Published in final edited form as:

Biometrics. 2011 September ; 67(3): 967–974. doi:10.1111/j.1541-0420.2010.01520.x.

Control of population stratification by correlation-selected principal components

Seunggeun Lee^{*}, Fred A. Wright^{**}, and Fei Zou^{***}

Department of Biostatistics, University of North Carolina, Chapel Hill, NC 27599, USA

Summary

In genome-wide association studies, population stratification is recognized as producing inflated type I error due to the inflation of test statistics. Principal component-based methods applied to genotypes provide information about population structure, and have been widely used to control for stratification. Here we explore the precise relationship between genotype principal components and inflation of association test statistics, thereby drawing a connection between principal component-based stratification control and the alternative approach of genomic control. Our results provide an inherent justification for the use of principal components, but call into question the popular practice of selecting principal components based on significance of eigenvalues alone. We propose a new approach, called EigenCorr, which selects principal components based on both their eigenvalues and their correlation with the (disease) phenotype. Our approach tends to select fewer principal components for stratification control than does testing of eigenvalues alone, providing substantial computational savings and improvements in power. Analyses of simulated and real data demonstrate the usefulness of the proposed approach.

Keywords

Genomic Control; GWAS; PCA; Population Stratification

1. Introduction

In tests of genetic association among unrelated individuals, it is recognized that population stratification can result in test statistics with inflated apparent significance, resulting in overall type I error that is far above the nominal level. The method of *genomic control* (Devlin and Roeder, 1999; Devlin et al., 2001) was among the first attempts to address this problem, and is straightforward. For (chi-square) statistics at numerous markers measuring association with phenotype, an estimate is obtained for the inflation of test statistics beyond that expected under the null hypothesis and assuming no stratification. Then the test statistics are all adjusted by the inflation factor. However, a typical genome-wide association scan (GWAS) tests a very large number of SNP markers, requiring a stringent significance threshold. In this setting, genomic control can fail to properly control the type I error (Devlin et al., 2004; Marchini et al., 2004), in part because variance inflation is not constant across the SNPs.

*slee@bios.unc.edu

**fwright@bios.unc.edu

***fzou@bios.unc.edu

Alternatively, the principal component (PC) approach (Price et al., 2006) uses PCs computed from all genotypes as covariates in phenotype-genotype regression or in stratified analyses. Results from numerous studies indicate that the PC values often reflect known substructure and ancestry (Price et al., 2008). A great advantage of the PC approach is its sensitivity, effectively adjusting test statistics for only those markers contributing to the stratification. However, several challenges remain for effective PC-based stratification control. The primary challenge lies in choosing which PCs to include as covariates. Clearly not all PCs can be included, as there are as many PCs as there are individuals under study. Price et al. (2006) originally suggested to use the 10 PCs with the highest eigenvalues. These investigators later proposed using the Tracy-Widom (TW) statistic (Patterson et al., 2006) to assess significance of eigenvalues to select PCs. However, this approach may detect a very large number of PCs as significant, with uncertain impact on the analysis. Moreover, the precise contribution of each PC to the overall type I error has not been established. As we shall see below, it is entirely possible for a relatively low-ranked PC to have a greater impact on type I error than a higher-ranked PC.

The paper is arranged as follows. In Section 2, we establish a relationship between PCs and the average of the test statistics. Based on this relationship, we propose a new method, EigenCorr, to select PCs based on their eigenvalues and their correlations with the phenotype. The explicit use of phenotypes in stratification control has been anticipated in previous work, e.g. Epstein et al. (2007) and Kimmel et al. (2007). However, EigenCorr provides a direct connection to the type I error inflation introduced by PCs. A straightforward generalization applies when only a subset of markers are used for stratification control. In Section 3, we demonstrate the usefulness of EigenCorr via simulation and real GWAS analysis. In Section 4, we conclude with a discussion of implications and future directions.

