RL-SKAT: An Exact and Efficient Score Test for Heritability and Set Tests

Regev Schweiger,^{*,1} Omer Weissbrod,[†] Elinor Rahmani,^{*} Martina Müller-Nurasyid,^{*,§,***} Sonja Kunze,^{††,##}
Christian Gieger,^{††,##} Melanie Waldenberger,^{**,††,##} Saharon Rosset,^{§§} and Eran Halperin^{***,†††,***}

^{*}Blavatnik School of Computer Science, Tel Aviv University, 6997801 Israel, [†]Department of Epidemiology, Harvard T.H. Chan School of Public Health, Boston, Massachusetts 02115, [‡]Institute of Genetic Epidemiology, ^{††}Institute of Epidemiology II, and ^{##}Research Unit of Molecular Epidemiology, Helmholtz Zentrum München, German Research Center for Environmental Health, 85764 Neuherberg, Germany, [§]Department of Medicine I, Ludwig-Maximilians-Universität, 80336 Munich, Germany, ^{**}German Centre for Cardiovascular Research, Partner Site Munich Heart Alliance, 80336 Munich, Germany, ^{§§}School of Mathematical Sciences, Department of Statistics, Tel Aviv University, 69978 Israel, and ^{***}Department of Computer Science, ^{†††}Department of Anesthesiology and Perioperative Medicine, and ^{***}Department of Human Genetics, University of California, Los Angeles, California 90095

ORCID IDs: 0000-0003-2450-8901 (R.S.); 0000-0002-9017-2070 (E.R.); 0000-0001-6986-9554 (C.G.); 0000-0003-0583-5093 (M.W.); 0000-0002-4458-9545 (S.R.); 0000-0002-2373-3691 (E.H.)

ABSTRACT Testing for the existence of variance components in linear mixed models is a fundamental task in many applicative fields. In statistical genetics, the score test has recently become instrumental in the task of testing an association between a set of genetic markers and a phenotype. With few markers, this amounts to set-based variance component tests, which attempt to increase power in association studies by aggregating weak individual effects. When the entire genome is considered, it allows testing for the heritability of a phenotype, defined as the proportion of phenotypic variance explained by genetics. In the popular score-based Sequence Kernel Association Test (SKAT) method, the assumed distribution of the score test statistic is uncalibrated in small samples, with a correction being computationally expensive. This may cause severe inflation or deflation of *P*-values, even when the null hypothesis is true. Here, we characterize the conditions under which this discrepancy holds, and show it may occur also in large real datasets, such as a dataset from the Wellcome Trust Case Control Consortium 2 ($n = 13,950$) study, and, in particular, when the individuals in the sample are unrelated. In these cases, the SKAT approximation tends to be highly overconservative and therefore underpowered. To address this limitation, we suggest an efficient method to calculate exact *P*-values for the score test in the case of a single variance component and a continuous response vector, which can speed up the analysis by orders of magnitude. Our results enable fast and accurate application of the score test in heritability and in set-based association tests. Our method is available in <http://github.com/cozygene/RL-SKAT>.

KEYWORDS statistical genetics; SKAT; heritability; set-tests

THE variance component model is a well-established statistical framework used in many scientific fields. Testing for an association between several explanatory variables and a univariate response produces a variety of useful applications. For example, in metagenomics, an association is tested between a phenotype (e.g., body mass index, blood glucose

levels, blood lipid levels, etc.) and the relative abundance counts of the measured species (Zhao *et al.* 2015).

In statistical genetics, testing for an association between a set of genetic markers and a phenotype, such as a disease or a trait, is a fundamental task. Since studies to detect genetic signals are often underpowered, even with large datasets becoming available, the common approach to help alleviate this issue is grouping together genetic markers and testing them jointly. Grouping genetic markers is commonly implemented under the framework of variance component models. In addition to association testing, this framework can be used to answer several questions, such as estimation of the underlying heritability of a phenotype (Kang *et al.* 2010); estimating the uncertainty of such estimation (Furlotte *et al.*

Copyright © 2017 by the Genetics Society of America

doi: <https://doi.org/10.1534/genetics.117.300395>

Manuscript received July 7, 2017; accepted for publication September 24, 2017; published Early Online October 12, 2017.

Supplemental material is available online at www.genetics.org/lookup/suppl/doi:10.1534/genetics.117.300395/-DC1.

[†]Corresponding author: Blavatnik School of Computer Science, Schreiber Building, Tel Aviv University, P.O. Box 39040, Ramat Aviv, Tel Aviv, 6997801 Israel. E-mail: schweiger@post.tau.ac.il

2014; Schweiger *et al.* 2016, 2017); phenotype prediction (Hayes *et al.* 2001), and more.

We consider two main scenarios in which such tests are performed: (i) a single phenotype, many sets of markers and (ii) many phenotypes, a single set of markers. Scenario (i) is common in set-testing, where relatively few markers are tested jointly. This is particularly useful in the case of rare variants, which are increasingly available for study using sequencing technologies, and which constitute a large part of human genetic variability. In such studies, a single phenotype is often tested against several sets of markers (for example, all rare variants in a single gene), because single-marker tests are often underpowered. Scenario (ii) occurs when studying heritability, defined as the proportion of phenotypic variance explained by genetics. Here, the tested markers are commonly the entire set of genotyped or sequenced single-nucleotide polymorphism (SNP) variants, or large portions of the genome (defined by, *e.g.*, chromosome or functional annotation), and they are often tested against many (*e.g.*, thousands) of phenotypes. Such phenotypes could be expression profiles of genes (Price *et al.* 2011; Wright *et al.* 2014; Lloyd-Jones *et al.* 2017), methylation levels across of various methylation sites in the DNA (Quon *et al.* 2013; Van Dongen *et al.* 2016) or neuroimaging measurements (Ganjgahi *et al.* 2015; Ge *et al.* 2015).

Within the variance components framework, a common approach for association testing is the score test. It is used, for example, for testing the heritability of morphometric measurements derived from brain structural MRI scans (Ge *et al.* 2015) and on fractional anisotropy measures in subjects from the Genetics of Brain Structure study (Ganjgahi *et al.* 2015).

The main popular alternative to the score test is the generalized likelihood ratio (LR) test, *e.g.*, as implemented by GCTA, a popular software package for heritability estimation (Yang *et al.* 2011). Both the score test and the LR test are based on properties of the likelihood function. The LR test statistic is calculated from the likelihood of the best fitting model across different heritability values, and from the likelihood of the model corresponding to no heritability. Conversely, the score test is based on the derivative of the likelihood function at the point corresponding to zero association, and testing if it is significantly nonzero. Compared with the LR test, the score test is often advantageous as it requires parameter estimation only for the null model, whereas the LR test requires parameter estimation for both the null and the alternative model. Additionally, the score test is the locally most powerful test; see Lippert *et al.* (2014) for a thorough comparison between the two tests, mainly in the context of set testing.

