it} \ i = 1, \dots, M\}$ of all hypotheses related to the phenotype t , addressing the existence of a genetic basis for the t th trait. Alternatively, one can focus on the families $\mathcal{F}_i = \{H_{it}, \ t = 1, \dots, P\}$ of all hypotheses involving SNP i , investigating the phenotypic effect of each genetic variant i . To these families we can associate global null hypotheses:

$H_{i\bullet} = \bigcap_{t=1}^P H_{it}$ signifies that variant i does not affect any trait, while $H_{\bullet t} = \bigcap_{i=1}^M H_{it}$ states that trait t is not influenced by any variant. Identifying a relevant family structure is important both because families are the appropriate universe for multiplicity adjustment and because they define discoveries. Ultimately this choice is study specific, but here we make one both in the interest of concreteness and to underscore a viewpoint that is often relevant. In most multiphenotype GWAS, scientists have solid reason to believe that the traits under investigation have a genetic underpinning, so rejecting $H_{\bullet t}$ would not represent an interesting discovery. In contrast, we expect most genetic variants to have no effect on any trait, so identifying those that are “functional” can arguably be considered the most important discovery of multiphenotype investigations. Consider, for example, eQTL studies: the discovery of SNPs that influence the expression of some genes is important as they are considered potential candidates for association with a variety of other medically relevant traits. Indeed, the reported results from multiphenotype GWAS tend to be organized in terms of associated variants. In what follows, then, we consider the hypotheses $\{H_{it} \ i = 1, \dots, M; \ t = 1, \dots, P\}$ as organized in M families \mathcal{F}_i defined by variants, and we identify the rejection of $H_{i\bullet}$ as an important discovery. We note that the grouping of *hypotheses* into families is not related to the dependence structure among the *tests* for these hypotheses. Rather, the information captured by hypotheses in the same family needs to be somewhat “exchangeable”: for example, rejecting any of the hypotheses H_{it} within the family \mathcal{F}_i has the same implication with respect to the hypothesis $H_{i\bullet}$. Once a decision has been made that the hypotheses under consideration can be grouped in different families, it becomes relevant and meaningful to talk about a variety of global error measures, as we are about to describe.

Material and Methods

Global Error Measures for Structured Hypotheses

We start by considering one simple example where we assume that we know the true status of the hypotheses and we can measure the realized false discovery proportion (FDP).

Table 1 presents a total of 40 hypotheses, relative to eight phenotypes and five variants, which define families \mathcal{F}_i ,

Table 1. Example of Structured Hypotheses

\mathcal{F}_1	\mathcal{F}_2	\mathcal{F}_3	\mathcal{F}_4	\mathcal{F}_5
$\star H_{11}$	H_{12}	$\star H_{13}$	H_{14}	H_{15}
H_{21}	H_{22}	$\star H_{23}$	H_{24}	H_{25}
$\star H_{31}$	H_{32}	$\star H_{33}$	H_{34}	H_{35}
H_{41}	$\star H_{42}$	$\star H_{43}$	H_{44}	H_{45}
H_{51}	H_{52}	$\star H_{53}$	H_{54}	H_{55}
$\star H_{61}$	H_{62}	$\star H_{63}$	H_{64}	H_{65}
H_{71}	H_{72}	H_{73}	H_{74}	H_{75}
H_{81}	H_{82}	H_{83}	H_{84}	H_{85}

Bold hypotheses are false null, and starred hypotheses correspond to rejections. The 40 Hypotheses H_{11}, \dots, H_{85} are Grouped into Families $\mathcal{F}_1, \dots, \mathcal{F}_5$.

Table 2. Relevant Error Rates and Measures of Power

Abbreviation	Description	Definition
gFDR	Global FDR	$E(V/\max\{R, 1\})$
FDR _{<i>i</i>}	FDR within family <i>i</i>	$E(V_i/\max\{R_i, 1\})$
aFDR	Average of within-family FDRs	$\frac{1}{M} \sum_{i=1}^M E(V_i/\max\{R_i, 1\})$
sFDR	Expected value of the average of within family FDPs across selected families	$E\left[\frac{1}{\max\{ \mathcal{S}(\mathbf{P}) , 1\}} \times \sum_{i \in \mathcal{S}(\mathbf{P})} (V_i/\max\{R_i, 1\})\right]$
vFDR	FDR for the discovery of variants (families)	$E(V^v/\max\{R^v, 1\})$
gFWER	Global familywise error rate	$E(I[V > 0])$
vFWER	FWER for the discovery of variants (families)	$E(I[V^v > 0])$
gPower	Global power	Proportion of false null $H_{i\bullet}$ rejected
vPower	Power to detect variants associated to at least one phenotype	Proportion of false null $H_{i\bullet}$ rejected

$i = 1, \dots, 5$. We use bold to indicate hypotheses that are false null (where signal/association is present) and asterisks to indicate hypotheses that are rejected. A variant is discovered if the corresponding family contains at least one rejected hypothesis. In Table 1 there are a total of 10 individual hypotheses rejected and two of these are true nulls: the global FDP equal to 2/10. Families \mathcal{F}_1 , \mathcal{F}_2 , and \mathcal{F}_3 are discovered, but all the hypotheses in \mathcal{F}_2 are true nulls: the proportion of false family discoveries is 1/3. The average FDP across all families is $0.2\bar{3} = (0 + 1 + 1/6 + 0 + 0)/5$; but if we focus only on families that have been discovered, the average FDP across selected families is $0.3\bar{8} = (0 + 1 + 1/6)/3$.

With this example in mind, we can define a variety of error rates. Let \mathbf{P} indicate the collection of P -values associated with all the individual hypotheses tested. Let $\mathcal{S}(\mathbf{P})$ be a selection procedure (which can depend on the observed P -values) that identifies interesting variants. Let R be the total number of rejections and V the total number of erroneous rejections across all hypotheses. Similarly, V_i and R_i count the false discoveries and total discoveries in family i . We say that variant i is discovered if the corresponding global null $H_{i\bullet}$ is rejected. We indicate with R^v and V^v , respectively, the total number of rejections and the total number of false discoveries among the M variant hypotheses $H_{i\bullet}$ s, each probing the role of a different variant i . In Table 2, we rely on these symbols to define the error rates of interest.

Given these alternatives, what error rate is relevant and important to control? The global FDR (gFDR) is a natural first choice as this is the error rate we would control if we had not identified a family structure among our hypotheses. Despite the appeal of its simplicity, there are caveats to be considered when targeting gFDR. As shown eloquently in Efron [2008], pooling hypotheses from multiple families that have different proportions of true nulls and controlling gFDR can result in rather suboptimal behavior: for families that contain none or very few false nulls, FDR_{*i*} will not be controlled at the desired level, while families with many false nulls will encounter a loss in power. If one targets FDR_{*i*} for each family separately, these difficulties are overcome but at the price of a large number of false discoveries: while average of within-family FDR (aFDR) would be controlled, gFDR and sFDR would not. In addition,

if we consider $H_{i\bullet}$ rejected as long as one of H_{it} , $t = 1, \dots, P$ is, it is important to note that neither of the two strategies above controls vFDR or sFDR.

