

Integrating Variant Functional Annotation Scores Have Varied Abilities To Improve Power Of Genome-Wide Association Studies

Jianhui Gao

University of Toronto

Oswaldo Espin-Garcia

University Health Network

Andrew D. Paterson

Hospital for Sick Children

Lei Sun (✉ sun@utstat.toronto.edu)

University of Toronto

Research Article

Keywords:

Posted Date: March 7th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1401415/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Integrating variant functional annotation scores have varied abilities to improve power of genome-wide association studies

Jianhui Gao¹, Osvaldo Espin-Garcia^{1,2}, Andrew D. Paterson^{1,3}, and Lei Sun^{1,4,*}

¹Division of Biostatistics, Dalla Lana School of Public Health, University of Toronto, Toronto, Canada

²Department of Biostatistics, Princess Margaret Cancer Centre, University Health Network, Toronto, Canada

³Program in Genetics and Genome Biology, The Hospital for Sick Children, Toronto, Canada

⁴Department of Statistical Sciences, Faculty of Arts and Science, University of Toronto, Toronto, Canada

*sun@utstat.toronto.edu

ABSTRACT

Functional annotations have the potential to increase power of genome-wide association studies (GWAS) by prioritizing variants according to their biological function, but this potential has not been well studied. We comprehensively evaluated all 1,132 traits in the UK Biobank whose SNP-heritability estimates were given "medium" or "high" labels by Neale's lab. For each trait, we integrated GWAS summary statistics of close to 8 million common variants (minor allele frequency > 1%) with either their 75 individual functional scores or their meta-scores, using three different data-integration methods. Overall, the number of new genome-wide significant findings after data-integration increases as a trait SNP-heritability estimate increases. However, there is a trade-off between new findings and loss of baseline GWAS findings, resulting in similar total numbers of significant findings between using GWAS alone and integrating GWAS with functional scores, across all 1,132 traits analyzed and all three data-integration methods considered. Our findings suggest that, even with the current biobank-level sample size, more informative functional scores and/or new data-integration methods are needed to further improve power of GWAS of common variants.

Introduction

In the last decade, genome-wide association studies (GWAS) have enabled the discovery and identification of thousands of genetic loci across a wide range of phenotypes [1]. However, despite their increasingly large sample sizes (e.g. $n > 100,000$) there is a need to improve the often modest power of GWAS, as effect sizes of causal variants are believed to be small for most complex human traits [2].

To increase power of GWAS, earlier work have leveraged linkage results or summary statistics from independent GWAS of the same or related traits [3, 4]. To integrate information across multiple GWAS, meta-analysis [5] and Fisher's method [6] are two standard and powerful approaches. For example, meta-analysis of summary statistics has been shown to be as powerful as mega-analysis of individual-level data, when there is no heterogeneity between the studies [7, 8]. On the other hand, Fisher's method is more robust to differential directions of effect by combining p-values from different studies.

Recently, it has been shown that variant functional annotations can predict the biological relevance of a variant [9, 10, 11, 12]. To overcome limitations such as incomparable metrics of measurement and differential ascertainment biases across different annotations, several authors have proposed methods to integrate multiple annotations into one single measure: a meta-score [13, 14, 15, 16, 17]. For instance, [16] combined more than 60 genomic features into one combined annotation dependent depletion (CADD) meta-score to provide a measure of the relative deleteriousness for each variant, while [17] developed Eigen, a functional meta-score of similar nature using an unsupervised spectral approach.

Despite the popularity of using these meta-scores for genomic studies [18, 19, 20], their potential for improving power of GWAS has not been well studied or understood. To integrate GWAS summary statistics with meta-scores, in addition to meta-analysis and Fisher's method, we also consider the weighted p-value approach [21] and the stratified false discovery rate (sFDR) control method [22], which extended the traditional FDR control methodology [23]. Both weighted p-value and sFDR have been used to leverage linkage evidence [24, 25], gene-expression data [26, 27] and pleiotropy [28] to increase power of GWAS. Here we use these data-integration methods to integrate CADD or Eigen functional meta-scores with GWAS summary statistics of 1,132 phenotypes from the UK Biobank data [29].

Integrating functional annotation scores with GWAS summary statistics has been previously studied. Recently, [30] proposed a modified weighted p-value-based method called FINDOR to leverage polygenic functional enrichment to improve

power of GWAS. To achieve this, FINDOR uses a stratified linkage disequilibrium (LD) score regression method [31] to compute the expected χ_1^2 statistic for each GWAS SNP, by regressing the observed GWAS χ_1^2 statistics of the tagging SNPs against their 75 functional annotation scores [32]. FINDOR then stratifies the GWAS SNPs into 100 equally-sized bins based on their expected GWAS χ_1^2 values and applies bin-specific weights to the corresponding GWAS p-values. An application of FINDOR by [30] to 27 traits, selected from the UK Biobank data [29], showed that the method was able to improve power of GWAS by identifying additional associated variants.

To answer the question of whether prioritizing variants according to their biological function could improve power of GWAS, our study here is different from the FINDOR evaluation of [30] in several ways. First, unlike FINDOR, we use methods that prioritize GWAS findings based on external information alone to minimize concern of over-fitting. That is, the weighting factor and stratification are determined based on the annotations alone, independent of the observed GWAS summary statistics. Second, we utilize existing meta-scores that are already calibrated and easier to implement in practice, instead of using many individual functional scores. Third, we focus on evaluating methods' robustness to the possibility of uninformative or even misleading functional annotations, because our understanding of the functionality of a genetic variant is incomplete and evolving. Finally, we comprehensively examine all 1,132 UK Biobank traits for which the confidence for their SNP-heritability estimates were considered medium to high by Benjamin Neale's lab from the Broad Institute (hereafter referred to as Nealelab; Web Resources). For each trait, we integrated GWAS summary statistics of close to 8 million common variants with their functional scores using three different data-integration methods: FINDOR with 75 individual functional scores, and weighted p-value and stratified false discovery control methods with CADD (or Eigen) meta-scores.

In addition to the large-scale UK Biobank application, we also conducted a large-scale simulation study using different study designs, from leveraging the observed genomic data combined with simulated genetic data or vice versa to using only simulated data. We also considered different and complementary performance measures, from the traditional family-wise error rate (FWER) to FDR, power, recall, precision, and relative efficiency. Finally, we sought to evaluate the functional annotation similarity between variants in linkage disequilibrium (LD), which has not been previously studied but an important consideration when integrating functional scores with GWAS.

Results

Method overview

Focusing on integrating functional scores with GWAS summary statistics to improve power of GWAS, we considered five data-integration methods, namely meta-analysis [5], Fisher's method [6], weighted p-value [21], sFDR [22], and FINDOR [30]. The meta-analysis and Fisher's method were only included in some of the simulation studies to demonstrate that, although commonly used in many other scientific settings, they are not suitable for integrating GWAS with genomic functional scores.

Prior to the large-scale UK Biobank application using methods for which we understand their performance properties, we conducted a large-scale simulation study using three complementary study designs: i) leveraging the observed functional annotations and integrating them with simulated GWAS, ii) simulating functional annotations and combining them with observed GWAS, and iii) using only simulated data. For an unbiased method evaluation, we also considered different performance measures, ranging from the traditional family-wise error rate to false discover control, and from the traditional power to recall, precision and relative efficiency.

Results of simulation design I, leveraging the observed genomic data

Here, simulated GWAS summary statistics, generated under the null of no association, were integrated with real functional annotations. The empirical family-wise error rates were estimated from 50,000 simulated replicates.

For the baseline analysis, using the null GWAS summary statistics alone, the empirical FWER is 0.0496 (Table S1). For the five different data-integration methods, the empirical rates are 0.0477, 0.0366, 0.0501, 0.0474, and 0.0537, respectively, for meta-analysis, Fisher's method, weight p-value, sFDR, and FINDOR, where FINDOR is the only method with slightly increased type I error rate. Although a method with an empirical FWER estimate outside [0.047, 0.053] can be considered inaccurate, overall all methods have reasonable type I error control in this setting.

Table S1 also provides a detailed account of the numbers of replicates, out of a total of 50,000 replicates, with at least one, two or three false findings for each of the methods; no method had more than three false findings per GWAS.

Results of simulation design II, leveraging the observed genetic data

Here, real UK Biobank GWAS summary statistics of the 1,132 traits were integrated with *permuted* CADD (or Eigen) meta-scores; FINDOR was not applicable here as previously explained. The performance measures used here are $Recall_t = TP_t/m_{1,t}$ and $Precision_t = 1 - FDR_t = TP_t/P_t$, where $m_{1,t}$ is the number of genome-wide significant GWAS findings prior to data-integration for trait t , and P_t and TP_t are, respectively, the numbers of total positives and true positives after data-integration.

The *Recall* results shown in Figure 1 confirm that meta-analysis and Fisher’s method are not suitable for integrating functional annotations with GWAS summary statistics. Across the 723 GWAS with at least one significant finding prior to data integration ($m_{1,t} > 0$), the [Q1, median, Q3] *Recall* rates are [50%, 66.67%, 73.34%] for meta-analysis and [70%, 84.23%, 92.15%] for Fisher’s method after integrating permuted CADD scores. In contrast, these values are [95.87%, 100%, 100%] for the weight p-value method and [100%, 100%, 100%] for the sFDR control. The *Precision* results in Figure 1 corroborate the findings based on *Recall*.

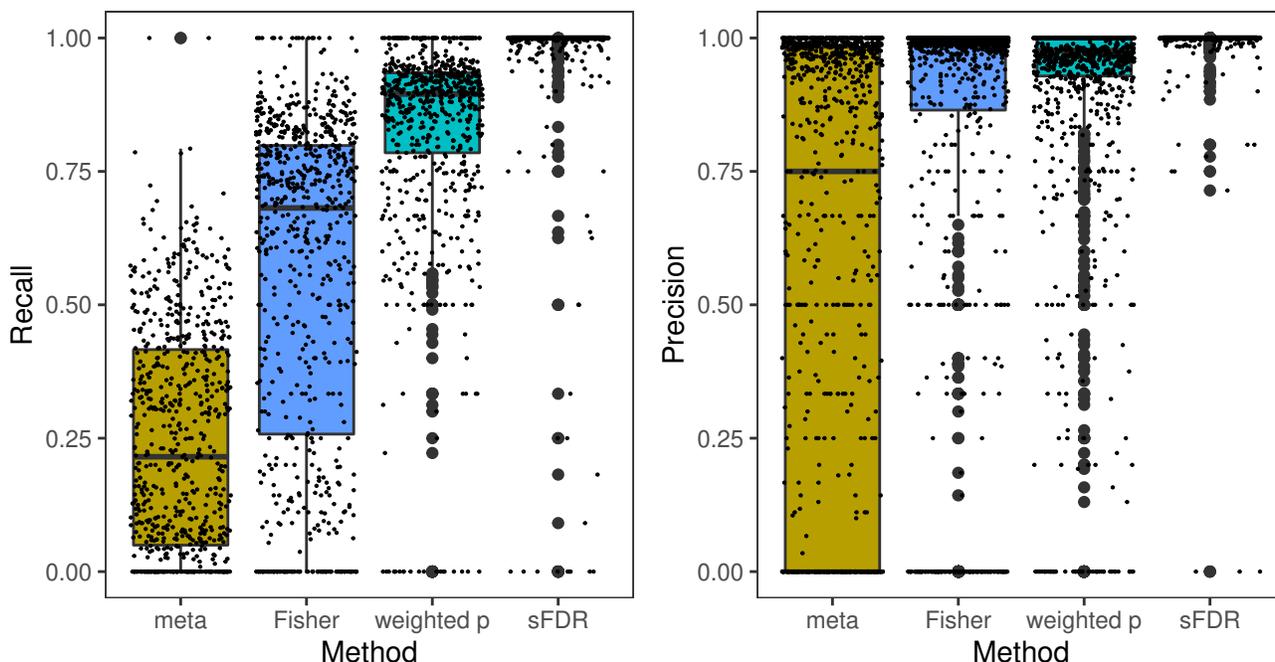


Figure 1. The *Recall* and *Precision* rates obtained from simulation study design II, integrating the 1,132 UK Biobank GWAS summary statistics with permuted CADD functional meta-scores, using meta-analysis, Fisher’s method, the weighted p-value approach, and the stratified FDR control. $Recall_t = TP_t/m_{1,t}$ and $Precision_t = 1 - FDR_t = TP_t/P_t$, where $m_{1,t}$ is the number of genome-wide significant independent loci prior to data-integration for trait t , and P_t and TP_t are the numbers of positives and true positives after data-integration; see Table S2 for additional results. Independent loci were defined using PLINK’s LDclumping algorithm with a 1 Mb window and an r^2 threshold of 0.1.

