A novel significance-based approach to two-stage hypothesis testing for discovery of gene-environment interactions

Eric S. Kawaguchi[a], Andre E. Kim[a], Juan Pablo Lewinger[a], and W. James Gauderman[a]
[a]Department of Population and Public Health Sciences, University of Southern California

Abstract

Two-step tests for gene-environment ($G \times E$) interactions, exploit marginal SNP effects to improve the power of genome-wide interaction scan (GWIS). They combine a screening step based on marginal effects with weighted hypothesis testing in the second step, to deliver greater power than single-step tests while preserving the genome-wide type I error. However, the presence of too many SNPs with detectable marginal effects on the trait of interest can make two-step tests suboptimal. We provide a novel framework that unifies aspects from two widely-adopted two-step hypothesis testing procedures. In step 1 we derive significance-based cutoffs to allocate tests into bins for step 2 $G \times E$ testing. Furthermore, we suggest a computationally efficient approach to take genetic correlation into account in the testing step by correcting for the effective number of independent tests within a bin. Simulation results demonstrate that this new two-step approach provides greater power over current methods under several scenarios. An application to a multi-study collaboration for understanding colorectal cancer (CRC) reveals a $G\times$Sex interaction located within the SMAD7 gene.

1 Introduction

Identifying gene-environment interactions ($G \times E$) is critical for understanding how health is affected by both an individual's diverse genetic background (G) and their exposure to environmental factors (E). Genome-wide interaction scans (GWIS) test $G \times E$ interactions one-SNP-at-a-time by modeling the genotype, an environmental exposure, and the corresponding interaction term. Testing of the $G \times E$ is based on the significance of the interaction term, typically from a logistic (for a binary trait), linear (for a quantitative trait), or Cox (1972) (for a survival trait) regression model. A standard one-step GWIS proceeds by testing each $G \times E$ interaction at significance level of $\alpha^* = 5 \times 10^{-8}$, common in genome-wide association studies (GWAS) or GWIS studies (Dudbridge and Gusnanto, 2008). However, the statistical power to detect an interaction in a one-step GWIS is generally much lower than the power for detecting a genetic marginal effect in a GWAS.

Two-step tests for GWIS have been proposed to improve the power of a $G \times E$ analysis while controlling the FWER for disease (Kooperberg and LeBlanc, 2008; Murcray et al., 2009, 2011; Hsu et al., 2012; Gauderman et al., 2013; Wang et al., 2021), quantitative traits (Paré et al., 2010; Zhang et al., 2016), and time-to-event traits (Kawaguchi et al., 2022). Independent information on $G \times E$ not captured by the standard $G \times E$ test, is used to perform an initial screening (Step 1) to prioritize SNPs that are more likely to be involved in an interaction. These SNPs are formally tested for an interaction (Step 2) under a modified

α^*, thus reducing the multiple testing burden (Kooperberg and LeBlanc, 2008; Murcray et al., 2009).

Two commonly-used screening statistics for case-control studies are 1) the marginal outcome-gene association statistic derived from modeling the outcome on each gene on the disease individually(Kooperberg and LeBlanc, 2008) and 2) the marginal exposure-gene association statistic, from modeling the relationship between each gene on the exposure(Murcray et al., 2009, 2011; Hsu et al., 2012; Gauderman et al., 2013). The marginal screen can also be used for quantitative (Zhang et al., 2016) and survival traits (Kawaguchi et al., 2022). A key requirement for validity of any two-step procedure is that the statistics used in Step 1 and Step 2 (Dai et al., 2012; Kawaguchi et al., 2022).

There are two widely-used procedures for prioritizing SNPs in Step-2 $G \times E$ testing after the Step-1 screening: subset (Kooperberg and LeBlanc, 2008; Murcray et al., 2009) and weighted hypothesis testing (Ionita-Laza et al., 2007). In subset testing, only the SNPs that pass a significance threshold based on the screening statistic are tested in Step 2 using a standard $G \times E$ test. The significance threshold in step 2 for $G \times E$ discovery is based on a Bonferroni correction based on the number of SNPs that pass the screening ($\alpha^* = \alpha/m$), which is much less stringent than the threshold used in a single step approach. A trade off for a relaxed threshold is that SNPs that do not pass the step 1 screening will not be tested. An alternative approach that does not rely on a pass/no pass hard rule is weighted hypothesis testing. Here, SNPs are allocated into bins based on the magnitude of the screening statistic. Each bin has a corresponding bin-wise error rate (BWER) such that the sum across all bins does not exceed α. Lower bins are allocated a larger fraction of α (see Section 2.1 for more detail), so that SNPs in those bins are tested at a more liberal significance threshold. SNPs in higher bins at much more stringent BWER. Unlike subset testing, every SNP is tested in Step 2 of the weighed approach; yet SNPs that are more likely to have an interaction, based on the screening statistic, will having a higher chance of being discovered. Power of the overall procedure depends on the size of each bin, which can be determined to achieve good power across a range of scenarios. Ionita-Laza et al. (2007) defined bin sizes based on a predetermined bin size scheme. Although weighted hypothesis testing is often more powerful than subset testing(Ionita-Laza et al., 2007; Gauderman et al., 2013), there is no universally most powerful approach.

The motivation behind two-step hypothesis testing is that in the presence of a true $G \times E$ interaction effect, one can typically expect there to be marginal G effects on the outcome, which makes the marginal outcome-gene statistic useful for screening/ranking SNPs in Step 1 (e.g being placed in bin 1, the bin with the largest BWER). However, this only proves useful if not too many SNPs have sizeable marginal effects but no $G \times E$ interaction. This is often not the case in a GWIS where known "hits" from a prior GWAS provides a set of SNPs for which there is a marginal outcome-gene effect. For example, more than 140 GWAS-significant (marginal-effect) loci have been previously identified for colorectal cancer (Huyghe et al., 2019) and the majority of these likely do not exhibit a $G \times E$ interaction. This will result in a phenomena we refer to as "bin overcrowding", where these SNPs will overcrowd the earlier bins due to having non-zero marginal effects. Thus, even if a true $G \times E$ effect admits a non-zero marginal effect, its competing against other known (or previously unknown) non-zero marginal effects which can force the true $G \times E$ effect to not be optimally tested (e.g. being placed/tested in a later bin with a stricter BWER). To avoid this loss of power, one can filter out these loci in advance. However, this

approach requires prior knowledge of loci with marginal effects, and it also removes these SNPs from consideration for $G \times E$ testing.

The contribution of this paper is the improvement of two-step GWIS testing in two ways. First, we propose a simple yet effective approach to prioritize tests in the screening step that eliminates the potential power loss due to the presence of SNPs with a marginal effect but no interaction effect. Second, we show how additional accounting for genetic correlation in the testing step can further increase power. We demonstrate via simulation that this new two-step testing enhancement yields greater power than its predecessors. We apply the approach to identify $G \times E$ interactions for colorectal cancer.

We describe current methods for testing $G \times E$ interactions using the two-step hypothesis testing paradigm and our proposed approach in Section 2. Simulation studies that compare the performance of two-step hypothesis testing procedures are given in Section 3 and an application to the FIGI study is provided in Section 4. Lastly, in Section 5, we provide concluding remarks, limitations, and areas of future research.