To illustrate these characteristics, we run a simulation with 300,000 hypotheses corresponding to $P=100$ phenotypes and $M=3,000$ variants. Families are defined by variants and contain only true null hypotheses, with the exception of 60 variants each associated to 25 phenotypes. P -values corresponding to the true null hypotheses are generated independently from a uniform distribution on the $[0, 1]$ interval. Test statistics for the false null hypotheses are generated independently from the $\mathcal{N}(2, \sigma^2)$ distribution, and the corresponding P -values are computed as the two-tailed P -values under the $\mathcal{N}(0, \sigma^2)$ distribution. Because larger values of the standard deviation σ make these two distributions more difficult to distinguish, we can interpret σ as the noise level. Figure 1 shows a set of global error measures as the noise level increases. We also provide two measures of power: gPower represents global power, and vPower represents power to detect variants associated to at least one phenotype. We compare three approaches for the analysis of the data sets: (a) the Benjamini-Hochberg (BH) method [Benjamini and Hochberg, 1995] applied to the pooled collection of all P -values with target level $q = 0.05$ for gFDR (“pooled BH”); (b) BH applied to each family separately with target level $q = 0.05$ for each FDR_{*i*} (“per family BH”); and (c) a hierarchical strategy we will discuss in the following section and include here for reference (“hierarchical BH”). Figure 1 illustrates how both (a) “pooled BH” and (b) “per family BH” control their target error rate (gFDR and aFDR, respectively), but not vFDR or sFDR. When (a) BH is applied to the entire collection of hypotheses, the false rejections are uniformly distributed across the true null hypotheses; in a context where many variants affect no phenotypes, this results in false variant discoveries. Furthermore, once we restrict attention to the families with at least one rejection, many have a within-family FDP close to 1: we do not have control of the error we make when declaring association between phenotypes and the selected SNPs.

If we apply BH in a per family manner (b), the aFDR is controlled: many families lead to no discoveries, resulting in a FDP equal to 0, which lowers the average FDR. However, the families associated with discovered variants tend to have very large FDP: neither sFDR or gFDR is controlled. From a certain point of view, applying BH to each family separately can be considered as ignoring the multiplicity due to different variants, so it is not surprising that vFDR and gFDR are quite high with this approach. In summary, (b) does not appear to be a viable strategy whenever M is large. We now introduce procedure (c) that overcomes this impasse.

Hierarchical Testing Procedure

Benjamini and Bogomolov [2014] describe how to control sFDR when families are selected according to a rather broad set of criteria. Here, we build upon their work and suggest selecting families so as to control the vFDR: this allows us to provide both guarantees on the discovered variants and

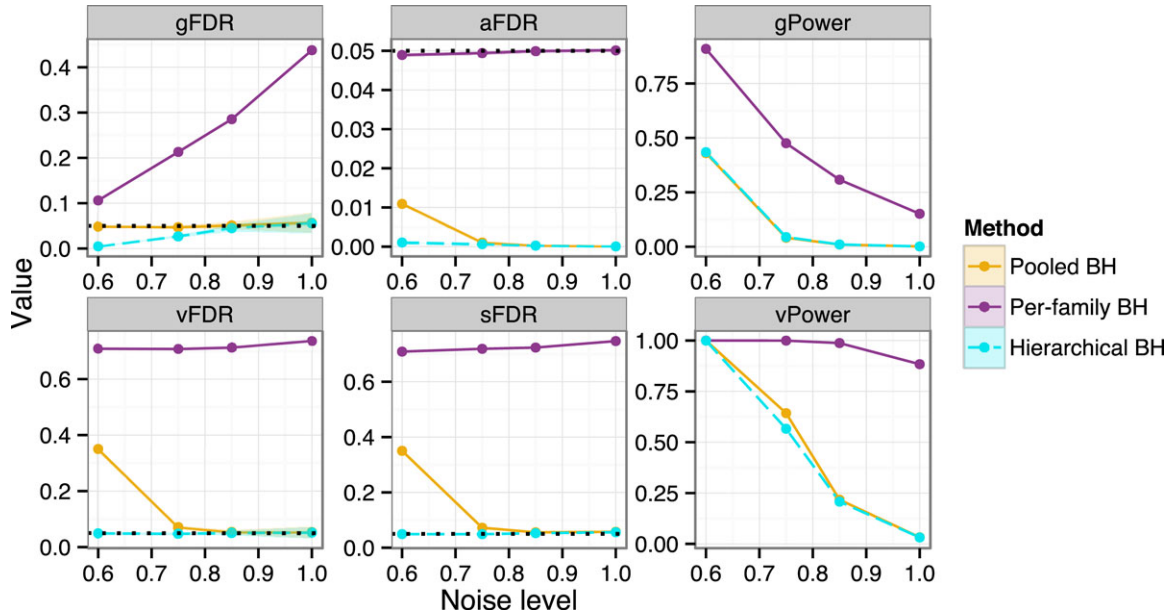


Figure 1. Global error achieved by different multiple comparisons controlling strategies as noise level increases. $M = 3,000$; $P = 100$; 60 variants are associated to 25 phenotypes and the rest have no association. The lines indicate the average and the shaded areas the standard error over 250 iterations.

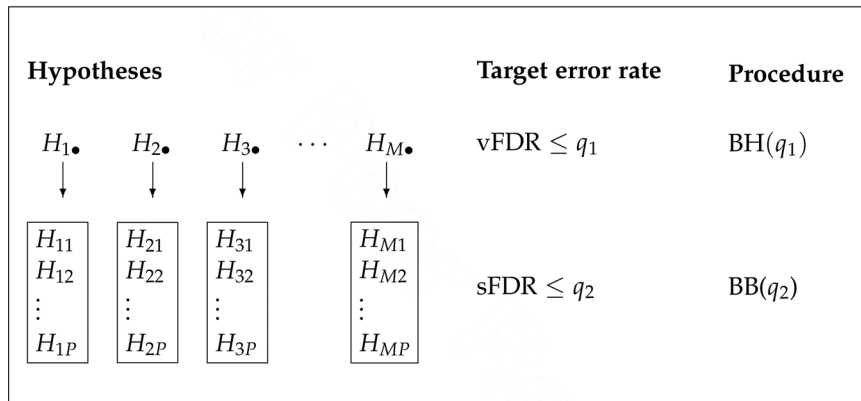


Figure 2. Hierarchical structure of hypotheses.

on the identification of the phenotypes they influence. To avoid unnecessary complexity, we assume that each family contains the same number of hypotheses, although this is not necessary.

We aim to control FDR on the collection of M global null hypotheses $H_{i\bullet} = \cap_{t=1}^P H_{it}$ $\{H_{i\bullet}, i = 1, \dots, M\}$ at level q_1 . Once a set of interesting families $\{\mathcal{F}_i, i \in \mathcal{S}\}$ has been identified by controlling the vFDR, we aim to control the sFDR, that is the average FDR on the selected families, at level q_2 . See Figure 2 for an illustration of the relevant hierarchical structure.

Testing is carried out on the basis of the P -values p_{it} obtained for each of the individual hypotheses H_{it} . The P -values for the intersection hypotheses $H_{i\bullet}$ are defined as the Simes's

P -values [Simes, 1986] for the respective families:

$$p_{i\bullet} = \min_{t=1, \dots, P} \frac{P p_{i(t)}}{t} \quad (1)$$

where $p_{i(t)}$ represents the t th ordered element of the vector $\{p_{it}, t = 1, \dots, P\}$. The hierarchical procedure is as follows:

Testing Procedure 1.

Stage 0 Use Simes's method to obtain P -values $p_{i\bullet}$ s for the intersection hypotheses $H_{i\bullet}$ s.

Stage 1 Apply BH to the collection of P -values $\{p_{i\bullet}, j = 1, \dots, M\}$ with an FDR target level q_1 . Let $\mathcal{S}(\mathbf{P})$

indicate the set of i corresponding to rejected hypotheses $H_{i\bullet}$.

Stage 2 Proceed to test the individual hypotheses H_{it} only in families \mathcal{F}_i with $i \in \mathcal{S}(\mathbf{P})$. Within such families, apply BH with target level $q_2 \times \frac{|\mathcal{S}(\mathbf{P})|}{M}$, the appropriate adjustment for the selection bias introduced in Stage 1.

Testing Procedure 1 guarantees vFDR control when the Simes's P -values are valid P -values for the intersection hypotheses and when BH applied to $\{p_{i\bullet}, i = 1, \dots, M\}$ controls FDR. It also guarantees control of sFDR when BH applied to each family \mathcal{F}_i controls FDR within the family and the P -values in each family are independent of the P -values in any other family, or when the pooled set of P -values satisfies a certain positive dependence property (see later for more details regarding the control of vFDR and sFDR of Testing Procedure 1 under dependence). Figure 1 illustrates how the hierarchical procedure controls vFDR and sFDR in the setting of the simulation described in the previous section. In the remainder of this paper, we will explore in some detail when conditions for Testing Procedure 1 to control its target error rate are satisfied and how applicable they are to the tests we encounter in GWAS with multiple phenotypes. First, however, some remarks are useful.

