Published in final edited form as:

Hum Genet. 2014 May ; 133(5): 639–650. doi:10.1007/s00439-013-1401-5.

Improving genetic risk prediction by leveraging pleiotropy

Cong Li,

Program in Computational Biology and Bioinformatics, Yale University, New Haven, Connecticut 06520, USA

Can Yang,

Department of Biostatistics, Yale School of Public Health, Yale University, New Haven, Connecticut 06520, USA, Department of Psychiatry, Yale University School of Medicine, New Haven, CT, USA

Joel Gelernter, and

Department of Psychiatry, Yale University School of Medicine, New Haven, CT, USA, VA CT Healthcare Center, Departments of Genetics and Neurobiology, Yale Univ. School of Medicine, West Haven, Connecticut 06516, USA

Hongyu Zhao

Department of Biostatistics, Yale School of Public Health, Yale University, New Haven, Connecticut 06520, USA, Program in Computational Biology and Bioinformatics, Yale University, New Haven, Connecticut 06520, USA

Cong Li: cong.li@yale.edu; Can Yang: can.yang@yale.edu; Joel Gelernter: joel.gelernter@yale.edu; Hongyu Zhao: hongyu.zhao@yale.edu

Abstract

An important task of human genetics studies is to predict accurately disease risks in individuals based on genetic markers, which allows for identifying individuals at high disease risks, and facilitating their disease treatment and prevention. Although hundreds of genome-wide association studies (GWAS) have been conducted on many complex human traits in recent years, there has been only limited success in translating these GWAS data into clinically useful risk prediction models. The predictive capability of GWAS data is largely bottlenecked by the available training sample size due to the presence of numerous variants carrying only small to modest effects. Recent studies have shown that different human traits may share common genetic bases. Therefore, an attractive strategy to increase the training sample size and hence improve the prediction accuracy is to integrate data from genetically correlated phenotypes. Yet the utility of genetic correlation in risk prediction has not been explored in the literature. In this paper, we analyzed GWAS data for bipolar and related disorders (BARD) and schizophrenia (SZ) with a bivariate ridge regression method, and found that jointly predicting the two phenotypes could substantially increase prediction accuracy as measured by the AUC (area under the *receiver operating characteristic curve*). We also found similar prediction accuracy improvements when we jointly analyzed GWAS data for Crohn's disease (CD) and ulcerative colitis (UC). The

Correspondence to: Hongyu Zhao, hongyu.zhao@yale.edu.

The authors declare that they have no conflict of interest.

empirical observations were substantiated through our comprehensive simulation studies, suggesting that a gain in prediction accuracy can be obtained by combining phenotypes with relatively high genetic correlations. Through both real data and simulation studies, we demonstrated pleiotropy can be leveraged as a valuable asset that opens up a new opportunity to improve genetic risk prediction in the future.

Keywords

Genetic risk prediction; Genome-wide association study; Pleiotropy; Ridge regression

1 Introduction

Predicting disease risks in individuals based on genetic markers, which is usually referred to as “genetic risk prediction” is an important task in human genetics studies [6]. Although hundreds of genome-wide association studies (GWAS) have been conducted and many thousands of genomic regions have been implicated for various human traits in recent years, these findings have not been translated into clinically useful risk prediction models based on genetic markers, which limits their potential impact on personalized disease prevention and treatment. Most genetic risk prediction models, including those used by direct-to-customer genetic testing companies, are constructed based on a few significant (usually validated) single nucleotide polymorphisms (SNPs). However, such SNPs typically account for only a small fraction of the total heritability and thus cannot provide satisfactory prediction accuracy [24]. A large fraction of the genetic variance is not accounted by such genetic risk predictions because this fraction is “missing” in the SNPwise GWAS results [22]. Several hypotheses have been put forward to explain this “missing heritability” [8], including epistasis [5,18], undetected CNVs [11], and rare variants [7], among others. However, it has also been suggested that much of the missing heritability is in fact hidden among the numerous common genetic variants carrying only small to modest effects [12,34,25]. Through a polygenic risk-score analysis on GWAS data of schizophrenia (SZ) and bipolar disorder (BP), Purcell et al [26] inferred that the genetic risk for SZ and BP may involve thousands of common SNPs individually of very small effect size. Using a linear mixed model, Yang et al [36] estimated that cumulatively, the whole-genome common SNP content could explain about 45% of human height variance. Lee et al [20] extended this method to complex human diseases and also found that a substantial fraction of the phenotypic variances could be explained by genotyped common SNPs. In this manuscript, we refer to the proportion of phenotypic variance explained by all genotyped common SNPs as the “chip heritability” [37], denoted as h^2 .

These results suggest that, instead of building genetic prediction models with only a few significant variants, a more sensible approach would be to impose less stringent criteria on variant selection or even to build prediction models using a whole genome variant set [4]. In this way, more variants with weak effects can be used, though at the expense of also including many variants that actually do not affect the phenotype one way or the other, or even act in the opposite effect direction. Following Makowsky et al [23], we refer to this class of methods as Whole Genome Prediction (WGP) methods. However, there is still a

significant gap between the prediction accuracy achieved and the genetic variance accounted by all common SNPs, even for the WGP methods [23]. These authors also investigated several parameters that affect prediction accuracy and showed that a substantial gain may be achieved by increasing the training sample size, and using training samples that are more related to the testing samples. Wray et al [35] also pointed out limited sample size as a bottleneck to achieve accurate genetic risk predictions.

These studies suggest that, in the presence of numerous weak-effect variants for most complex traits one of the most effective ways to develop a more accurate risk prediction model is to increase the size of the training sample. But subject recruitment may be difficult and expensive. Alternatively, if another phenotype shares some genetic variants that affect the primary phenotype of interest, it might be possible to incorporate the available data for this correlated phenotype in risk model development. In this way, we can increase the effective training sample size with the caveat that the additional samples are not *directly* related to the primary phenotype of interest. Appropriate statistical methods are needed to analyze jointly these distinct yet related data sets.

In fact, there is accumulating evidence suggesting that different complex human traits are genetically correlated, i.e. multiple traits share common genetic bases, which is also formally known as “pleiotropy”. In a systematic analysis of the open-access NHGRI catalog, 17% of the trait-associated genes and 5% of the trait-associated SNPs showed pleiotropic effects [28]. Vattikuti et al [33] used a bivariate linear mixed model to analyze the Atherosclerosis Risk in Communities GWAS and found significant genetic correlations between several metabolic syndrome traits, including body-mass index, waist-to-hip ratio, systolic blood pressure, fasting glucose, fasting insulin, fasting triglycerides, and fasting high-density lipoprotein. Lee et al [21] extended this bivariate linear mixed model so that it could deal with binary traits, e.g. presence or absence of a disease. Andreassen et al [1] applied a “pleiotropic enrichment” method on GWAS data of schizophrenia and cardiovascular-disease and showed that the power to detect schizophrenia-associated common variants can be improved by exploiting the pleiotropy between these two phenotypes. More recently, a study on genome-wide SNP data for five psychiatric disorders in 33,332 cases and 27,888 controls identified four significant loci ($P < 5 \times 10^{-8}$) affecting multiple disorders, including two genes encoding two L-type voltage-gated calcium channel subunits, *CACNA1C* and *CACNB2* [29]. Results from the large scale Collaborative Oncological Gene-environment Study also highlighted the existence of “carcinogenic pleiotropy”, i.e. the overlap between loci that confer genetic susceptibility to multiple types of tumor [27].