The Sequence Kernel Association Test (SKAT) (Wu *et al.* 2011) has become the standard score-based test in statistical genetics and in metagenomics (Zhao *et al.* 2015), in large part due to its computational tractability. One of its merits is that it does not rely on the asymptotic distribution of the score test statistic, instead specifying a nonasymptotic distribution for the statistic under the null hypothesis of no asso-

ciation. However, it has been observed that this distribution may be inaccurate. In the SKAT-O extension (Lee *et al.* 2012), a resampling-based moment-matching correction is suggested. An adaptive permutation testing procedure is suggested in Hasegawa *et al.* (2016). Chen *et al.* (2016) provide a method for calculating exact *P*-values; however, their method may be significantly slower than that of SKAT, as it requires the eigendecomposition of a full rank square matrix, whose computational complexity is typically cubic in the sample size, for each distinct response variable (*e.g.*, phenotype), or each set of explanatory variables (*e.g.*, SNP set). Finally, in these works, it is reported that this discrepancy occurs mainly in studies having a small sample size, and it is currently unclear to which extent the *P*-values of SKAT are calibrated for large sample sizes.

Here, we undertake a thorough analysis of the null distribution of the score test statistic, and its discrepancy under the SKAT approximation. We suggest a practical way to quantify this discrepancy, and show that such discrepancies may occur even at large sample sizes. We show that a discrepancy is expected when the number of markers is comparable to or larger than the number of individuals, and when the individuals are relatively unrelated. In particular, in addition to such inaccuracies occurring in tests of sets of rare-variants in small samples, we conclude that they may also occur in large scale heritability studies. We further suggest a computational method, Recalibrated Lightweight SKAT (RL-SKAT), that allows exact *P*-value computation while maintaining computation time as in SKAT; in particular, for multiple phenotypes tested against the same marker set, only a single eigendecomposition is required. Finally, we demonstrate and validate our results on two real datasets, a large dataset from the Wellcome Trust Case Control Consortium 2 (International Multiple Sclerosis Genetics Consortium *et al.* 2011) (WTCCC2) study and the Cooperative health research in the Region of Augsburg (KORA) study (Holle *et al.* 2005) dataset.

Materials and Methods

We begin by reviewing the score test, as defined by the SKAT method (Wu *et al.* 2011) [see also the Supplementary Information in Lippert *et al.* (2014) for an excellent review]. We focus here on continuous phenotypes, and on the case of a single variance component; for other cases, see the *Discussion* below.

The variance components model

We consider the following standard variance components model [see Searle *et al.* (2009) for a detailed review]. Let n be the number of observations, and $\mathbf{y}$ be a $n \times 1$ vector of responses. Let $\mathbf{X}$ be a $n \times p$ design matrix of p covariates, associated with fixed effects (possibly including an intercept vector $\mathbf{1}_n$ as a first column, as well as other covariates) and let $\boldsymbol{\beta}$ be a $p \times 1$ vector of fixed effects. Finally, let $\mathbf{K}$ be a kernel matrix, which, in a kernel-based method such as SKAT, can be taken to be any symmetric positive-definite matrix that

encodes similarity between individuals. Then, $\mathbf{y}$ is assumed to follow:

$$\mathbf{y} \sim \mathcal{N}(\mathbf{X}\boldsymbol{\beta}, \sigma_g^2 \mathbf{K} + \sigma_e^2 \mathbf{I}_n), \quad (1)$$

The fixed effects $\boldsymbol{\beta}$ and the coefficients σ_g^2 and σ_e^2 are the parameters of the model.

In the context of statistical genetics, $\mathbf{y}$ is a vector of phenotype measurements for each individual, and $\mathbf{X}$ is a matrix of covariates (often including an intercept, sex, age, *etc.*). Let $\mathbf{Z}$ be a $n \times m$ standardized (*i.e.*, columns have zero mean and unit variance) genotype matrix containing the m SNPs we test. The common choice for $\mathbf{K}$ is a weighted dot product of the genetic markers (Yang *et al.* 2010); formally, define $\mathbf{K} = \mathbf{Z}\mathbf{W}\mathbf{Z}^\top$, where $\mathbf{W}$ is a non-negative $m \times m$ diagonal matrix assigning a weight per SNP. A standard choice is the uniform $\mathbf{W}_{i,i} = 1/m$ [see Wu *et al.* (2011) for a discussion]. The narrow-sense heritability due to genotyped common SNPs is defined as the proportion of total variance explained by genetic factors (Visscher *et al.* 2008):

$$h^2 = \frac{\sigma_g^2}{\sigma_g^2 + \sigma_e^2}. \quad (2)$$

The score test

Under the above model, evaluating whether the tested covariates influence the response, while adjusting for additional covariates, corresponds to testing the null hypothesis $\sigma_g^2 = 0$. SKAT tests this hypothesis with a variance component score test in the corresponding mixed model. Specifically, the score statistic in the single-kernel case is obtained from the derivative of the restricted likelihood, discarding terms that are constant with respect to $\mathbf{y}$ (Lippert *et al.* 2014):

$$Q(\mathbf{y}) = \mathbf{y}^\top \mathbf{S} \mathbf{K} \mathbf{S} \mathbf{y} \quad (3)$$

where $\mathbf{S} = \mathbf{I}_n - \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$ is the projection matrix to the subspace orthogonal to the covariates $\mathbf{X}$. For clarity of presentation, we will divide the statistic by σ_e^2 . Then,

Proposition 1: Let $\{\phi_i\}$ be the eigenvalues of $\mathbf{S} \mathbf{K} \mathbf{S}^\top$ and be $\chi_{1,i}^2$ are i.i.d. random variables distributed chi-square with one degree of freedom. Then,

$$Q/\sigma_e^2 \sim \sum_{i=1}^n \phi_i \chi_{1,i}^2. \quad (4)$$

The proof of Proposition 1, as well as all proofs below, are deferred to the Supplemental Material in [File S1](#).