The results here consistently show the sensitivity issue of meta-analysis and Fisher’s methods, and they confirm that sFDR is more robust than the weighted p-value approach, which was demonstrated by [25] when integrating linkage results with GWAS. Results stratified by the four types of traits analyzed, nonsig, nominal, z4, and z7 (Figure S10), counting significant SNPs instead of loci (Figure S11), or using permuted Eigen scores (Figure S12) led to the same conclusion.

Permuting the meta-scores provides random functional annotations, independent of the GWAS summary statistics, for type I error evaluation. However, as noted earlier, it is of value to examine annotation similarities between SNPs in linkage disequilibrium. A similarity measure, $s_{i,j}^2$, was introduced to be compared with the LD measure, $r_{i,j}^2$. Results in Figure S7 show that there is no clear concordance between the two measures. A closer examination of $s_{i,j}^2$ and $r_{i,j}^2$ for two randomly selected regions is shown in Figure S8, and the contrast between variant-specific CADD meta-score and LD score across the genome is shown Figure S9. Both figures led to the same conclusion that functional scores of SNPs in strong LD are not necessarily similar.

Results of simulation design III, varying the informativeness of genomic information

Here, both the GWAS summary statistics and the additional information available for data-integration were simulated, with varying degree of informativeness, including uninformative or possibly misleading annotation scores. The performance measures here are power and relative efficiency (RE), where RE was defined as one minus (the average ranks of the truly associated SNP after data-integration) divided by (their average baseline ranks using GWAS data alone).

The RE results in Figure 2 are consistent with those from simulation study design II. While meta-analysis and Fisher’s method work well when the additional information is completely informative (i.e. the two data resources are homogeneous with

each other), they are not suitable data-integration methods for this study setting.

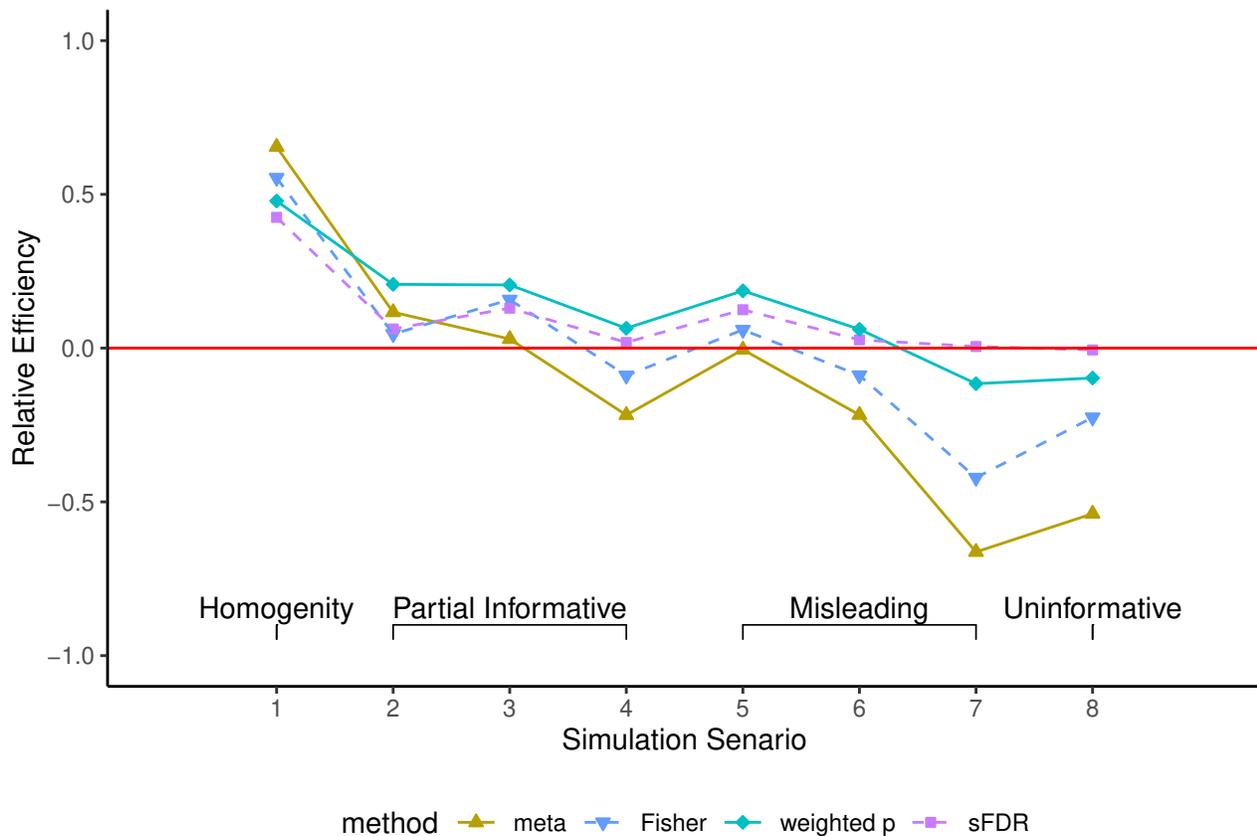


Figure 2. The relative efficiency (*RE*) obtained from simulation study design III, integrating simulated GWAS summary statistics with simulated additional information with varying degrees of informativeness, using meta-analysis, Fisher’s method, the weighted p-value approach, and the stratified FDR control. There were 10,000 independent SNPs, among which 100 were truly associated whose summary statistics were drawn from $N(3, 1)$; the rest from $N(0, 1)$. For the additional information available for data integration, the details of the eight simulation scenarios are provided in the text and illustrated in Figure S13. *RE* is one minus (the average ranks of the truly associated SNP after data-integration) divided (by their average base-line ranks using GWAS data alone), averaged across 1000 simulation replicates.

The *RE* results are consistent with the power results in Figure S14, across different rejection rules including controlling FWER at 5%, rejecting top 100 ranked SNPs, and controlling FDR at 5% or 20%. Consistent results were also observed when μ_1 was varied from 0.1 to 4 to represent different power scenarios of a GWAS (Figures S15 to S20). Thus, meta-analysis and Fisher’s method were excluded from the application study.

Results from integrating functional annotations to improve power of UK Biobank GWAS

The UK Biobank GWAS summary statistics from Nealelab (between each of the 7,895,174 common SNPs and each of the 1,132 UK Biobank traits) were integrated with variant’s functional annotation scores using three data-integration methods, where FINDOR used 75 individual annotation scores, and weighted p-value and sFDR used CADD (or Eigen) meta-scores as described earlier. The total number of independent, significant loci detected at the 5×10^{-8} level, as well as *Recall* and *New Discoveries* were calculated.

No striking improvements across the 1,132 traits

Figure 3 shows the distributions of the total number of independent, significant loci identified by using GWAS alone (as a baseline, the first box-plot within each sub-figure) or after applying FINDOR, weight p-value and sFDR data-integration methods, stratified by the four types of traits analyzed (nonsig, nominal, z4, and z7).

Overall, integrating the existing functional annotations with the UK Biobank GWAS association statistics did not lead to