2 Methods

Consider a gene-environment study with a (continuous, binary, or time-to-event) trait/outcome Y_i, an environmental exposure of interest E_i, and M SNPs measured or imputed for each of the N subjects, with $\mathbf{G}_i = (G_{i1}, G_{i2}, \ldots, G_{iM})$ for $i = 1, \ldots, N$ representing the genotypes. To perform a GWIS, we assume M tests of $G \times E$ interaction with test statistics $\{T_j\}_{j=1}^M$ and corresponding p-values $\{p_j\}_{j=1}^M$ are computed. If for example the trait Y is quantitative, a standard one-step GWIS generates the $G \times E$ interaction one-at-a-time by using the following model:

$$E(Y|G_j, E) = \beta_{0j} + \beta_{G_j} G_j + \beta_E E + \beta_{G_j \times E}(G_j \times E) \tag{1}$$

for each of the M SNPs. For ease of exposition, we did not include subject-level covariates in the model, but in practice adjustment covariates like sex, age, and principal components capturing genetic ancestry are critical for preventing spurious associations . Each T_j corresponds to the test statistic for testing the null hypothesis $H_0 : \beta_{G_j \times E} = 0$. An adjustment for multiple comparisons is applied to preserve the family-wise Type I error rate (FWER) at a prespecified significance level α (e.g., $\alpha^* = 5 \times 10^{-8}$ or $\alpha^* = \alpha/M$). However, this correction will lead to low power in detecting a $G \times E$ interaction.

2.1 Two-step tests

Two-step methods have been developed to improve the power of GWIS(Kooperberg and LeBlanc, 2008; Murcray et al., 2009; Ionita-Laza et al., 2007). The test of interaction is also based on T_j (Step 2); however, the significance level at which each of the M SNPs is tested depends on an initial screening step (Step 1) based on a screening statistics S_j computed for each SNP. A common choice of screening statistic is based on the marginal effect of a SNP on the trait, as the presence of an interaction will typically induce a marginal G effect on the outcome (Kooperberg and LeBlanc, 2008). For a quantitative trait, the marginal outcome-gene association can be modeled using

$$E(Y|G_j) = \mu_{0j} + \mu_{G_j} G_j, \tag{2}$$

and S_j is the test statistic corresponding to $H_0 : \mu_{G_j} = 0$. Model 2 has the form typically used to identify SNPs associated with the outcome in a standard GWAS.

Two approaches to two-step testing are being widely used: subset testing (Kooperberg and LeBlanc, 2008) and weighted hypothesis testing (Ionita-Laza et al., 2007). Subset testing: In subset testing, each of the M screening statistics is compared to a prespecified significance threshold α_0. Let $\mathcal{A}$ be the collection (subset) of indices of SNPs that pass the Step 1 screen (i.e. SNPs for which S_j is statistically significant at the α_0 level) with $|\mathcal{A}| = m$, where $|A|$ represents the cardinality of set A. Then, for each $k \in \mathcal{A}$, the test statistic p-value p_k is compared against the Bonferroni-corrected significance level $\alpha^* = \alpha/m$, which is a less stringent significance level than in the standard single-step GWIS. Note that tests where $k \notin \mathcal{A}$ will never be tested – or equivalently tested against a significance level of $\alpha^* = 0$. This ensures that the overall Type I error rate is retained at α. Weighted testing: Instead of only testing a subset of hypotheses based on some dichotomous screening procedure, one can test all M test statistics in Step 2 using a weighted hypothesis test (Ionita-Laza et al., 2007) where the significance level assigned to each Step 2 SNP is based on the ordered (largest to smallest) values of $|S_j|$ from Step 1. The rationale behind the weighted screening approach is that more promising SNPs - as measured by the screening statistic – are tested at a less stringent significance level. To control the type I error, the SNPs are assigned into B bins according to their M screening statistics such that all SNPs in a bin b are tested at the significance level α_b, where $\alpha_1 + \alpha_2 + \ldots + \alpha_B \leq \alpha$. A common choice is $\alpha_b = \alpha/2^{b+1}$; then, for sufficiently large B,

$$\alpha_1 + \alpha_2 + \ldots + \alpha_B = \sum_{b=0}^{B} \frac{\alpha}{2^{b+1}} \approx \alpha.$$

By partitioning α in such a way, the overall Type I error rate is controlled at α while allocating a greater fraction of α to the top bins (i.e. those with the most promising SNPs). A Bonferroni-like correction can be used to preserve the bin-wise error rate (BWER) by dividing α_b by $|\mathcal{A}_b|$, where $\mathcal{A}_b$ is the set of indices for tests in bin b (size of the bin). We are now left with the non-trivial task of deciding how many SNPs should be allocated to each bin. Using a predetermined initial bin size, B_0, Ionita-Laza et al. (2007) suggested binning tests such that $|\mathcal{A}_b| = 2^{b-1}B_0$. For example, the SNPs with the B_0 largest values of $|S_j|$ are placed in $\mathcal{A}_1$ and tested against a significance threshold of α_1/B_0, the next $2B_0$ SNPs are tested in bin 2 ($\mathcal{A}_2$) at level $\alpha_2/2B_0$ and so forth. SNPs in the earlier bins (e.g. $\mathcal{A}_1, \mathcal{A}_2$) are tested at a less stringent threshold than in the standard one-step GWIS, while SNPs in the latter bins are tested at a more stringent significance level than in a one-step GWIS. Since the binning of SNPs can be done in several ways, we refer to the latter as *rank-based (RB)* binning to contrast with significance-based binning that we propose. Two-step RB-weighted tests are in general, more powerful than subset testing (Gauderman et al., 2013). Although the particular choice of α_0 (subset testing) and B_0 (RB-weighted testing), does not affect the FWER they greatly affect power.(Lewinger et al., 2013) Rule-of-thumb values for α_0 and B_0 have been proposed (Kooperberg and LeBlanc, 2008; Murcray et al., 2009; Ionita-Laza et al., 2007).

2.2 Proposal # 1: Significance-based weighting in Step 1

Note that subset testing can be seen as a unique case of RB-weighted testing with two bins with $\alpha_1 = \alpha$, $\mathcal{A}_1 = \mathcal{A}$, $\alpha_2 = 0$, and $\mathcal{A}_2 = \mathcal{A}^c$ where $\mathcal{A}$ is the subset of indices of SNPs that have a Step 1 p-value $< \alpha_0$. Unlike RB-weighted testing, where the size of bin 1 is set *a priori*, here the size of the bin is determined by the number of p-values in Step 1 that fall within the interval $(0, \alpha_0)$. Rather than creating two bins based off of one p-value (e.g. α_0), one can allocate tests into B bins using a series of significance-level cutoffs.

Assuming that the null hypothesis of no marginal association holds for all M SNPs, each p_{S_j} is uniformly distributed in the interval $[0, 1]$. In expectation, B_0 of the marginal outcome-gene screening statistics should have a p-value less than B_0/M, $2B_0$ of them should be in the interval $(B_0/M, 3B_0/M)$, and so forth. However, this relationship only holds under the null hypothesis of no outcome-gene association. In practice, this assumption may not hold and will lead to a more-than-expected number of p-values to lie in the top bins. This "overcrowding" of the top bins (e.g., $\mathcal{A}_1, \mathcal{A}_2$) will reduce the power of the RB-weighted screening approach since bin sizes are capped and SNPs with true $G \times E$ effects may be pushed to a bin in which the interaction will be tested using a stricter threshold. As we will show in our numerical studies, the performance of RB-weighted testing is dependent on both the strength and number of non-null outcome-gene associations, which typically is unknown *a priori*.