The exact distribution of the score test statistic

The above derivation is exact whenever σ_e^2 is known. However, in practice, σ_e^2 is not known and needs to be estimated from the data; most often, from the single response vector we are testing. In practice, σ_e^2 is replaced with its restricted maximum likelihood (REML) estimate. The REML estimate is

simply the corrected mean of the squared entries of the phenotype, after regressing out the covariates and using $\mathbf{S}^\top \mathbf{S} = \mathbf{S}$:

$$\hat{\sigma}_e^2(\mathbf{y}) = \frac{\|\mathbf{S}\mathbf{y}\|^2}{n-p} = \frac{\mathbf{y}^\top \mathbf{S} \mathbf{y}}{n-p}. \quad (5)$$

We note that sometimes the ML estimate $\mathbf{y}^\top \mathbf{S} \mathbf{y}/n$ is used, or just $\mathbf{y}^\top \mathbf{S} \mathbf{y}$; as this only introduces a multiplicative constant, we use the unbiased REML estimate for simplicity of presentation later. The statistic Q and $\hat{\sigma}_e^2$, are, in fact, dependent random variables. Therefore, the assumed distribution of $Q/\hat{\sigma}_e^2$ (described in Proposition 1) does not hold when substituting σ_e^2 with its estimate, $\hat{\sigma}_e^2$. In Zhang and Lin (2003), Liu *et al.* (2007, 2008), and Wu *et al.* (2011), this substitution is justified by the claim that the (restricted) ML estimator $\hat{\sigma}_e^2$ is consistent, and may therefore be substituted by its true value for a large enough sample size, n . However, this argument does not take into consideration the dependency between Q and $\hat{\sigma}_e^2$. Also, as shown below, this distribution might not hold in realistic settings. In Chen *et al.* (2016), this discrepancy is reported for small samples, and an exact distribution is derived for the statistic $Q/\hat{\sigma}_e^2$, and for any $n, \mathbf{K}$ and $\mathbf{X}$, which we review here:

Proposition 2: The distribution of $Q/\hat{\sigma}_e^2$ may be modeled as a ratio of quadratic forms of normal variables. In particular, if $\mathbf{z} \sim \mathcal{N}(\mathbf{0}_n, \mathbf{I}_n)$, then

$$\frac{Q}{\hat{\sigma}_e^2} \stackrel{d}{=} (n-p) \cdot \frac{\mathbf{z}^\top \mathbf{S} \mathbf{K} \mathbf{S} \mathbf{z}}{\mathbf{z}^\top \mathbf{S} \mathbf{z}} \quad (6)$$

Assessing the discrepancy

While noted in the literature (Zhao *et al.* 2015; Chen *et al.* 2016), the above discrepancy is reported for small samples only. However, as we show now, it may occur also when the number of individuals is large. We give a qualitative measure for when to expect large discrepancies between the asymptotic approximation of a weighted mixture of chi-squares and the exact distribution.

In the Supplemental Material in [File S1](#), it is shown that the distributions of Q/σ_e^2 and $Q/\hat{\sigma}_e^2$ have the same means, but that $\text{Var}(Q/\sigma_e^2) > \text{Var}(Q/\hat{\sigma}_e^2)$, *i.e.*, the latter has a smaller variance. We can further quantify the ratio between the variances as an indicator to the discrepancy between the distributions.

Proposition 3: Denote the eigenvalues of $\mathbf{S} \mathbf{K} \mathbf{S}$ by $\phi_1, \dots, \phi_n$, and note that there are at most $n-p$ nonzero eigenvalues ϕ_i . Denote the first two sample moments of the eigenvalues by $\bar{\phi} = \sum_{i=1}^n \phi_i / (n-p)$ and $\bar{\phi}^2 = \sum_{i=1}^n \phi_i^2 / (n-p)$. Denote the empirical variance of the eigenvalues by $\sigma^2(\phi) = \bar{\phi}^2 - (\bar{\phi})^2$. Then,

$$R := \frac{\text{Var}(Q/\sigma_e^2)}{\text{Var}(Q/\hat{\sigma}_e^2)} = \frac{n-p+2}{n-p} \cdot \left(\left(\frac{\sigma(\phi)}{\bar{\phi}} \right)^{-2} + 1 \right) \quad (7)$$

The expression $\sigma(\phi)/\bar{\phi}$ is the (sample) coefficient of variation (CV) of the eigenvalues—a unitless, relative measure of their

Table 1 Performance summary

Scenario	Algorithm	Exact?	Preprocessing	Calculating $Q/\hat{\sigma}_e^2$	Calculating P -value
Heritability	SKAT	Approximate	$O(np^2 + n^2p + n^3)$	$O(n^2)$	$O(n)$
	MiRKAT	Exact	$O(np^2 + n^2p)$	$O(n^2)$	$O(n^3)$
	RL-SKAT	Exact	$O(np^2 + n^2p + n^3)$	$O(n^2)$	$O(n)$
Set-testing	SKAT	Approximate	$O(np^2 + nmp + nm^2)$	$O(n(m+p))$	$O(n)$
	MiRKAT	Exact	$O(np^2 + nmp)$	$O(n(m+p))$	$O(n^3)$
	RL-SKAT	Exact	$O(np^2 + nmp + n(m+p)^2)$	$O(n(m+p))$	$O(n)$

Comparison of the different approaches for P -value calculation discussed. RL-SKAT achieves accuracy while remaining computationally efficient.

dispersion. Therefore, the ratio becomes larger when the CV is smaller. Also, as noted above, since the approximation wrongly ignores the dependency between the statistic Q and $\hat{\sigma}_e^2$, we expect the discrepancy to grow larger as the correlation between Q and $\hat{\sigma}_e^2$ increases. We therefore examine this correlation as an additional measure of this discrepancy.

Proposition 4: Let $\sigma(\phi)/\bar{\phi}$ be the CV of the eigenvalues as above. Then,

$$\text{Corr}(Q, \hat{\sigma}_e^2) = \left(\left(\frac{\sigma(\phi)}{\bar{\phi}} \right)^2 + 1 \right)^{-1/2} \quad (8)$$

This again demonstrates that CV affects discrepancy—the correlation becomes stronger when the CV is smaller. When $CV \ll 1$, for example, when $\mathbf{K} \approx \mathbf{I}_n$, we have $R \gg 1$ and $R \approx (n-p+2)/(n-p) \cdot (1+n/m)/(n/m)$. Conversely, when $CV \gg 1$, we have $R \approx 1$ and $\text{Corr}(Q, \hat{\sigma}_e^2) \approx 1/CV$. This also gives the variance ratio as the function of the correlation as

$$R = \frac{n-p+2}{n-p} \cdot \frac{1}{1 - \text{Corr}^2(Q, \hat{\sigma}_e^2)}. \quad (9)$$

To summarize, the discrepancy is strong when the eigenvalues are more uniformly dispersed, and is weak when they have large variability. The dispersion of the eigenvalues of a kinship matrix has been previously shown to be related to the uncertainty in estimation of heritability: In Visscher and Goddard (2015), it is shown that the asymptotic variance of the heritability REML estimator decreases with the variance of the entries of the kinship matrix, and with the variance of the eigenvalues. In Schweiger *et al.* (2016), this result is shown without assumptions of asymptotics.