- In Stage 0, we suggested using Simes's P -value for three reasons: it can be easily constructed from the single hypothesis P -values; it is robust to most common types of dependence between the test statistics in the family [Sarkar, 1998; Hochberg, 2006]; and, finally, its combination with BH leads to consistent results between stages, as will be discussed in more detail later. However, other choices are possible and might be more effective in specific situations. For example, when the tests across phenotypes can be considered independent, it might be advantageous to combine P -values using Fisher's rule [Fisher, 1932]: this might lead to the identification of SNPs that have a very modest effect on multiple phenotypes, so that their influence can only be gathered by combining these effects. If appropriate distributional assumptions are satisfied, another choice might be the higher criticism statistic [Donoho and Jin, 2004]. Finally, one might obtain a P -value $p_{j\bullet}$ for the intersection hypothesis by means other than the combination of the P -values for individual hypotheses. For example, one can use a reverse regression approach as in O'Reilly et al. [2012], in which a regression is fit for each genetic variant treating the full set of phenotypes as the predictors and the SNP genotype as an ordinal response.
- Stage 1 focuses on the discovery of interesting families, which correspond to genetic variants associated with variability in phenotypes: a multiplicity adjustment that controls the desired error rate on $\{H_{i\bullet}, i = 1, \dots, M\}$ needs to be in place. For FDR control we rely on BH, which has been shown to perform well under the types of dependence across markers present in the GWAS setting [Schadt et al., 2003]. The more conservative Benjamini-Yekutieli procedure

[Benjamini and Yekutieli, 2001], with its theoretical guarantees, is also possible. Some might prefer to control FWER at this level via a Bonferroni procedure: this would be in keeping with the criteria routinely adopted in GWASs. In the simulations that follow, we explore the properties of this approach as well.

- Stage 2 identifies phenotypes associated with interesting SNPs. It rests on the results in Benjamini and Bogomolov [2014]: to control the average error rate across the selected families at level q_2 , one has to perform a multiplicity adjustment within each family at a more stringent level $q_2 \times \frac{|\mathcal{S}(\mathbf{P})|}{M}$ to account for the selection effect. Again, this result is more general than implied in Testing Procedure 1. For example, one might want to control the average FWER across selected families: this would be possible by using Bonferroni at the appropriate level. It is useful to observe the interplay of selection penalty and Bonferroni correction. If only one family is selected, the threshold for significance is $\frac{q_2}{M}$, the same that would result from applying Bonferroni to the entire collection of hypotheses. If all families are selected, the threshold for significance is simply $\frac{q_2}{P}$, and there is no price for multiplicity across families. When more than one family is selected, the threshold is between these two. In general, it can be shown that controlling the average FWER across selected families is more liberal than controlling global FWER. It is not possible to make such a general statement with respect to FDR, but it remains true that the hierarchical procedure has the potential of increasing power by reducing the multiple comparisons burden via relevant selection of which hypotheses to test.
- Testing Procedure 1 controls sFDR in Stage 2 by controlling FDR within each selected family at a more stringent level. One interesting aspect of this approach is that BH is applied to each selected family separately: this allows for adaptivity to the family-specific proportion of true nulls, overcoming one of the limitations of BH applied to the entire collection of hypotheses.
- Stages 1 and 2 are governed by two separate testing procedures. Generally speaking, this could imply that the set of discoveries in the two steps are not in perfect correspondence: one could reject the intersection null hypothesis corresponding to a variant, but not reject any of the single hypotheses on the association between that variant and the individual phenotypes. The setup of Testing Procedure 1—where P -values for the intersection hypotheses are obtained with Simes's rule and Stages 1 and 2 use BH—assures that this is not the case whenever $q_1 \leq q_2$: as long as the global null corresponding to one variant is rejected, this variant is declared to be associated with at least one phenotype.

Results

Simulations with Independent Tests

To illustrate the operating characteristics of the hierarchical procedure, we rely first on simulations with all tests