2. Materials and Methods

Let g_{ij} be the genotype of SNP i and individual j , where $i = 1, \dots, M$ and $j = 1, \dots, N$. We

define a normalized genotype x_{ij} as $x_{ij} = (g_{ij} - \bar{g}_i) / \sqrt{\sum_{j=1}^N (g_{ij} - \bar{g}_i)^2 / N}$, where

$\bar{g}_i = \sum_{j=1}^N g_{ij} / N$. Let $\mathbf{X}$ be the resulting $M \times N$ normalized genotype matrix, and $\mathbf{x}_i$ the i th row of $\mathbf{X}$. We have $\sum_{j=1}^N x_{ij} = 0$ and $\sum_{j=1}^N x_{ij}^2 = 1$. For mathematical precision in later development, the normalization used here is slightly different from that used in Price et al. (2006), but produces nearly identical PCs. From the singular value decomposition (SVD) we obtain $\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{P}^T$, where $\mathbf{D}$ is an $M \times N$ diagonal matrix of ordered singular values with j th diagonal element d_j , $\mathbf{U}$ is an $M \times N$ loading matrix, and $\mathbf{P}$ is the $N \times N$ normalized principal component matrix. Let $\mathbf{p}_j$ be the j th column of $\mathbf{P}$, (i.e., the j th PC), for $j \in \{1, \dots, N\}$. Note that $\mathbf{p}_j^T \mathbf{p}_j = 1$, $\mathbf{p}_j^T \mathbf{p}_k = 0$ for $k \neq j$, and $\mathbf{p}_j^T \mathbf{1} = 0$ for $j \in \{1, \dots, N-1\}$ where $\mathbf{1} = \{1, \dots, 1\}$. Finally, we use $\mathbf{y}$ to denote the vector of N phenotypes.

THEOREM 1: Let $\gamma_j = \mathbf{p}_j^T \mathbf{y}$. We have $\sum_{i=1}^M (\mathbf{x}_i^T \mathbf{y})^2 = \sum_{j=1}^N \gamma_j^2 \lambda_j$, where $\lambda_j = d_j^2$ is the j th eigenvalue of $\mathbf{X}^T \mathbf{X}$.

The proof is given in Web Appendix A. As γ_j is an inner product between the (normalized) $\mathbf{p}_j$ and $\mathbf{y}$, it is easy to show that

$$\gamma_j = \sqrt{\sum_{k=1}^N (y_k - \bar{y})^2} \times \text{corr}(\mathbf{p}_{\cdot j}, \mathbf{y}), \quad (1)$$

where “corr” is the Pearson correlation coefficient. Thus γ_j is proportional to the correlation between the phenotype $\mathbf{y}$ and the j th PC. We emphasize that the correlation is a *sample* quantity, observable from the data. Similarly, each term $\mathbf{x}_i^T \mathbf{y}$ in the equality is proportional to the correlation between the genotype at SNP i and the phenotype. The importance of the result lies in the explicit connection between these M genotype-phenotype correlations to the N PC-phenotype correlations.

2.1 Relationship between Genomic Control and Principal Components

Here we obtain explicit results for the relevant test statistics applied in association mapping. Theorem 1 technically applies to score test statistics. However, we later demonstrate that the results apply to other common choices of test statistic.

Quantitative Traits—For continuous quantitative phenotype Y , we assume a simple linear regression model at each SNP i : $y_i = \beta_{0i} + \beta_{1i}x_{ij} + \epsilon_j$, $\epsilon_j \sim N(0, \sigma^2)$. To test an association of the SNP i with the phenotype, we can use the following score test statistic:

$$S_i = \frac{(\mathbf{x}_i^T \mathbf{y})^2}{\sum_{j=1}^N (y_j - \bar{y})^2 / N}. \quad (2)$$

By Theorem 1,

$$\frac{1}{M} \sum_{i=1}^M S_i = \left\{ M \sum_{j=1}^N (y_j - \bar{y})^2 / N \right\}^{-1} \sum_{j=1}^N \gamma_j^2 \lambda_j = \frac{N}{M} \sum_{j=1}^N \text{corr}^2(\mathbf{p}_{\cdot j}, \mathbf{y}) \lambda_j. \quad (3)$$

The observed mean of all score test statistics across the M SNPs is thus proportional to the sum of the squared PC-phenotype correlations multiplied by their respective eigenvalues.

In a justification of genomic control, it has been argued that under certain models of population stratification, the inflation of test statistics should be similar across all “null” SNPs (Devlin et al., 2001). This and related work (Devlin and Roeder, 1999) compared the sample median of test statistics to the chi-square median value, for an estimated inflation factor $\hat{\tau} = \text{median}(S)/0.456$, in order to be robust to outlying test statistics which presumably correspond to “alternative” SNPs. Other work (Reich and Goldstein, 2001; Devlin et al.,

2004) has used the sample mean $\hat{\tau} = \bar{S} = (1/M) \sum_{i=1}^M S_i$ directly. In practice, the median and mean typically give similar results, as the proportion of alternative SNPs is typically small.

Using either approach, for each i a new statistic $S'_i = S_i / \hat{\tau}$ is then compared to χ^2_1 .