Examples: We now employ Propositions 3 and 4 to analyze several interesting examples in a genetic context. For simplicity, in the following, we use $\mathbf{X} = \mathbf{0}$, so that $p = 0$ and $\mathbf{S} = \mathbf{I}_n$.

Completely unrelated cohort: Suppose the cohort contains completely unrelated individuals; then, $\mathbf{K} = \mathbf{I}_n$. Thus, $\phi_1 = \dots = \phi_n = 1$, so $R = \infty$, $\text{Corr}(Q, \hat{\sigma}_e^2) = 1$, and $Q/\hat{\sigma}_e^2$ is the constant n . Compare this to the case where σ_e^2 is known; then, it can be easily seen that $Q/\sigma_e^2 \sim \chi_n^2$. Therefore, the mean is the same but the variance vanishes completely.

Rank-one kinship matrix: Consider the case of a simple burden test (Lee *et al.* 2012): if we assume the random effects $\mathbf{s}$ of all SNPs are identical, the burden test becomes equivalent

to the score test with $\mathbf{K} = \mathbf{u}\mathbf{u}^\top$, where $\mathbf{u} = \mathbf{Z}\mathbf{1}_m$. Alternatively, consider the extreme case, where all the individuals are identical: $\mathbf{K} = \mathbf{1}\mathbf{1}^\top$ (while unlikely in human, this could be approximately true in studies of plants, yeast, *etc.*). In both these cases, there is a single nonzero eigenvalue: $\phi_2 = \dots = \phi_n = 0$, which gives $R \approx 1$ and $\text{Corr}(Q, \hat{\sigma}_e^2) = (\phi_1/n)/\sqrt{\phi_1^2/n} = 1/\sqrt{n}$; that is, with large enough sample size, we expect the correlation to be effectively zero, and the SKAT mixture approximation to hold well.

A full rank kinship matrix: Assume the matrix $\mathbf{Z}$ contains $m > n$ SNPs in linkage equilibrium, where each column was mean-centered and normalized to have unit variance. Choosing the linear kernel $\mathbf{K} = \mathbf{Z}\mathbf{Z}^\top/m$, we follow Patterson *et al.* (2006) in modeling $\mathbf{Z}$ as a matrix of random standard normal variables, from which it follows that $\mathbf{K}$ is a Wishart matrix. The limit distribution of the density of the eigenvalues of $\mathbf{K}$ is specified by the Marčenko-Pastur distribution (Marcenko and Pastur 1967), with its first two moments known to be 1 and $1 + n/m$. Under this approximation, $\bar{\phi} \approx 1$, $\phi^2 \approx 1 + n/m$, $\sigma^2(\phi) \approx n/m$, $R \approx (n-p+2)/(n-p) \cdot (1+n/m)/(n/m)$ and $\text{Corr}(Q, \hat{\sigma}_e^2) \approx 1/\sqrt{1+n/m}$. When $m \gg n$, as is often the case, $R \gg 1$ and $\text{Corr}(Q, \hat{\sigma}_e^2) \approx 1$. This shows that, for a

Figure 1 Statistic distribution. Results of the WTCCC2 data analysis, presented by quantile-quantile plots of the $-\log_{10}(p)$ -values for heritability significance of 10,000 random phenotypes drawn under the null distribution. Significant deviation from the black line indicates a deflation arising from an inaccurate null distribution. Calculation under the assumption of a weighted mixture of chi-square distributions, gives deflated P -values and potentially creating false negatives. Using the correct distribution, as implemented in RL-SKAT, results in calibrated P -values.

Figure 2 Heritability study. Histograms of the P -values of the studied phenotypes in the KORA dataset, as calculated by the accurate method (left) and the inaccurate method (right). Histograms are shown in log-scale, and are capped at $p = 10^{-8}$ for clarity of presentation. SKAT tends to severely deflate P -values, which are small according to the accurate calculation, leading to a severe loss of power.

large class of kinship matrices, we would expect the SKAT mixture approximation to hold poorly.

SNP set: Now, consider the case of set-testing, where $\mathbf{Z}$ is a normalized matrix of $m < n$ SNPs in linkage equilibrium. Following the modeling above, we have again $R \approx (n - p + 2)/(n - p) \cdot (1 + n/m)/(n/m)$ and $\text{Corr}(Q, \hat{\sigma}_e^2) \approx 1/\sqrt{n/m}$; when $m \ll n$, $R \approx 1$ and $\text{Corr}(Q, \hat{\sigma}_e^2) \approx \sqrt{m/n} \ll 1$, and thus expecting a good approximation by the mixture. This perhaps shows why the SKAT mixture approximation was considered good in the context of set-tests, when few variants or a large sample is considered. This also shows why, in small samples, the mixture is expected to be a poor approximation.

Calculating P -values

We now describe how to efficiently calculate P -values for the distribution of the statistic $r = Q(\mathbf{y})/\hat{\sigma}_e^2(\mathbf{y})$ calculated from the data; that is, given an observed statistic r , what is $\Pr(Q/\hat{\sigma}_e^2 > r)$ under the null? We review the result in Chen *et al.* (2016):

Proposition 5: Let r be the observed value of the statistic. Denote by $\alpha_1^{(r)}, \dots, \alpha_n^{(r)}$ the eigenvalues of $\mathbf{SKS} - r/(n - p) \cdot \mathbf{S}$. Then,

$$\Pr\left(\frac{Q}{\hat{\sigma}_e^2} > r\right) = \Pr\left(\sum_{i=1}^n \alpha_i^{(r)} \chi_{1,i}^2 > 0\right) \quad (10)$$

where $\chi_{1,i}^2$ are i.i.d. random variables distributed chi-square with one degree of freedom.

However, this condition requires us to calculate the eigenvalues of $\mathbf{SKS} - r/(n - p) \cdot \mathbf{S}$ for each new value r , which, naively, has a complexity of $O(n^3)$. We consider two scenarios where this is problematic. First, in many heritability studies, we wish to test the heritability of many (e.g., thousands) of phenotypes, all relative to the same kernel or kinship matrix (see above). For each phenotype $\mathbf{y}_1, \dots, \mathbf{y}_N$, we calculate its score test statistic r_i . For P -value calculation, we need to compute the eigendecomposition of $\mathbf{SKS} - r_i/(n - p) \cdot \mathbf{S}$ for each observed statistic r_i , which is a significant computational burden.