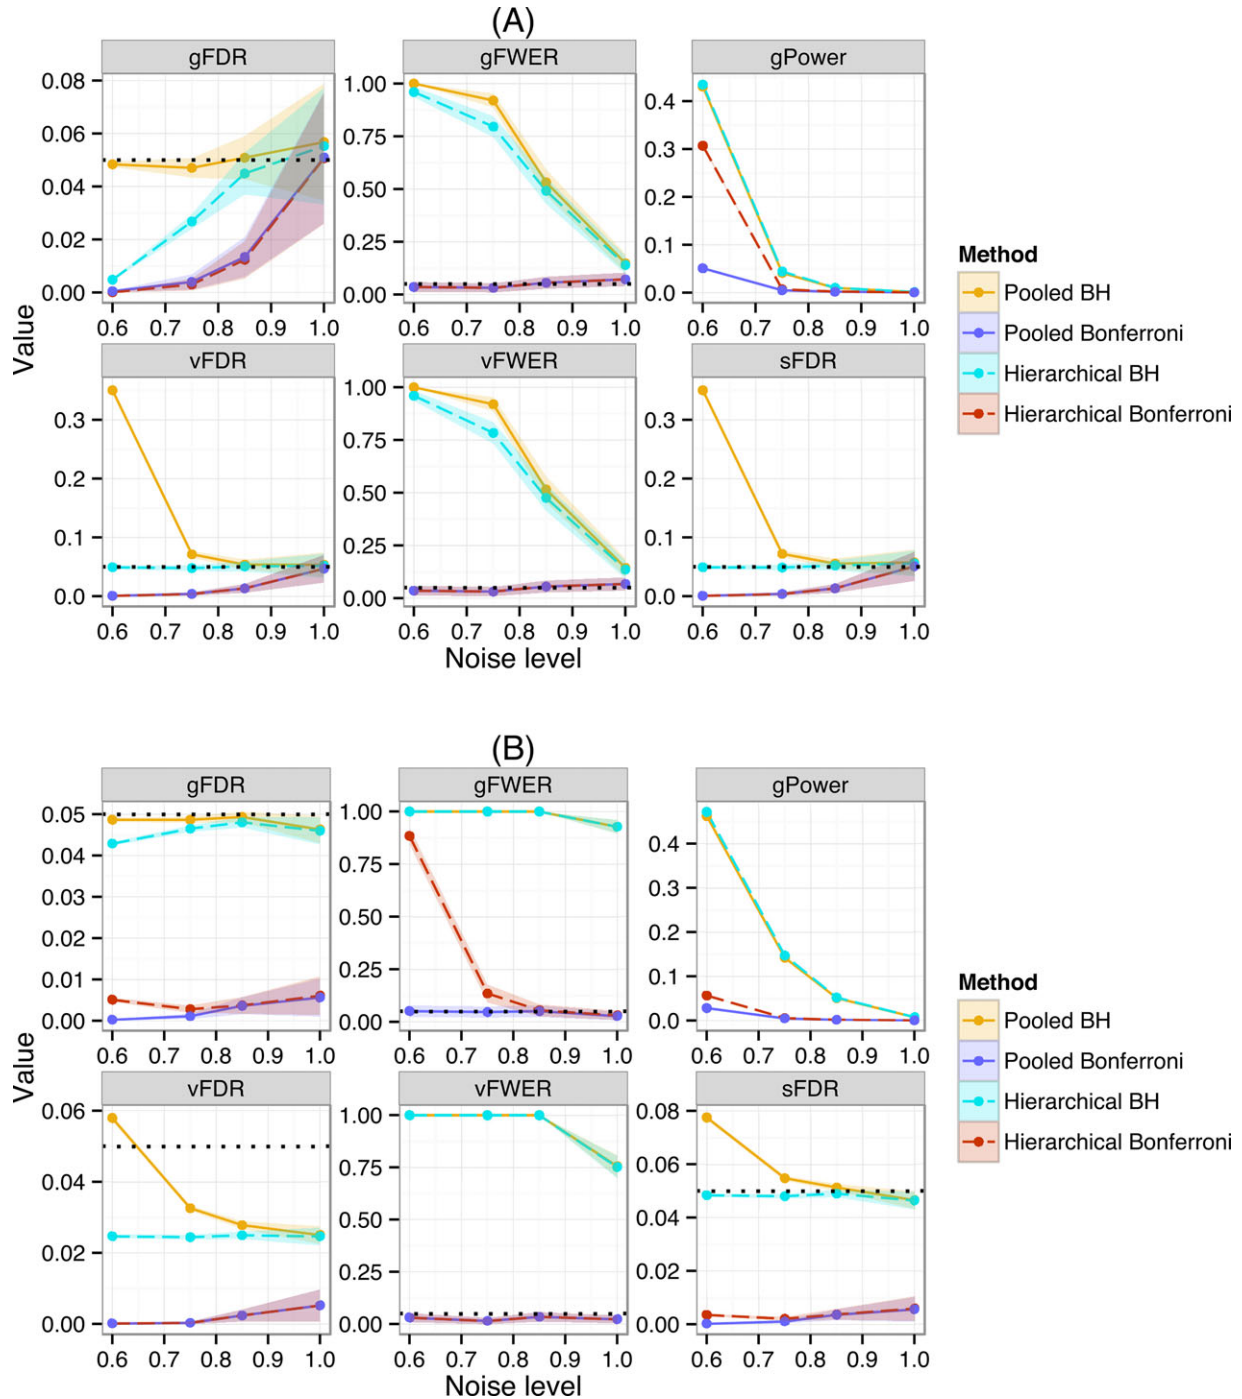


Figure 3. Error rates and power for four multiple-testing strategies. $M = 3,000$, $P = 100$ and test statistics are independent. In (A) 60 variants are associated with 25 traits each and in (B) 1,500 variants are associated with five phenotypes each. The solid lines show the average, the shaded areas represent the standard error over 250 iterations, and the dotted horizontal lines mark the level 0.05.

independent. Exploration of typical GWAS dependence will be discussed in the next section. Figure 3 summarizes the results of two scenarios: $M=3,000$, $P=100$ and in (A) 60 variants are associated with 25 phenotypes (as in Fig. 1), while in (B) 1,500 variants are associated with five phenotypes. P -values were generated as for Figure 1. Four strategies are compared: (a) gFDR control with BH (“pooled BH”); (b) Bonfer-

roni targeting gFWER (“pooled Bonferroni”); (c) Testing Procedure 1 (“hierarchical BH”); (d) hierarchical testing targeting vFWER, via Bonferroni applied on the Simes’s P -values, and sFDR (“hierarchical Bonferroni”). The target for all error rates is 0.05.

All procedures control their respective targeted error rates, and the two hierarchical procedures also control gFDR. The

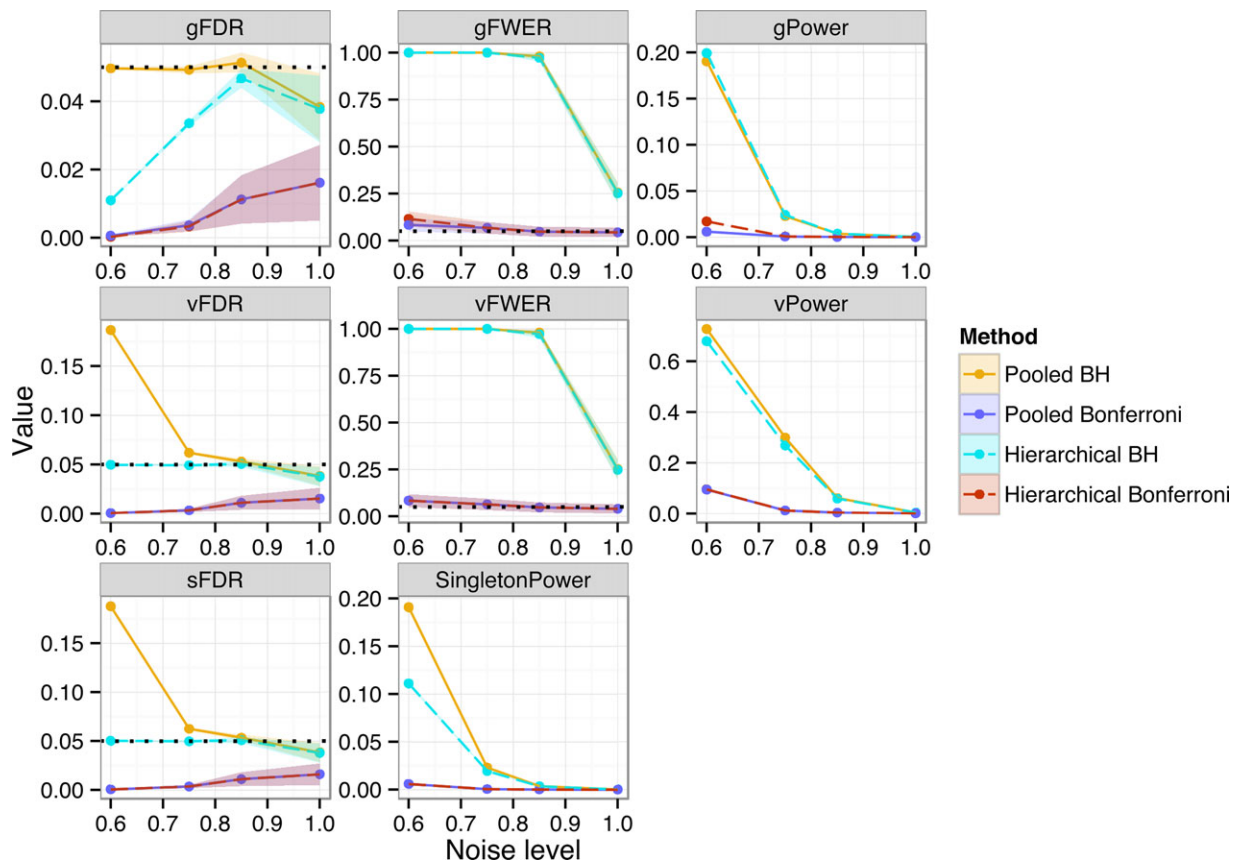


Figure 4. Error rates and power for four multiple-testing strategies. $M = 100,000$, $P = 100$, and 1,000 variants are associated with 25 traits and 500 variants are associated with one trait each. The lines show the average, the shaded areas represent the standard error over 250 iterations, and the dotted horizontal lines mark level 0.05.

power of the hierarchical procedure that controls vFDR is comparable to that of applying BH to the entire collection of hypotheses, and the power of the procedure that targets vFWER is comparable to or better than that of Bonferroni on the entire collection. The hierarchical procedures show an advantage when the families with non-null hypotheses are a small subset of the total families. In such cases, BH applied to the pooled collection of P -values fails to control vFDR and sFDR. This is precisely the situation we expect to hold in GWAS: only a small proportion of SNPs are associated to any phenotype. The substantial increase in power of “hierarchical Bonferroni” over “pooled Bonferroni” in (A) is due to the adaptivity of BH to the proportion of false null hypotheses in the families: when an SNP is selected, which has effects on multiple phenotypes, it becomes easier to detect these associations.