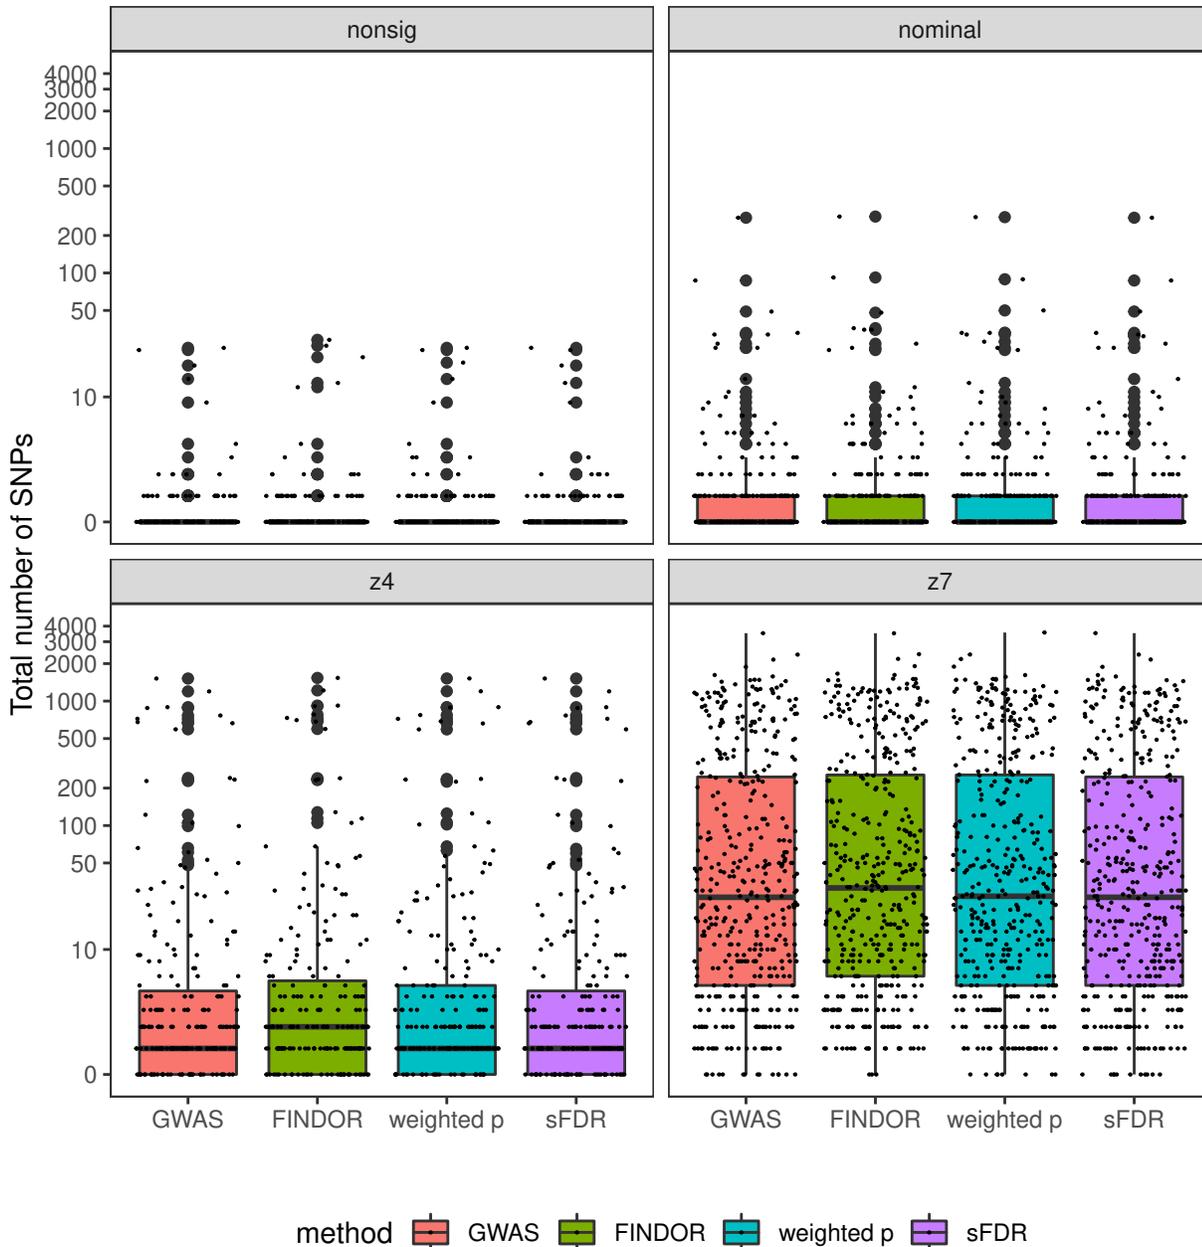


Figure 3. The total numbers of genome-wide significant independent loci of the UK Biobank GWAS application study, before and after data-integration with functional annotations, stratified by the four phenotype categories. In each figure, the total number of significant loci identified based on the UK Biobank GWAS data alone serves as a baseline. The GWAS baseline box-plot is followed by the box-plots for the total numbers of significant loci after integrating the UK Biobank GWAS summary statistics with functional annotations using FINDOR (using 75 individual annotation scores), and the weighted p-value and stratified FDR control methods (each using the CADD meta-score), analyzing 7,895,174 variants for each of the 1,132 UK Biobank traits. The 1,132 traits were rated by Nealelab as having medium to high confidence for their heritability estimates, and they fall into four categories: nonsig (182 traits; heritability testing p-value $p > 0.05$), nominal (277 traits; $p < 0.05$), z4 (235 traits; $p < 3.17 \times 10^{-5}$), and z7 (438 traits; $p < 1.28 \times 10^{-12}$). Independent loci were defined using PLINK's LDclumping algorithm with a 1 Mb window and an r^2 threshold of 0.1.

striking improvements, irrespective of the data-integration method. Results of using Eigen (Figure S21) or counting SNPs instead of independent loci (Figure S22) are characteristically similar.

The overall limited improvement is also evident from Table 1. For example, among the 1,132 traits, prior to data-integration 772 have at least one genome-significant, independent loci. After data-integration these numbers are 738, 746 and 717 by, respectively, FINDOR, weighted p-value and sFDR; the counts stratified by the four trait categories are also provided in Table 1. Similarly, 337 traits have more than ten significant loci prior to data-integration, and the numbers are 353, 346 and 337 post-data-integration by the three method.

	GWAS alone	After data-integration with functional annotation scores		
		75 individual scores	CADD meta-score	
		FINDOR	Weighted p-value	Stratified FDR
# of traits with >0; 5; 10 significant loci				
all traits	772; 402; 337	738; 420; 353	746; 408; 346	717; 403; 337
nonsig	36; 5; 4	39; 5; 5	40; 5; 4	36; 5; 4
nominal	110; 18,11	102; 21; 11	113; 21; 11	109; 18; 11
z4	160; 55; 42	169; 59; 45	172; 56; 43	157; 56; 42
z7	416; 324; 280	428; 335; 292	421; 326; 288	415; 324; 280
[Q1, Median, Q3] of # of significant loci across traits				
all traits	[0, 2, 18]	[0, 2, 21]	[0, 2, 20]	[0, 2, 18]
nonsig	[0, 0, 0]	[0, 0, 0]	[0, 0, 0]	[0, 0, 0]
nominal	[0, 0, 1]	[0, 0, 1]	[0, 0, 1]	[0, 0, 1]
z4	[0, 1, 5]	[0, 2, 6]	[0, 1, 5]	[0, 1, 5]
z7	[5, 27, 246]	[6, 32, 255]	[5, 27, 255]	[5, 27, 246]
# of traits with >0; 5; 10 <i>New Discoveries</i>				
all traits	NA	553; 227; 165	472; 180; 139	89; 0; 0
nonsig	NA	20; 0; 0	15; 0; 0	0; 0; 0
nominal	NA	50; 3; 1	37; 1; 0	3; 0; 0
z4	NA	103; 22; 11	77; 16; 10	9; 0; 0
z7	NA	380; 202; 153	343; 163; 130	77; 0; 0
[Q1, Median, Q3] of # of <i>New Discoveries</i> across traits				
all traits	NA	[0, 0, 3]	[0, 0, 2]	[0, 0, 0]
nonsig	NA	[0, 0, 0]	[0, 0, 0]	[0, 0, 0]
nominal	NA	[0, 0, 0]	[0, 0, 0]	[0, 0, 0]
z4	NA	[0, 0, 1]	[0, 0, 1]	[0, 0, 0]
z7	NA	[1, 4, 19]	[1, 3, 14]	[0, 0, 0]

Table 1. Results of the UK Biobank application study, before and after data-integration with functional annotations, stratified by the four phenotype categories. The three data-integration methods integrated the UK Biobank GWAS summary statistics with functional annotations using FINDOR (using 75 individual annotation scores), and the weighted p-value and stratified FDR control methods (each using the Eigen meta-score), analyzing 7,895,174 variants for each of the 1,132 UK Biobank traits. The 1,132 traits were rated by Nealelab having medium to high confidence for their heritability estimates, and they fall into four categories: nonsig (182 traits; heritability testing $p > 0.05$), nominal (277 traits; $p < 0.05$), z4 (235 traits; $p < 3.17 \times 10^{-5}$), and z7 (438 traits; $p < 1.28 \times 10^{-12}$). See Table S1 for additional results.

Further, the intersection of significant loci between methods displayed in Figure S23 shows that out of a total of 59,764 significant loci identified in the z7 category, 46,631 (78%) were common across the three data-integration methods and the baseline GWAS alone. Additionally, Figure S28 shows that the total numbers of significant loci after data-integration are similar to those based on GWAS alone, for all three data-integration methods and across all 1,132 traits.

New Discoveries for the 182 traits in the nonsig category

Although the ground-truth is unknown in application studies, the *New Discoveries* for traits in the nonsig category may be considered as false positives, as their SNP-heritability testing p-values were > 0.05 and the inferences were given "medium" or "high" confidence by Nealelab; traits in this category include, for example, "Fizzy drink intake", "Apple intake", "Time spent doing moderate physical activity", and "Work hours".

Among the 182 traits in the nonsig category, 20, 15 and 0 traits had at least one *New Discoveries* after data-integration using, respectively, FINDOR, and weighted p-value and sFDR when using CADD (Table 1 and Figure 4). Weighted p-value and sFDR using Eigen led to 9 and 1 traits with at least one *New Discoveries* (Table S2). Reassuringly, no data-integration

methods led to more than five *New Discoveries* for any of the 182 traits in the nonsig category.

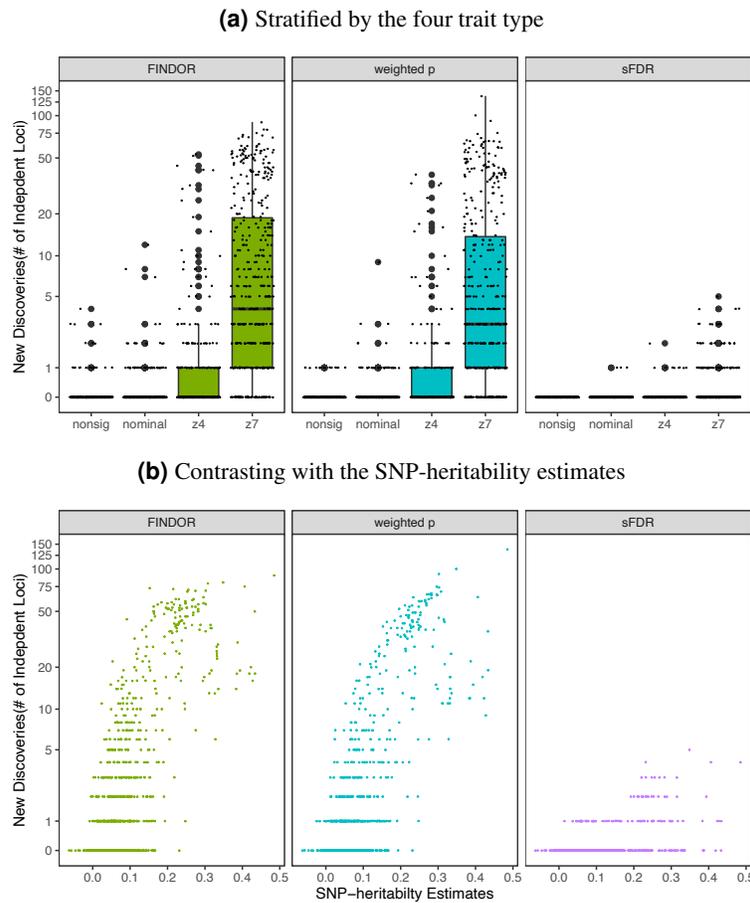


Figure 4. The total numbers of genome-wide significant independent loci of the UK Biobank GWAS application study, before and after data-integration with functional annotations, stratified by the four phenotype categories. In each figure, the total number of significant loci identified based on the UK Biobank GWAS data alone serves as a baseline. The GWAS baseline box-plot is followed by the box-plots for the total numbers of significant loci after integrating the UK Biobank GWAS summary statistics with functional annotations using FINDOR (using 75 individual annotation scores), and the weighted p-value and stratified FDR control methods (each using the CADD meta-score), analyzing 7,895,174 variants for each of the 1,132 UK Biobank traits. The 1,132 traits were rated by Nealelab as having medium to high confidence for their heritability estimates, and they fall into four categories: nonsig (182 traits; heritability testing p-value $p > 0.05$), nominal (277 traits; $p < 0.05$), z4 (235 traits; $p < 3.17 \times 10^{-5}$), and z7 (438 traits; $p < 1.28 \times 10^{-12}$). Independent loci were defined using PLINK’s LDclumping algorithm with a 1 Mb window and an r^2 threshold of 0.1.

New Discoveries for the 438 traits in the z7 category

The number of *New Discoveries* increases as a trait’s SNP-heritability estimate increases (Figure 4(B)), and there are increased numbers of *New Discoveries* for the 438 traits in the z7 category (Table 1). This is consistent with the results of [30] who studied 27 highly heritable traits in the UK Biobank, which we replicated here.

Among the 438 traits in the z7 category, 380, 343 and 77 traits had at least one *New Discoveries* after data-integration using, respectively, FINDOR, and weighted p-value and sFDR when using CADD (Table 1 and Figure 4). Weighted p-value and sFDR using Eigen led to characteristically similar results (Table S2 and Figure S31).