Motivated by the idea of overcrowding, we propose an alternative binning approach, henceforth referred to as significance-based (SB) weighted testing, that prioritizes SNPs based on the significance levels of the p-values corresponding to each screening statistic (p_{S_j}). More specifically, defining $\tau = (\tau_1, \tau_2, \ldots, \tau_{B+1})$ as the set of significance level cutoffs, the collection of tests in bin b is $\mathcal{A}_b = \{j : p_{S_j} \in [\tau_b, \tau_{b+1})\}$. Thus, SB-weighted testing can be viewed as a hybrid between both subset and RB-weighted testing. While τ can be defined arbitrarily, we propose to set $\tau = (0, B_0/M, 3B_0/M, 7B_0/M, \ldots, 1)$ which, in expectation, corresponds to binnings that are identical to RB-weighted testing. However, bin sizes are not capped since p-values are used to determine bin allocation.

2.3 Proposal #2: Accounting for correlation among tests in Step 2

While two-step hypothesis testing generally provides increased power over the standard one-step GWIS, a Bonferroni-type correction is still used to control the FWER in Step 2. This adjustment tends to be overly conservative in situations where correlation is present as it is the case in association studies due to linkage disequilibrium (LD). Accounting for genetic correlation can provide additional power gains.

Permutation tests are seen as the gold standard by permuting the data in a way that simulates the null hypothesis while simultaneously maintaining the original correlation structure. However, in large association studies, permutation-based approaches are computationally prohibitive. Furthermore, it is not obvious how to extend permutation tests in the $G \times E$ setting due to the hierarchical structure of the model. Several, more efficient, methods to account for correlation between tests have been proposed. For example, Conneely and Boehnke (2007) developed a method (p_{ACT}) that attains the accuracy of permutation or simulation-based tests in much less computation time. Their approach, however, requires valid estimation of the covariance matrix, which is computationally prohibitive if

applied genome wide. Another alternative is to replace the denominator of the Bonferroni correction, the total number of tests within a bin, with an estimate of the effective number of independent tests (M^*). Cheverud et al. (1983), among others, have shown that the overall correlation among traits in a set can be captured by the variance of the eigenvalues derived from their correlation matrix (i.e. higher correlation among traits will lead to higher eigenvalue variance).

Let $\{\lambda_k\}_{k=1}^K$ be the set of eigenvalues obtained by a principal component analysis (PCA) from the pair-wise SNP correlation matrix (i.e. LD matrix), arranged in decreasing order. Gao et al. (2008) proposed a simple and fast correction (simple$M$) that well approximates permutation-based corrections in both simulated and real data based on these eigenvalues

$$M^* = \min_x \left\{ \frac{\sum_{i=1}^x \lambda_i}{\sum_{k=1}^K \lambda_j} > C \right\}, \tag{3}$$

where C determines the percentage cutoff of explained variation. Gao et al. (2008) suggest setting $C = 0.995$ so that M^* corresponds to the minimum number of eigenvalues needed to explain at least 99.5% of the variation in the data.

Applying this correction genome wide is infeasible as it requires, first, the calculation of an $M \times M$ correlation matrix and, second, an eigendecomposition (e.g. PCA) to derive the eigenvalues. We instead apply the correction to the set of SNPs within each bin. For example, if we let $\mathbf{G}$ be the $N \times M$ matrix of SNP dosages for N subjects, then we calculate M_b^* for bin b using the correlation matrix of $\mathbf{G}_{\mathcal{A}_b}$, the sub-matrix of $\mathbf{G}$ that corresponds to the indices in $\mathcal{A}_b$. Note that, by design, $\mathcal{A}_b$ grows larger, in expectation, as b increases. Thus the calculation of M^* for the larger bins may still be computationally demanding. In practice, as shown in Section 3, we restrict the computation of M_b^* for the first seven bins and set the significance threshold for the later bins to be 0 since the threshold to be declared statistically significant is very stringent (e.g. BWER < 0.0004) for these bins.

3 Simulation studies

We compare the performance of our two-step significance-based (SB) weighted testing procedure to 1) the standard one-step GWIS, and 2) the two-step rank-based (RB) weighted testing procedure. For both two-step methods (rank and significance based), we compare the standard Bonferroni correction to the LD-adjusted Bonferroni correction using the simple$M$ approach (Gao et al., 2008) to account for LD within bins. As recommended by Gao et al. (2008), we set the tuning parameter $C = 0.995$.

Quantitative traits are simulated according to the following linear model:

$$\mathbf{y} = \mathbf{G}\boldsymbol{\beta} + E\gamma + (\mathbf{G} \times E)\boldsymbol{\omega} + \boldsymbol{\epsilon},$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}_M(\mathbf{0}, \sigma_\epsilon^2 I)$ for some $\sigma_\epsilon^2 > 0$, $\mathbf{G} = [\mathbf{G}_1, \mathbf{G}_2, \ldots]$ is an $N \times M$ genotype matrix with $\mathbf{G}_j$ being the jth block of $N \times 50$ simulated allele counts, E is the exposure variable (assumed to be binary) with $\Pr(E = 1) = 0.3$, $(\mathbf{G} \times E)$ corresponds to the $N \times M$ matrix of gene-environment interaction terms and $(\boldsymbol{\beta}, \gamma, \boldsymbol{\omega})$ are the respective parameters of interest. Each $\mathbf{G}_j$ is simulated based on sampled minor allele frequencies (MAFs) and LD-matrices from the 1000 Genomes Project. We refer readers to the Appendix for more information on the simulation setup.

We designate the 25th SNP within block 1 ($\mathbf{G}_1$) as having a true $G \times E$ effect on the outcome (i.e. the $G \times E$ locus). The value of the parameters β_{25}, γ, and ω_{25} were set to achieve a predetermined R^2 for each term: $R^2_{G_{25}} = 0.01$, $R^2_E = 0.005$ and $R^2_{G_{25} \times E} = 0.01$. Given a minor allele frequency of 0.231 for the $G \times E$ locus and setting $Var(\mathbf{y}) = 1$, yields $\beta_{25} \approx 0.06$, $\gamma \approx -0.01$ and $\omega_{25} \approx 0.37$. We also simulate additional "G-only" effects (G-only loci), SNPs that are only marginally associated with the outcome but have no $G \times E$ effect. Each G-only loci was placed in a different block so that these SNPs, and SNPs in LD, are expected to be prioritized in Step 1 and thus potentially affect power in identifying the true $G \times E$ effect due to bin overcrowding. For simplicity, we set the expected R^2 ($R^2_{G_0}$) for each of the G-only effects to be the same and vary the number of G-only effects (n_{G_0}).