To summarize, the results above provide a direct relationship between the mean version of the genomic control quantity $\hat{\tau}$ (left-hand side of (3)) and the PC-phenotype correlations and eigenvalues. This relationship is more than a simple curiosity. While it is known that distinct subpopulations can be represented using PCs (Price et al., 2008), we are not aware that a natural relationship has been previously established between the PCs and the testing

procedures. Moreover, (3) is exact (not based on expectations), holding for any $\mathbf{X}$ and $\mathbf{y}$, regardless of the underlying population substructure and the proportion of alternative SNPs. Thus the right-hand side of (3) is subject to the same sampling variation as $\bar{S}$. We also note that, to the extent that increases in $\bar{S}$ above 1 determine type I error inflation, the equation specifically highlights the terms $\text{corr}^2(\mathbf{p}_j, \mathbf{y}) \lambda_j$ as contributors to this inflation. Principal components contributing meaningfully to this inflation must have appreciable values for both λ_j and $\text{corr}^2(\mathbf{p}_j, \mathbf{y})$. Due to sampling variation, PCs will exhibit sample $\text{corr}^2(\mathbf{p}_j, \mathbf{y}) > 0$ for each j , even if the PCs are truly uncorrelated with the population from which $\mathbf{y}$ is drawn. The eigenvalues $\lambda_1, \dots, \lambda_{N-1}$ are also non-zero. Thus we must consider the magnitude of the terms, as well as sampling variation. Our general approach in the later sections will be to (i) re-rank the PCs by the terms $\text{corr}^2(\mathbf{p}_j, \mathbf{y}) \lambda_j$, (ii) test for the statistical significance of each of the terms, and (iii) control for stratification using only those PCs with significant terms.

Case-Control Traits—We now establish that the relationships described above apply to case-control studies, with $Y = 0$ and $Y = 1$ corresponding to controls and cases, respectively. We use the logistic regression model for SNP i : $\log [P(Y = 1) / \{1 - P(Y = 1)\}] = \beta_{0i} + \beta_{1i}x_{ij}$, which is conditional on the sampling scheme. With N_1 cases and N_0 controls, the score test

statistic is $S_i = (\mathbf{x}_i^T \mathbf{y})^2 / \{(N_0 N_1) / N^2\}$. It is simple to show that $(N_0 N_1) / N^2 = \sum_{j=1}^N (y_j - \bar{y})^2 / N$, and so by comparison with (2), we see that (3) directly applies.

We have now demonstrated a direct connection between the genomic control inflation factor and the PCs, but the two correction approaches are fundamentally different. In genomic control, the inflation of test statistics is effectively assumed to be constant across all null SNPs. However, it can be shown that regression analyses with the top PCs used as covariates is equivalent to adjusting the inflation effect of each SNP separately (see Web Appendix B for details).

2.2 EigenCorr : An Eigenvalue and Correlation-Based PC Selection Procedure

Using the result that the effect of $\mathbf{p}_j$ on $\bar{S}$ is proportional to $\gamma_j^2 \lambda_j$, we propose to select PCs based on the $\gamma_j^2 \lambda_j$, which we call the *EigenCorr* scores. To determine the impact of a given PC, we describe two procedures, which differ in their underlying assumptions.

1) EigenCorr1—We adopt the null hypothesis that the population correlation of the PCs and phenotypes is zero and that there is no population substructure. We are able to directly estimate the null distribution of EigenCorr scores using the Tracy-Widom approximation (Johnstone, 2001; Patterson et al., 2006) and the Fisher z-transformation applied to sample correlations (Fisher, 1921). That is, $\text{corr}(\mathbf{p}_j, \mathbf{y})$ approximately follows the distribution of $(e^{2Z} - 1) / (e^{2Z} + 1)$, where $Z \sim N(0, 1/(N - 3))$. Using these approximations to the distributions of (independent) γ_j^2 and λ_j , we obtain null distributions for $\gamma_j^2 \lambda_j$ by simulation, resulting in p -values for each EigenCorr score. The process proceeds sequentially. We simulate a random λ_1^* from the distribution of $(\xi T + \mu) \sum_{k=1}^{N-1} \lambda_k / (N - 1)$, where T is a Tracy-Widom random variable, and ξ and μ are as described in Web Appendix C. We then simulate γ_1^* using the Fisher z-transformation to obtain p -values for $\gamma_1^2 \lambda_1$. After excluding the first eigenvalue, we set $N = N - 1$, recompute μ and ξ , and follow the same procedure sequentially to obtain p -values for the remaining EigenCorr scores. Significant PCs are selected by the p -values, acknowledging multiple comparisons using the Benjamini-Hochberg false discovery rate (FDR) procedure (Benjamini and Hochberg, 1995).