These findings are exciting because they imply that genetic correlation is prevalent among complex human diseases and hence leveraging the genetic correlations between phenotypes might be a promising strategy to improve genetic risk prediction. Although genetic correlations have been extensively studied for association analyses [19,17], little attention has been paid to their utility in genetic risk prediction. In this paper, we propose to use a bivariate ridge regression method to leverage the genetic correlation between two diseases in genetic risk prediction. We analyzed real GWAS data sets for two pairs of related common diseases. We performed a comprehensive simulation study on the utility of genetic

correlation by investigating the gain of prediction accuracy as a function of the strength of genetic correlation between two traits. We also examined the effects of several other parameters such as the “chip heritability” h^2 , the training sample size and the number/proportion of causal SNPs.

2 Methods

In this study, four dbGap data sets were analyzed. We did not obtain the consent from the participants, because we downloaded anonymous data from the database. For two of the dbGap datasets, we obtained the approval of the institutional review board. For the other two data sets, institutional review board approval is not required to access the data sets.

Data Description

We downloaded a GWAS data set of bipolar and related disorders (BARD) and a GWAS data set of schizophrenia (SZ) from the dbGaP database (<https://dbgap.ncbi.nlm.nih.gov>). All the subjects were genotyped on the Affymetrix Genome-Wide Human SNP Array 6.0 platform. See http://www.ncbi.nlm.nih.gov/projects/gap/cgi-bin/study.cgi?study_id=phs000017.v3.p1 and http://www.ncbi.nlm.nih.gov/projects/gap/cgi-bin/study.cgi?study_id=phs000021.v3.p2 for more details. We merged the two data sets after removing redundant subjects and collected the genotype data for subjects with only European ancestry. The merged data set consists of 653 affected subjects for BARD, 1170 affected subjects for SZ and 1403 controls. SNPs with minor allele frequency (MAF) smaller than 0.05 and missing rate greater than 0.01 were excluded. The remaining missing genotypes were randomly drawn from binomial distributions based on the allele frequencies at the given loci. SNPs that failed the Hardy-Weinberg Equilibrium test ($P < 0.0001$) in either BARD, SZ or control group were also excluded. We also performed linkage-disequilibrium pruning so that every pair of SNPs within a 50-SNP window had an R-squared value no greater than 0.8. After these procedures, 298,604 SNPs remained.

For the second pair of diseases, we downloaded a GWAS data set of Crohn's disease (CD) and a GWAS data set of ulcerative colitis (UC). The subjects in the CD data set were genotyped on the ILLUMINA HumanHap300v1.1 platform. See http://www.ncbi.nlm.nih.gov/projects/gap/cgi-bin/study.cgi?study_id=phs000130.v1.p1 for more details. UC subjects were genotyped on either the ILLUMINA HumanHap300v2 or the HumanHap550v3 platform. See http://www.ncbi.nlm.nih.gov/projects/gap/cgi-bin/study.cgi?study_id=phs000345.v1.p1 for more details. Since the SNPs on the HumanHap300v1.1 platform are a subset of the SNPs on the other two platforms, we retained only the HumanHap300v1.1 SNPs in all the data sets for further analysis. In the CD data set, only subjects with non-Jewish, European ancestry were included in analysis, which include of 515 unaffected subjects and 513 affected subjects. The UC data set consists of 902 non-Hispanic, European ancestry, affected individuals. To have a genetically matched control data set for UC, we downloaded all the Caucasian subjects that have genotypes from the HumanHap300v1.1 SNPs from the ILLUMINA iControl database (<http://www.illumina.com/science/icontroldb.ilmn>). We performed principal component analysis on the iControl samples and the UC samples. For each UC sample, we selected an iControl sample with the smallest Euclidean distance between their first 20 principal components.

SNPs that were overlapped between all these data sets were used for analysis after excluding SNPs with minor allele frequency smaller than 0.05 and missing rate greater than 0.01. The remaining missing genotypes were randomly drawn from binomial distributions based on the allele frequencies at the given loci. SNPs that failed the Hardy-Weinberg Equilibrium test ($P < 0.0001$) in either data set were also excluded. We also performed linkage-disequilibrium pruning so that every pair of SNPs within a 50-SNP window had an R-squared value no greater than 0.8. After these procedures, 241,649 SNPs remained.

Bivariate Ridge Regression

To take advantage of the pleiotropy, we used a bivariate ridge regression method to predict jointly the status for two diseases. In fact, the bivariate (or multivariate) ridge regression was motivated as a generalization of univariate ridge regression to account for the “across regression” correlations when multiple linear regressions are simultaneously considered [2,14]. Similar to the relationship between the univariate ridge regression and the univariate linear mixed model, the bivariate ridge regression is also closely related to the bivariate linear mixed model, which has been implemented to estimate the genetic correlation between phenotypes recently [21]. Here we highlight this connection while we go through the formulation of bivariate ridge regression.

Consider the following bivariate linear mixed model [31]:

$$\begin{aligned} \mathbf{y}^{(1)} &= \mathbf{X}^{(1)}\beta^{(1)} + \mathbf{g}^{(1)} + \mathbf{e}^{(1)}, \\ \mathbf{y}^{(2)} &= \mathbf{X}^{(2)}\beta^{(2)} + \mathbf{g}^{(2)} + \mathbf{e}^{(2)} \end{aligned} \quad (1)$$

where $\mathbf{y}^{(1)}$ and $\mathbf{y}^{(2)}$ are the two vectors of two phenotypes measured on n_1 and n_2 individuals, respectively. In this study, $\mathbf{X}^{(1)}$ and $\mathbf{X}^{(2)}$ consist of only a column of ones respectively, indicating that only the intercept is included as fixed effect. In this case, $\beta^{(1)}$ and $\beta^{(2)}$ are simply the means of the two phenotypes. In the rest of this paper, we assume that $\mathbf{y}^{(1)}$ and $\mathbf{y}^{(2)}$ are already mean-subtracted for convenience. $\mathbf{g}^{(1)}$ and $\mathbf{g}^{(2)}$ are the genetic values for the two phenotypes. $\mathbf{e}^{(1)}$ and $\mathbf{e}^{(2)}$ are the residuals. The genetic values $\mathbf{g}^{(1)}$ and $\mathbf{g}^{(2)}$ are treated as random effects:

$$\mathbf{g}^{(1)} = \mathbf{G}^{(1)}\mathbf{u}^{(1)}, \quad \mathbf{g}^{(2)} = \mathbf{G}^{(2)}\mathbf{u}^{(2)} \quad (2)$$

$$\begin{bmatrix} \mathbf{u}_j^{(1)} \\ \mathbf{u}_j^{(2)} \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{g}}) \quad (3)$$

$$\Sigma_{\mathbf{g}} = \begin{bmatrix} \sigma_{g_1}^2 & \rho_g \sigma_{g_1} \sigma_{g_2} \\ \rho_g \sigma_{g_1} \sigma_{g_2} & \sigma_{g_2}^2 \end{bmatrix} \quad (4)$$

where $\mathbf{G}^{(t)}$ is the standardized genotype matrix for a total of p SNPs of the training samples for the t -th phenotype. Specifically, let B and b be the two alleles at the j -th locus and f_j be the frequency of the B allele, then $\mathbf{G}_{ij}^{(t)}$ takes a value of