The testing statistics, T_j, are based on the following hypotheses $H_0 : \omega_j = 0$ for $j = 1, \ldots, M$, where ω_j is estimated based on the following one-SNP-at-a-time model

$$\mathbf{y} = \mu_0 + g_j\beta_j + E\gamma + (g_j \times E)\omega_j + \boldsymbol{\epsilon},$$

where g_j denotes the jth SNP ($j = 1, \ldots, M$). For both two-step procedures (RB and SB-weighted), the screening statistics used in Step 1 are based on the test statistics corresponding to the marginal outcome-gene association (Kooperberg and LeBlanc, 2008) based on

$$\mathbf{y} = \tilde{\beta}_0 + G_j\tilde{\beta}_j + \tilde{\boldsymbol{\epsilon}} \quad j = 1, \ldots M.$$

The initial bin size B_0 was set to 5 SNPs for the RB-weighted approach and to an expectation of 5 SNPs assuming no genetic effects for the SB-weighted approach. For the latter, this corresponds to $\tau = (0, 5/25000, 15/25000, 35/25000, \ldots, 1)$. Under this scheme, we expect both RB and SB-weighted hypothesis testing to have comparable performance for detecting $G \times E$ loci when G-only effects are weak or absent. The FWER was set to $\alpha = 0.05$. As aforementioned, calculation of the effective number of independent tests for both two-step approaches will be computationally demanding for bins with larger number of SNPs. To avoid this, we restrict computation of M^* to the first seven bins and set the significance threshold in the later bins to 0, since the threshold to be labeled statistically significant is very stringent for these bins.

Since SNPs in $\mathbf{G}_1$, the block with the causal $G \times E$ locus, are correlated, power is calculated as the number of times we reject the null hypothesis for any of the SNPs in $\mathbf{G}_1$ at the corresponding significance level. FWER is defined as rejecting the null hypothesis $H_0 : \omega_j = 0$ for $j = 51, \ldots, M$ at the corresponding significance level. Furthermore, we recorded the ranking of the $G \times E$ locus in terms of its Step 1 statistic as well as the smallest ranking of any loci in $\mathbf{G}_1$ in terms of its Step 1 statistic. Results are averaged over 5,000 Monte Carlo replications.

In our first set of simulations, we set $N = 2000$, $M = 25000$, and $R^2_{G_0} = R^2_{G_{25}} = 0.01$ so that the marginal effects of the $G \times E$ locus and G-only loci are comparable (Figure 1 Panel A). Power to detect the true $G \times E$ effect ω_{25} with the standard one-step GWIS is approximately 45%. When no G-only effects are present ($n_{G_0} = 0$), both approaches (RB and SB-weighted) have comparable power substantially higher than the standard one-step GWIS. The comparability of RB and SB-weighted testing in this scenario is expected since the binning of tests for both approaches should be nearly identical with no additional G-only effects. The median ranking of the $G \times E$ loci in terms of its Step 1 statistic was 3; whereas the median for the smallest ranking among loci in $\mathbf{G}_1$ was 1. However, power

for RB-weighted testing is sensitive to the number of G-only effects, dropping from $\approx 85\%$ when only 10 G-only effects are present to $\approx 64\%$ when 80 are present. This decrease in power is expected since the G-only loci, and their LD-regions, are being ranked higher than the $G \times E$ loci, and hence pushing the $G \times E$ loci into bins with more stringent significance thresholds in the testing step. In this scenario the median ranking for the $G \times E$ locus and median ranking for the smallest ranked locus among loci in $\mathbf{G}_1$ were 89 and 106, respectively. The loss in power for SB-weighted testing is less dramatic, which suggests that it is more robust to bin flooding than RB-weighted testing. Furthermore, adjusting for LD-based correlations in Step 2 provides a modest increase in power across all scenarios.

Unlike our first set of simulations, it is highly unlikely that over $> 80\%$ of the variation can be explained by the data in typical GWAS/GWIS studies. We include another set of simulations (Figure 1 Panel B) where we double the sample size N to 4000 and halve the expected R^2 for each factor (i.e. $R^2_{G_{25}} = 0.005$, $R^2_E = 0.0025$ and $R^2_{G_{25}\times E} = 0.005$, and $R^2_{G_0} = 0.005$). Our results are consistent to what we have seen above which suggests that the power benefit of SB-weighted testing, coupled with the simple$M$ procedure, can be achieved under more realistic effect sizes to what is commonly detected in GWAS and GWIS studies. Estimated FWER is presented in Figure S1 for $N = 2,000$ and $N = 4,000$. Compared to the standard Bonferroni correction (solid line), the simple$M$ procedure provides FWER that are closer to nominal although still slightly convservative.

In addition to increasing the number of marginal effects, we also evaluate both weighted testing methods when we vary both the size and magnitude of the G-only SNPs (Table 1). In this set of experiments, we fixed the % of explained variation explained by the G-only SNPs to 40% and vary n_{G_0} from $10, 20, 40$ to 80 such that the expected R^2 for each G-only SNP is $R^2_{G_0} = 0.04, 0.02, 0.01$ and 0.005, respectively. We keep $R^2_{G_{25}} = 0.01$ for the causal $G \times E$ locus so that we capture scenarios where the 'competition' in prioritizing the causal locus ranges from weak to strong. When $R^2_{G_0}$ is large relative to $R^2_{G_{25}}$, the test for the causal locus has little-to-no chance of being in the top bins and will thus always be tested at a more stringent significance threshold. Power for the standard one-step GWIS is unaffected by the magnitude and number of the G-only SNPs. The difference in power is largest between the SB and RB-weighted testing methods when a few number of 'strong' G-only SNPs ($n_{G_0} = 10$, $R^2_{G_0} = 0.04$) are present (median rank of the Step 1 statistic for the $G \times E$ locus = 61). In this scenario, RB-weighted testing has around 70% power of detecting the $G \times E$ locus whereas SB-weighted testing has power closer to 80%. On the other hand, power between both RB and SB-weighted testing are comparable when more but weaker G-only SNPs are present (around 80% power when $n_{G_0} = 80$ and $R^2_{G_0} = 0.005$; median rank of the Step 1 statistic for the $G \times E$ locus = 21). Thus, both magnitude of effect and the number of non-null screening statistics can greatly affect power for RB-weighted hypothesis testing. Conversely, SB-weighted screening is more robust to the size and magnitude of the G-only SNPs, with power hovering around $80 - 81\%$ in all four scenarios. We see similar trends in terms of FWER (Table S1).

It has been shown that the choice of B_0 affects power for weighted testing (Ionita-Laza et al., 2007; Lewinger et al., 2013). We investigate the performance of both RB and SB-weighted hypothesis testing under various initial bin sizes (Figure 2). For this simulation, we assume 80 G-only SNPs with $R^2_{G_0} = 0.01$ are present ($R^2_{G_{25}} = 0.01$, $R^2_E = 0.005$ and $R^2_{G_{25}\times E} = 0.01$) with $N = 2,000$ and $M = 25,000$. We see that RB-weighted screening, is sensitive to B_0, with an increase in power as B_0 increases. The increase in power can