Given that the relative advantages of the procedures we are considering depend on the number of families and the number of true null hypotheses they contain, we run a simulation with dimensions that should resemble that of a GWAS involving multiple traits: 100,000 SNPs and 100 phenotypes. In Figure 4 most of the families contain only true null hypotheses, except for 1,000 variants that are associated with 25 phenotypes and 500 variants that are

associated with one phenotype each. This last type of family is included both to account for phenotype-specific effects and to evaluate the possible loss of power in detecting these variants for the hierarchical strategy: in addition to the global power (gPower), we report power to detect variants (vPower) and power to detect variants that affect only one phenotype (SingletonPower). As expected, simply applying BH to the entire collection of hypotheses results in a substantial increase of the vFDR and sFDR, with no substantial power advantage. Indeed, the overall power is better for the hierarchical strategy, even if this encounters a loss of power to detect SNPs that are associated with only one phenotype. Although these simulations were based on independent test statistics, we obtained similar results when the test statistics within each family were correlated, as is the case when there is dependence among phenotypes. Specifically, when the test statistics within each family were sampled from a multivariate normal distribution with a covariance matrix inducing spatial correlation, the measured error rates and relative performance of the different error control strategies were consistent with those shown in Figure 3. For more details on this simulation, see http://web.stanford.edu/~cbp/Peterson_multi_trait_supplement.pdf. In the next section, we explore the effects of more complex dependence including

correlation across the predictor variables on the hierarchical procedure.

GWAS Dependence Structure

The markers typed in GWAS are typically chosen to span the entire genome at a high density. SNPs in the same neighborhood are not independent, but in linkage disequilibrium. This redundancy assures that the typed markers can effectively act as proxies for untyped variants and is one of the sources of dependency relevant for our study.

To understand other departures from independence, it is useful to look at the relationship between phenotypes and genotypes and the methods with which these are analyzed. In its simplest form, the data-generating model considered by geneticists to link each phenotype t to genotypes is $y_{it} = \mathbf{x}'_i \boldsymbol{\beta} + \epsilon_i$, where ϵ_i are uncorrelated and i indicates subjects. The coefficient vector $\boldsymbol{\beta}$ is thought to be sparse (i.e., with a small proportion of nonzero elements) or effectively sparse in the sense that a small portion of the coefficients have appreciable size. When considering multiple phenotypes and n subjects, this translates into

$$\mathbf{Y} = \mathbf{X}\mathbf{B} + \mathbf{E}, \quad (2)$$

where $\mathbf{Y}_{n \times P}$, $\mathbf{X}_{n \times M}$, $\mathbf{B}_{M \times P}$, and $\mathbf{E}_{n \times P}$ are matrices containing phenotypes, genotypes, coefficients, and error terms, respectively. Although most of the rows of \mathbf{B} are full of zeros, some rows are expected to contain more than one nonzero element, corresponding to genomic locations that influence multiple phenotype (pleiotropy): the resulting phenotypes are not independent, even when the elements of the error matrix are *iid*.

GWAS data are generally analyzed using a collection of univariate regressions linking each phenotype t to one genetic variant i :

$$\hat{Y}_{[i,t]} = \hat{\alpha} + \mathbf{x}_{[i,t]} \hat{\beta}_{it} + \hat{\mathbf{E}}_{[i,t]}, \quad (3)$$

and the hypothesis H_{it} translates into $H : \beta_{it} = 0$, tested with the standard t -statistics. Clearly, the discrepancy between even the theoretical model (2) and the regression (3) used for analysis leads to a number of consequences. For example, as the error terms $\hat{\mathbf{E}}_{[i,t]}$ cannot be expected to be uncorrelated across individuals, linear mixed models are often used in single phenotype analysis [Kang et al., 2010]. Moreover, the combination of spatial dependence existing across SNPs and the univariate testing approach (3) induces spatial structure among both the test statistics and the hypotheses. Consider the case of a complete null where the phenotypes under study have no genetic underpinning. If by random chance one variant appears to have some explanatory power for one phenotype, the P -values of neighboring SNPs will also tend to have lower values—this is dependence among the test statistics. Consider now a data-generating model (2) where variant i has a coefficient different from zero while its neighbors do not. With respect to model (2) H_{it} is false and the H_{it} for neighboring SNPs l are true. However, once we decide to look at the data through the lenses of (3), the

hypotheses H_{it} are redefined to mean the lack of any association between SNP l and phenotype t and—as long as SNP l can act as a reasonable proxy for one of the causal variants— H_{it} is false. We expect clusters of null hypotheses corresponding to neighboring SNPs to be false or true together. Indeed, in GWAS studies it is common to find a number of nearby variants significantly associated with the trait: this is interpreted as evidence for the presence of one or more causal variants in the specific genomic region. Looking at multiple phenotypes that might share genetic determinants adds another layer to this phenomenon.

On the one hand, dependence between test statistics can be problematic for multiplicity adjustment strategies. The Bonferroni approach controls FWER even if tests are dependent; the BH procedure, instead, is guaranteed to control FDR under independence or positive regression dependence on a subset (PRDS) [Benjamini and Yekutieli, 2001], even if it has been empirically observed to provide FDR control under broader conditions. When the BH procedure controls FDR under the dependence of the P -values within each family and the P -values in each family are independent of the P -values in any other family, the Testing Procedure 1 controls vFDR and sFDR. Provided that certain overall positive dependence properties hold, these error rates remain controlled when the P -values across the families are not independent. In particular, when the pooled set of P -values is PRDS, sFDR is controlled (see Theorem 3 in Benjamini and Bogomolov [2014]; note that this is the same condition needed for pooled BH to control gFDR). In addition, it can be concluded from the simulation results of Benjamini and Heller [2008] that when $\{p_{it}, i = 1, \dots, M\}$ are PRDS for each $t \in \{1, \dots, P\}$, and when $\{p_{it}, t = 1, \dots, P\}$ are PRDS for each $i \in \{1, \dots, M\}$, vFDR is controlled.

On the other hand, the fact that tested hypotheses H_{it} are defined with respect to (3) rather than the data generative model (2) makes it challenging to evaluate the error made by a multiple-testing procedure: if we use (2) as ground truth, we expect many false rejections that really do not correspond to a mistake with reference to (3). In order to avoid this problem, we consider all the hypotheses relative to variants that are within 1 Mb and have correlation of magnitude at least 0.2 to a causal variant in the generative model as correctly rejected.

For the simulations below, we use genotype data obtained from 1966 Northern Finland Birth Cohort (NFBC) Sabatti et al. [2009]. We exclude copy number variants and markers with P -values for Hardy-Weinberg equilibrium below 1×10^{-5} , with minor allele frequency (MAF) less than 0.01, or with call rate less than 95%. This screening results in $M = 334,103$ SNPs on $n = 5,402$ subjects. We code SNPs by minor allele count and impute missing genotypes by average variant allele count. We simulate $P = 100$ traits. In each iteration, we select 130 SNPs at random and use them to generate phenotypes, as follows: the first 10 SNPs affect 50 phenotypes, the next 10 affect 25, the next 10 affect 10, and the final 100 each affect five phenotypes, always chosen at random. In this set up, each trait reflects the contribution of 13.5 SNPs on average. The more than 300,000 SNPs remaining have no