Additionally, FINDOR and weighted p-value led to more than ten *New Discoveries* for, respectively, 153 and 130 traits in the z7 category. However, the two methods also led to loss of significant loci that were present in the baseline GWAS, resulting in similar total numbers of significant loci before and after data-integration (Figure S29). In general, FINDOR and weighted p-value methods yielded similar performance, which is somewhat expected as FINDOR applies the weighted p-value principle. Both methods have noticeably more findings than sFDR. This is also expected given the trade-off between power and robustness, which we explore further.

Trade-off between New Discoveries and Recall

Figure 5(A) shows the *Recall* for the 337 traits with more than ten GWAS signals prior to data-integration ($m_{1,t} > 10$), as the stability of a *Recall* estimate depends on $m_{1,t}$. For the 795 traits with $m_{1,t} \leq 10$, instead of showing $P_t/m_{1,t}$, Figure 5(B) contrasts P_t with $m_{1,t}$; see Figure S24(A) for *Recall* of all traits and Figure S24(B) for *Recall* versus SNP-heritability estimates.

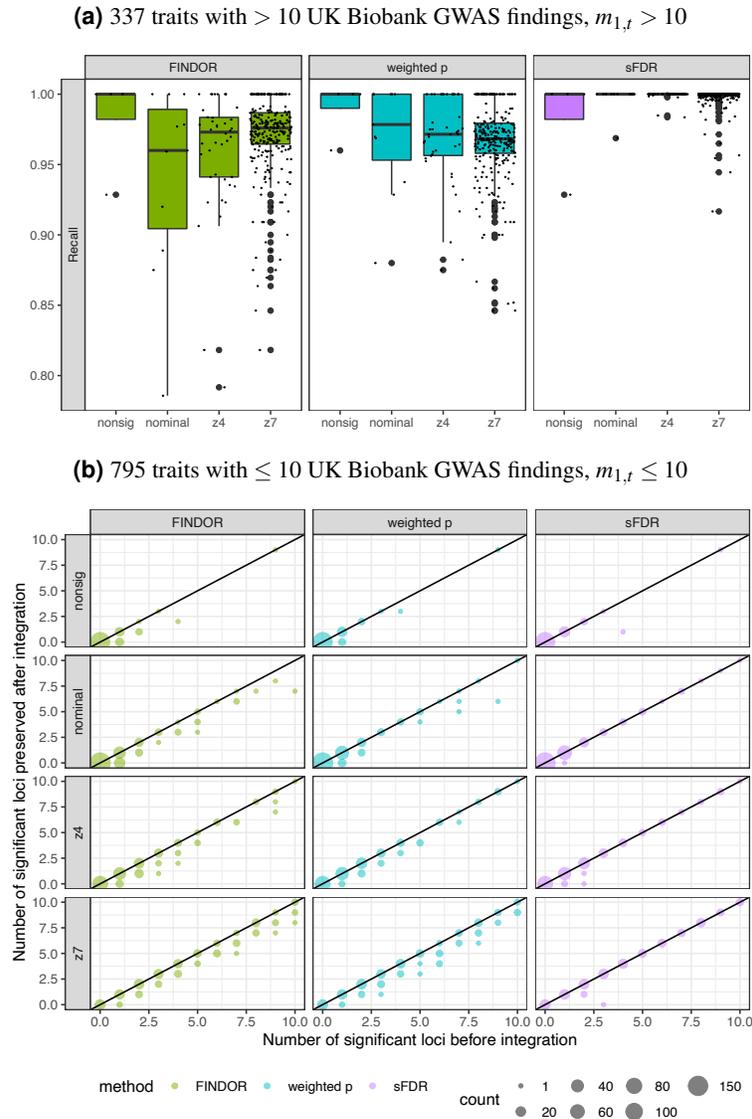


Figure 5. Results of the UK Biobank GWAS application study, before and after data-integration with functional annotations, stratified by the four phenotype categories. (A) $Recall_t = TP_t/m_{1,t}$, where $m_{1,t}$ is the number of genome-wide significant independent loci prior to data-integration for trait t , and TP_t is the number of true positives after data-integration. *Recall* estimation is not stable when $m_{1,t}$ is small so for $m_{1,t} \leq 10$, (B) contrasts the number of significant loci preserved after data-integration with $m_{1,t}$. The three data-integration methods integrated the UK Biobank GWAS summary statistics with functional annotations using FINDOR (using 75 individual annotation scores), and the weighted p-value and stratified FDR control methods (each using the CADD meta-score), analyzing 7,895,174 variants for each of the 1,132 UK Biobank traits. The 1,132 traits were rated by Nealelab having medium to high confidence for their heritability estimates, and they fall into four categories: nonsig (182 traits; heritability testing $p > 0.05$), nominal (277 traits; $p < 0.05$), z4 (235 traits; $p < 3.17 \times 10^{-5}$), and z7 (438 traits; $p < 1.28 \times 10^{-12}$). Independent loci were defined using PLINK’s LDclumping algorithm with a 1 Mb window and an r^2 threshold of 0.1.

It is clear that sFDR has better *Recall* than either FINDOR or weighted p-value for the traits in the nominal, z4 and z7 categories. For example, for the 280 traits in the z7 category with at least one significant finding prior to data-integration, the

median *Recall* is 97.6%, 96.8% and 100%, respectively, for the FINDOR, weighted p-value and sFDR methods (Table S2).

The trade-off between power and robustness is also supported by results in Figure S25 (*Recall* of the 402 traits with $m_{1,t} > 5$), and other supplementary figures including Figures S26 (P_t versus $m_{1,t}$ for traits with $m_{1,t} \leq 50$), Figure S27 (traits with $m_{1,t} \leq 10$) and Figure S28 (for all traits), as well as Figure S30 (*Recall* of weighted p-value and sFDR using Eigen, instead of CADD). It is also clear that, *Recall* increases as a trait SNP-heritability estimate increases (Figure S24(B)), which was also observed for *New Discoveries* (Figure 4(B)).

Discussion

There has been much discussion about the value of integrating functional annotations into genetic association studies [33, 34, 35], but theoretical evaluation and large-scale application to test this hypothesis has been limited [30]. In addition to conducting comprehensive simulation studies, we performed a large-scale application study of all 1,132 traits in the UK Biobank, for which the SNP-heritability estimates were given "medium" or "high" labels by Nealelab. For each trait, we integrated GWAS summary statistics of close to 8 million common variants with their functional scores using three different data-integration methods: FINDOR with 75 individual functional scores, and weighted p-value and stratified false discovery control methods with CADD (or Eigen) meta-scores.

We observed that, although the numbers of new genome-wide significant findings after data-integration increase as trait SNP-heritability estimates increase, there is a trade-off between new findings and loss of the original GWAS findings. This resulted in similar total numbers of significant findings between using GWAS alone and integrating GWAS with functional scores, across all 1,132 traits analyzed and all three data-integration methods considered. A closer examination of method performance and trait heritability revealed that all methods performed better (more *New Discoveries* and higher *Recall*) for traits with higher estimates of SNP-heritability (Figures 4(B) and S24(B)).

Our study used CADD and Eigen as the functional meta-score available for data-integration using weighted p-value and sFDR. To the best of our knowledge, CADD was the first meta-score in the literature and Eigen was the first to use unsupervised learning approach, and both meta-scores have been shown to be superior to other scores in some genomic studies [16, 17]. However, the recent work by [18] has proposed annotation-PCs, an alternative that warrants further investigation.

In one of our simulation studies, we permuted CADD (and Eigen) to provide a set of meta-scores that are independent of the GWAS summary statistics of the UK Biobank data. Although this approach is valid for examining type I error control, we were intrigued by the question of whether meta-scores of SNPs in strong linkage disequilibrium are similar. To answer this question, we defined $s_{i,j}^2 = 1 - |CADD_i - CADD_j| / (CADD_i + CADD_j)$ as the functional similarity measure between SNPs i and j . Interestingly, there was no clear concordance between $s_{i,j}^2$ and r^2 , the LD measure for genotype similarity (Figures S7 and S8). Additionally, there was no relationship between CADD and LD score (Figure S9).

Throughout this paper, we have used the default tuning parameter values, $\beta = 2$ for the weighted p-value approach and $k = 2$ for the sFDR method. We did not tune the parameters to select values that lead to the 'best' results, for which valid result interpretation requires adjustment for the inherent data-dredging or selective inference. The choice of different β and k values, however, has an effect on method performance. Figure S32 shows the results of the full analysis of the UK Biobank application study. For the weighted p-value approach, the default $\beta = 2$ led to the highest number of *New Discoveries*, but at the same time it resulted in the lowest *Recall* rate. For the sFDR method, $k = 10$ or 20 lead to an increased number of *New Discoveries* as compared with the default $k = 2$, at the cost of slightly reduced *Recall* rates. Thus, unlike the previous linkage and GWAS integration setting, the default value of $k = 2$ for sFDR appears to be sub-optimal for integrating functional meta-score with GWAS.

Among the three data-integration methods, FINDOR and weighted p-value behaved more similarly to each other than with sFDR, where FINDOR and weighted p-value led to more new discoveries for traits considered heritable (the traits in the nominal, z4 and z7 categories). However, sFDR is the most robust with zero new findings ("false positives") for the traits in the nonsig category, and it also had the highest *Recall* and *Precision*.

Conclusions

Overall, the number of new genome-wide significant discoveries is limited across the 1,132 UK Biobank traits analyzed and the three data-integration methods investigated. This suggests the need for more informative functional scores and/or new data integration methods to further improve power of GWAS through leveraging variant functional annotations. It is important to note that this conclusion applies to bi-allelic common autosomal variants with MAF greater than 1%, which may not be generalizable to rare variants.

Methods

The integration methods: meta-analysis, Fisher’s method, weighted p-value, and stratified false discovery rate (sFDR) control

Notation and set-up

Let z_i and p_i be the association test statistic and its corresponding p-value for SNP i , $i = 1, \dots, m$, from a genome-wide association study, the primary data of interest. Without loss of generality, we assume z_i follows $N(0, 1)$, the standard normal distribution, under the null hypothesis of no association between the SNP and the GWAS trait under the study.

Let $z_{i,add}$ and $p_{i,add}$ be additional information available for the SNP, based on data *independent* of z_i and p_i from the GWAS. Note that $z_{i,add}$ may or may not be normally distributed depending on the application setting, e.g. $z_{i,add}$ can be the CADD [16] or Eigen [17] functional meta-score available for SNP i , which will be the focus of this study.

Meta-analysis and Fisher’s method

For the meta-analysis approach, we first assume the best-case scenario where $z_{i,add}$ is normally distributed. We then use the inverse variance approach [36] to integrate z_i and $z_{i,add}$, $Z_i^{meta} = (\sqrt{1/v_i}z_i + \sqrt{1/v_{i,add}}z_{i,add})/\sqrt{1/v_i + 1/v_{i,add}}$, where the weights depend on v_i and $v_{i,add}$, the variance estimates associated with, respectively, the GWAS and the additional study available for data integration. Under the null hypothesis of no association *and* assuming the functional meta-score is uninformative, Z_i^{meta} is $N(0, 1)$ distributed.

Fisher’s method combines p-values instead of the test statistics, $Z_i^{Fisher} = -2(\log(p_i) + \log(p_{i,add}))$. Fisher’s method is omnibus to directions of effect, and as a result it can be more powerful than meta-analysis when signs of z_i and $z_{i,add}$ differ. Under the null that both p_i and $p_{i,add}$ are independently $\text{Unif}(0, 1)$ distributed, Z_i^{Fisher} is $\chi_{df=4}^2$ distributed.