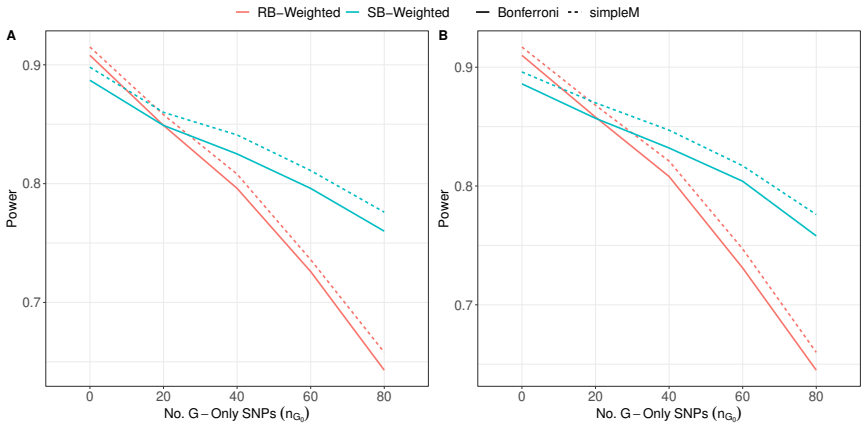

Figure 1: Estimated power when n_{G_0} G-only effects are present ($n_{G_0} \in \{10, 20, 40, 80\}$) that each explain $R^2_{G_0} \times 100\%$ of the variation. RB-Weighted: Rank-based weighted hypothesis testing using with initial bin size B_0 in Step 1; SB-Weighted: Significance-based weighted hypothesis testing using $\tau = (0, B_0/25000, 3B_0/25000, \ldots, 1)$ as the p-value cutoffs in Step 1. Bonferroni: Standard Bonferroni correction within bin; simpleM: The simpleM procedure proposed by Gao et al. (2008) with $C = 0.995$. Results are averaged over 5000 simulations. Panel A: $R^2_{G_{25}} = R^2_{G_{25}\times E} = R^2_{G_0} = 0.01, R^2_E = 0.005, N = 2,000$; Panel B: $R^2_{G_{25}} = R^2_{G_{25}\times E} = R^2_{G_0} = 0.005, R^2_E = 0.0025, N = 4,000$.

be explained by the fact that we allow more tests to be included in each bin. Thus, even when comparably-sized G-only effects are competing with the causal $G \times E$ locus for bin placement, the probability of being in a smaller bin, and thus tested against a less stringent BWER, is higher. However, as others have explored, power can be sensitive to the choice of B_0, and larger values of B_0 can correspond to lower power when no G-only effects are present(Lewinger et al., 2013; Gauderman et al., 2013). In contrast, SB-weighted screening is relatively robust to the choice of B_0 using in creating the significance thresholds.

4 Application to colorectal cancer

We applied our proposed two-step approach to a genome-wide scan of biological sex by gene interactions on the risk of colorectal cancer (CRC). The FIGI (Functionally Informed Gene-Environment Interaction) study is a multi-institutional collaborative effort to identify novel $G \times E$ interactions for colorectal cancer (CRC). Details of the FIGI study have been previously published (Schumacher et al., 2015; Schmit et al., 2019; Huyghe et al., 2019). In brief, epidemiological and genotype data were pooled from 61 cohort and case-control CRC studies from 3 large consortia - the Genetics and Epidemiology of Colorectal Cancer Consortium (GECCO), the Colorectal Transdisciplinary Study (CORECT), and Colon Cancer Family Registry (CCFR). Analyses include only individuals with complete exposure and covariate information, and were limited to individuals of European ancestry as determined by self-reported race and clustering of principal components with 1000 Genomes EUR sample. Details on genotyping, quality control, data collection and harmonization have been described previously (Hutter et al., 2012; Huyghe et al., 2019). Our analysis included $N = 89,304$ individuals (40,647 cases and 48,657 controls) and $M = 7,809,725$

Table 1: Estimated power when n_{G_0} G-only effects are present ($n_{G_0} \in \{10, 20, 40, 80\}$) and the total amount of variation explained is fixed at 40% ($R^2_{G_0} = 0.4/n_{G_0}$). One-step GWIS: The standard one-step GWIS. RB-Weighted: Rank-based weighted hypothesis testing proposed by Ionita-Laza et al. (2007) with $B_0 = 5$; SB-Weighted: Our proposed significance-based weighted hypothesis testing with $\tau = (0, 5/25000, 15/25000, \ldots, 1)$ as the p-value cutoffs. Bonferroni: Standard Bonferroni correction within bin; simple$M$: The simpleM procedure proposed by Gao et al. (2008) with $C = 0.995$. Results are averaged over 5000 simulations.

$n_{G_0} =$	10	20	40	80
$R^2_{G_0} =$	0.04	0.02	0.01	0.005
One-step GWIS	0.439	0.437	0.443	0.442
RB-Weighted				
Bonferroni	0.694	0.670	0.726	0.801
simple$M$	0.713	0.688	0.736	0.812
SB-Weighted				
Bonferroni	0.814	0.797	0.796	0.812
simple$M$	0.829	0.811	0.812	0.826

SNPs.

Logistic regression models were used to model SNP allelic dosages (G) and sex (E) on case/control status (y) genome-wide. The marginal outcome-gene association, regressing y on G genome-wide, was used as the screening statistic in Step 1. All models included the following set of adjustment covariates - age, study/genotyping batch, and ancestry as defined by the first three principal components. Analyses using similar adjustment covariates were used in Jordahl et al. (2022).

Similar to the simulation study, bin significance threshold cutoffs for SB-weighted testing were based on the following cutoffs - $(0, 5/M, 15/M, 35/M, \ldots, 1)$. Bin flooding is apparent as bin 1, contains 3,795 SNPs (Figure 3) wth a direct G-CRC association p-value $< 5/M$. Under the RB-weighted testing approach with $B_0 = 5$, these SNPs would fill the first 8 bins. For testing in the second step, we use the simple$M$ correction to estimate the effective number of independent tests. For bin 1, the estimated effective number of independent tests was 423. Thus, each of the 3,795 $G \times E$ tests are tested against a significance threshold of 0.025/423. One SNP (18:46458950:G:A) passed the threshold in bin 1. This SNP is located near the SMAD7 gene, which has been previously shown as a potential marker of colorectal cancer (Thompson et al., 2009; Jiang et al., 2013; Stolfi et al., 2014; Alidoust et al., 2022). No $G \times E$ effects were discovered using either the standard GWIS or RB-weighted two-step approaches. Under an initial bin size of $B_0 = 5$, SNP 18:46458950:G:A landed in the fifth bin, which corresponds to a significance threshold that is slightly more stringent than 0.025 / 423.

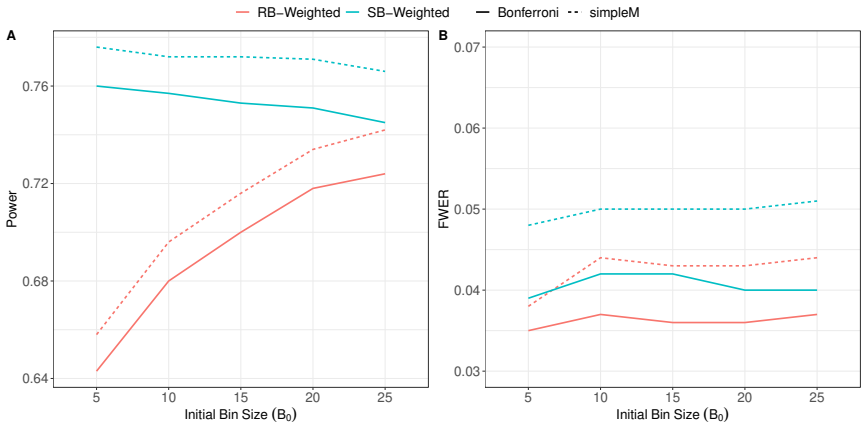

Figure 2: Estimated power and FWER as a function of the initial bin size B_0. 80 G-only effects are present with $R^2_G = 0.01$. RB-Weighted: Rank-based weighted hypothesis testing using with initial bin size B_0 in Step 1; SB-Weighted: Significance-based weighted hypothesis testing using $\tau = (0, B_0/25000, 3B_0/25000, \ldots, 1)$ as the p-value cutoffs in Step 1. Bonferroni: Standard Bonferroni correction within bin or subset; simpleM: The simpleM procedure proposed by Gao et al. (2008) with $C = 0.995$. Results are averaged over 5000 simulations.