A second problematic scenario is of an association study of a single phenotype with many sets of SNPs, e.g., rare variants. Choosing a weighted linear kernel as in SKAT (Wu *et al.* 2011), we have $\mathbf{K}_i = \mathbf{Z}_i \mathbf{W}_i \mathbf{Z}_i^T$ for each set. As $\mathbf{K}_i$ changes with

each test, in principle, we need to perform a costly $O(n^3)$ eigendecomposition for each matrix $\mathbf{K}_i$. However, a significant computational saving is gained due to the fact that the nonzero eigenvalues of $\mathbf{SK}_i \mathbf{S} = \mathbf{S} \mathbf{Z}_i \mathbf{W}_i \mathbf{Z}_i^T \mathbf{S}$ are the same as those of $\mathbf{W}_i^{1/2} \mathbf{Z}_i^T \mathbf{S} \mathbf{Z}_i \mathbf{W}_i^{1/2}$, which is an $m \times m$ matrix (Lippert *et al.* 2014). As the number of tested SNPs m is often small, calculating the eigenvalues of this matrix instead is significantly faster, taking only $O(m^3)$, with matrix construction taking only $O(n(m + p)^2)$ (see Lippert *et al.* 2014). However, with the exact approach, we need to calculate the eigenvalues of $\mathbf{SK}_i \mathbf{S} - r_i/(n - p) \cdot \mathbf{S}$ instead of $\mathbf{SK}_i \mathbf{S}$. Even when $\mathbf{K}_i$ is low rank, the matrix $\mathbf{SK}_i \mathbf{S} - r_i/(n - p) \cdot \mathbf{S}$ may be close to full rank, so another approach is needed.

The following characterizes the eigenvalues of $\mathbf{SKS} - r/(n - p) \cdot \mathbf{S}$ given the eigenvalues of $\mathbf{SKS}$:

Proposition 6: Let r be the observed score test statistic. Denote by $\phi_1, \dots, \phi_n$ the eigenvalues of $\mathbf{SKS}$. Denote the column space of a matrix $\mathbf{A}$ by $\text{col}(\mathbf{A})$, its null space by $\text{ker}(\mathbf{A})$. Then,

$$\Pr\left(\frac{Q}{\hat{\sigma}_e^2} > r\right) = \Pr\left(\sum_{i=1}^k \left(\phi_i - \frac{r}{n - p}\right) \chi_{i,1}^2 - \sum_{i=k+1}^{k+q} \frac{r}{n - p} \cdot \chi_{i,1}^2 > 0\right) \quad (11)$$

Figure 3 Power study. The power of the accurate approach and SKAT is shown for P -value threshold of $p = 0.05$, for the KORA dataset, on 10,000 simulated phenotypes with varying degrees of true underlying heritability. SKAT is seen to be severely underpowered.

Table 2 Benchmarks

Scenario	Algorithm	Exact?	Preprocessing	Calculating $Q/\hat{\sigma}_e^2$	Calculating P -value
Heritability	SKAT	Approximate	3 sec	0.3 sec	5 sec
	MiRKAT	Exact	0.2 sec	0.3 sec	37 min
	RL-SKAT	Exact	9 sec	0.3 sec	5 sec
Set-testing	SKAT	Approximate	45 sec	2 sec	1 sec
	MiRKAT	Exact	5 sec	2 sec	43 min
	RL-SKAT	Exact	50 sec	2 sec	4 sec

Benchmark of the performance of different approaches for P -value calculation, applied to the KORA dataset.

where $k = \text{rank}(\mathbf{SKS})$ is the number of nonzero eigenvalues ϕ_i , $q = \dim(\ker(\mathbf{SKS}) \cap \text{col}(\mathbf{S}))$, and $\chi_{1,i}^2$ are i.i.d. random variables distributed chi-square with one degree of freedom, $i = 1, \dots, k + q$.

Proposition 6 shows that calculating the P -value amounts to evaluating the cumulative distribution function of a certain weighted mixture of chi-square distribution at 0. This can be done rapidly using the Davies method (Davies 1980), which is based on the numerical inversion of the characteristic function and runs in $O(n)$ complexity, or using other methods (Duchesne and De Micheaux 2010).

It remains to calculate k and q . Naively, this can be done in $O(n^3)$, for example by calculating the singular value decomposition (SVD) of $\mathbf{SKS}$ and $\mathbf{S}$ to get k and to obtain vector bases for $\ker(\mathbf{SKS})$ and $\text{col}(\mathbf{S})$, and by calculating the SVD of a matrix whose columns are the two vector bases to obtain q . When the same kernel is used with many phenotypes, it is a single preprocessing step. However, when the number of SNPs used to construct the kernel and the number of covariates are small, these quantities can be calculated much faster:

Proposition 7: Suppose $\mathbf{K} = \mathbf{ZWZ}^T$, and let $k = \text{rank}(\mathbf{SKS})$ and $q = \dim(\ker(\mathbf{SKS}) \cap \text{col}(\mathbf{S}))$. Then, k and q can be calculated in complexity $O(n(m+p)^2)$.

Most commonly, $k = \min(m, n) - 1$. When the number of SNPs m and the number of covariates p are small, the computational saving is substantial.

Data availability

This study makes use of data generated by the Wellcome Trust Case Control Consortium. A full list of the investigators who contributed to the generation of the data is available from www.wtccc.org.uk. The data used in this manuscript were obtained via KORA.PASST (<https://epi.helmholtz-muenchen.de/>) with the following variables: KORA F4 Illumina Human-Methylation450K BeadChip array, BMIQ normalization KORA

F4 Affymetrix 6.0 SNP Array; imputed (HapMap2 reference panel). Access to the data may be obtained by request to KORA.

Results

Performance summary

We summarize the results described in *Materials and Methods* above in Table 1 and in Algorithms 1 and 2. We compare our method, RL-SKAT, with the SKAT formulation and the correction of Chen *et al.* (2016) using the naive implementation of Proposition 5, as implemented by the MiRKAT software package (Zhao *et al.* 2015). The two scenarios discussed are those of a heritability study (same $\mathbf{K}$ with many responses $\mathbf{y}_i$) and SNP set-testing (many low rank $\mathbf{K}_i$). In all methods, a preprocessing step of calculating $\mathbf{X}^\dagger$ and $\{\phi_i\}$ is required. In a heritability study, calculating the statistic $Q/\hat{\sigma}_e^2$ amounts to evaluating two quadratic forms in $O(n^2)$. Compared to RL-SKAT, MiRKAT requires a full $O(n^3)$ eigendecomposition for each $\mathbf{y}_i$. For a set-testing study, these quadratic forms can be calculated in $O(n(m+p))$ due to the low rank of $\mathbf{K}_i$. Again, MiRKAT requires a full $O(n^3)$ eigendecomposition, compared to the $O(n(m+p)^2)$ procedure described in Proposition 7.