$-2f_j/\sqrt{2f_j(1-f_j)p}$, $(1-2f_j)/\sqrt{2f_j(1-f_j)p}$ or $2(1-f_j)/\sqrt{2f_j(1-f_j)p}$ if the genotype of the i -th individual at the j -th locus is bb , Bb or BB , respectively. $\mathbf{u}_j^{(1)}$ and $\mathbf{u}_j^{(2)}$ are the random effects of the j -th locus for the two phenotypes and ρ_g measures the strength of genetic correlation between the two phenotypes. We also assume the residuals follow a multivariate normal distribution:

$$\begin{bmatrix} \mathbf{e}^{(1)} \\ \mathbf{e}^{(2)} \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \sum_{\mathbf{e}}) \quad (5)$$

$$\sum_{\mathbf{e}} = \begin{bmatrix} \sigma_{e_1}^2 \mathbf{I}_{n_1} & 0 \\ 0 & \sigma_{e_2}^2 \mathbf{I}_{n_2} \end{bmatrix}. \quad (6)$$

In Lee et al. [21], the random effects $\mathbf{u}^{(1)}$, $\mathbf{u}^{(2)}$, $\mathbf{e}^{(1)}$ and $\mathbf{e}^{(2)}$ were integrated out and the average information restricted maximum likelihood (AI-REML) algorithm [13] was used to estimate the variance parameters $\sigma_{g_1}^2, \sigma_{g_2}^2, \rho_g, \sigma_{e_1}^2, \sigma_{e_2}^2$. ρ_g measures the genetic correlation between the two phenotypes whereas the other four variance parameters can be used to calculate the “chip” heritability or “variance explained by SNPs” for the two phenotypes. However, our focus here is not to estimate the “chip heritability” or the genetic correlation but to predict the phenotypes. Therefore we are more interested in estimating the random effects: $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$. Given a certain set of variance parameters $\sigma_{g_1}^2, \sigma_{g_2}^2, \rho_g, \sigma_{e_1}^2$ and $\sigma_{e_2}^2$, the posterior means of the random effects $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ can be written out as:

$$\begin{bmatrix} \hat{\mathbf{u}}^{(1)} \\ \hat{\mathbf{u}}^{(2)} \end{bmatrix} = ((\sum_{\mathbf{g}} \otimes \mathbf{I}_p)^{-1} + \mathbf{G}' \sum_{\mathbf{e}}^{-1} \mathbf{G})^{-1} \mathbf{G}' \sum_{\mathbf{e}}^{-1} \begin{bmatrix} \mathbf{y}^{(1)} \\ \mathbf{y}^{(2)} \end{bmatrix} \quad (7)$$

where $\mathbf{G} = \begin{bmatrix} \mathbf{G}^{(1)} & \mathbf{0} \\ \mathbf{0} & \mathbf{G}^{(2)} \end{bmatrix}$ and $\otimes$ stands for the Kronecker product. After some rearrangements, we can see that this is essentially equivalent to the solution to a bivariate ridge regression given a certain set of regularization parameters:

$$\begin{bmatrix} \hat{\mathbf{u}}^{(1)} \\ \hat{\mathbf{u}}^{(2)} \end{bmatrix} = (\mathbf{K} \otimes \mathbf{I}_p + \mathbf{G}' \mathbf{G})^{-1} \mathbf{G}' \begin{bmatrix} \mathbf{y}^{(1)} \\ \mathbf{y}^{(2)} \end{bmatrix} \quad (8)$$

where $\mathbf{K} = (1 - \rho_g^2)^{-1} \begin{bmatrix} \sigma_{e_1}^2 / \sigma_{g_1}^2 & -\rho_g \sigma_{e_2}^2 / (\sigma_{g_1} \sigma_{g_2}) \\ -\rho_g \sigma_{e_1}^2 / (\sigma_{g_1} \sigma_{g_2}) & \sigma_{e_2}^2 / \sigma_{g_2}^2 \end{bmatrix}$ is the “ridge matrix” as defined in Brown and Zidek [2].

We note that directly calculating Eq. (8) is infeasible because it involves inversion of a $2p$ by $2p$ matrix where p is total number of SNPs, which could be in the order of 10^6 in a typical GWAS. Therefore we used the matrix inversion lemma to rearrange it as:

$$\begin{bmatrix} \hat{\mathbf{u}}^{(1)} \\ \hat{\mathbf{u}}^{(2)} \end{bmatrix} = (\mathbf{K}^{-1} \otimes \mathbf{I}_p) \mathbf{G}' [\mathbf{I}_{(n_1+n_2)} + \mathbf{G}(\mathbf{K}^{-1} \otimes \mathbf{I}_p) \mathbf{G}']^{-1} \begin{bmatrix} \mathbf{y}^{(1)} \\ \mathbf{y}^{(2)} \end{bmatrix} \quad (9)$$

where $\mathbf{K}^{-1} = \begin{bmatrix} \sigma_{g_1}^2 / \sigma_{e_1}^2 & \rho_g \sigma_{g_1} \sigma_{g_2} / \sigma_{e_1}^2 \\ \rho_g \sigma_{g_1} \sigma_{g_2} / \sigma_{e_2}^2 & \sigma_{g_2}^2 / \sigma_{e_2}^2 \end{bmatrix}$. Now we only need to evaluate the inverse of a matrix of dimension $(n_1 + n_2)$ by $(n_1 + n_2)$, which is usually in the order of 10^3 . Given the estimates of the random effects $\hat{\mathbf{u}}^{(1)}$ and $\hat{\mathbf{u}}^{(2)}$ and two sets of validation individuals with standardized genotype matrix $\mathbf{G}_v^{(1)}$ and $\mathbf{G}_v^{(2)}$, the predicted values for the two phenotypes are given by:

$$\hat{\mathbf{y}}^{(1)} = \mathbf{G}_v^{(1)} \hat{\mathbf{u}}^{(1)}, \hat{\mathbf{y}}^{(2)} = \mathbf{G}_v^{(2)} \hat{\mathbf{u}}^{(2)} \quad (10)$$