2) EigenCorr2—In EigenCorr1, we assumed no population substructure. Although this assumption underlies Tracy-Widom testing (Patterson et al., 2006), we can relax the assumption, recognizing that significant eigenvalues alone are not sufficient to produce inflation of type I error. For EigenCorr2, we assume only that the PCs are uncorrelated with the population phenotype distribution. Here we treat the λ_j values as fixed, and use the Fisher z approximation to compute p -values for high values of γ_j^2 . Although EigenCorr2 is simpler than EigenCorr1, and is shown to perform well in later simulations, both approaches may have value in different situations. In particular, EigenCorr1 may have an advantage in situations where few eigenvalues are truly significant. For either approach, our experience indicates that a relatively small number of PCs will be chosen for stratification control, which is desirable for both computational and statistical simplicity.

In the current practice of PC-based stratification control, investigators are often concerned that the inclusion of SNPs in high linkage disequilibrium can produce misleading results. Thus many investigators “thin” out SNPs so that only a subset with lower correlations is used to generate PCs (e.g. Fellay et al., 2007), which is a special case of the *shrinkage PC* approach (Zou et al., 2010). We have shown (Web Appendix F) that EigenCorr procedures can be applied, but to the EigenCorr scores based on the weighted PCs. Further, simulations show that even when genomic control is performed using all SNPs, but PCs are calculated using a thinned set of SNPs, the approximate relationship still holds (Web Appendix G).

3. Simulations and Real Data Analysis

We investigated the performance of the proposed EigenCorr approach in applications to simulated data and two real GWAS datasets.

3.1 Simulation Studies

Simulation 1—We simulated 1000 samples from 5 subpopulations with 20,000 uncorrelated SNPs, with 210 samples from each of the first four subpopulations, and the remaining 160 samples from subpopulation 5. For each SNP, the overall minor allele frequency (MAF) was uniform from 0.05 to 0.5, and F_{st} uniform from 0.01 to 0.04. From these values, the MAF for each subpopulation was generated according to the Balding-Nichols model (Balding and Nichols, 1995). PC analysis showed that the top 4 PCs were significant according to the TW test, with $p_{.4}$ specifically distinguishing subpopulation 5 from the others. To simulate population stratification, we generated a disease phenotype from a logistic model with $\log(\text{odds ratio})=1.6$ between subpopulation 5 and the remaining samples. Therefore, increases in type I error resulting from population stratification arise entirely from the differing disease prevalence in the subpopulations. Figure 1 shows eigenvalues and EigenCorr scores of the first 10 PCs. The TW test selected the top 4 PCs as significant at $p < 0.01$, since its selection is entirely eigenvalue based, while only $p_{.4}$ was identified, correctly, by EigenCorr. This simulation is illustrative of the intended advantage of EigenCorr scores.

Simulation 2 simulations based on a real dataset—To investigate type I error and power for PC-based methods, we simulated phenotypes based on a real schizophrenia GWAS from the GAIN consortium [Version 2, Accession number: phs000021.v2.p1] (Sanders et al., 2008). In this manner we intended to reflect the genetic complexity encountered in real studies. The General Research Use (GRU) African American data consists of 1904 samples with 845,814 SNPs, and was downloaded from dbGap at NCBI (ncbi.nlm.nih.gov). After several data filtering steps (See Web Appendix D), we obtained a final data set with 1835 samples and 96,346 SNPs. TW testing identified 91 significant PCs with $p < 0.01$, which also corresponded to FDR control at 0.1. In contrast, the informal (but

widely used) “scree plot” method of simply inspecting eigenvalues would result in choosing two PCs (See Web Figure 1). A different version of the TW test is implemented in the GEM software (Luca et al., 2008), referred to hereafter as TW-GEM. TW-GEM selected 12 PCs at FDR level 0.1. In contrast to the TW test implemented in EigenSoft, TW-GEM estimates the effective number of markers only once and uses it for all PCs when computing TW statistics.