We now demonstrate our results on two datasets: a dataset from the Wellcome Trust Case Control Consortium 2 (International Multiple Sclerosis Genetics Consortium *et al.* 2011) (WTCCC2) study and the Cooperative health research in the Region of Augsburg (KORA) study (Holle *et al.* 2005). A full description of data preprocessing is given in the Supplemental Material in File S1.

A simulation study using WTCCC2 data

We first analyze data with real genotypes from the WTCCC2 Multiple Sclerosis dataset, and simulated phenotypes. We used the same data processing described in Yang *et al.*

Algorithm 1 RL-SKAT for heritability

procedure PREPROCESSING($\mathbf{X}, \mathbf{K}$)	▷ Preprocessing step, done once
Calculate $\mathbf{X}^\dagger = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$	▷ $O(np^2)$
Calculate $\mathbf{SKS}$ using $\mathbf{S} = \mathbf{I} - \mathbf{XX}^\dagger$	▷ $O(n^2p)$
Calculate $\phi_1, \dots, \phi_n$, the eigenvalues of $\mathbf{SKS}$	▷ $O(n^3)$
Extract $k = \text{rank}(\mathbf{SKS})$	▷ $O(1)$
Calculate $q = \dim(\ker(\mathbf{SKS}) \cap \text{col}(\mathbf{S}))$ using Proposition 7	▷ $O(n(n+p)^2)$
procedure TEST($\mathbf{y}$)	▷ Calculate p-value for a single phenotype $\mathbf{y}$
Calculate the score: $r := Q/\hat{\sigma}_e^2 = (n-p) \cdot \mathbf{y}^T \mathbf{SKS} \mathbf{y} / \mathbf{y}^T \mathbf{S} \mathbf{y}$	▷ $O(n^2)$
Calculate $\{\alpha_i^{(r)}\}$ as in Propositions 5 and 6	▷ $O(n)$
Calculate the p-value $p = \Pr(\sum_{i=1}^n \alpha_i^{(r)} \chi_{1,i}^2 > 0)$ using the Davies method	▷ $O(n)$

Algorithm 2 RL-SKAT for set-tests

```

procedure PREPROCESSING( $\mathbf{X}, \mathbf{ZW}^{1/2}$ )
    Calculate  $\mathbf{X}^\dagger = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$ 
    Calculate  $\mathbf{SZW}^{1/2}$  using  $\mathbf{S} = \mathbf{I} - \mathbf{XX}^\dagger$ 
    Calculate  $\varphi_1, \dots, \varphi_n$  as the squares of the singular values of  $\mathbf{SZW}^{1/2}$ 
    Extract  $k = \text{rank}(\mathbf{SKS})$ 
    Calculate  $q = \dim(\ker(\mathbf{SKS}) \cap \text{col}(\mathbf{S}))$  using Proposition 7
procedure TEST( $\mathbf{y}$ )
    Calculate the score:  $r := Q/\hat{\sigma}_e^2 = (n-p) \cdot \mathbf{y}^\top \mathbf{SKS} \mathbf{y} / \mathbf{y}^\top \mathbf{S} \mathbf{y}$ , using  $\mathbf{K} = \mathbf{ZWZ}^\top$ 
    Calculate  $\{\alpha_i^{(r)}\}$  as in Propositions 5 and 6
    Calculate the p-value  $p = \Pr(\sum_{i=1}^n \alpha_i^{(r)} \chi_{1,i}^2 > 0)$  using the Davies method

```

▷ Preprocessing step, done once

▷ $O(np^2)$

▷ $O(nmp)$

▷ $O(nm^2)$

▷ $O(1)$

▷ $O(n(m+p)^2)$

▷ Calculate p-value for a single phenotype $\mathbf{y}$

▷ $O(n(m+p))$

▷ $O(n)$

▷ $O(n)$

(2014), resulting in $m = 360,556$ SNPs for $n = 13,950$ individuals. We constructed the kinship matrix by a standard, uniformly weighted linear kernel. We sought to demonstrate the discrepancy between the true null distribution and the chi-square weighted mixture distribution. Following Proposition 4, we calculated the correlation to be 0.886 and variance ratio to be $R = 4.69$, indicating that a large discrepancy is possibly expected. To verify this, we simulated 10,000 random phenotypes, where each phenotype is a vector of i.i.d. standard normal variables. We tested whether the variance component is significantly > 0 , and calculated their P -values under the assumptions of either of the two distributions. In Figure 1, we show the quantile-quantile plots for the two sets of P -values. As evidenced, using the SKAT mixture distribution results in a severe deflation of small P -values, while using the correct distribution as in Proposition 1 results in an accurate P -value distribution. This shows that even for large sample sizes ($n = 13,950$), such a discrepancy is possible.

Testing for heritable methylation sites in the KORA dataset

The longitudinal KORA study consists of whole-blood methylation levels and genotypes of $n = 1799$ individuals. The phenotype is the proportion of methylated samples at a specific site, averaged across DNA samples of an individual. The study consists of independent population-based subjects from the general population living in the region of Augsburg, southern Germany (Holle *et al.* 2005). Whole-blood samples of the KORA F4 study were used as described elsewhere (Pfeiffer *et al.* 2015). In summary, a total of 431,366 methylation site phenotypes, and 657,103 SNPs, were available for analysis. The correlation as in Proposition 4 is 0.976 and the variance ratio is $R = 22.01$, indicating again that a large discrepancy is expected. We performed a heritability study of multiple phenotypes with the same kinship matrix, by testing the heritability of the $N = 43,140$ methylation sites on chromosome 1. As it is common for a methylation site to be correlated with its surrounding SNPs (Gibbs *et al.* 2010; Zhang *et al.* 2010; Bell *et al.* 2011), we avoided such *cis* effects by using a kinship matrix constructed from the $m = 604,170$ SNPs on all chromosomes other than 1. The kinship matrix is constructed by a standard, uniformly weighted linear kernel. For covariates, we used $\mathbf{X}$ consisting only

of an intercept vector. Again, we calculated P -values under the assumption of the two distributions. We note that it has been shown that some methylation site profiles often display significant heritability, while others do not; thus, both significant and insignificant P -values are expected (Rahmani *et al.* 2017).