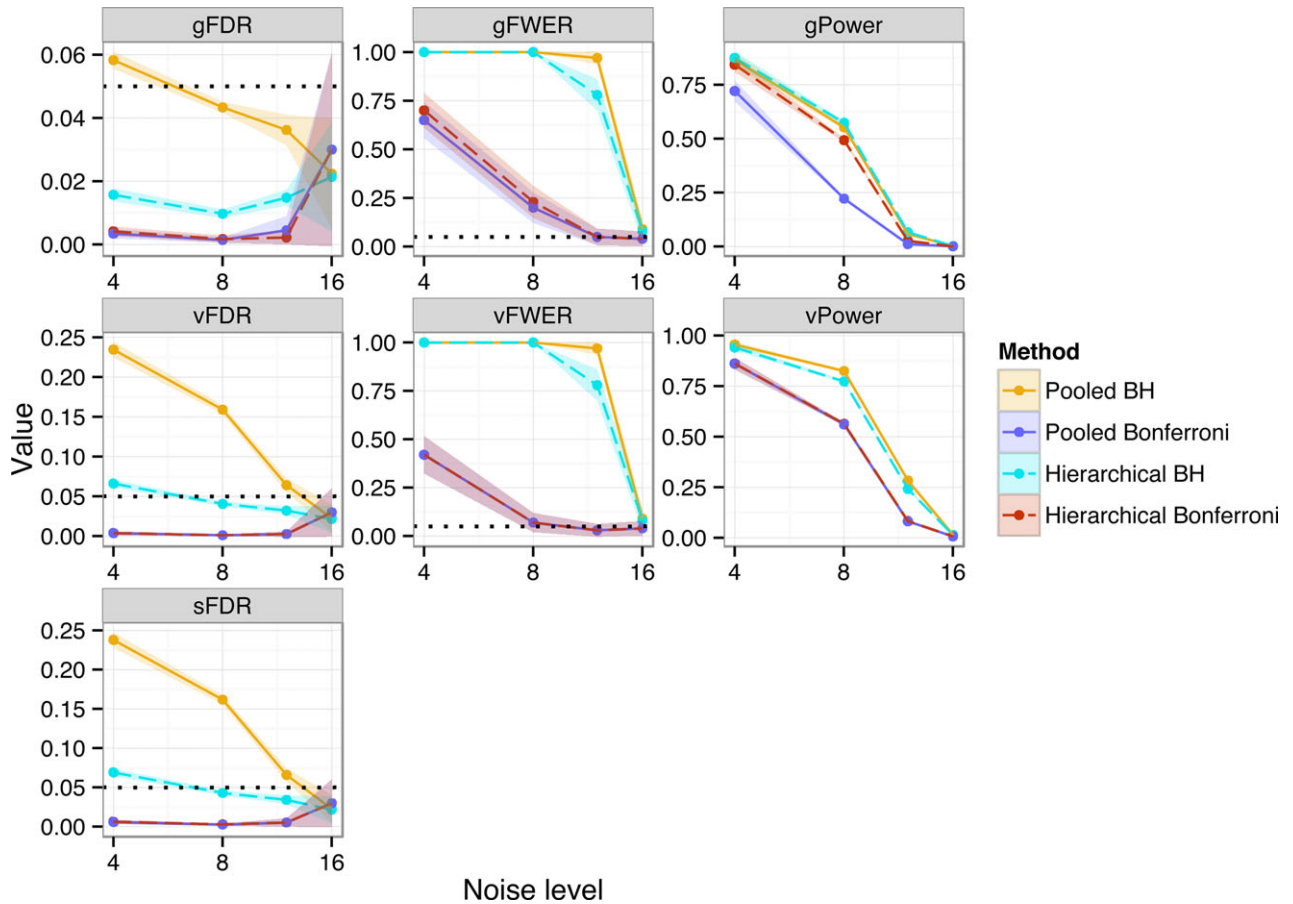


Figure 5. Error rates and power for four multiple-testing strategies applied to simulated data starting from real genotypes. The lines show the average, the shaded areas report the standard error over 100 iterations, and the dotted horizontal lines mark the 0.05 level.

functional role. To generate the simulated traits, we follow the linear model in equation (2), where \mathbf{B}_{it} is 1 in presence of an association between variant i and trait t and 0 otherwise. In summary, the test statistics generated as part of this simulation will exhibit dependence due to two sources: spatial correlation between neighboring SNPs and shared genetic background among phenotypes.

Due to the large number of hypotheses under consideration, we rely on *MatrixEQTL* [Shabalin, 2012] to allow efficient computation of the P -values of association. This software, originally designed for the analysis of eQTL data, utilizes large matrix operations to increase computational speed and has the option to reduce the required memory by saving only P -values beneath a given threshold. As long as this threshold is above the P -value cutoff for selection under all error control methods, this shortcut does not affect the results. In applying *MatrixEQTL*, we use a threshold of 5×10^{-4} for saving output and include the first five principal components of the genotype data as covariates to adjust for the effects of population structure.

Under varying levels of noise σ , we compare four adjustment strategies studied before. When analyzing the results, we consider a discovery a true positive if it lies within 1 Mb

and has correlation at least 0.2 to the truly causal SNP. The results, given in Figure 5, show that even with this allowance, there are still settings where some of the methods under consideration fail to control their target error rates. In particular, pooled Bonferroni fails to control gFWER and hierarchical Bonferroni fails to control vFWER for settings with higher levels of power. In addition, gFDR is somewhat above 0.05 for pooled BH and vFDR exceeds 0.05 for hierarchical BH in the setting with highest power. Rather than a failure of the multiple comparisons procedure, this is to be attributed to the confusion induced by the use of model (3) to analyze data generated with model (2); when we rerun the analysis using phenotypes adjusted for the effects of variables omitted by the univariate model, these errors appear appropriately controlled. FWER is more sensitive to these misspecification errors simply because one single mistake is enough to raise the realized FWE to 1; in contrast, as long as these mistakes are accompanied by a number of true discoveries, the realized FDP will only be marginally inflated. Focusing on the performance of hierarchical methods compared, we again conclude that they appear to control their targeted error rates whenever the corresponding pooled approach controls gFDR or gFWER.

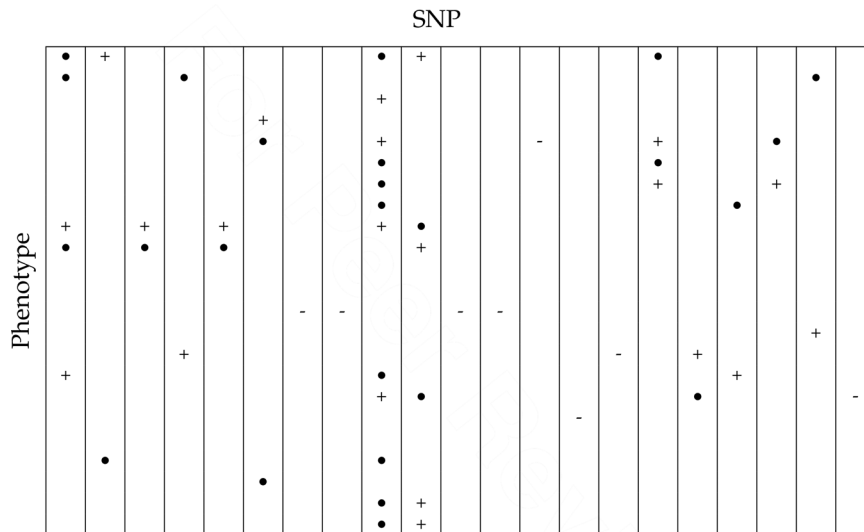


Figure 6. Families with differential results from pooled BH and hierarchical BH for the case study. Discoveries under both methods are marked with •, discoveries made only under pooled BH are marked with “-”, and discoveries made only under hierarchical BH are marked with “+.”

Case Study: Flowering in *Arabidopsis thaliana*

We use Testing Procedure 1 to reanalyze data on the genetic basis of flowering phenotypes in *A. thaliana* [Atwell et al., 2010] online at <https://cynin.gmi.oew.ac.at/home/resources/atpolydb/genomic-polymorphism-data-in-arabidopsis-thaliana>. Although the original study includes 109 different traits, we focus on 23 phenotypes related to flowering including days to flowering under different conditions, plant diameter at flowering, number of leaves at flowering, etc.; the results in Atwell et al. [2010] indicate that a shared genetic basis is likely for at least some of these traits. Genotypes are available for 199 inbred lines at 216,130 SNPs.

To obtain P -values of association, we follow the steps described in Atwell et al. [2010]: exclude SNPs with a MAF ≤ 0.1 , transform certain phenotypes to the log scale, and fit the variance components model implemented in Kang et al. [2008], which allows us to account for population structure. The original analysis underscored the difficulties of identifying true positives only on the basis of statistical considerations and did not attempt formal multiplicity adjustment. Although these challenges clearly still stand, here we compare the results of applying BH across the full set of P -values targeting gFDR at level 0.05, with those of Testing Procedure 1 targeting vFDR and sFDR, each at level 0.05. This means that for the hierarchical procedure, we have $M = 216, 130$ families corresponding to SNPs, each consisting of 23 hypotheses.