We used $p_{.1}$, $p_{.2}$, $p_{.5}$, and $p_{.10}$ to generate association of strata with a simulated phenotype. For quantitative traits, under the null hypothesis for SNP association, we simulated phenotypes according to $y_j = \eta_1 p_{j1} + \eta_2 p_{j2} + \eta_5 p_{j5} + \eta_{10} p_{j10} + \epsilon_j$. Here the ϵ_j were generated from $N(0, 1)$, and $\eta_1, \eta_2, \eta_5, \eta_{10}$ generated as independent and identically normally distributed so as to contribute the half of the variability of y . Note that for all simulations, the genotypes were fixed, but the phenotypes were simulated prospectively, producing variation in the PCs selected by EigenCorr. On average, EigenCorr1 and EigenCorr2 selected 3.52 and 3.51 PCs respectively out of the first 200 PCs at FDR level 0.1. These PCs mostly overlapped and reflected the true PC stratification. Similarly, for a dichotomous trait, we simulated case-control status from the model

$$\log \left[P(Y_j=1) / \{1 - P(Y_j=1)\} \right] = \eta_1 p_{j1} + \eta_2 p_{j2} + \eta_5 p_{j5} + \eta_{10} p_{j10},$$

where the coefficients $\eta_1, \eta_2, \eta_5, \eta_{10}$ were randomly generated from a normal distribution such that the variance of the log odds ratio equaled 4.0.

For the alternative hypothesis, given the genotype x at a causal SNP, quantitative traits were generated from $y_j = \beta x + \eta_1 p_{j1} + \eta_2 p_{j2} + \eta_5 p_{j5} + \eta_{10} p_{j10} + \epsilon_j$, where $\beta = 0.15$, and dichotomous traits were generated from

$$\log \left[P(Y_j=1) / \{1 - P(Y_j=1)\} \right] = \beta x + \eta_1 p_{j1} + \eta_2 p_{j2} + \eta_5 p_{j5} + \eta_{10} p_{j10},$$

with $\beta = \log(1.6) = 0.47$. The same η distributions used in the null models were applied.

To investigate the type I error and power, we ran 10,000 simulations for each model. In each null simulation a single SNP was chosen at random for investigation and p-value computation. Table 1 provides the empirical type I errors. We compared seven different approaches for PC selection: 1) no PCs for adjustment; 2) the four PCs that reflected the true population stratification; 3) the first two PCs, selected by scree plot examination; 4) the 91 PCs selected by the TW test; 5) the 12 PCs selected by TW-GEM; 6) the PCs selected by EigenCorr1, and 7) the PCs selected by EigenCorr2. The “true PC” approach can be viewed as a gold standard. *P*-values were computed from the score statistics for the additive genotype effect. The results from likelihood ratio testing were similar (See Web Table 1).

From the table we can see that the “no-adjustment” and “scree plot” methods result in improper control of type I error. As expected, use of the four known confounding PCs properly controlled type I error. The 91 top PCs from the TW test controlled the type I error well for the quantitative trait, although the model was highly over-parameterized. However, these 91 PCs results in somewhat inflated type I error for the dichotomous trait. In practice, investigators might be reluctant to fit so many covariates, but the absence of a principled alternate procedure based solely on eigenvalues makes it difficult to prescribe an alternative, when so many eigenvalues are clearly significant. The fewer PCs selected by the TW-GEM test controlled type I error in these data, as did the PCs from both EigenCorr methods. The statistical power of the models chosen by TW-GEM and both EigenCorr procedures were

comparable to the gold standard of known PCs, while the TW procedure resulted in somewhat reduced power.

The average estimates of the genetic effect β are shown in Web Table 2. For the quantitative traits, all seven methods gave essentially unbiased estimates. The estimates from the logistic regression, however, were biased downward with adjustment by zero PCs or by the two scree-based PCs. Estimates were upwardly biased when using the 91 TW PCs. The presence of bias and poorly-controlled type I error are well-known features of logistic regression if a large number of unnecessary (null) covariates are included in the analysis (Lubin, 1981).

To investigate the family-wise error rate (FWER) under more stringent testing thresholds, we used the same setup to run 1000 null whole-genome scans using all 810,264 SNPs. Here the minimum GWAS p-value in each simulation was compared to prespecified thresholds (Table 2). The threshold of 6.15×10^{-8} , corresponds to Bonferroni adjustment for FWER=0.05. The precise FWER-controlling threshold is difficult to predetermine, due to the effects of SNP correlation and the use of asymptotic p-values at extreme thresholds. Nonetheless, for each threshold, the "known PC" situation can serve as a gold-standard for comparison. Table 2 shows that the large number of PCs from the TW test inflates the FWER for dichotomous traits, doubling or tripling the error compared to using known PCs. The "no adjustment" and scree inspection methods failed to control FWER for both continuous and dichotomous traits, but because too few PCs are used. For the simulation setup, we find that EigenCorr1, EigenCorr2 and TW-GEM offer reasonable FWER control. However, the real-data example below shows that for some datasets TW-GEM can fail to detect PCs which have a substantial impact on type I error.