In Figure 2 we show the histograms of the $\log_{10}$ of the P -value of all the considered phenotypes. The two histograms are indeed very different; P -values calculated using the inaccurate SKAT mixture distribution indicate that the heritability of almost all sites is considered insignificant; for example, using a Bonferroni threshold of $0.05 \cdot 1/43140 \approx 10^{-6}$, only 8/43,140 sites are significant. In light of the results above, it is reasonable to suspect that P -values of many heritable phenotypes are deflated, thus causing false negatives. The P -values distribution has a peak at ~ 0.5 , likely an artifact of the inaccurate calculation method. In comparison, P -values calculated by RL-SKAT do not exhibit such a peak. They are significantly smaller, and using the same Bonferroni threshold, we now find 319/43,140 significant sites. Indeed, a simulated power study of both approaches under varying degrees of a true underlying heritability validates that the inaccurate approach results in a severe decrease in power (Figure 3), which has been reported in the literature (Uemoto *et al.* 2013). As a point of reference, we compared the power of RL-SKAT with that of the popular LR test approach, and found they have similar power (see the Supplemental Material in File S1). We conclude that in this dataset, using the SKAT distribution for P -value calculation is highly problematic.

Benchmarks

Finally, we benchmarked the methods discussed here on the KORA dataset under the two above scenarios, on a 64G, 2.2 GHz Linux workstation, using our implementation in the Python language. We verified that the relevant part of our implementation is equivalent to MiRKAT and has a very similar running time. For the scenario of heritability testing, we calculated the P -values of 1000 phenotypes with the kinship matrix. For the scenario of set testing, we used 1000 sets of 100 SNPs each. The results are summarized in Table 2; as expected, the computational savings are very significant, achieving a speedup of more than two orders of magnitude. We expect the speedup to be even more significant for larger datasets.

Discussion

In summary, we have shown that the distribution suggested by SKAT to the score test statistic may be very inaccurate. Unlike previous studies, which have noted this discrepancy only in small sample sizes, we have shown that it might occur in large studies as well. We have proposed a computational method to accurately calculate P -values without compromising computational time. Finally, we demonstrated our findings in two datasets.

The exact calculation of P -values can be applied to other variants of the score test; for example, the SKAT-O (Lee *et al.* 2012) test seeks to find an optimal combination of burden tests and nonburden tests, which amounts to the score test with a certain kernel.

In this work, we focused on the case of a single kernel, and on a continuous phenotype. The extension of this work to multiple kernels (e.g., corresponding to several sets of SNPs) or to binary phenotypes (e.g., case/control studies) is non-trivial, as the null distribution cannot be modeled as a ratio of quadratic forms (see, e.g., Wang 2016; Wu *et al.* 2016). It therefore remains a subject for future work.

We believe that the prominence of likelihood-ratio based tests in heritability studies might stem from the statistical issues discussed above (see, for example, Uemoto *et al.* 2013), where SKAT was found to be significantly less powerful. It is our hope that this paper would facilitate the use of score tests in heritability studies in the future.

Acknowledgments

R.S. is supported by the Colton Family Foundation. This study was supported in part by a fellowship from the Edmond J. Safra Center for Bioinformatics at Tel Aviv University to R.S. and E.R. E.H. was partially supported by NSF grant 1705197. Funding for the project was provided by the Wellcome Trust under award 076113. The KORA study was initiated and financed by the Helmholtz Zentrum München, German Research Center for Environmental Health, which is funded by the German Federal Ministry of Education and Research and by the State of Bavaria. Furthermore, KORA research was supported within the Munich Center of Health Sciences, Ludwig-Maximilians-Universität, as part of LMUinnovativ.

Literature Cited

Bell, J. T., A. A. Pai, J. K. Pickrell, D. J. Gaffney, R. Pique-Regi *et al.*, 2011 DNA methylation patterns associate with genetic and gene expression variation in hapmap cell lines. *Genome Biol.* 12: R10.
 Chen, J., W. Chen, N. Zhao, M. C. Wu, and D. J. Schaid, 2016 Small sample kernel association tests for human genetic and microbiome association studies. *Genet. Epidemiol.* 40: 5–19.
 Davies, R. B., 1980 Algorithm AS 155: the distribution of a linear combination of χ^2 random variables. *J. R. Stat. Soc. Ser. C Appl. Stat.* 29: 323–333.