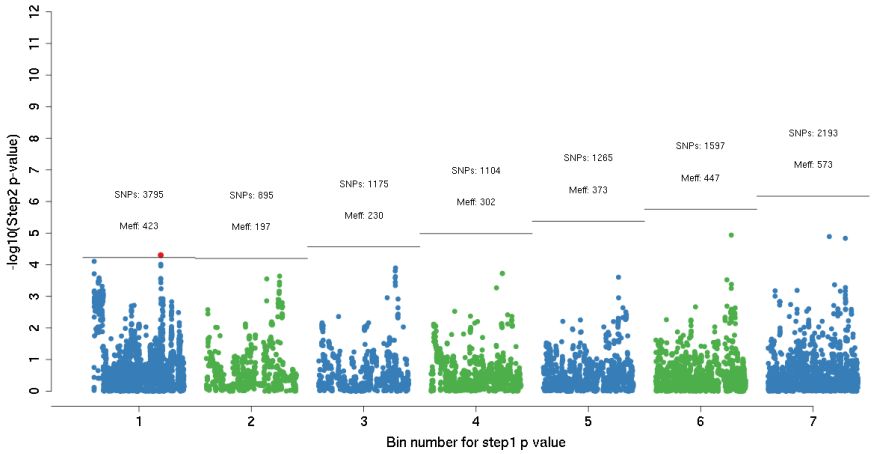

Figure 3: Results from the G-by-sex interaction scan using data provided by the FIGI consortium ($N = 89,304$, $M = 7,809,725$). x-axis: Bins are based on the marginal outcome-gene association statistic (e.g. SNPs that have a Step 1 statistic $< 5/M$ are included in bin 1). y-axis: p-value of the $G \times E$ association provided by the GWIS (on the $-\log_{10}$ scale). Number of SNPs in each bin as well as the effective number of independent SNPs using the simple$M$ approach are included. Horizontal line indicates the threshold the Step 2 p-value must cross to be statistically significant. Only SNPs in the first 7 bins are shown in this figure.

5 Discussion

Two-step tests generally provide greater power than standard one-step approaches for genome-wide $G \times E$ discovery (Gauderman et al., 2017). We propose a novel approach, significance-based (SB) weighted hypothesis testing, that aims to address the shortcomings of its predecessors. In Step 1, we allocate SNPs into bins based on the significance of the screening statistic rather than on their ranking. Then, in step 2, to account for SNP-SNP correlations due to LD, we control the bin-wise error rate based on an estimate of the effective number of independent tests. We show that SB-weighted testing is comparable to RB-weighted hypothesis testing when 1) no marginal G-only effects are present or 2) weak marginal G-only effect are present (i.e. little-to-no skew in the p-value distribution) and outperforms RB-weighted hypothesis testing when G-only loci flooding of the top bins exists. In addition, using an estimate of the effective number tests (simple$M$) to account for LD-based correlation among SNPs provides additional power for either of the two-step methods. We demonstrated our SB-weighted approach to identify novel gene-sex interactions for colorectal cancer using data provided by the FIGI consortium.

In our simulation study, we also show that power for the RB-weighted hypothesis test is sensitive to both the number and magnitude of the step 1 screening statistics as well as the initial bin size. This corresponds to the number of marginal genetic effects that are associated with the outcome if the outcome-gene association statistic proposed by Kooperberg and LeBlanc (2008) is used in Step 1. These factors, among others, will negatively affect the prioritization of the causal ($G \times E$) loci in the testing step since the number of tests per bin are capped *a priori*. Alternatively, the use of significance cutoffs, as opposed to fixed bin sizes, in SB weighted hypothesis testing procedure is robust to both the number and magnitude of the step 1 statistics and has greater power over weighted testing when there is competition in bin prioritization. Adjusting for the effective number of tests using the simple$M$ method of Gao et al. (2008) yields additional modest, improvements in power while providing close-to-nominal coverage.

Our two-step hypothesis testing procedure identified a new locus on SMAD7 that interacts with biological sex to modulate CRC risk. Further studies should examine the potential mechanisms through which this newly discovered locus impacts CRC risk. We note that $G\times$ Sex interaction for this locus would not have been found using either the standard one-step GWIS or the RB-weighted testing approach.

We envision several directions to further explore SB-weighted testing. First, the marginal outcome-gene association was used as the screening statistic in Step 1. However, SB-weighted testing should be valid using any screening statistic in Step 1 as long as its independent of the testing statistic in Step 2 (Dai et al., 2012). Further investigation into the performance of SB-weighted testing is warranted under different phenotypes, sampling designs, and screening statistics. The binning of tests was performed by developing significance-level based cutoffs of the p-values in Step 1. Our rationale for this approach was to, in expectation, develop bins that were similar to the rank-based binnings proposed by Ionita-Laza et al. (2007). However, one can bin tests arbitrarily based on the p-value distribution of the Step 1 screening statistics or on the distribution of the screening statistics themselves. To this end, we have not yet considered alternative cutoffs for binning tests.

LD was taken into account through estimating the effective number of independent tests. While we explored the performance of the simple$M$ procedure proposed by Gao

et al. (2008), several other estimators have been suggested (Cheverud, 2001; Li and Ji, 2005; Galwey, 2009; Nyholt, 2004; Li et al., 2012). In addition, these estimates do not take into account the distribution of the test statistics. Alternative approaches to adjust p-values within bins (or the testing subset) can be adopted. Conneely and Boehnke (2007) propose p_{ACT}, p-value adjusted for correlated tests, which compares the observed test statistics directly with their asymptotic distribution through numerical integration. While genetic correlation within bins have been considered, between bin LD can be present as well. Methods to account for both within and between bin genetic correlation may further improve power (at the expense of computational efficiency) and should be investigated further. Alternatively, one may instead control the false discovery rate (Benjamini and Hochberg, 1995), the expected proportion of discoveries that are false, rather than the FWER. Future research is needed into developing FDR-based two-step hypothesis tests.

The current two-step hypothesis testing paradigm can be greatly improved by 1) incorporating a more robust binning procedure for the screening step (Step 1) and 2) taking into account LD when adjusting the BWER. Our proposal suggests to bin test using a significance-based weighting scheme and to adjust the BWER by using the simple$M$ procedure. While this pair improved power without inflating the FWER across several simulation scenarios and identified a novel loci previously not found using conventional methods, other combinations should be explored.

Acknowledgements

The work was partially supported by NIH grants T32ES013678, P01CA196569, R01CA201407, and P30ES007048. These awards had no influence over the experimental design, data analysis or interpretation, or writing the manuscript.

References

Alidoust, M., L. Hamzehzadeh, A. Khorshid Shamshiri, F. Afzaljavan, M. A. Kerachian, A. Fanipakdel, S. A. Aledavood, A. Allahyari, A. Bari, H. Moosanen Mozaffari, et al. (2022). Association of smad7 genetic markers and haplotypes with colorectal cancer risk. *BMC medical genomics 15*(1), 1–9.

Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological) 57*(1), 289–300.

Cheverud, J. M. (2001). A simple correction for multiple comparisons in interval mapping genome scans. *Heredity 87*(1), 52–58.