3.2 Real Data Analysis

We next applied the EigenCorr methods to two schizophrenia GWAS studies: (i) the CATIE dataset of Sullivan et al. (2008); (ii) the GAIN consortium dataset described earlier, using the actual study phenotypes. The CATIE analysis is described in detail here, with the GAIN analysis described in Web Appendix E.

3.3 The CATIE Schizophrenia Data

CATIE Schizophrenia GWAS data was obtained from the NIMH Genetic Repository (nimhgenetics.org), with 1,492 samples (741 cases and 751 controls) and 495,172 SNPs. After applying the same data filtering steps described in Web Appendix D, 1,439 samples remained, with 71,985 thinned SNPs for PC analysis. Web Figure 5 presents sample eigenvalues, their correlations with disease status, and the corresponding EigenCorr scores. Inspection of the scree plot suggested use of the top two PCs. At FDR = 0.1, the TW test selected 45 PCs. In contrast, TW-GEM selected only the first PC. Of the top 200 PCs, both EigenCorr1 and EigenCorr2 selected three PCs, with two of them overlapping.

We applied all the PC-selection approaches in Simulation 2, except the "true PC" approach, as the true confounding PCs are of course unknown. Figure 2 shows the $-\log_{10}(p\text{-values})$ (observed vs. expected) from the various approaches. The TW-GEM shows a large deviation from the diagonal, indicating that adjustment by only the first PC was not sufficient to control for stratification. In contrast, the plots for EigenCorr1 and EigenCorr2 suggest proper type I error control. PC 2 has the largest EigenCorr score, with a large eigenvalue and high correlation with the case-control phenotype. In addition, the EigenCorr methods successfully find a short list of SNPs that meet genome-wide significance.

4. Discussion

We have shown that the average inflation of test statistics caused by population stratification can be expressed in terms of genotype PCs, according to their eigenvalues and their correlations with the phenotype. PCs that are uncorrelated with the phenotype have a negligible effect on the inflation of test statistics. The explicit connection of the PCs to genomic control provides insight into the advantages of PC-based methods. Importantly, this natural motivation for PC-based control does not require any assumptions about the nature of population structure. In addition, the results show that PCs computed from genotypes are “natural” variables to consider, even though discrete genotypes are far removed from the normality assumptions that underlie classical multivariate analysis. In addition to the statistical advantages of EigenCorr PC selection, the small number of PCs typically selected has a distinct computational advantage when applied at large scales, for example with eQTL analysis.

For simulations and real data, we have shown that the TW test selects too many PCs. In practice, investigators typically choose few PCs, but informal methods such as scree plot inspection can select too few. Similarly, the TW-GEM test clearly selected too few PCs for CATIE, because it does not approach PC testing sequentially, and can be less powerful for detecting important but lower-ranked PCs. We believe the EigenCorr approaches offer a principled alternative, and tend to correctly select the few most important PCs.

In our analyses, EigenCorr1 and EigenCorr2 produced similar results. EigenCorr1 offers a potential conceptual advantage, in that the EigenCorr score is used directly as a statistic, and lower-ranked PCs are given less emphasis. However, when it is clear that many eigenvalues are highly significant, the null PC assumption may not be realistic, and EigenCorr2 may be preferred.

Other investigators have used the phenotype to identify or construct genotype-based covariates. Epstein et al. (2007) described a general stratification score approach, using partial least squares (PLS) of phenotype on a number of markers. Lee et al. (2008) pointed out that PLS can result in overfitting to the phenotype and reduce power. However, Epstein et al. (2007) and the rejoinder (Epstein et al., 2008) provide important clarification that desirable stratification control should explain some phenotype variability. This fact was also recognized by Kimmel et al. (2007), who described initial correction of phenotype by a small number of PC clusters before association testing. Zhao et al. (2009) described a propensity score using (genetic) covariates to predict genotype at a test locus before inclusion in a phenotype-genotype model. This approach, although implemented with relatively few genetic covariates (serving a role analogous to our PCs), potentially avoids overfitting in selecting covariates. However, it is potentially susceptible to inclusion of an excessive number of genetic covariates, which can present a problem for logistic regression.