Duchesne, P., and P. L. De Micheaux, 2010 Computing the distribution of quadratic forms: further comparisons between the Liu-Tang-Zhang approximation and exact methods. *Comput. Stat. Data Anal.* 54: 858–862.
 Furlotte, N. A., D. Heckerman, and C. Lippert, 2014 Quantifying the uncertainty in heritability. *J. Hum. Genet.* 59: 269–275.
 Ganjgahi, H., A. M. Winkler, D. C. Glahn, J. Blangero, P. Kochunov *et al.*, 2015 Fast and powerful heritability inference for family-based neuroimaging studies. *Neuroimage* 115: 256–268.
 Ge, T., T. E. Nichols, P. H. Lee, A. J. Holmes, J. L. Roffman *et al.*, 2015 Massively expedited genome-wide heritability analysis (MEGHA). *Proc. Natl. Acad. Sci. USA* 112: 2479–2484.
 Gibbs, J. R., M. P. van der Brug, D. G. Hernandez, B. J. Traynor, M. A. Nalls *et al.*, 2010 Abundant quantitative trait loci exist for DNA methylation and gene expression in human brain. *PLoS Genet.* 6: e1000952.
 Hasegawa, T., K. Kojima, Y. Kawai, K. Misawa, T. Mimori *et al.*, 2016 AP-SKAT: highly-efficient genome-wide rare variant association test. *BMC Genomics* 17: 745.
 Hayes, B., M. Goddard, and M. E. Goddard, 2001 Prediction of total genetic value using genome-wide dense marker maps. *Genetics* 157: 1819–1829.
 Holle, R., M. Happich, H. Löwel, and H. Wichmann MONICA/KORA Study Group, 2005 KORA - a research platform for population based health research. *Gesundheitswesen* 67: S19–S25.
 International Multiple Sclerosis Genetics Consortium Wellcome Trust Case Control Consortium 2 Sawcer, S., G. Hellenthal, M. Pirinen, C. C. Spencer, N. A. Patsopoulos *et al.*, 2011 Genetic risk and a primary role for cell-mediated immune mechanisms in multiple sclerosis. *Nature* 476: 214–219.
 Kang, H. M., J. H. Sul, S. K. Service, N. A. Zaitlen, S.-y. Kong *et al.*, 2010 Variance component model to account for sample structure in genome-wide association studies. *Nat. Genet.* 42: 348–354.
 Lee, S., M. J. Emond, M. J. Bamshad, K. C. Barnes, M. J. Rieder *et al.*, 2012 Optimal unified approach for rare-variant association testing with application to small-sample case-control whole-exome sequencing studies. *Am. J. Hum. Genet.* 91: 224–237.
 Lippert, C., J. Xiang, D. Horta, C. Widmer, C. Kadie *et al.*, 2014 Greater power and computational efficiency for kernel-based association testing of sets of genetic variants. *Bioinformatics* 30: 3206–3214.
 Liu, D., X. Lin, and D. Ghosh, 2007 Semiparametric regression of multidimensional genetic pathway data: least-squares kernel machines and linear mixed models. *Biometrics* 63: 1079–1088.
 Liu, D., D. Ghosh, and X. Lin, 2008 Estimation and testing for the effect of a genetic pathway on a disease outcome using logistic kernel machine regression via logistic mixed models. *BMC Bioinformatics* 9: 292.
 Lloyd-Jones, L. R., A. Holloway, A. McRae, J. Yang, K. Small *et al.*, 2017 The genetic architecture of gene expression in peripheral blood. *Am. J. Hum. Genet.* 100: 228–237.
 Marcenko, V. A., and L. A. Pastur, 1967 Distribution of eigenvalues for some sets of random matrices. *Math. USSR-Sbornik* 1: 457.
 Patterson, N., A. L. Price, and D. Reich, 2006 Population structure and eigenanalysis. *PLoS Genet.* 2: e190.
 Pfeifferm, L., S. Wahl, L. C. Pilling, E. Reischl, J. K. Sandling *et al.*, 2015 DNA methylation of lipid-related genes affects blood lipid levels. *Circ. Cardiovasc. Genet.* 8: 334–342.
 Price, A. L., A. Helgason, G. Thorleifsson, S. A. McCarroll, A. Kong *et al.*, 2011 Single-tissue and cross-tissue heritability of gene expression via identity-by-descent in related or unrelated individuals. *PLoS Genet.* 7: e1001317.
 Quon, G., C. Lippert, D. Heckerman, and J. Listgarten, 2013 Patterns of methylation heritability in a genome-wide analysis of four brain regions. *Nucleic Acids Res.* 41: 2095–2104.

- Rahmani, E., L. Shenhav, R. Schweiger, P. Yousefi, K. Huen *et al.*, 2017 Genome-wide methylation data mirror ancestry information. *Epigenetics Chromatin* 10: 1.
- Schweiger, R., S. Kaufman, R. Laaksonen, M. E. Kleber, W. März *et al.*, 2016 Fast and accurate construction of confidence intervals for heritability. *Am. J. Hum. Genet.* 98: 1181–1192.
- Schweiger, R., E. Fisher, E. Rahmani, L. Shenhav, S. Rosset *et al.*, 2017 Using stochastic approximation techniques to efficiently construct confidence intervals for heritability, pp. 241–256 in *International Conference on Research in Computational Molecular Biology*. Springer, Berlin.
- Searle, S. R., G. Casella, and C. E. McCulloch, 2009 *Variance Components*, Vol. 391. John Wiley & Sons, Hoboken, NJ.
- Uemoto, Y., R. Pong-Wong, P. Navarro, V. Vitart, C. Hayward *et al.*, 2013 The power of regional heritability analysis for rare and common variant detection: simulations and application to eye biometrical traits. *Front. Genet.* 4: 232.
- Van Dongen, J., M. G. Nivard, G. Willemsen, J.-J. Hottenga, Q. Helmer *et al.*, 2016 Genetic and environmental influences interact with age and sex in shaping the human methylome. *Nat. Commun.* 7: 11115.
- Visscher, P. M., and M. E. Goddard, 2015 A general unified framework to assess the sampling variance of heritability estimates using pedigree or marker-based relationships. *Genetics* 199: 223–232.
- Visscher, P. M., W. G. Hill, and N. R. Wray, 2008 Heritability in the genomics era—concepts and misconceptions. *Nat. Rev. Genet.* 9: 255–266.
- Wang, K., 2016 Boosting the power of the sequence kernel association test by properly estimating its null distribution. *Am. J. Hum. Genet.* 99: 104–114.
- Wright, F. A., P. F. Sullivan, A. I. Brooks, F. Zou, W. Sun *et al.*, 2014 Heritability and genomics of gene expression in peripheral blood. *Nat. Genet.* 46: 430–437.
- Wu, B., W. Guan, and J. S. Pankow, 2016 On efficient and accurate calculation of significance p-values for sequence kernel association testing of variant set. *Ann. Hum. Genet.* 80: 123–135.
- Wu, M. C., S. Lee, T. Cai, Y. Li, M. Boehnke *et al.*, 2011 Rare-variant association testing for sequencing data with the sequence kernel association test. *Am. J. Hum. Genet.* 89: 82–93.
- Yang, J., B. Benyamin, B. P. McEvoy, S. Gordon, A. K. Henders *et al.*, 2010 Common SNPs explain a large proportion of the heritability for human height. *Nat. Genet.* 42: 565–569.
- Yang, J., S. H. Lee, M. E. Goddard, and P. M. Visscher, 2011 GCTA: a tool for genome-wide complex trait analysis. *Am. J. Hum. Genet.* 88: 76–82.
- Yang, J., N. A. Zaitlen, M. E. Goddard, P. M. Visscher, and A. L. Price, 2014 Advantages and pitfalls in the application of mixed-model association methods. *Nat. Genet.* 46: 100–106.
- Zhang, D., and X. Lin, 2003 Hypothesis testing in semiparametric additive mixed models. *Biostatistics* 4: 57–74.
- Zhang, D., L. Cheng, J. A. Badner, C. Chen, Q. Chen *et al.*, 2010 Genetic control of individual differences in gene-specific methylation in human brain. *Am. J. Hum. Genet.* 86: 411–419.
- Zhao, N., J. Chen, I. M. Carroll, T. Ringel-Kulka, M. P. Epstein *et al.*, 2015 Testing in microbiome-profiling studies with MiRKAT, the microbiome regression-based kernel association test. *Am. J. Hum. Genet.* 96: 797–807.

Communicating editor: C. Sabatti