The solution of the bivariate ridge regression depends on four regularization parameters. To alleviate the computational burden of tuning the regularization parameters, we impose the constraint that $\sigma_{e_1} = \sigma_{e_2}$, which leads to a symmetric “ridge matrix” $\mathbf{K}$ and hence a

symmetric $\mathbf{K}^{-1}$. Then we can reparameterize $\mathbf{K}^{-1}$ as $\begin{bmatrix} \lambda_1 & \rho_g \sqrt{\lambda_1 \lambda_2} \\ \rho_g \sqrt{\lambda_1 \lambda_2} & \lambda_2 \end{bmatrix}$. As a result, we only need to tune three regularization parameters λ_1 , λ_2 and ρ_g . λ_1 and λ_2 control the shrinkage level of genetic effects for the two phenotypes whereas ρ_g controls the correlation between the genetic effects between the two phenotypes. In order to find the optimal regularization parameters, we perform grid searches with λ_1 and λ_2 ranging from a very small number λ_{min} to a very large number λ_{max} and $\rho_g \in [0, 1)$. In the real data analysis, the parameter values that give the highest cross-validation mean AUC for each phenotype were chosen. In the simulation studies, the parameter values that give the highest AUC in the validation data for each phenotype were chosen.

We classify the validation individuals into affected ones and unaffected ones by dichotomizing the predicted values with a grid of threshold values ranging from the largest predicted value to the smallest predicted value and obtain the *receiver operating characteristic* curve by evaluating the sensitivity and specificity at each threshold value. Then we evaluate the area under *receiver operating characteristic* curve (AUC) as a measure of the prediction accuracy.

In the bivariate ridge regression, we used cross-validation to tune the regularization parameters instead of trying to estimate the parameters from the data as in [21]. There are primarily two reasons: 1) the AI-REML algorithm that Lee et al. used in [21] sometimes fails to converge, for example on the CD-UC data set. 2) The estimated parameters are based on normality assumptions and henceforth may not be optimal for predicting binary phenotypes.

Simulation Study

In order to examine the relationship between the gain of predictive power and the level of genetic correlation between two diseases, we performed the following simulation studies

based on the classical liability threshold model [9] to simulate the case-control data. Specifically, given a desired sample size N , the total number of SNPs p , the proportion of cases in the case-control data P and the disease prevalence Q , genotypes of a cohort of at least NP/Q individuals were generated as follows: first, the MAF of p SNPs were uniformly drawn from $[0.05, 0.5]$; then the genotypes (the number of minor allele copies) for the NP/Q individuals at each SNP were drawn from a binomial distribution of size two and the probability of success being the MAF of the corresponding SNP.

After the genotypes were generated, m causal SNPs were randomly chosen with each causal SNP carrying a per-minor-allele effect drawn from a normal distribution with zero mean and

variance of $\frac{h^2}{(1-h^2)f_j(1-f_j)m}$ where h^2 is the desired level of variance explained by all SNPs on the liability scale and f_j is the MAF of the corresponding causal SNP. Then the environmental effect on the liability scale for each individual was independently drawn from a standard normal distribution (zero mean and unit variance). Then, we obtained the liability for each individual by adding up the genetic effects conferred by all causal SNPs and environmental effect. Once the liabilities were obtained, individuals with liabilities greater than the $1 - Q$ quantile were classified as cases and the others classified as controls. Then NP cases and $N(1 - P)$ controls were randomly drawn from the cohort. When simulating data for two diseases simultaneously, the two data set were simulated from two disjoint cohorts. The MAF for each SNP was the same for the two cohorts. A total of m' causal SNPs were chosen to be shared between the two diseases to mimic the shared genetic basis between them. We assume shared causal SNPs have a correlation of 0.8 between their liability scale effect sizes in the two diseases.

In our simulation studies, we set the sample size N at 1000, 1500, 2000 or 2500. The number of causal SNPs m was chosen to be 500, 1000, 1500 or 2,000. P was fixed at 0.5 and Q was fixed at 0.05. We assumed the two diseases had the same number of causal SNPs. We varied the proportion of the causal SNPs that were shared between the two diseases $\gamma = m'/m$ from 0 to 1 to model different levels of genetic correlation and investigated the gain of predictive power under different settings. We also investigated the effects of unequal sample size and unequal h^2 between the two diseases.

3 Results

Results from Real Data

We analyzed GWAS data for bipolar and related disorders (BARD) and schizophrenia (SZ) that we downloaded from the dbGaP database (<https://dbgap.ncbi.nlm.nih.gov>). After pre-processing (see Methods), the combined data set consisted of 3,226 individuals (653 cases for BARD, 1,170 cases for SZ and 1,403 controls) with genotypes at 298,604 SNPs. In order to organize two case-control data sets for the two diseases, we randomly partitioned the controls into two disjoint sets of individuals and assigned each set to one of the diseases. We partitioned the control individuals so that the proportions of cases for the two data sets were approximately equal, i.e. we assigned 505 controls to BARD and 898 controls to SZ. The random partitioning was repeated 50 times.

We then used the bivariate ridge regression method [2] described above to predict jointly the disease status for BARD and SZ using this data set. Prediction performance was evaluated through five-fold cross-validation, i.e., the data were randomly partitioned into five equal-sized folds and four of them were used as the training set and the other one was used as the validation set each time. We calculated the mean AUC (area under the *receiver-operating-characteristic* curve) for the five folds to measure prediction performance. For comparison, we also evaluated prediction performance of three other prediction methods: univariate ridge regression, SVM (support vector machine) [10] with linear kernel, and LASSO (least absolute shrinkage and selection operator) [32] that treat each disease separately (Figure 1). Note that the bivariate ridge regression method we used is a direct generalization of the univariate ridge regression (see Methods). Therefore the utility of pleiotropy can be best demonstrated through the comparison between the two approaches. Indeed, the bivariate ridge regression method performed better than the univariate ridge regression for both BARD and SZ. We also note that the gain of AUC of BARD (0.033) was larger than that of SZ (0.012), possibly as a result of the sample size difference between the two diseases and a stronger genetic signals for SZ. SVM performed worse than univariate ridge regression in BARD, but slightly outperformed both bivariate and univariate ridge regression in SZ. LASSO, which prefers a sparse genetic architecture with a small number of major (strong effect) SNPs, did not perform well for either disease.

To evaluate whether the improvement was the result of the genetic correlation between BARD and SZ or an artificial effect of the bivariate ridge regression method, we randomly selected 25%, 50%, 75% and all of the SNPs and shuffled their identities in the BARD samples while keeping the SNP identities unchanged in the SZ samples. This led to reduced genetic correlation between BARD and SZ and consequently reduced improvement of prediction accuracy for the bivariate ridge regression (Figure 2). In particular, when all the SNPs were shuffled, the performance of bivariate ridge regression was almost the same as univariate ridge regression, confirming that the gain of prediction accuracy of the bivariate ridge regression indeed came from specific SNPs, and thus the genetic correlation between the two phenotypes. To further demonstrate this, we also analyzed the WTCCC (Wellcome Trust Case Control Consortium) type 2 diabetes case-control samples [3], combining them with the BARD and SZ samples. The prediction accuracy was almost the same between univariate ridge regression and bivariate ridge regression (Figure S1), suggesting that the prediction accuracy can only be improved by bivariate ridge regression when genuine genetic correlation exists between two phenotypes.