We view the EigenCorr approach as generally similar to that advocated by Epstein et al. (2007), using PC-based phenotype adjustment in the form of regression covariates. However, the EigenCorr procedure, with a fixed number of PCs to choose from, is intentionally less flexible than procedures such as PLS, and FDR-based testing provides a natural penalty against overfitting. Moreover, in contrast to other procedures, the EigenCorr motivation and approach is explicitly connected to the source of test statistic inflation. The fact that genotype must also be associated with population stratum in order to create confounding is also implicit in EigenCorr, because each informative marker has an influence on, and will be associated with, at least one eigenvector. We believe that EigenCorr offers an efficient filter to identify the confounding variables of greatest influence.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Support was provided in part by NIH R01GM074175, EPA R832720, and a Gillings Innovation Laboratory award. The authors gratefully acknowledge the comments and suggestions of the reviewers and editors, which greatly improved the manuscript. GAIN Dataset were obtained from the GAIN Database found at <http://view.ncbi.nlm.nih.gov/dbgap-controlled-through/dbGaP> accession number phs000021.v2.p1. Samples and associated phenotype data for the Linking Genome-Wide Association Study of Schizophrenia were provided by P. Gejman. The principal investigators of the CATIE trial were Jeffrey A. Lieberman, M.D., T. Scott Stroup, M.D., M.P.H., and Joseph P. McEvoy, M.D. The CATIE trial was funded by a grant from the National Institute of Mental Health (N01 MH900001) along with MH074027 (PI PF Sullivan). Genotyping was funded by Eli Lilly and Company.

References

- Balding D, Nichols R. A method for quantifying differentiation between populations at multi-allelic loci and its implications for investigating identity and paternity. *Genetica*. 1995; 96:3–12. [PubMed: 7607457]
- Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society*. 1995; 57:289–300. Series B (Methodological)
- Devlin B, Bacanu S, Roeder K. Genomic Control to the extreme. *Nature Genetics*. 2004; 36:1129–1130. [PubMed: 15514657]
- Devlin B, Roeder K. Genomic Control for Association Studies. *Biometrics*. 1999; 55:997–1004. [PubMed: 11315092]
- Devlin B, Roeder K, Wasserman L. Genomic Control, a New Approach to Genetic-Based Association Studies. *Theoretical Population Biology*. 2001; 60:155–166. [PubMed: 11855950]
- Epstein M, Allen A, Satten G. A simple and improved correction for population stratification in case-control studies. *The American Journal of Human Genetics*. 2007; 80:921–930.
- Epstein M, Allen A, Satten G. Response to lee et al. *The American Journal of Human Genetics*. 2008; 82:526–528.
- Fellay J, Shianna K, Ge D, Colombo S, Ledergerber B, Weale M, Zhang K, Gumbs C, Castagna A, Cossarizza A, et al. A Whole-Genome Association Study of Major Determinants for Host Control of HIV-1. *Science*. 2007; 317:944–947. [PubMed: 17641165]
- Fisher R. On the probable error of a coefficient of correlation deduced from a small sample. *Metron*. 1921; 1:3–32.
- Johnstone I. On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics*. 2001; 29:295–327.
- Kimmel G, Jordan M, Halperin E, Shamir R, Karp R. A randomization test for controlling population stratification in whole-genome association studies. *The American Journal of Human Genetics*. 2007; 81:895–905.
- Lee S, Sullivan P, Zou F, Wright F. Comment on a simple and improved correction for population stratification. *The American Journal of Human Genetics*. 2008; 82:524–526.
- Lubin J. An empirical evaluation of the use of conditional and unconditional likelihoods for case-control data. *Biometrika*. 1981; 68:567–571.
- Luca D, Ringquist S, Klei L, Lee A, Gieger C, Wichmann H, Schreiber S, Krawczak M, Lu Y, Styche A, et al. On the use of general control samples for genome-wide association studies: genetic matching highlights causal variants. *The American Journal of Human Genetics*. 2008; 82:453–463.
- Marchini J, Cardon L, Phillips M, Donnelly P. The effects of human population structure on large genetic association studies. *Nature Genetics*. 2004; 36:512–517. [PubMed: 15052271]
- Patterson N, Price A, Reich D. Population structure and eigenanalysis. *PLoS Genet*. 2006; 2:e190. [PubMed: 17194218]