Hierarchical BH identifies 131 variants versus the 139 of pooled BH, reflecting a tighter standard for variant discovery. At the same time, hierarchical BH increases global power over pooled BH, resulting in a total of 174 discoveries versus 161: an increase of 8%. The variants that pooled BH discovers in excess of hierarchical BH are declared associated to one phenotype only. There are 7% fewer such SNPs according to the hierarchical procedure. Figure 6 presents

variants with different results under the two methods: eight SNPs discovered by pooled BH as associated with only one phenotype are not selected by hierarchical BH, while several SNPs discovered under pooled BH are associated to a larger number of phenotypes by hierarchical BH. For example, the SNP in column 1 of Figure 6 corresponds to a particular location in the short vegetative phase (SVP) gene, that is known to be involved in flowering and associated to two additional phenotypes under the hierarchical method.

Discussion

Contemporary genomic investigations result in testing very large number of hypotheses, making it vital to adopt appropriate strategies for multiplicity adjustment: the risk of lack of reproducibility of results is too high to be overlooked. When the collection of tested hypotheses has some structure, discoveries often occur at multiple levels and reports typically do not focus on the rejection of hypotheses at the finest scales. In the hope of increasing both power and interpretability, scientists often attempt to outline an overall picture with statements that are supported by groups of hypotheses. We considered one example of such situations: in GWASs concerning a large number of phenotypes the primary object of inference is often the identification of variants that are associated to any trait.

The simulations presented make clear that in these settings it is necessary to identify what is to be considered a discovery and to perform a multiplicity adjustment that allows one to control measures of global error defined on the discoveries of interest. By adapting the work in Benjamini and Bogomolov [2014], we outline one such strategy and explore its performance and relative advantages in the context of GWAS studies involving multiple phenotypes.

Our hierarchical strategy aims at (a) identifying SNPs that affect some phenotypes (while controlling errors at this level) and (b) detecting which phenotypes are influenced by such SNPs (controlling the average error measure across selected SNPs). Aim (a) focuses on the discovery of SNPs with any functional effects, while aim (b) focuses on the association of these SNPs to the specific phenotypes they influence. Following the terminology defined in Table 2, the measures of global error controlled in stages (a) and (b) are v FDR and s FDR, respectively. We consider two error measures: FDR and FWER. We show that while our strategy achieves these goals, applying FDR controlling rules (as BH) on the entire collection of hypotheses (“pooled BH”) does not control the FDR of the discoveries in (a) and (b): whenever the reporting of results emphasizes these, other multiplicity adjustments need to be in place. On the other hand, the “hierarchical BH” procedure is not guaranteed to control the global FDR (gFDR) in general, but it effectively appears to do so in the situations we simulated. Applying Bonferroni to the pooled collection of hypotheses does control FWER for the discoveries in (a) and s FDR for the discoveries in (b), but it is excessively conservative if these are the target error rates. Conversely, the “hierarchical Bonferroni” strategy does not control global FWER.

To complete this summary of results, we shall make a few remarks. First, while the application to GWAS studies has motivated us and guided the exposition of material as well as some specific implementation choices, it is important to note that Testing Procedure 1 is applicable to much broader settings. It simply rests on the possibility of organizing the entire collection of tested hypotheses in groups of separate families, each probing a meaningful scientific question.

Second, it is worth noting that the hierarchical strategy represents one example of valid selective inference. More and more, as the modalities of data collection become increasingly comprehensive rather than targeted, scientists tend to “look at the data first and ask questions later.” In other words, initial exploratory analyses are used to identify possible meaningful patterns and formulate precise hypotheses for formal statistical testing. When this is the case, however, the traditional rules for determining significance are inappropriate and procedures that account for the selection effects are called for. The work of Benjamini and Bogomolov [2014] that we adapt here is an important step in this direction.

Moving on to the specific implications for multiphenotype GWAS, the results of our simulations using actual genotypes contribute to the debate on whether to choose FDR or FWER as targeted error rate. The combination of correlation between SNPs and misspecification of the linear model that is routinely used in GWAS applications can result in the rejection of hypotheses of no association between an SNP and a phenotype even when the SNP has no causal effect and is reasonably far from any causal variants. In procedures that target FDR control, these “false” rejections are accompanied by a number of correct ones and their effect on the error rate is modest. Conversely, the presence of even one such wrong rejection equates the realized FWE to one: this makes it very hard to control FWER in situations other than global null.

Because of the disparities in targeted error rates, it is difficult to contrast the power of the hierarchical and pooled strategies as this comparison is most meaningful across procedures that guarantee the same error level. However, it is of practical relevance to contrast the number and characteristics of true findings that a researcher can expect when adopting the pooled and the hierarchical procedure targeting the respective error rates at the 0.05 level. Both the BH strategies appear to control global FDR and our simulations indicate that overall power is quite similar: the pooled approach discovers more SNPs that truly affect a single phenotype and the hierarchical approach discovers more SNPs that affect multiple phenotypes. The same trend is evident in the real-data analysis. Note that the FDR among SNPs that are declared associated with one phenotype by the pooled BH strategy can be very high. Both Bonferroni strategies control the FWER of SNP discoveries and the expected value of the average FDP for SNP-phenotype associations across selected SNPs: the hierarchical approach (which does not control global FWER) has greater power, once again thanks to the increased discovery of SNPs associated to multiple phenotypes.

Although we have not discussed this so far, it will not have escaped the attentive reader that the hierarchical procedure we propose can be applied in meta-analyses of GWAS studies of the same trait. In this setting, one typically has independence across studies and multiple powerful choices of P -value for the global nulls are available in Stage 0. The contribution of the hierarchical procedure in this context is in Stage 2, where studies with significant association are identified.

A remark is in order with reference to the application of the proposed approach to multiphenotype GWAS studies. Although we found that our method is robust to correlation among phenotypes given independent predictors, we have not fully explored the results of dependence across phenotypes due to environmental components in a realistic GWAS setting. Consider eQTL studies where the traits are measurements of expression levels of multiple genes: it has been repeatedly observed that experimental batch effects can result in strong dependence between traits. If such correlation between phenotypes is present, it would be crucial to account for it in the method of analysis used to define P -values. In absence of this, it is quite possible that some of the environmental effects might be accidentally correlated with the genotype value of some of the SNPs in the study resulting in a number of false positives that would be exacerbated by the hierarchical approaches. Indeed, the procedures we outlined here are valid as long as the P -values used as input are accurate; obtaining such P -values is clearly of paramount importance. There have been several recent proposals that rely on the linear mixed model framework to handle correlation across phenotypes in the context of GWAS with population structure [Joo et al., 2015; Korte et al., 2012; Zhou et al., 2014]; these could provide a useful alternative to P -values obtained via standard linear regression models when correlation across traits is strong. Alternatively, the issue of dependence among phenotypes could be addressed by constructing independent traits using a technique such as principal component analysis (PCA). In fact, there are

many possible methods that could be used; an advantage of the proposed hierarchical testing procedure is that it is quite flexible and can be applied given association P -values from whichever approach is preferred for a particular context.

Finally, we want to point out that, while our focus has been the problem of multiple comparison in multitrait GWAS studies that start with obtaining P -values for the hypotheses of association between each SNP and each phenotype, there are other approaches to this problem. Specifically, one can explicitly model the dependence between phenotypes and use statistical approaches that identify loci underlying these multivariate traits. The linkage literature has some interesting examples [Williams et al., 1999], and steps along these directions have been also documented in association studies [Flutre et al., 2013; O'Reilly et al., 2012]. One caveat to keep in mind is that in some cases explicit multivariate models can be very computationally intensive and become impractical when the number of phenotypes is quite large. When, however, a more powerful test for the hypotheses $H_{i\bullet}$ can be obtained with these approaches, we recommend using them in Stage 0 of our procedure.

Software

The hierarchical error control methods described in this paper have been implemented as a part of the TreeQTL R package [Peterson, 2015], available online at <http://bioinformatics.org/treeqtl>.