To investigate the effect of sample size, we randomly sampled 50% and 75% individuals from the original data and performed both univariate and bivariate ridge regressions. As expected, the prediction accuracy increased as the sample size increased (Figure 3). Moreover, the prediction accuracy of bivariate ridge regression for BARD with 50% samples even outperformed univariate ridge regression with 100% samples. As for SZ, the prediction of univariate ridge regression with 100% samples was close to bivariate ridge regression with 75% samples. This difference may also be due the sample size difference between the two diseases.

We did the same analyses on a GWAS data set for Crohn's disease (CD) and a GWAS data set for ulcerative colitis (UC) obtained from the dbGap database. The CD data set consists of 515 unaffected subjects and 513 affected subjects with non-Jewish European ancestry and the UC data set consists of 902 affected subjects with non-Jewish European ancestry. Because there are no control samples for the UC data set, we obtained 902 samples from the Illumina iControl database (<http://www.illumina.com/science/icontroldb.ilmn>) that were genetically matched with the UC samples to minimize the confounding effect of potential population stratification (see Methods). After preprocessing (see Methods), we obtained a data set consisting of 513 affected subjects for CD, 515 unaffected subjects for CD, 902 affected subjects for UC and 902 unaffected subjects for UC with genotypes for 241,649 SNPs. As opposed to the BARD and SZ data set, the control samples for CD and UC were fixed instead of being randomly partitioned.

Prediction accuracies of different methods on the CD-UC data set were also evaluated through five-fold cross-validation, with the cross-validation repeated 50 times. The results are shown in Figure 4. Similar to the results for the BARD-SZ data, bivariate ridge regression performed better than univariate ridge regression in both diseases. Moreover, the increase of AUC for CD, which has a smaller sample size, was larger than that of UC (0.027 versus 0.015). The prediction accuracy of SVM was close to that of univariate ridge regression but was outperformed by bivariate ridge regression in both diseases.

Similar to the BARD-SZ data set, we performed random shuffling of the identities of the SNPs in the CD samples while keeping their identities in the UC samples. The results are shown in Figure 5. As expected, the improvement of prediction accuracy of bivariate ridge regression also decreases as the fraction of SNPs increases and eventually diminishes when all SNPs are shuffled.

We also examined the prediction accuracy of bivariate and univariate ridge regression on subsamples of the CD-UC data set. The results are shown in Figure 6. For CD, the prediction of univariate ridge regression with 100% samples was close to bivariate ridge regression with 50% samples. Whereas for UC, the prediction of univariate ridge regression with 100% samples was close to bivariate ridge regression with 75% samples.

Simulation Study

We simulated data sets for two genetically correlated diseases to evaluate the utility of genetic correlation in genetic risk prediction. We simulated a total of $p = 20,000$ SNPs. We let the sample size N to be 1000, 1500, 2000 or 2500. The number of causal SNPs m was also chosen to be 500, 1000, 1500 or 2000. We assumed the two traits had the same number of causal SNPs. We varied the proportion of the causal SNPs that were shared between the two diseases γ from 0 to 1 to mimic different levels of genetic correlation. We assumed that the shared causal SNPs had a correlation of 0.8 between their effect sizes on the two diseases. Our empirical results suggest that the realized genetic correlation is approximately 0.8γ under this setup (results not shown). We simulated two h^2 levels (on liability scale, see [20]): 0.3 and 0.6. Disease prevalence was set to be 0.05 and equal numbers of cases and controls were drawn from the simulated population. The simulations were repeated for 25 times in each scenario. In each repeat, a training set and a validation set were independently

generated and were used to evaluate the prediction accuracies of bivariate ridge regression and univariate ridge regression.

We first considered the scenario where both diseases had the same number of samples and the same h^2 level. The results for $m = 1000$, $N = 1000$ and 2000 are shown in Figure 7 and more results for other scenarios are shown in Figures S2–S9. As expected, the prediction accuracy for bivariate and univariate ridge regressions both increased as sample size increased. When the genetic correlation between the two diseases was zero, the prediction accuracy of bivariate ridge regression was almost the same as univariate ridge regression. As the genetic correlation increased, the AUC for bivariate ridge regression also increased and the improvement over univariate ridge regression became quite noticeable as long as $\gamma \geq 0.5$, which corresponds to a genetic correlation of approximately 0.4.

We also examined the case when the sample sizes (Figure 6 and Figure S10) or the h^2 levels (Figure 7 and Figure S11) were different between the two traits. When the sample sizes were different but the h^2 levels were the same between the two diseases, the prediction accuracy of the disease with the smaller sample size improved more than that of the disease with the larger sample size by joint modeling, which echoes with the results from the real data analysis. When the h^2 level were different but the sample sizes were the same between the two diseases, the prediction accuracy of the disease with the lower h^2 level improved more than that of the disease with the higher h^2 level.

In Figures S6–S9 we examined the effect of m on the prediction accuracy of univariate and bivariate ridge regressions and we found that these two methods were relatively insensitive to the change of m .

4 Discussion

Genetic risk prediction is a challenging task due to the existence of numerous weak-effect variants. Risk prediction is also bottlenecked by the available sample size of GWAS. A promising yet under-appreciated strategy to increase the effective training sample size, and thus to improve prediction accuracy, is to integrate information from diseases that are genetically correlated. In this paper, we performed a systematic investigation of the utility of genetic correlation in genetic risk prediction using a bivariate ridge regression method. In the *real data* analysis, we demonstrated that gain of predictive power can be achieved by using the genetic correlation between phenotypes. In the *simulation* study, the results confirmed our findings in the real data, as well as offering us insights into the relationship between gain of prediction accuracy and the level of genetic correlation between two phenotypes. These results can provide an important guide for researchers to make productive use of pleiotropy while developing disease risk prediction models.

In both the real data and simulation studies, we found that the prediction accuracy of the disease with the smaller sample size benefited more from the joint modeling than that of the disease with the larger sample size. A possible reason is that when jointly modeling the two data sets, the larger data set contains more information or statistical evidence about the genetic effects shared by the two diseases than the smaller data set, and therefore more information can be borrowed from the larger data set to the smaller data set than the

opposite way. In the simulation studies, we also observed that when the sample sizes were the same but the h^2 levels were different, joint modeling offered more benefit to the low-heritability disease than that of the high-heritability disease. A similar explanation can be made: the high-heritability disease data set can provide stronger statistical evidence about the genetic effects shared by the two diseases than the low-heritability disease. Note that the two explanations are based on the simulation setup that the number of associated SNPs is the same for the two diseases, all the SNPs are independent, and the genetic effects are randomly distributed among the causal SNPs.