- Price A, Butler J, Patterson N, Capelli C, Pascali V, Scarnicci F, Ruiz-Linares A, Groop L, Saetta A, Korkolopoulou P, et al. Discerning the ancestry of European Americans in genetic association studies. *PLoS Genet.* 2008; 4:e236. [PubMed: 18208327]
- Price A, Patterson N, Plenge R, Weinblatt M, Shadick N, Reich D. Principal components analysis corrects for stratification in genome-wide association studies. *Nature Genetics.* 2006; 38:904–909. [PubMed: 16862161]
- Reich D, Goldstein D. Detecting association in a case-control study while correcting for population stratification. *Genetic Epidemiology.* 2001; 20:4–16. [PubMed: 11119293]
- Sanders A, Duan J, Levinson D, Shi J, He D, Hou C, Burrell G, Rice J, Nertney D, Olincy A, et al. No significant association of 14 candidate genes with schizophrenia in a large European ancestry sample: implications for psychiatric genetics. *American Journal of Psychiatry.* 2008; 165:497–506. [PubMed: 18198266]
- Sullivan P, Lin D, Tzeng J, van den Oord E, Perkins D, Stroup T, Wagner M, Lee S, Wright F, Zou F, et al. Genomewide association for schizophrenia in the CATIE study: results of stage 1. *Molecular psychiatry.* 2008; 13:570–584. [PubMed: 18347602]
- Zhao H, Rebbeck T, Mitra N. A propensity score approach to correction for bias due to population stratification using genetic and non-genetic factors. *Genetic epidemiology.* 2009; 33:679–690. [PubMed: 19353632]
- Zou F, Lee S, Knowles M, Wright F. Quantification of Population Structure Using Correlated SNPs by Shrinkage Principal Components. *Human heredity.* 2010; 70:9–22. [PubMed: 20413978]

Figure 1. Illustration of the EigenCorr scores for simulation1. The left panel presents the first 10 eigenvalues and the right panel presents the first 10 EigenCorr scores. The importance of PC4 is clearly indicated by the EigenCorr score.

Figure 2. Plots of observed vs. expected p -values ($-\log_{10}$ scale) for the CATIE data. The dashed lines indicate 95% prediction bands.

Table 1

Performance of the methods for 10,000 GWAS simulations, evaluated at a null or alternative SNP. Values in the table represent Type I error (for the null simulations) or Power (for the alternative simulations) from the score test. The simulation setups are described in “Simulations and Real Data Analysis”.

Nominal Significance α	No Adjustment	Known Confounding PCs	Scree Method	TW	TW-GEM	EigenCorr1	EigenCorr2
Quantitative Trait							
NULL							
0.05	0.1661	0.0501	0.0789	0.0499	0.0497	0.0508	0.0508
10^{-2}	0.0782	0.0104	0.0229	0.0106	0.0099	0.0104	0.0105
Alternative							
10^{-2}	0.6871	0.7221	0.7110	0.6983	0.7205	0.7201	0.7191
10^{-4}	0.3901	0.3743	0.3777	0.3388	0.3713	0.3708	0.3703
10^{-6}	0.1877	0.1369	0.1487	0.1155	0.1345	0.1346	0.1334
Case Control Trait							
NULL							
0.05	0.1642	0.0490	0.0704	0.0561	0.0483	0.0492	0.0500
10^{-2}	0.0770	0.0095	0.0188	0.0118	0.0099	0.0096	0.0097
Alternative							
10^{-2}	0.7260	0.8532	0.7835	0.8464	0.8528	0.8524	0.8519
10^{-4}	0.4593	0.6315	0.5152	0.6193	0.6269	0.6294	0.6293
10^{-6}	0.2611	0.3943	0.2840	0.3832	0.3909	0.3917	0.3915

Table 2

Family Wise Error rates (FWER) for the minimum score statistic p-values in 1,000 simulated whole genome scans, with population structure as described in the text. Values represent FWER when applying significance level α to each of the 810, 264 SNPs. A significance level $\alpha = 6.17 \times 10^{-8}$ corresponds to Bonferroni control of FWER ≤ 0.05 .

Nominal Significance α	No Adjustment	Known Confounding PCs	Scree Method	TW	TW-GEM	EigenCorr1	EigenCorr2
Quantitative Trait							
10^{-6}	0.9870	0.4460	0.8600	0.4460	0.4570	0.4450	0.4450
10^{-7}	0.9470	0.0620	0.5600	0.0680	0.0620	0.0660	0.0660
6.17×10^{-8}	0.9230	0.0380	0.5150	0.0400	0.0420	0.0400	0.0400
Case Control Trait							
10^{-6}	0.9970	0.4850	0.9120	0.7140	0.4840	0.4910	0.4910
10^{-7}	0.9690	0.0640	0.6800	0.1350	0.0620	0.0630	0.0630
6.17×10^{-8}	0.9620	0.0290	0.6410	0.0850	0.0390	0.0330	0.0330