Although meta-analysis and Fisher’s method are applicable in many scientific settings, their applications to genetic association studies are typically restricted to combining evidence from multiple GWAS of the same phenotype and from the same population. This is because the statistical power of meta-analysis (and Fisher’s method) relies on the assumption of homogeneity beyond direction of effect [37]. In practice, given two families of multiple tests, the underlying compositions of the null and alternative hypotheses may differ, unless the two studies used the same study design including phenotype definition, genotyping platform, environmental exposure, and study population [38]. When the truly associated SNPs do not completely overlap between the different studies, using random-effect (instead of fixed-effect meta-analysis) does not guarantee improved power, because it violates the assumption that the effect sizes come from the *same* distribution.

The use of meta-analysis and Fisher’s method is also questionable when z_i and $z_{i,add}$ from the two studies offer different types of information. In essence, the use of weights $\sqrt{1/v_i}$ and $\sqrt{1/v_{i,add}}$ notwithstanding, meta-analysis and Fisher’s method implicitly assume z_i and $z_{i,add}$ carry ‘exchangeable’ information. For our study, however, z_i is the genetic association summary statistic, while $z_{i,add}$ is the genomic annotation meta-score. Thus, meta-analysis and Fisher’s method are likely to be sub-optimal for the purpose of this study. However, for completeness we include the two classical data-integration methods in our initial method evaluation.

For a practical implementation of meta-analysis when $z_{i,add}$ is the CADD or Eigen meta-score, we use equal weights as the sample size of a functional study is unclear. Further, we use the inverse normal transformation to re-scale $z_{i,add}$ while keeping the sign of the re-scaled $z_{i,add}$ to be the same as z_i , creating the best-case scenario for the meta-analysis. Similarly, for a practical implementation of Fisher’s method, we use a rank-based transformation and let $p_{i,add} = (\text{rank of } z_{i,add}/m)$, which is also related the phred-scaled CADD and Eigen scores which we discuss later.

The weighted p-value approach

Unlike meta-analysis and Fisher’s method, which assume z_i and $z_{i,add}$ carry similar information, the weighted p-value approach [21] treats z_i and $z_{i,add}$ differently. That is, the method considers z_i and p_i as the primary data of interest, and it transforms $z_{i,add}$ to w_i , a weight to be applied to p_i . Thus, the weighted p-value approach is an attractive method for this study setting, where the primary data are GWAS summary statistics, and the additional information available are genomic functional scores derived *independently* from the GWAS of interest.

For a valid weighted p-value implementation, the w_i ’s must satisfy two conditions: $w_i \geq 0$ and $\bar{w} = \sum w_i/m = 1$ [21]. To transform $z_{i,add}$ to w_i , [24] studied two possible weighting schemes: exponential, $w_i = m(\exp(\beta \times z_{i,add})/\sum_i \exp(\beta \times z_{i,add}))$, and cumulative,

$$w_i = m \frac{\Phi(z_{i,add} - \beta)}{\sum_i \Phi(z_{i,add} - \beta)}, \quad (1)$$

where Φ is the cumulative distribution function of the standard normal. In either case,

$$p_{i,weighted} = \min\left\{\frac{p_i}{w_i}, 1\right\}. \quad (2)$$

Here we choose the cumulative weighting scheme, with the recommended default value of $\beta = 2$ [24]. This is because the exponential weighting scheme is highly sensitive to large values of $z_{i,add}$, which is the case here; functional meta-scores can be as large as 80 [16].

Stratified false discovery rate (sFDR) control

Unlike the weighted p-value approach that up- or down-weights each SNP according to its external information $z_{i,add}$, the sFDR method separates the GWAS SNPs into different groups based on $z_{i,add}$, which can be categorical or continuous [22]. When $z_{i,add}$ is continuous, it has been shown that categorizing $z_{i,add}$ does not necessarily result in loss of power, as the additional information available are unlikely to be precisely informative [25]. In addition, sFDR is robust to the situation when $z_{i,add}$ is uninformative (i.e. random) or possibly misleading.

To implement sFDR in our setting where $z_{i,add}$ is the continuous functional meta-score, without loss of generality, we first stratify GWAS SNPs into two groups based on whether their meta-scores are among the top five percent or not, *irrespective* of z_i and p_i , the GWAS summary statistics. (The choice of the number of groups and thresholds, however, is subjective, similar to choosing the weighting scheme and β value for the weighted p-value approach above.) As a result, there are two groups of GWAS SNPs, where group 1 contains 5% of the GWAS SNPs with the highest functional meta-scores and group 2 contains the remaining SNPs. It is worth emphasizing that group 1 is only presumed to be the high-priority group, as the stratification is based on genomic $z_{i,add}$ alone, independent of the GWAS z_i or p_i .

We then apply FDR control, separately, to the two groups of GWAS p-values, but using the same pre-specified FDR $\gamma\%$ level. Following the sFDR method of [22], for each group of SNPs we first transform their GWAS p-values, p_i 's, to q-values, q_i 's [39], and we then reject the SNPs with $q_i < \gamma\%$; this sFDR procedure controls the overall FDR at the $\gamma\%$ level. Although sFDR does not explicitly use weights, group-specific weights can be derived [25].

Let m^k be the number of SNPs in group k , and let $\pi_0^{(k)}$ be the proportion of truly associated SNPs in the group. Within each group, we obtain q-values recursively [39], $q_i = \min\{\hat{\pi}_0^{(k)} m^{(k)} p_{(i)}/i, q_{i+1}\}$, where $p_{(1)} \leq \dots \leq p_{(i)} \leq \dots \leq p_{(m)}$ are the ordered GWAS p-values, and the procedure starts from $q_{(m)} = \hat{\pi}_0 p_{(m)}$. To obtain $\hat{\pi}_0$, we choose the commonly used conservative estimate [40], $\hat{\pi}_0 = \{\text{the number of SNPs with } p_i > 0.5\} / \{0.5m^{(k)}\}$.

After rejecting SNPs with $q_i < \gamma\%$ separately for each group of SNPs, let $\alpha^{(k)}$ be the maximum GWAS p-values among the rejected SNPs for group k , the group-specific weight is,

$$w^{(k)} = m \frac{\alpha^{(k)}}{\sum_k \alpha^{(k)} m^{(k)}}. \quad (3)$$

We can then obtain sFDR weighted p-values,

$$p_{i,sFDR} = \min\left\{\frac{p_i}{w^{(k)}}, 1\right\}. \quad (4)$$

If group 1 has no rejections at the pre-specified FDR $\gamma\%$ level, we set $w^{(1)} = 0$ and $w^{(2)} = m/m^{(2)}$. Similarly, if group 2 has no rejections, $w^{(1)} = m/m^{(1)}$ and $w^{(2)} = 0$. If both groups have no rejections at the $\gamma\%$ level, then $w^{(1)} = w^{(2)} = 1$. That is, the study is reduced to the unweighted case.

The sFDR group-specific weights, $w^{(k)}$'s, satisfy the constraints imposed by the weighted p-value approach [21], and they have been shown to be a robust version of the SNP-specific w_i 's [25]. If the additional information is truly informative, $w^{(1)} > 1$ while $w^{(2)} < 1$, and they can be considered as dichotomized w_i 's of the weighted p-value approach. In that case, the weighted p-value approach is more powerful than sFDR. On the other hand, if the information is just random noise, $w^{(1)} \approx w^{(2)} \approx 1$ for sFDR, while the weighted p-value method still up- or down-weights the GWAS p-values according to the SNP-specific w_i 's, which are proportional to the observed $z_{i,add}$'s. In the event of misleading information, $w^{(1)} < 1$ while $w^{(2)} > 1$ even though group 1 was presumed to be the high-priority group. Thus, sFDR is robust to uninformative or even misleading added information.

The UK Biobank GWAS summary statistics for 1,132 complex traits

We obtained the UK Biobank GWAS round 2 summary statistics from Nealelab (Web Resources). Nealelab performed association studies for 4,236 complex traits using regression model with additively coded genotype, as well as age, sex and the first 20 principal components as covariates. For each of these traits, Nealelab also applied the LD-score regression method [41] to estimate the SNP-heritability, which ranges from 0% to 48%.

In addition to testing if the SNP-heritability is 0%, Nealelab also provided a confidence level ("low", "medium" or "high") to the heritability inference for each trait. Thus, we restricted our analysis to the 1,132 traits denoted with "medium" or "high" confidence labels, which were primarily based on the effective sample sizes $> 20,000$. Of the 1,132 traits analyzed, 531 are continuous and 601 are binary traits. For a binary trait, the effective sample size depends on the number of cases or controls;

see Figure S1 for a histogram of the case rates for the 601 binary traits. Nealelab then classified these 1,132 traits into four categories: nonsig (182 traits with SNP-heritability testing p-value $p > 0.05$), nominal (277 traits; $p < 0.05$), z4 (235 traits; $p < 3.17 \times 10^{-5}$), and z7 (438 traits; $p < 1.28 \times 10^{-12}$), where the nonsig category can serve a negative control for the purpose of this study. Figure S2 contrasts the heritability h_g^2 estimates of the 1,132 traits with their SNP-heritability testing z-values.

The UK Biobank GWAS results of Nealelab were derived from $n = 361,194$ individuals of white-British ancestry and 10.9 million variants that passed a set of quality control (QC) steps; see Web Resources for the detailed QC steps performed by Nealelab. Our data-integration analysis focused on $m = 7,895,174$ common bi-allelic autosomal SNPs. We excluded indel variants because their functional meta-scores are unavailable. We additionally excluded X-chromosomal variants because their functional annotations are not always available and the association testing may not be optimal [42]. Lastly, we excluded SNPs with minor allele frequency (MAF) less than 1%, as joint analysis of multiple rare variants simultaneously [43] is beyond the scope of our study.

CADD and Eigen functional meta-scores

We obtained the CADD meta-scores (v1.6), using the CADD tool [44], and the Eigen meta-scores (v1.0), using the ANNOVAR tool [45], for all the 7,895,174 common, bi-allelic autosomal SNPs.

In addition to the raw CADD meta-scores, the CADD tool also made available rank-based scores called phred scores,

$$-10\log_{10}(\text{ranks of the raw scores/the total number SNPs});$$

the phred-scaled scores are positive and have better interpretation compared to the raw scores. For example, a phred score of 10 or greater indicates that the SNP is predicted to be among the top 10% most deleterious variants of the human genome, while a phred score 20 or greater implies top 1% most deleterious.

For consistency between CADD and Eigen, we similarly obtained phred-scaled Eigen scores. Figure S3 in Supplementary Information shows the histograms of CADD and Eigen phred-scaled scores; each is expected to be $2.17\chi_2^2$ distributed, because a phred-scaled score is Unif(0,1) distributed, and as a result $-2\log(\text{Unif}(0,1))$ is χ_2^2 distributed; hereafter scores mean phred-scaled scores unless specified otherwise.

Because Eigen scores were calculated using an unsupervised learning approach, in contrast to CADD scores inferred using labeled data, we also compared these two scores genome-wide (Figure S4) and across four different consequence categories (Figure S5): missense, non-coding, synonymous, and protein truncating variants (PTV). Variants in the missense and PTV categories tend to have higher CADD than Eigen scores, while variants in the non-coding and synonymous categories tend to have higher Eigen than CADD scores. However, overall the two meta-scores are consistent and led to qualitatively comparable data-integration results, which we discuss next.