To our knowledge, only Hartley et al [15,16] have studied pleiotropy in genetic risk prediction to date. However, our study is distinct from theirs in two main aspects. First, Hartley et al. focused on methodological development, whereas our study emphasizes the utility of genetic correlation using real GWAS data sets. In fact, our study is the first to demonstrate a substantial contribution of genetic correlation to prediction accuracy on real data. Second, their method requires that all the phenotypes of interest be observed in all individuals, which limits its applications. In contrast, we consider the case where two phenotypes are observed on two disjoint sets of individuals. Not only does this allow more GWAS data sets to be integrated, we can expect much more improvement of the predictive ability since the sample size is effectively increased by combining the data for the two phenotypes whereas there is no increase in the sample size in the scenario considered by Hartley and colleagues. In addition, even when we have phenotype information from multiple traits for each study subject, the bivariate mixed effect model may not appropriately accommodate specific, sometimes complex, ascertainment schemes used to collect the GWAS samples. This may lead to potential biases and reduced prediction accuracy, e.g. causing some shared genetic effects between the two traits to be cancelled out.

Solovieff et al [30] identified several sources of spurious pleiotropy. One is sample ascertainment bias. In our real data examples, each “case” individual is affected by only one disease and the controls are not shared by the two diseases. Therefore, it is unlikely that our study is affected by ascertainment bias. Another source is misdiagnosis, e.g. people with schizophrenia may be misclassified as bipolar disorder. However they also noted that the misclassification rate has to be very high to generate a substantial genetic correlation, which suggests the gain of prediction accuracy in our results may still primarily be a result of the true genetic correlation between the two diseases.

The method that we used to predict jointly two phenotypes is a bivariate generalization to the ridge regression. Although the comparison between bivariate ridge regression and univariate suggests that pleiotropy can indeed substantially contribute to genetic risk prediction, we do note that there is also a need to develop better methods for joint prediction of multiple phenotypes. For example, SVM sometimes does better than univariate ridge regression. Development of such methods that jointly predict multiple phenotypes is an important future task. In this study, we demonstrated the utility of pleiotropy through two pairs of diseases that are known to share a lot of common genetic bases. Another important future task is to comprehensively examine the genetic correlations between other complex human diseases and their utility in risk prediction.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This study was supported by NIH grants R01 AA11330, AA017535, DA030976, GM59507, VA Cooperative Studies Program 572, and a fellowship from China Scholarship Council. We also thank Yale University High Performance Computing Center (funded by NIH RR19895) for the computation resource and data storage.

Funding support for the Whole Genome Association Study of Bipolar Disorder was provided by the National Institute of Mental Health (NIMH) and the genotyping of samples was provided through the Genetic Association Information Network (GAIN). The datasets used for the analyses described in this manuscript were obtained from the database of Genotypes and Phenotypes (dbGaP) found at <http://www.ncbi.nlm.nih.gov/gap> through dbGaP accession number phs000017.v3.p1. Samples and associated phenotype data for the Collaborative Genomic Study of Bipolar Disorder were provided by the The NIMH Genetics Initiative for Bipolar Disorder. Data and biomaterials were collected in four projects that participated in NIMH Bipolar Disorder Genetics Initiative. From 1991–98, the Principal Investigators and Co-Investigators were: Indiana University, Indianapolis, IN, U01 MH46282, John Nurnberger, M.D., Ph.D., Marvin Miller, M.D., and Elizabeth Bowman, M.D.; Washington University, St. Louis, MO, U01 MH46280, Theodore Reich, M.D., Allison Goate, Ph.D., and John Rice, Ph.D.; Johns Hopkins University, Baltimore, MD U01 MH46274, J. Raymond DePaulo, Jr., M.D., Sylvia Simpson, M.D., MPH, and Colin Stine, Ph.D.; NIMH Intramural Research Program, Clinical Neurogenetics Branch, Bethesda, MD, Elliot Gershon, M.D., Diane Kazuba, B.A., and Elizabeth Maxwell, M.S.W. Data and bio-materials were collected as part of ten projects that participated in the NIMH Bipolar Disorder Genetics Initiative. From 1999–03, the Principal Investigators and Co-Investigators were: Indiana University, Indianapolis, IN, R01 MH59545, John Nurnberger, M.D., Ph.D., Marvin J. Miller, M.D., Elizabeth S. Bowman, M.D., N. Leela Rau, M.D., P. Ryan Moe, M.D., Nalini Samavedy, M.D., Rif El-Mallakh, M.D. (at University of Louisville), Husseini Manji, M.D. (at Wayne State University), Debra A. Glitz, M.D. (at Wayne State University), Eric T. Meyer, M.S., Carrie Smiley, R.N., Tatiana Foroud, Ph.D., Leah Flury, M.S., Danielle M. Dick, Ph.D., Howard Edenberg, Ph.D.; Washington University, St. Louis, MO, R01 MH059534, John Rice, Ph.D., Theodore Reich, M.D., Allison Goate, Ph.D., Laura Bierut, M.D.; Johns Hopkins University, Baltimore, MD, R01 MH59533, Melvin McInnis M.D., J. Raymond DePaulo, Jr., M.D., Dean F. MacKinnon, M.D., Francis M. Mondimore, M.D., James B. Potash, M.D., Peter P. Zandi, Ph.D., Dimitrios Avramopoulos, and Jennifer Payne; University of Pennsylvania, PA, R01 MH59553, Wade Berrettini M.D., Ph.D.; University of California at Irvine, CA, R01 MH60068, William Byerley M.D., and Mark Vawter M.D.; University of Iowa, IA, R01 MH059548, William Coryell M.D., and Raymond Crowe M.D.; University of Chicago, IL, R01 MH59535, Elliot Gershon, M.D., Judith Badner Ph.D., Francis McMahon M.D., Chunyu Liu Ph.D., Alan Sanders M.D., Maria Caserta, Steven Dinwiddie M.D., Tu Nguyen, Donna Harakal; University of California at San Diego, CA, R01 MH59567, John Kelsoe, M.D., Rebecca McKinney, B.A.; Rush University, IL, R01 MH059556, William Scheftner M.D., Howard M. Kravitz, D.O., M.P.H., Diana Marta, B.S., Annette Vaughn-Brown, MSN, RN, and Laurie Bederow, MA; NIMH Intramural Research Program, Bethesda, MD, 1Z01MH002810-01, Francis J. McMahon, M.D., Layla Kassem, PsyD, Sevilla Detera-Wadleigh, Ph.D., Lisa Austin, Ph.D., Dennis L. Murphy, M.D.