Cheverud, J. M., J. Rutledge, and W. R. Atchley (1983). Quantitative genetics of development: genetic correlations among age-specific trait values and the evolution of ontogeny. *Evolution*, 895–905.

Conneely, K. N. and M. Boehnke (2007). So many correlated tests, so little time! rapid adjustment of p values for multiple correlated tests. *The American Journal of Human Genetics 81*(6), 1158–1168.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological) 34*(2), 187–202.

Dai, J. Y., C. Kooperberg, M. Leblanc, and R. L. Prentice (2012). Two-stage testing procedures with independent filtering for genome-wide gene-environment interaction. *Biometrika 99*(4), 929–944.

Dudbridge, F. and A. Gusnanto (2008). Estimation of significance thresholds for genomewide association scans. *Genetic Epidemiology: The Official Publication of the International Genetic Epidemiology Society 32*(3), 227–234.

Galwey, N. W. (2009). A new measure of the effective number of tests, a practical tool for comparing families of non-independent significance tests. *Genetic Epidemiology: The Official Publication of the International Genetic Epidemiology Society 33*(7), 559–568.

Gao, X., J. Starmer, and E. R. Martin (2008). A multiple testing correction method for genetic association studies using correlated single nucleotide polymorphisms. *Genetic Epidemiology: The Official Publication of the International Genetic Epidemiology Society 32*(4), 361–369.

Gauderman, W. J., B. Mukherjee, H. Aschard, L. Hsu, J. P. Lewinger, C. J. Patel, J. S. Witte, C. Amos, C. G. Tai, D. Conti, et al. (2017). Update on the state of the science for analytical methods for gene-environment interactions. *American journal of epidemiology 186*(7), 762–770.

Gauderman, W. J., P. Zhang, J. L. Morrison, and J. P. Lewinger (2013). Finding novel genes by testing g× e interactions in a genome-wide association study. *Genetic epidemiology 37*(6), 603–613.

Hsu, L., S. Jiao, J. Y. Dai, C. Hutter, U. Peters, and C. Kooperberg (2012). Powerful cocktail methods for detecting genome-wide gene-environment interaction. *Genetic Epidemiology 36*(3), 183–194.

Hutter, C. M., J. Chang-Claude, M. L. Slattery, B. M. Pflugeisen, Y. Lin, D. Duggan, H. Nan, M. Lemire, J. Rangrej, J. C. Figueiredo, et al. (2012). Characterization of gene–environment interactions for colorectal cancer susceptibility loci. *Cancer research 72*(8), 2036–2044.

Huyghe, J. R., S. A. Bien, T. A. Harrison, H. M. Kang, S. Chen, S. L. Schmit, D. V. Conti, C. Qu, J. Jeon, C. K. Edlund, et al. (2019). Discovery of common and rare genetic risk variants for colorectal cancer. *Nature genetics 51*(1), 76–87.

Ionita-Laza, I., M. B. McQueen, N. M. Laird, and C. Lange (2007). Genomewide weighted hypothesis testing in family-based association studies, with an application to a 100k scan. *The American Journal of Human Genetics 81*(3), 607–614.

Jiang, X., J. E. Castelao, D. Vandenberg, A. Carracedo, C. M. Redondo, D. V. Conti, J. P. Paredes Cotoré, J. D. Potter, P. A. Newcomb, M. N. Passarelli, et al. (2013). Genetic variations in smad7 are associated with colorectal cancer risk in the colon cancer family registry. *PloS one 8*(4), e60464.

Jordahl, K. M., A. Shcherbina, A. E. Kim, Y.-R. Su, Y. Lin, J. Wang, C. Qu, D. Albanes, V. Arndt, J. W. Baurley, et al. (2022). Beyond gwas of colorectal cancer: Evidence of interaction with alcohol consumption and putative causal variant for the 10q24. 2 region. *Cancer Epidemiology, Biomarkers & Prevention 31*(5), 1077–1089.

Kawaguchi, E. S., G. Li, J. P. Lewinger, and W. J. Gauderman (2022). Two-step hypothesis testing to detect gene-environment interactions in a genome-wide scan with a survival endpoint. *Statistics in Medicine*.

Kooperberg, C. and M. LeBlanc (2008). Increasing the power of identifying gene× gene interactions in genome-wide association studies. *Genetic Epidemiology: The Official Publication of the International Genetic Epidemiology Society 32*(3), 255–263.

Lewinger, J. P., J. L. Morrison, D. C. Thomas, C. E. Murcray, D. V. Conti, D. Li, and W. J. Gauderman (2013). Efficient two-step testing of gene-gene interactions in genome-wide association studies. *Genetic epidemiology 37*(5), 440–451.

Li, J. and L. Ji (2005). Adjusting multiple testing in multilocus analyses using the eigenvalues of a correlation matrix. *Heredity 95*(3), 221–227.

Li, M.-X., J. M. Yeung, S. S. Cherny, and P. C. Sham (2012). Evaluating the effective numbers of independent tests and significant p-value thresholds in commercial genotyping arrays and public imputation reference datasets. *Human genetics 131*(5), 747–756.

Murcray, C. E., J. P. Lewinger, D. V. Conti, D. C. Thomas, and W. J. Gauderman (2011). Sample size requirements to detect gene-environment interactions in genome-wide association studies. *Genetic epidemiology 35*(3), 201–210.

Murcray, C. E., J. P. Lewinger, and W. J. Gauderman (2009). Gene-environment interaction in genome-wide association studies. *American journal of epidemiology 169*(2), 219–226.

Nyholt, D. R. (2004). A simple correction for multiple testing for single-nucleotide polymorphisms in linkage disequilibrium with each other. *The American Journal of Human Genetics 74*(4), 765–769.

Paré, G., N. R. Cook, P. M. Ridker, and D. I. Chasman (2010). On the use of variance per genotype as a tool to identify quantitative trait interaction effects: a report from the women's genome health study. *PLoS Genet 6*(6), e1000981.

Schmit, S. L., C. K. Edlund, F. R. Schumacher, J. Gong, T. A. Harrison, J. R. Huyghe, C. Qu, M. Melas, D. J. Van Den Berg, H. Wang, et al. (2019). Novel common genetic susceptibility loci for colorectal cancer. *JNCI: Journal of the National Cancer Institute 111*(2), 146–157.

Schumacher, F. R., S. L. Schmit, S. Jiao, C. K. Edlund, H. Wang, B. Zhang, L. Hsu, S.-C. Huang, C. P. Fischer, J. F. Harju, et al. (2015). Genome-wide association study of colorectal cancer identifies six new susceptibility loci. *Nature communications 6*(1), 1–7.

Stolfi, C., V. De Simone, A. Colantoni, E. Franze, E. Ribichini, M. Fantini, R. Caruso, I. Monteleone, G. Sica, P. Sileri, et al. (2014). A functional role for smad7 in sustaining colon cancer cell growth and survival. *Cell death & disease 5*(2), e1073–e1073.

Thompson, C. L., S. J. Plummer, L. S. Acheson, T. C. Tucker, G. Casey, and L. Li (2009). Association of common genetic variants in smad7 and risk of colon cancer. *Carcinogenesis 30*(6), 982–986.

Wang, J., A. Patel, J. M. Wason, and P. J. Newcombe (2021). Two-stage penalized regression screening to detect biomarker-treatment interactions in randomized clinical trials. *Biometrics*.