Simulation study design I, leveraging the observed genomic data

Here we used the real CADD and Eigen functional meta-scores, combined with simulated GWAS summary statistics, to evaluate type I error control of the data-integration methods examined.

Simulated GWAS summary statistics under the null of no association combined with real functional annotation scores

To simulate GWAS summary statistics that contain realistic LD patterns, we utilized the publicly available genotype data of the 1000 Genomes Project [46]. Independent of the observed genotype data, we simulated trait values, from $N(0,1)$, for 1,756 individuals from the 1000 Genomes Project who are unrelated to each other [47]. We examined 422,923 autosomal, bi-allelic and common (MAF > 5%) SNPs that a) passed the quality control conducted by [47], b) have CADD and Eigen meta-scores available, and c) have the 75 annotations used by FINDOR.

We then obtained GWAS summary statistics for the 422,923 SNPs by regressing the trait values of the 1,756 individuals on their additively coded genotypes. Because the trait values were randomly generated, independent of the genotypes and populations, the resulting GWAS z_i 's are $N(0,1)$ distributed and p_i 's Unif(0,1) distributed, as expected under the null of no association; the histograms of z_i 's and p_i 's from one randomly selected simulation run are shown in Figure S6.

Finally, we integrated the GWAS summary statistics with their corresponding CADD (or Eigen) meta-scores using the four methods, meta-analysis, Fisher's method, weighted p-value, and sFDR control as described above. Although [30] showed that FINDOR calibrates well when a GWAS consists of a mixture of null and associated SNPs, we also examined the performance of FINDOR in this setting when all GWAS SNPs are under the null hypothesis of no association. We applied the FINDOR tool using the same set of LDscores and the 75 annotations [32] that were used by [30] for their study; see Web Resources.

Method evaluation: family-wise error rate (FWER)

For each simulation replicate (i.e. a GWAS simulated under the null of no association combined with real functional scores), we obtained the number of false positives using the conservative Bonferroni corrected significance level, $\alpha = 0.05/422923 = 1.2 \times 10^{-7}$. We repeated the simulation, independently, 50,000 times, and calculated the family-wise error rate as the proportion

of the number of replicates with at least one significant finding. Assuming the true FWER is 0.05, we expect the FWER estimate obtained from the 50,000 independent simulation replicates to have a standard error of $\sqrt{0.05 \times 0.95/50000} \approx 0.001$. Thus, a method with an empirical FWER outside [0.047, 0.053] can be considered inaccurate.

Simulation study design II, leveraging the observed genetic data

Here we combined the observed UK Biobank GWAS summary statistics with *permuted* CADD (or Eigen) scores to evaluate robustness of a method to random annotation scores. Prior to the permutation, we examined the similarity of functional annotations between SNPs in linkage disequilibrium.

Permuted functional annotation scores combined with real GWAS summary statistics

Permutation does not preserve the potential correlation between functional scores of nearby SNPs. However, for the purpose of evaluating type I error control, it provides a valid set of annotation scores that are independent of the GWAS summary statistics. Nevertheless, we examined if SNPs in strong LD have similar annotation scores, as this has not been previously studied.

Using CADD as an example, let $CADD_i$ and $CADD_j$ be the annotation scores of SNPs i and j , respectively. We first defined a pair-wise similarity measure as $s_{i,j}^2 = 1 - |CADD_i - CADD_j| / (CADD_i + CADD_j)$. The measure $s_{i,j}^2$ is bounded between 0 and 1, where 1 means two scores are identical whereas a value close to 0 suggests a lack of similarity. We then contrasted $s_{i,j}^2$ with $r_{i,j}^2$, the traditional LD measure of genotype similarity between two SNPs.

After we permuted the functional scores of the 7,895,174 common, bi-allelic autosomal SNPs, for each of the 1,132 traits of the UK Biobank data, we integrated the GWAS summary statistics with the permuted annotation scores, using meta-analysis, Fisher's method, weighted p-value, and sFDR control. We were not able to evaluate FINDOR here, because FINDOR implements the LD Scoring (LDSC) tool [41] and the validity of using LDSC for permuted annotation is not clear.

Method evaluation: Recall, Precision and FDR

Before data integration, we first used $\alpha = 5 \times 10^{-8}$ [48] to identify genome-wide significance findings, $m_{1,t}$, for each trait t , $t = 1, \dots, 1,132$, based on the UK Biobank summary statistics alone. For the purpose of this simulation study, we treated $m_{1,t}$ as the total number of truly associated SNPs to be discovered after data-integration for trait t . In addition to counting the number of significant SNPs per GWAS, we also counted the number of independent, significant loci. We first defined independent loci using the LDclumping algorithm of PLINK (v1.07) [49], with a sliding window of 1 Mb and a LD r^2 threshold of 0.1 as per standard practice. We then considered a locus significant if it contained at least one genome-wide significant SNP.

After integrating the UK Biobank GWAS summary statistics with *permuted* functional scores for each trait t , we used the same $\alpha = 5 \times 10^{-8}$ to identify genome-wide significance findings (SNPs or loci as defined above), denoted as P_t . Among the P_t positives, we defined false positives, FP_t , as the new findings that were not part of $m_{1,t}$, because the information used here for data integration were random noise. Similarly, we defined $TP_t = P_t - FP_t$ as the number of true positives for trait t .

Finally, we defined and calculated recall, precision and false discovery rate by

$$Recall_t = \frac{TP_t}{m_{1,t}}, Precision_t = \frac{TP_t}{P_t} \text{ and } FDR_t = 1 - Precision_t = \frac{FP_t}{P_t}.$$

Recall is conceptually the same as *Power*, defined later for our simulation studies where the ground truth is known and $m_{1,t}$ SNPs were simulated as truly associated SNPs. We calculated $Recall_t$ only when $m_{1,t} > 0$. That is, for the 409 out of 1,132 traits with no GWAS significant findings *before* data-integration (i.e. $m_{1,t} = 0$) we did not calculate $Recall_t$. Regardless of whether $m_{1,t} = 0$ or not, for traits with no significant findings *after* data-integration (i.e. $P_t = 0$) we conservatively defined $Precision_t = 1$ and $FDR_t = 0$.

Simulation study design III, varying the informativeness of genomic information

To further investigate method performance in the presence of completely informative, partially informative, uninformative, or even misleading added information, we performed an additional set of simulation studies. Although linkage disequilibrium is an important aspect of GWAS, given the simulation study designs I and II and our findings in the results section, the simulation studies here focused on independent SNPs to delineate other potentially influencing factors.

Without loss of generality, we assumed the total number of SNPs $m = 10,000$, among which the first $m_1 = 100$ SNPs are truly associated. The corresponding summary statistics z_i 's were drawn, independently, from $N(\mu_1, 1)$ for the m_1 associated SNPs, and from $N(0, 1)$ for the remaining null SNPs. The top left plot in Figure S13 shows the Manhattan plot for one simulated GWAS replicate with $\mu_1 = 3$; we also varied μ_1 from 0.1 to 4 to represent different power scenarios of a GWAS.

We then assumed $z_{i,add}$'s as the additional information available, which were drawn, independently, from $N(\mu_{add}, 1)$ for the m_{add} SNPs and from $N(0, 1)$ for the remaining SNPs. Importantly, the locations of the m_{add} SNPs may differ from those of the m_1 associated SNPs. That is, the additional information available for a truly associated SNP may be random noise. On the other

hand, for a null SNP with no association (i.e. z_i drawn from $N(\mu_1, 1)$), its $z_{i,add}$ could be drawn from $N(\mu_{add}, 1)$, representing misleading information. We also varied μ_{add} , which may or may not be the same as μ_1 .

Using $m_1 = 100$ and $\mu_1 = 3$ as an example for the GWAS component, we considered the following eight scenarios for the additional information available for data integration (Figure S13), which fall into four categories.

- Category I is completely informative (homogeneity): (1) $m_{add} = 100$, $\mu_{add} = 3$, and locations of the m_{add} SNPs perfectly match those of m_1 GWAS truly associated SNPs.
- Category II is partially informative: (2) $m_{add} = 100$ and $\mu_{add} = 1.5$; (3) $m_{add} = 50$ and $\mu_{add} = 3$; (4) $m_{add} = 50$, $\mu_{add} = 1.5$, and all m_{add} SNPs coincide with (some of) the m_1 SNPs.
- Category III is (partially or completely) misleading: (5) $m_{add} = 100$ and $\mu_{add} = 3$; (6) $m_{add} = 100$ and $\mu_{add} = 1.5$, but in both scenarios only 50 out of the m_{add} SNPs coincide with 50 of the m_1 SNPs. And (7) $m_{add} = 100$ and $\mu_{add} = 3$, but none of the m_{add} SNPs coincide with the m_1 SNPs.
- Category IV is uninformative: (8) $m_{add} = 0$ and $\mu_{add} = 0$. That is, the additional information available is white noise.

For each of the eight scenarios, we simulated 1,000 data replicates, independently of each other. For each replicate, we applied the four data-integration methods that are suitable for this simulation study, namely meta-analysis, Fisher's method, weighted p-value, and sFDR control. Finally, we evaluated the methods using various performance measures, which we describe below.

Method evaluation: power and relative efficiency (RE)

We first used the Bonferroni corrected threshold to declare significance, p-value $< 0.05/10000 = 5 \times 10^{-6}$. Let P_{rep_i} be the number of positives for each of the 1000 simulation replicates after data-integration, we defined power as

$$Power_{rep_i} = \frac{P_{rep_i}}{m_1},$$

the proportion of the truly associated GWAS SNPs that were found after data-integration, which is similar to $Recall_i$ defined earlier in simulation study design I.

As the Bonferroni approach can be conservative, we explored two alternative decision rules: fixed-region and fixed-FDR rejections. The fixed-region rule rejected the top k SNPs (e.g. $k = 100$), while the fixed-FDR rule rejected SNPs by controlling FDR at $\gamma\%$ level (e.g. $\gamma\% = 5\%$). For each rejection rule, we then calculated power of a method as described above.

Finally, we considered ranked-based relative efficiency as a performance measure. To this end, we first ranked all the truly associated m_1 SNPs based on the GWAS summary statistics alone, denoted as $R_{baseline}$. After data-integration, we use R_{method} to denote the ranks of the m_1 SNPs based on their Z^{meta} , Z^{Fisher} , $p_{weighted}$, and p_{sFDR} values. Finally, after averaging $R_{baseline}$ and R_{method} across the m_1 SNPs and across the 1,000 simulated replicates, we defined relative efficiency as

$$RE_{method} = 1 - \frac{\bar{R}_{method}}{\bar{R}_{baseline}}.$$

A positive RE_{method} value means the truly associated m_1 SNPs are ranked higher, on average, after data-integration using the method; a RE_{method} value of zero means that the data-integration method did not improve performance; and a negative RE_{method} value suggests that the data-integration effort was counter-productive.

Integrating UK Biobank GWAS summary statistics with functional annotations

We studied all 1,132 UK Biobank traits for which the confidence for their heritability inference was considered "medium" or "high" by Nealelab. For each trait, we analyzed the 7,895,174 autosomal SNPs that are bi-allelic and common (MAF $> 5\%$), integrating their GWAS summary statistics with the CADD (or Eigen) meta-scores using the weighted p-value and sFDR methods. We excluded meta-analysis and Fisher's method from the analysis here, because severe robustness issues (to partially informative, uninformative, or misleading $z_{i,add}$) were found in simulation studies; see results for details. For comparison, we also applied FINDOR using the set of 75 publicly available annotations recommended by the authors; see Web Resources.