Funding support for the Genome-Wide Association of Schizophrenia Study was provided by the National Institute of Mental Health (R01 MH67257, R01 MH59588, R01 MH59571, R01 MH59565, R01 MH59587, R01 MH60870, R01 MH59566, R01 MH59586, R01 MH61675, R01 MH60879, R01 MH81800, U01 MH46276, U01 MH46289, U01 MH46318, U01 MH79469, and U01 MH79470) and the genotyping of samples was provided through the Genetic Association Information Network (GAIN). The datasets used for the analyses described in this manuscript were obtained from the database of Genotypes and Phenotypes (dbGaP) found at <http://www.ncbi.nlm.nih.gov/gap> through dbGaP accession number phs000021.v3.p2. Samples and associated phenotype data for the Genome-Wide Association of Schizophrenia Study were provided by the Molecular Genetics of Schizophrenia Collaboration (PI: Pablo V. Gejman, Evanston Northwestern Healthcare (ENH) and Northwestern University, Evanston, IL, USA).

The NIDDK IBD Genetics Consortium Crohn's Disease Genome-Wide Association Study was conducted by the NIDDK IBD Genetics Consortium Crohn's Disease Genome-Wide Association Study Investigators and supported by the National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK). The data and samples from the NIDDK IBD Genetics Consortium Crohn's Disease Genome-Wide Association Study reported here were supplied by the NIDDK Central Repositories. The datasets used for the analyses described in this manuscript were obtained from the database of Genotypes and Phenotypes (dbGaP) found at <http://www.ncbi.nlm.nih.gov/gap> through dbGaP accession number phs000130.v1.p1. This manuscript was not prepared in collaboration with Investigators of the NIDDK IBD Genetics Consortium Crohn's Disease Genome-Wide Association Study and does not necessarily reflect the opinions or views of the NIDDK IBD Genetics Consortium Crohn's Disease Genome-Wide Association Study, the NIDDK Central Repositories, or the NIDDK.

The NIDDK IBD Genetics Consortium Ulcerative Colitis Genome-Wide Association Study was conducted by the NIDDK IBD Genetics Consortium Ulcerative Colitis Genome-Wide Association Study Investigators and supported by the National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK). The data and samples from the NIDDK IBD Genetics Consortium Ulcerative Colitis Genome-Wide Association Study reported here were supplied by the NIDDK Central Repositories. The datasets used for the analyses described in this manuscript were obtained from the database of Genotypes and Phenotypes (dbGaP) found at <http://www.ncbi.nlm.nih.gov/gap> through dbGaP accession number phs000345.v1.p1. This manuscript was not prepared in collaboration with Investigators of the NIDDK IBD Genetics Consortium Ulcerative Colitis Genome-Wide Association Study and does not necessarily reflect the opinions or views of the NIDDK IBD Genetics Consortium Ulcerative Colitis Genome-Wide Association Study, the NIDDK Central Repositories, or the NIDDK.

References

1. Andreassen OA, Djurovic S, Thompson WK, Schork AJ, Kendler KS, O'Donovan MC, Rujescu D, Werge T, van de Bunt M, Morris AP, et al. Improved detection of common variants associated with schizophrenia by leveraging pleiotropy with cardiovascular-disease risk factors. *The American Journal of Human Genetics*. 2013; 92(2):197–209.
2. Brown PJ, Zidek JV. Adaptive multivariate ridge regression. *The Annals of Statistics*. 1980; 8(1): 64–74.
3. Burton PR, Clayton DG, Cardon LR, Craddock N, Deloukas P, Duncanson A, Kwiatkowski DP, McCarthy MI, Ouwehand WH, Samani NJ, et al. Genome-wide association study of 14,000 cases of seven common diseases and 3,000 shared controls. *Nature*. 2007; 447(7145):661–678. [PubMed: 17554300]
4. de los Campos G, Gianola D, Allison D. Predicting genetic predisposition in humans: the promise of whole-genome markers. *Nature Reviews Genetics*. 2010; 11(12):880–886.
5. Clarke AJ, Cooper DN. GWAS: heritability missing in action? *European Journal of Human Genetics*. 2010; 18(8):859–861. [PubMed: 20234388]
6. Collins FS, McKusick VA. Implications of the Human Genome Project for medical science. *JAMA: the journal of the American Medical Association*. 2001; 285(5):540–544. [PubMed: 11176855]
7. Dickson SP, Wang K, Krantz I, Hakonarson H, Goldstein DB. Rare variants create synthetic genome-wide associations. *PLoS Biology*. 2010; 8(1):e1000, 294.
8. Eichler EE, Flint J, Gibson G, Kong A, Leal SM, Moore JH, Nadeau JH. Missing heritability and strategies for finding the underlying causes of complex disease. *Nature Reviews Genetics*. 2010; 11(6):446–450.
9. Falconer DS. The inheritance of liability to certain diseases, estimated from the incidence among relatives. *Annals of Human Genetics*. 1965; 29(1):51–76.
10. Fan RE, Chang KW, Hsieh CJ, Wang XR, Lin CJ. Liblinear: A library for large linear classification. *The Journal of Machine Learning Research*. 2008; 9:1871–1874.
11. Forer L, Schönherr S, Weissensteiner H, Haider F, Kluckner T, Gieger C, Wichmann HE, Specht G, Kronenberg F, Kloss-Brandstätter A. CONAN: copy number variation analysis software for genome-wide association studies. *BMC Bioinformatics*. 2010; 11(1):318. [PubMed: 20546565]
12. Gibson G, et al. Hints of hidden heritability in GWAS. *Nature Genetics*. 2010; 42(7):558–560. [PubMed: 20581876]
13. Gilmour AR, Thompson R, Cullis BR. Average information REML: an efficient algorithm for variance parameter estimation in linear mixed models. *Biometrics*. 1995; 51(4):1440–1450.
14. Haitovsky Y. On multivariate ridge regression. *Biometrika*. 1987; 74(3):563–570.
15. Hartley SW, Monti S, Liu CT, Steinberg MH, Sebastiani P. Bayesian methods for multivariate modeling of pleiotropic SNP associations and genetic risk prediction. *Frontiers in Genetics*. 2012; 3
16. Hartley SW, Sebastiani P. PleioGRiP: genetic risk prediction with pleiotropy. *Bioinformatics*. 2013; 29(8):1086–1088. [PubMed: 23419378]
17. Huang J, Johnson AD, O'Donnell CJ. PRIME: a method for characterization and evaluation of pleiotropic regions from multiple genome-wide association studies. *Bioinformatics*. 2011; 27(9): 1201–1206. [PubMed: 21398673]