Acknowledgments

The Northern Finland Birth Cohort 1966 (NFBC1966) data used in the article were obtained via dbGaP, accession number phs000276.v1.p1. C. Peterson was supported by a CEHG fellowship and by NIH grant MH101782; M. Bogomolov was partially supported by Israel Science Foundation grant no. 1112/14; Y. Benjamini was partially supported by NIH HG006695; C. Sabatti was partially supported by NIH MH101782 and HG006695. We are grateful to these sponsors and to H. Tombropoulos for editorial assistance.

References

Atwell S. 2010. Online data for Atwell et al. 2010. <https://cynin.gmi.oeaw.ac.at/home/resources/atpolydb/genomic-polymorphism-data-in-arabidopsis-thaliana> [Accessed: 2015].

Atwell S, Huang YS, Vilhjálmsson BJ, Willems G, Horton M, Li Y, Meng D, Platt A, Tarone AM, Hu TT, and others. 2010. Genome-wide association study of 107 phenotypes in *Arabidopsis thaliana* inbred lines. *Nature* 465:627–631.

Benjamini Y, Bogomolov M. 2014. Selective inference on multiple families of hypotheses. *J R Stat Soc B* 76:297–318.

Benjamini Y, Heller R. 2008. Screening for partial conjunction hypotheses. *Biometrics* 64:1215–1222.

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

Benjamini Y, Yekutieli D. 2001. The control of the false discovery rate in multiple testing under dependency. *Ann Stat* 29:1165–1188.

Benjamini Y, Yekutieli D. 2005. Quantitative trait loci analysis using the false discovery rate. *Genetics* 171:783–790.

Brem RB, Yvert G, Clinton R, Kruglyak L. 2002. Genetic dissection of transcriptional regulation in budding yeast. *Science* 296:752–755.

Cheung VG, Spielman RS, Ewens KG, Weber TM, Burdick JT. 2005. Mapping determinants of human gene expression by regional and genome-wide association. *Nature* 437:1365–1369.

Donoho D, Jin J. 2004. Higher criticism for detecting sparse heterogeneous mixtures. *Ann Stat* 32:962–994.

Efron B. 2008. Simultaneous inference: when should hypothesis testing problems be combined? *Ann Appl Stat* 2:197–223.

Efron B, Tibshirani R, Storey JD, Tusher V. 2001. Empirical Bayes analysis of a microarray experiment. *J Am Stat Assoc* 96:1151–1160.

Elston RC, Lange K. 1975. The prior probability of autosomal linkage. *Ann Hum Genet* 38:341–350.

Feingold E, Brown PO, Siegmund D. 1993. Gaussian models for genetic linkage analysis using complete high-resolution maps of identity by descent. *Am J Hum Genet* 53:234–251.

Fisher RA. 1932. *Statistical Methods for Research Workers*. Edinburgh: Oliver and Boyd.

Flutre T, Wen X, Pritchard J, Stephens M. 2013. A statistical framework for joint eQTL analysis in multiple tissues. *PLoS Genet* 9:e1003486. DOI: 10.1371/journal.pgen.1003486

Foss EJ, Radulovic D, Shaffer SA, Ruderfer DM, Bedalov A, Goodlett DR, Kruglyak L. 2007. Genetic basis of proteome variation in yeast. *Nat Genet* 39:1369–1375.

Hochberg Y, Hommel G. 2006. Simes' test of multiple hypotheses. In: Kotz S, Read CB, Balakrishnan N, Vidakovic B, editors. *Encyclopedia of Statistical Sciences*. Wiley, New York, pp. 1–5.

Illig T, Gieger C, Zhai G, Römisch-Margl W, Wang-Sattler R, Prehn C, Altmäier E, Kastenmüller G, Kato BS, Mewes H, and others. 2010. A genome-wide perspective of genetic variation in human metabolism. *Nat Genet* 42:137–141.

Joo JW, Kang EY, Org E, Furlotte N, Parks B, Lusis AJ, Eskin E. 2015. Efficient and accurate multiple-phenotypes regression method for high dimensional data considering population structure. In: Przytycka T, editor. *Research in Computational Molecular Biology*. Springer, pp. 136–153.

Kang HM, Sul JH, Service SK, Zaitlen NA, Kong S, Freimer NB, Sabatti C, Eskin E. 2010. Variance component model to account for sample structure in genome-wide association studies. *Nat Genet* 42:348–354.

Kang HM, Zaitlen NA, Wade CM, Kirby A, Heckerman D, Daly MJ, Eskin E. 2008. Efficient control of population structure in model organism association mapping. *Genetics* 178:1709–1723.

Keurentjes JJ, Fu J, Ric De Vos CH, Lommen A, Hall RD, Bino RJ, van der Plas LH, Jansen RC, Vreugdenhil D, Koornneef M. 2006. The genetics of plant metabolism. *Nat Genet* 38:842–849.

Korte A, Vilhjálmsson BJ, Segura V, Platt A, Long Q, Nordborg M. 2012. A mixed-model approach for genome-wide association studies of correlated traits in structured populations. *Nat Genet* 44:1066–1071.

Lander E, Kruglyak L. 1995. Genetic dissection of complex traits: guidelines for interpreting and reporting linkage results. *Nat Genet* 11:241–247.

Million Veteran Program. <http://www.research.va.gov/MVP/> [Accessed: 2015].

Morton NE. 1955. Sequential tests for the detection of linkage. *Am J Hum Genet* 7:277–318.

O'Reilly PF, Hoggart CJ, Pomyen Y, Calboli FC, Elliott P, Jarvelin M, Coin LJ. 2012. MultiPhen: joint model of multiple phenotypes can increase discovery in GWAS. *PLoS One* 7:e34861.

Peterson CB, Bogomolov M, Benjamini Y, Sabatti C. 2015. TreeQTL: hierarchical error control for eQTL findings. *bioRxiv*. doi: <http://dx.doi.org/10.1101/021170>.

Reiner A, Yekutieli D, Benjamini Y. 2003. Identifying differentially expressed genes using false discovery rate controlling procedures. *Bioinformatics* 19:368–375. Research Program on Genes, Environment, and Health. <http://www.dor.kaiser.org/external/DORExternal/rpgeh/> [Accessed: 2015].

Sabatti C, Service SK, Hartikainen A, Pouta A, Ripatti S, Brodsky J, Jones CG, Zaitlen NA, Varilo T, Kaakinen M, and others. 2009. Genome-wide association analysis of metabolic traits in a birth cohort from a founder population. *Nat Genet* 41:35–46.

Sarkar SK. 1998. Some probability inequalities for ordered MTP2 random variables: a proof of the Simes conjecture. *Ann Stat* 26:494–504.

Schadt EE, Monks SA, Drake TA, Lusis AJ, Che N, Colinao V, Ruff TG, Milligan SB, Lamb JR, Cavet G, and others. 2003. Genetics of gene expression surveyed in maize, mouse and man. *Nature* 422:297–302.

Shabalina AA. 2012. Matrix eQTL: ultra fast eQTL analysis via large matrix operations. *Bioinformatics* 28:1353–1358.

Simes RJ. 1986. An improved Bonferroni procedure for multiple tests of significance. *Biometrika* 73:751–754.

Stein JL, Hua X, Lee S, Ho AJ, Leow AD, Toga AW, Saykin AJ, Shen L, Foroud T, Pankratz N, and others. 2010. Voxelwise genome-wide association study (vGWAS). *Neuroimage* 53:1160–1174.

Williams J, Van Eerdewegh P, Almasly L, Blangero J. 1999. Joint multipoint linkage analysis of multivariate qualitative and quantitative traits. I. Likelihood formulation and simulation results. *Am J Hum Genet* 65:1134–1147.

Wu L, Candille SI, Choi Y, Xie D, Jiang L, Li-Pook-Than J, Tang H, Snyder M. 2013. Variation and genetic control of protein abundance in humans. *Nature* 499:79–82.

Zhou X, Stephens M. 2014. Efficient multivariate linear mixed model algorithms for genome-wide association studies. *Nat Methods* 11:407–409.