To summarize the application results, we first counted the numbers of independent, significant loci (at the 5×10^{-8} level) identified before and after data-integration for each of the 1,132 traits, stratified by the four trait categories (nonsig, nominal, z4, and z7). We then calculated $Recall$, the proportion of the initial GWAS findings that were retained after data-integration, as previously defined for the simulation studies. Finally, we used *New Discoveries* to represent the number of new genome-wide significant findings at the 5×10^{-8} level.

Availability of data and materials

All codes used for data analyses and simulation studies are open-resource at <https://github.com/jianhuig/Integrate-gwas/#readme>.

Data used in this work are GWAS summary statistics and functional annotation scores, which are all publicly available:

UK Biobank GWAS summary statistics from Nealelab, <http://www.nealelab.is/uk-biobank>

UK Biobank SNP-heritability estimates from Nealelab, https://nealelab.github.io/UKBB_ldsc/

The 1,000 Genome Projects, <http://tcag.ca/tools/1000genomes.html>

CADD(v1.6) , <https://cadd.gs.washington.edu>

Eigen (v1.0) through ANNOVAR software, <http://annovar.openbioinformatics.org/en/latest/user-guide/filter/#eigen-score-annotations>

FINDOR, <https://github.com/gkichaev/FINDOR>

References

1. Visscher, P.M., Wray, N.R., Zhang, Q., Sklar, P., McCarthy, M.I., Brown, M.A., Yang, J.: 10 Years of GWAS Discovery: Biology, Function, and Translation. *American Journal of Human Genetics* **101**(1), 5–22 (2017). doi:10.1016/j.ajhg.2017.06.005
2. Spencer, C.C.A., Su, Z., Donnelly, P., Marchini, J.: Designing Genome-Wide Association Studies: Sample Size, Power, Imputation, and the Choice of Genotyping Chip. *PLoS Genetics* **5**(5) (2009). doi:10.1371/journal.pgen.1000477. Accessed 2020-12-07
3. Ott, J., Wang, J., Leal, S.M.: Genetic linkage analysis in the age of whole-genome sequencing. *Nature Reviews Genetics* **16**(5), 275–284 (2015). doi:10.1038/nrg3908. Number: 5 Publisher: Nature Publishing Group. Accessed 2020-12-09
4. Buniello, A., MacArthur, J.A.L., Cerezo, M., Harris, L.W., Hayhurst, J., Malangone, C., McMahon, A., Morales, J., Mountjoy, E., Sollis, E., Suveges, D., Vrousitou, O., Whetzel, P.L., Amode, R., Guillen, J.A., Riat, H.S., Trevanion, S.J., Hall, P., Junkins, H., Flicek, P., Burdett, T., Hindorf, L.A., Cunningham, F., Parkinson, H.: The NHGRI-EBI GWAS Catalog of published genome-wide association studies, targeted arrays and summary statistics 2019. *Nucleic Acids Research* **47**(D1), 1005–1012 (2019). doi:10.1093/nar/gky1120
5. Cochran, W.G.: The Combination of Estimates from Different Experiments. *Biometrics* **10**(1), 101–129 (1954). doi:10.2307/3001666. Publisher: [Wiley, International Biometric Society]. Accessed 2020-06-17
6. Fisher, R.A.: *Statistical Methods for Research Workers*. Edinburgh, Oliver and Boyd, ??? (1938). <http://archive.org/details/statisticalmethoe7fish> Accessed 2021-01-23
7. Lin, D.Y., Zeng, D.: Meta-Analysis of Genome-Wide Association Studies: No Efficiency Gain in Using Individual Participant Data. *Genetic epidemiology* **34**(1) (2010). doi:10.1002/gepi.20435. Accessed 2020-04-04
8. Sung, Y.J., Schwander, K., Arnett, D.K., Kardia, S.L.R., Rankinen, T., Bouchard, C., Boerwinkle, E., Hunt, S.C., Rao, D.C.: An empirical comparison of meta-analysis and mega-analysis of individual participant data for identifying gene-environment interactions. *Genetic epidemiology* **38**(4), 369–378 (2014). doi:10.1002/gepi.21800. Accessed 2021-06-20
9. Kundaje, A., Meuleman, W., Ernst, J., Bilenky, M., Yen, A., Heravi-Moussavi, A., Kheradpour, P., Zhang, Z., Wang, J., Ziller, M.J., Amin, V., Whitaker, J.W., Schultz, M.D., Ward, L.D., Sarkar, A., Quon, G., Sandstrom, R.S., Eaton, M.L., Wu, Y.-C., Pfening, A.R., Wang, X., Claussnitzer, M., Yaping Liu, Coarfa, C., Alan Harris, R., Shores, N., Epstein, C.B., Gjonneska, E., Leung, D., Xie, W., David Hawkins, R., Lister, R., Hong, C., Gascard, P., Mungall, A.J., Moore, R., Chuah, E., Tam, A., Canfield, T.K., Scott Hansen, R., Kaul, R., Sabo, P.J., Bansal, M.S., Carles, A., Dixon, J.R., Farh, K.-H., Feizi, S., Karlic, R., Kim, A.-R., Kulkarni, A., Li, D., Lowdon, R., Elliott, G., Mercer, T.R., Neph, S.J., Onuchic, V., Polak, P., Rajagopal, N., Ray, P., Sallari, R.C., Siebenthal, K.T., Sinnott-Armstrong, N.A., Stevens, M., Thurman, R.E., Wu, J., Zhang, B., Zhou, X., Beaudet, A.E., Boyer, L.A., Jager, P.L.D., Farnham, P.J., Fisher, S.J., Haussler, D., Jones, S.J.M., Li, W., Marra, M.A., McManus, M.T., Sunyaev, S., Thomson, J.A., Tlsty, T.D., Tsai, L.-H., Wang, W., Waterland, R.A., Zhang, M.Q., Chadwick, L.H., Bernstein, B.E., Costello, J.F., Ecker, J.R., Hirst, M., Meissner, A., Milosavljevic, A., Ren, B., Stamatoyannopoulos, J.A., Wang, T., Kellis, M.: Integrative analysis of 111 reference human epigenomes. *Nature* **518**(7539), 317–330 (2015). doi:10.1038/nature14248. Number: 7539 Publisher: Nature Publishing Group. Accessed 2020-12-10
10. Dunham, I., *et al.*: An integrated encyclopedia of DNA elements in the human genome. *Nature* **489**(7414), 57–74 (2012). doi:10.1038/nature11247. Number: 7414 Publisher: Nature Publishing Group. Accessed 2020-12-10
11. Davydov, E.V., Goode, D.L., Sirota, M., Cooper, G.M., Sidow, A., Batzoglou, S.: Identifying a High Fraction of the Human Genome to be under Selective Constraint Using GERP++. *PLOS Computational Biology* **6**(12), 1001025 (2010). doi:10.1371/journal.pcbi.1001025. Publisher: Public Library of Science. Accessed 2020-12-10
12. Adzhubei, I.A., Schmidt, S., Peshkin, L., Ramensky, V.E., Gerasimova, A., Bork, P., Kondrashov, A.S., Sunyaev, S.R.: A method and server for predicting damaging missense mutations. *Nature Methods* **7**(4), 248–249 (2010). doi:10.1038/nmeth0410-248. Accessed 2020-04-04

13. Shihab, H.A., Rogers, M.F., Gough, J., Mort, M., Cooper, D.N., Day, I.N.M., Gaunt, T.R., Campbell, C.: An integrative approach to predicting the functional effects of non-coding and coding sequence variation. *Bioinformatics* **31**(10), 1536–1543 (2015). doi:10.1093/bioinformatics/btv009. Accessed 2020-12-10
14. Lu, Q., Hu, Y., Sun, J., Cheng, Y., Cheung, K.-H., Zhao, H.: A Statistical Framework to Predict Functional Non-Coding Regions in the Human Genome Through Integrated Analysis of Annotation Data. *Scientific Reports* **5** (2015). doi:10.1038/srep10576. Accessed 2020-12-10
15. Ritchie, G.R.S., Dunham, I., Zeggini, E., Flicek, P.: Functional annotation of noncoding sequence variants. *Nature Methods* **11**(3), 294–296 (2014). doi:10.1038/nmeth.2832. Number: 3 Publisher: Nature Publishing Group. Accessed 2020-12-10
16. Kircher, M., Witten, D.M., Jain, P., O’Roak, B.J., Cooper, G.M., Shendure, J.: A general framework for estimating the relative pathogenicity of human genetic variants. *Nature genetics* **46**(3), 310–315 (2014). doi:10.1038/ng.2892. Accessed 2020-07-09
17. Ionita-laza, I., McCallum, K., Xu, B., Buxbaum, J.: A SPECTRAL APPROACH INTEGRATING FUNCTIONAL GENOMIC ANNOTATIONS FOR CODING AND NONCODING VARIANTS. *Nature genetics* **48**(2), 214–220 (2016). doi:10.1038/ng.3477. Accessed 2020-03-12
18. Li, X., Li, Z., Zhou, H., Gaynor, S.M., Liu, Y., Chen, H., Sun, R., Dey, R., Arnett, D.K., Aslibekyan, S., Ballantyne, C.M., Bielak, L.F., Blangero, J., Boerwinkle, E., Bowden, D.W., Broome, J.G., Conomos, M.P., Correa, A., Cupples, L.A., Curran, J.E., Freedman, B.I., Guo, X., Hindy, G., Irvin, M.R., Kardia, S.L.R., Kathiresan, S., Khan, A.T., Kooperberg, C.L., Laurie, C.C., Liu, X.S., Mahaney, M.C., Manichaikul, A.W., Martin, L.W., Mathias, R.A., McGarvey, S.T., Mitchell, B.D., Montasser, M.E., Moore, J.E., Morrison, A.C., O’Connell, J.R., Palmer, N.D., Pampana, A., Peralta, J.M., Peyser, P.A., Psaty, B.M., Redline, S., Rice, K.M., Rich, S.S., Smith, J.A., Tiwari, H.K., Tsai, M.Y., Vasani, R.S., Wang, F.F., Weeks, D.E., Weng, Z., Wilson, J.G., Yanek, L.R., Neale, B.M., Shamil R. Sunyaev, Abecasis, G.R., Rotter, J.I., Willer, C.J., Peloso, G.M., Natarajan, P., Lin, X.: Dynamic incorporation of multiple in silico functional annotations empowers rare variant association analysis of large whole-genome sequencing studies at scale. *Nature Genetics* **52**(9), 969–983 (2020). doi:10.1038/s41588-020-0676-4. Number: 9 Publisher: Nature Publishing Group. Accessed 2021-01-23
19. Liang, J., Cade, B.E., He, K.Y., Wang, H., Lee, J., Sofer, T., Williams, S., Li, R., Chen, H., Gottlieb, D.J., Evans, D.S., Guo, X., Gharib, S.A., Hale, L., Hillman, D.R., Lutsey, P.L., Mukherjee, S., Ochs-Balcom, H.M., Palmer, L.J., Rhodes, J., Purcell, S., Patel, S.R., Saxena, R., Stone, K.L., Tang, W., Tranah, G.J., Boerwinkle, E., Lin, X., Liu, Y., Psaty, B.M., Vasani, R.S., Cho, M.H., Manichaikul, A., Silverman, E.K., Barr, R.G., Rich, S.S., Rotter, J.I., Wilson, J.G., Redline, S., Zhu, X.: Sequencing Analysis at 8p23 Identifies Multiple Rare Variants in *DLC1* Associated with Sleep-Related Oxyhemoglobin Saturation Level. *The American Journal of Human Genetics* **105**(5), 1057–1068 (2019). doi:10.1016/j.ajhg.2019.10.002. Accessed 2021-01-23
20. Pereira, S.V.-N., Ribeiro, J.D., Ribeiro, A.F., Bertuzzo, C.S., Marson, F.A.L.: Novel, rare and common pathogenic variants in the *CFTR* gene screened by high-throughput sequencing technology and predicted by in silico tools. *Scientific Reports* **9**(1), 6234 (2019). doi:10.1038/s41598-019-42404-6. Number: 1 Publisher: Nature Publishing Group. Accessed 2021-01-23
21. Genovese, C.R., Roeder, K., Wasserman, L.: False Discovery Control with p-Value Weighting. *Biometrika* **93**(3), 509–524 (2006). Publisher: [Oxford University Press, Biometrika Trust]. Accessed 2021-01-13
22. Sun, L., Craiu, R.V., Paterson, A.D., Bull, S.B.: Stratified false discovery control for large-scale hypothesis testing with application to genome-wide association studies. *Genetic Epidemiology* **30**(6), 519–530 (2006). doi:10.1002/gepi.20164. Accessed 2021-01-26
23. Benjamini, Y., Hochberg, Y.: 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 (1995). Publisher: [Royal Statistical Society, Wiley]. Accessed 2020-04-04
24. Roeder, K., Bacanu, S.-A., Wasserman, L., Devlin, B.: Using linkage genome scans to improve power of association in genome scans. *American Journal of Human Genetics* **78**(2), 243–252 (2006). doi:10.1086/500026
25. Yoo, Y.J., Bull, S.B., Paterson, A.D., Waggott, D., Sun, L.: Were Genome-Wide Linkage Studies a Waste of Time? Exploiting Candidate Regions Within Genome-Wide Association Studies. *Genetic epidemiology* **34**(2), 107–118 (2010). doi:10.1002/gepi.20438. Accessed 2020-04-04
26. Li, L., Kabesch, M., Bouzigon, E., Demenais, F., Farrall, M., Moffatt, M.F., Lin, X., Liang, L.: Using eQTL weights to improve power for genome-wide association studies: a genetic study of childhood asthma. *Frontiers in Genetics* **4** (2013). doi:10.3389/fgene.2013.00103. Publisher: Frontiers. Accessed 2020-10-09
27. Keel, B.N., Snelling, W.M., Lindholm-Perry, A.K., Oliver, W.T., Kuehn, L.A., Rohrer, G.A.: Using SNP Weights Derived From Gene Expression Modules to Improve GWAS Power for Feed Efficiency in Pigs. *Frontiers in Genetics* **10** (2020). doi:10.3389/fgene.2019.01339. Publisher: Frontiers. Accessed 2020-12-10
28. Andreassen, O.A., Thompson, W.K., Schork, A.J., Ripke, S., Mattingsdal, M., Kelsoe, J.R., Kendler, K.S., O’Donovan,