18. Huebinger RM, Garner HR, Barber RC. Pathway genetic load allows simultaneous evaluation of multiple genetic associations. *Burns*. 2010; 36(6):787–792. [PubMed: 20381257]
19. Korte A, Vilhjálmsson BJ, Segura V, Platt A, Long Q, Nordborg M. A mixed-model approach for genome-wide association studies of correlated traits in structured populations. *Nature Genetics*. 2012; 44(9):1066–1071. [PubMed: 22902788]
20. Lee S, Wray N, Goddard M, Visscher P. Estimating missing heritability for disease from genome-wide association studies. *The American Journal of Human Genetics*. 2011; 88(3):294–305.
21. Lee S, Yang J, Goddard M, Visscher P, Wray N. Estimation of pleiotropy between complex diseases using single-nucleotide polymorphism-derived genomic relationships and restricted maximum likelihood. *Bioinformatics*. 2012; 28(19):2540–2542. [PubMed: 22843982]
22. Maher B. The case of the missing heritability. *Nature*. 2008; 456(7218):18–21. [PubMed: 18987709]
23. Makowsky R, Pajewski NM, Klimentidis YC, Vazquez AI, Duarte CW, Allison DB, de los Campos G. Beyond missing heritability: prediction of complex traits. *PLoS Genetics*. 2011; 7(4):e1002, 051.
24. Manolio TA, Collins FS, Cox NJ, Goldstein DB, Hindorff LA, Hunter DJ, McCarthy MI, Ramos EM, Cardon LR, Chakravarti A, et al. Finding the missing heritability of complex diseases. *Nature*. 2009; 461(7265):747–753. [PubMed: 19812666]
25. Park JH, Wacholder S, Gail MH, Peters U, Jacobs KB, Chanock SJ, Chatterjee N. Estimation of effect size distribution from genome-wide association studies and implications for future discoveries. *Nature Genetics*. 2010; 42(7):570–575. [PubMed: 20562874]
26. Purcell SM, Wray NR, Stone JL, Visscher PM, O'Donovan MC, Sullivan PF, Sklar P, Ruderfer DM, McQuillin A, Morris DW, et al. Common polygenic variation contributes to risk of schizophrenia and bipolar disorder. *Nature*. 2009; 460(7256):748–752. [PubMed: 19571811]
27. Sakoda LC, Jorgenson E, Witte JS. Turning of COGS moves forward findings for hormonally mediated cancers. *Nature Genetics*. 2013; 45(4):345–348. [PubMed: 23535722]
28. Sivakumaran S, Agakov F, Theodoratou E, Prendergast JG, Zgaga L, Manolio T, Rudan I, McKeigue P, Wilson JF, Campbell H. Abundant pleiotropy in human complex diseases and traits. *The American Journal of Human Genetics*. 2011; 89(5):607–618.
29. Smoller JW, Craddock N, Kendler K, Lee PH, Neale BM, Nurnberger JI, Ripke S, Santangelo S, Sullivan PF, et al. Identification of risk loci with shared effects on five major psychiatric disorders: a genome-wide analysis. *Lancet*. 2013; 381(9875):1371–1379. [PubMed: 23453885]
30. Solovieff N, Cotsapas C, Lee PH, Purcell SM, Smoller JW. Pleiotropy in complex traits: challenges and strategies. *Nature Reviews Genetics*. 2013; 14(7):483–495.
31. Thompson R. The estimation of variance and covariance components with an application when records are subject to culling. *Biometrics*. 1973; 29(3):527–550.
32. Tibshirani R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B (Methodological)*. 1996; 58(1):267–288.
33. Vattikuti S, Guo J, Chow CC. Heritability and Genetic Correlations Explained by Common SNPs for Metabolic Syndrome Traits. *PLoS Genetics*. 2012; 8(3):e1002, 637.
34. Visscher P, Brown M, McCarthy M, Yang J. Five years of GWAS discovery. *American Journal of Human Genetics*. 2012; 90(1):7. [PubMed: 22243964]
35. Wray NR, Yang J, Hayes BJ, Price AL, Goddard ME, Visscher PM. Pitfalls of predicting complex traits from SNPs. *Nature Reviews Genetics*. 2013; 14(7):507–515.
36. Yang J, Benyamin B, McEvoy B, Gordon S, Henders A, Nyholt D, Madden P, Heath A, Martin N, Montgomery G, et al. Common SNPs explain a large proportion of the heritability for human height. *Nature Genetics*. 2010; 42(7):565–569. [PubMed: 20562875]
37. Zhou X, Carbonetto P, Stephens M. Polygenic modeling with Bayesian sparse linear mixed models. *PLoS Genetics*. 2013; 9(2):e1003, 264.

Fig. 1.

Prediction accuracy of different methods on the BARD-SZ data. “BVR”: bivariate ridge regression; “UVR”: univariate ridge regression. The numbers in the brackets are the mean AUCs achieved by each method in the 50 repeats.

Fig. 2.

Prediction accuracy of the bivariate ridge regression after shuffling the SNP identities and of ridge regression. Red plots represent the results of bivariate ridge regression and blue plots represent those of ridge regression. The percentage below each red plot represents the fraction of the SNPs that were shuffled.

Fig. 3. The prediction accuracy of bivariate and univariate ridge regression for BARD and SZ with subsamples.

Fig. 4. Prediction accuracy of different methods on the CD-UC data. “BVR” and “UVR” are defined as in Figure 1. The numbers in the brackets are the mean AUCs achieved by each method in the 50 repeats.

Fig. 5. Prediction accuracy of the bivariate ridge regression after shuffling the SNP identities and of ridge regression. Red plots represent the results of bivariate ridge regression and blue plots represent those of the ridge regression. The percentage below each red plot represents the fraction of the SNPs that were shuffled.

Fig. 6.
The prediction accuracy of bivariate and univariate ridge regression for CD and UC with subsamples.

Fig. 7.

Simulation results for the case when the two diseases have equal sample sizes and h^2 levels and $m = 1000$. “BVR” and “UVR” are defined as in Figure 1. Two h^2 levels (0.3 and 0.6) and two sample sizes (1000 and 2000) were simulated. The proportion of shared causal SNPs, γ was varied from 0 to 1 with an increment of 0.25. The numbers below the UVR box plots are the sample sizes. Following the UVR box plots are the box plots representing the results of BVR with the same sample sizes at different γ values (below the BVR box plots).

Fig. 8.

Simulation results for the case when the two diseases have unequal sample sizes and equal h^2 levels and $m = 1000$. “BVR” and “UVR” are defined as in Figure 1. One of the diseases has 2000 samples and the other has 1000 samples. Two h^2 levels (0.6 and 0.3) were simulated. The proportion of shared causal SNPs, γ was varied from 0 to 1 with an increment of 0.25. The numbers below the UVR box plots are the h^2 levels. Following the UVR box plots are the results of BVR with the same h^2 levels at different γ values (below the BVR box plots).

Fig. 9.

Simulation results for the case when the two diseases have equal sample sizes and unequal h^2 levels and $m = 1000$. “BVR” and “UVR” are defined as in Figure 1. One of the diseases has $h^2 = 0.6$ and the other has $h^2 = 0.3$. Two sample sizes (2000 and 1000) were simulated. The proportion of shared causal SNPs, γ was varied from 0 to 1 with an increment of 0.25. The numbers below the UVR box plots are the sample sizes. Following the UVR box plots are the results of BVR with the same sample sizes at different γ values (below the BVR box plots).