Zhang, P., J. P. Lewinger, D. Conti, J. L. Morrison, and W. J. Gauderman (2016). Detecting gene–environment interactions for a quantitative trait in a genome-wide association study. *Genetic epidemiology 40*(5), 394–403.

S1 Appendix: Additional figures and tables

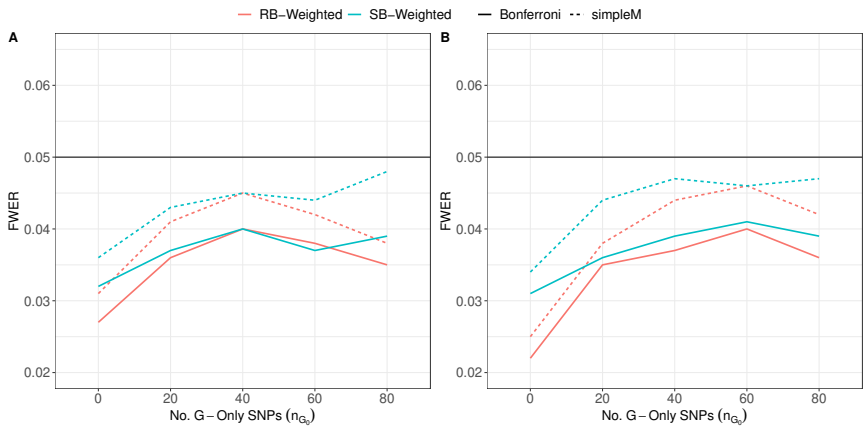

Figure S1: Estimated FWER when n_{G_0} G-only effects are present ($n_{G_0} \in \{10, 20, 40, 80\}$) each explaining $R^2_{G_0} \times 100\%$ of the variation in the quantitative trait. RB-Weighted: Rank-based two-step weighted testing with initial bin size 5 in Step 1; SB-Weighted: Significance-based weighted hypothesis testing with $\tau = (0, 5/25000, 15/25000, \ldots, 1)$ as the p-value cutoffs in Step 1. Bonferroni: Standard Bonferroni correction within bin; simpleM: The simpleM procedure proposed by Gao et al. (2008) with $C = 0.995$. Results are averaged over 5000 simulations. Panel A: $R^2_{G_{25}} = R^2_{G_{25} \times E} = R^2_{G_0} = 0.01, R^2_E = 0.005, N = 2,000$; Panel B: $R^2_{G_{25}} = R^2_{G_{25} \times E} = R^2_{G_0} = 0.005, R^2_E = 0.0025, N = 4,000$.

Table S1: Estimated FWER when n_{G_0} G-only effects are present ($n_{G_0} \in \{10, 20, 40, 80\}$) and the total amount of variation explained is fixed at 40% ($R^2_{G_0} = 0.4/n_{G_0}$). Standard: Standard one-step GWIS; Weighted: Weighted hypothesis testing using $B_0 = 5$ as the initial bin size in Step 1; EBR: Expectation-based ranked testing using $\tau = (0, 5/25000, 15/25000, \ldots, 1)$ as the p-value cutoffs in Step 1. Bonf: Standard Bonferroni correction within bin; Gao: Using the effective number of tests defined by Gao et al. with $C = 0.995$; Li & Ji: Using the effective number of tests defined by Li & Ji. Results are averaged over 5000 simulations.

$n_{G_0} =$	10	20	40	80
$R^2_{G_0} =$	0.04	0.02	0.01	0.005
One Step GWIS	0.055	0.056	0.054	0.057
RB-Weighted				
Bonferroni	0.044	0.040	0.035	0.036
simple$M$	0.045	0.045	0.039	0.039
SB-Weighted				
Bonferroni	0.037	0.037	0.036	0.036
simple$M$	0.045	0.045	0.043	0.044

S2 Appendix: Simulation setup

We simulate quantitative traits according to the following linear model:

$$\mathbf{y} = \mathbf{G}\boldsymbol{\beta} + E\gamma + (\mathbf{G} \times E)\boldsymbol{\omega} + \boldsymbol{\epsilon},$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}_M(\mathbf{0}, \sigma^2_\epsilon I)$ with $\sigma^2_\epsilon > 0$, $\mathbf{G}$ is an $N \times M$ genotype matrix, E is a binary exposure with $\Pr(E = 1) = 0.3$, $(\mathbf{G} \times E)$ corresponds to the $N \times M$ matrix gene-environment interaction terms and $(\boldsymbol{\beta}, \gamma, \boldsymbol{\omega})$ are the respective parameters of interest. σ^2_ϵ was chosen to explain the remaining variance in the outcome after accounting for the variance explained by the causal locus, the exposure, and the G-only loci across the different scenarios outlined in the main text.

The $M = 25,0000$ genotypes are simulated in blocks ($\mathbf{G} = [\mathbf{G}_1, \mathbf{G}_2, \ldots, \mathbf{G}_{500}]$) such that each block consists of 50 SNPs drawn from a mean zero multivariate normal distribution with variance-covariance matrix based on the LD pattern derived from a sampled region of the 1000 Genomes Project. The normal variates are then trichotomized into genotypes based on the 1000 Genomes Project derived MAFs assuming Hardy-Weinberg equilibrium. Thus, genotypes are correlated within a block but independent across blocks.

In other words, let $\mathbf{V}_j$ be the 50×50 LD matrix and $\mathbf{f}_j = (f_{1,1}, \ldots f_{1,50})$ be the corresponding vector of minor allele frequencies (MAF) of the 50 SNPs in region $j = 1, \ldots, 50$. To simulate the first block of 50 SNPs, let $\mathbf{X}_1 \sim \mathcal{N}_{50}(\mathbf{0}, \mathbf{V}_1)$ where $\mathbf{X}_1 = (x_{1,1}, \ldots, x_{1,50})^T$. Let $\mathbf{G}_1 = (g_{1,1}, \ldots, g_{1,50})^T$ denote the genotype matrix for the first block , and $a^i_{1,1}$ as the ith element of a generic vector $a_{1,j}$, for $i = 1, \ldots N$ and $j = 1, \ldots 50$. We can now define the i-th element of $g_{1,j}$ as

$$g_{1,j}^i = \begin{cases} 0 & \text{if } x_{1,j}^i < \Phi(f_{1,j}^2) \\ 1 & \text{if } \Phi(f_{1,j}^2) \leq x_{1,j}^i < \Phi(f_{1,j}^2 + 2f_{1,j}(1 - f_{1,j})) \\ 2 & \text{if } \Phi(f_{1,j}^2 + 2f_{1,j}(1 - f_{1,j})) \leq x_{1,j}^i \end{cases}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. We repeat this process to obtain $\mathbf{G}_j$ ($j = 1, \ldots, 500$). To avoid storing 500 unique values of $\mathbf{V}_j$ and $\mathbf{f}_j$, we only store 50 unique values (randomly sampled regions from the 1000 Genomes Project), and recycled them such that the LD and MAF pair ($\{\mathbf{V}_j, \mathbf{f}_j\}_{j=1}^{50}$) used for $\mathbf{G}_1, \ldots, \mathbf{G}_{50}$ are different but the same between $\mathbf{G}_1, \mathbf{G}_{51}, \mathbf{G}_{101}, \ldots$ (i.e. $\mathbf{V}_1, \mathbf{f}_1$).