- M.C., Rujescu, D., Werge, T., Sklar, P., Psychiatric Genomics Consortium (PGC), Bipolar Disorder and Schizophrenia Working Groups, Roddey, J.C., Chen, C.-H., McEvoy, L., Desikan, R.S., Djurovic, S., Dale, A.M.: Improved detection of common variants associated with schizophrenia and bipolar disorder using pleiotropy-informed conditional false discovery rate. *PLoS genetics* **9**(4), 1003455 (2013). doi:10.1371/journal.pgen.1003455
29. Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., Liu, B., Matthews, P., Ong, G., Pell, J., Silman, A., Young, A., Sprosen, T., Peakman, T., Collins, R.: UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age. *PLoS Medicine* **12**(3) (2015). doi:10.1371/journal.pmed.1001779. Accessed 2021-06-01
 30. Kichaev, G., Bhatia, G., Loh, P.-R., Gazal, S., Burch, K., Freund, M.K., Schoech, A., Pasaniuc, B., Price, A.L.: Leveraging Polygenic Functional Enrichment to Improve GWAS Power. *The American Journal of Human Genetics* **104**(1), 65–75 (2019). doi:10.1016/j.ajhg.2018.11.008. Publisher: Elsevier. Accessed 2020-11-17
 31. Finucane, H.K., Bulik-Sullivan, B., Gusev, A., Trynka, G., Reshef, Y., Loh, P.-R., Anttila, V., Xu, H., Zang, C., Farh, K., Ripke, S., Day, F.R., Purcell, S., Stahl, E., Lindstrom, S., Perry, J.R.B., Okada, Y., Raychaudhuri, S., Daly, M.J., Patterson, N., Neale, B.M., Price, A.L.: Partitioning heritability by functional annotation using genome-wide association summary statistics. *Nature Genetics* **47**(11), 1228–1235 (2015). doi:10.1038/ng.3404. Number: 11 Publisher: Nature Publishing Group. Accessed 2021-01-11
 32. Gazal, S., Finucane, H.K., Furlotte, N.A., Loh, P.-R., Palamara, P.F., Liu, X., Schoech, A., Bulik-Sullivan, B., Neale, B.M., Gusev, A., Price, A.L.: Linkage disequilibrium–dependent architecture of human complex traits shows action of negative selection. *Nature Genetics* **49**(10), 1421–1427 (2017). doi:10.1038/ng.3954. Number: 10 Publisher: Nature Publishing Group. Accessed 2021-06-03
 33. Visscher, P.M., Wray, N.R., Zhang, Q., Sklar, P., McCarthy, M.I., Brown, M.A., Yang, J.: 10 years of gwas discovery: biology, function, and translation. *The American Journal of Human Genetics* **101**(1), 5–22 (2017)
 34. Li, Y., Haug, S., Schlosser, P., Teumer, A., Tin, A., Pattaro, C., Köttgen, A., Wuttke, M.: Integration of gwas summary statistics and gene expression reveals target cell types underlying kidney function traits. *Journal of the American Society of Nephrology* **31**(10), 2326–2340 (2020). doi:10.1681/ASN.2020010051. <https://jasn.asnjournals.org/content/31/10/2326.full.pdf>
 35. Uffelmann, E., Huang, Q.Q., Munung, N.S., de Vries, J., Okada, Y., Martin, A.R., Martin, H.C., Lappalainen, T., Posthuma, D.: Genome-wide association studies. *Nature Reviews Methods Primers* **1**(1), 1–21 (2021)
 36. Hedges, L.V., Vevea, J.L.: Fixed- and random-effects models in meta-analysis. *Psychological Methods* **3**(4), 486–504 (1998). doi:10.1037/1082-989X.3.4.486. Num Pages: 486-504 Place: Washington, US Publisher: American Psychological Association (US). Accessed 2021-05-17
 37. Thompson, S.G.: Why sources of heterogeneity in meta-analysis should be investigated. *BMJ : British Medical Journal* **309**(6965), 1351–1355 (1994). Accessed 2020-12-09
 38. Begum, F., Ghosh, D., Tseng, G.C., Feingold, E.: Comprehensive literature review and statistical considerations for GWAS meta-analysis. *Nucleic Acids Research* **40**(9), 3777–3784 (2012). doi:10.1093/nar/gkr1255. Accessed 2020-12-09
 39. Storey, J.D.: A direct approach to false discovery rates. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **64**(3), 479–498 (2002). doi:10.1111/1467-9868.00346. Accessed 2020-06-17
 40. Storey, J.D., Tibshirani, R.: Statistical significance for genomewide studies. *Proceedings of the National Academy of Sciences of the United States of America* **100**(16), 9440–9445 (2003). doi:10.1073/pnas.1530509100
 41. Bulik-Sullivan, B.K., Loh, P.-R., Finucane, H.K., Ripke, S., Yang, J., Patterson, N., Daly, M.J., Price, A.L., Neale, B.M.: LD Score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nature Genetics* **47**(3), 291–295 (2015). doi:10.1038/ng.3211. Number: 3 Publisher: Nature Publishing Group. Accessed 2021-06-02
 42. Chen, B., Craiu, R.V., Strug, L.J., Sun, L.: The x factor: A robust and powerful approach to x-chromosome-inclusive whole-genome association studies. *Genetic epidemiology* **45**(7), 694–709 (2021)
 43. Derkach, A., Lawless, J.F., Sun, L.: Pooled Association Tests for Rare Genetic Variants: A Review and Some New Results. *Statistical Science* **29**(2), 302–321 (2014). doi:10.1214/13-STS456. Publisher: Institute of Mathematical Statistics. Accessed 2021-06-03
 44. Rentzsch, P., Witten, D., Cooper, G.M., Shendure, J., Kircher, M.: CADD: predicting the deleteriousness of variants throughout the human genome. *Nucleic Acids Research* **47**(D1), 886–894 (2019). doi:10.1093/nar/gky1016. Publisher: Oxford Academic. Accessed 2020-10-12
 45. Wang, K., Li, M., Hakonarson, H.: ANNOVAR: functional annotation of genetic variants from high-throughput sequencing data. *Nucleic Acids Research* **38**(16), 164–164 (2010). doi:10.1093/nar/gkq603. Accessed 2020-06-22
 46. Auton, A., *et al.*: A global reference for human genetic variation. *Nature* **526**(7571), 68–74 (2015). doi:10.1038/nature15393. Number: 7571 Publisher: Nature Publishing Group. Accessed 2021-06-03
 47. Roslin, N.M., Weili, L., Paterson, A.D., Strug, L.J.: Quality control analysis of the 1000 Genomes Project Omni2.5

genotypes. bioRxiv, 078600 (2016). doi:10.1101/078600. Publisher: Cold Spring Harbor Laboratory Section: New Results. Accessed 2021-06-03

48. Dudbridge, F., Gusnanto, A.: Estimation of significance thresholds for genomewide association scans. *Genetic Epidemiology* **32**(3), 227–234 (2008). doi:10.1002/gepi.20297. Accessed 2021-06-03
49. Purcell, S., Neale, B., Todd-Brown, K., Thomas, L., Ferreira, M.A.R., Bender, D., Maller, J., Sklar, P., de Bakker, P.I.W., Daly, M.J., Sham, P.C.: PLINK: a tool set for whole-genome association and population-based linkage analyses. *American Journal of Human Genetics* **81**(3), 559–575 (2007). doi:10.1086/519795

Acknowledgements

We sincerely thank Dr. Angelo Canty for helpful discussions.

Author contributions statement

The authors confirm contribution to the paper as follows: draft manuscript preparation and data analysis: G.J. All authors contributed to the writing and editing of the manuscript.

Additional information

Funding

This research was funded by the Natural Sciences and Engineering Research Council of Canada (NSERC, RGPIN-04934 and RGPAS-522594 to LS), the Canadian Institutes of Health Research (CIHR, MOP-310732 to LS), and the University of Toronto McLaughlin Centre Accelerator Grants in Genomic Medicine (MC-2019-15 to LS and ADP).

Competing interests

The authors declare no competing interests.

Additional Files

Addition file 1 – Supplementary tables

Supplementary Tables S1 and S2.

Addition file 2 – Supplementary figures

Supplementary Figures S1-S32.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [suppfigure.pdf](#)
- [supptable.pdf](#)