

UCLA

UCLA Previously Published Works

Title

OpenMendel: a cooperative programming project for statistical genetics

Permalink

<https://escholarship.org/uc/item/06k1386p>

Journal

Human Genetics, 139(1)

ISSN

0340-6717

Authors

Zhou, Hua

Sinsheimer, Janet S

Bates, Douglas M

et al.

Publication Date

2020

DOI

10.1007/s00439-019-02001-z

Peer reviewed

OPENMENDEL: A Cooperative Programming Project for Statistical Genetics

Hua Zhou^{1,2} * Janet S. Sinsheimer^{1,2} † Douglas M. Bates ‡ Benjamin B. Chu §
 Christopher A. German* Sarah S. Ji* Kevin L. Keys ¶ Juhyun Kim* Seyoon Ko ||
 Gordon D. Mosher ** Jeanette C. Papp † Eric M. Sobel † Jing Zhai †† Jin J. Zhou ††
 Kenneth Lange² §

¹ These authors contributed equally.

² To whom correspondence should be addressed.

Email: huazhou@ucla.edu, jsinshei@g.ucla.edu, klange@ucla.edu

Phone: +1 (310) 794-7835, +1 (310) 825-8002, +1 (310) 206-8076

February 15, 2019

Abstract

Statistical methods for genomewide association studies (GWAS) continue to improve. However, the increasing volume and variety of genetic and genomic data make computational speed and ease of data manipulation mandatory in future software. In our view, a collaborative effort of statistical geneticists is required to develop open source software targeted to genetic epidemiology. Our attempt to meet this need is called the OPENMENDELproject (<https://openmendel.github.io>). It aims to (1) enable interactive and reproducible analyses with informative intermediate results, (2) scale to big data analytics, (3) embrace parallel and distributed computing, (4) adapt to rapid hardware evolution, (5) allow cloud computing, (6) allow integration of varied genetic data types, and (7) foster easy communication between clinicians, geneticists, statisticians, and computer scientists. This article reviews and makes recommendations to the genetic epidemiology community in the context of the OPENMENDEL project.

Keywords: Statistical Genomics, GWAS, Computational Statistics, Open Source, Collaborative Programming

* Department of Biostatistics, UCLA Fielding School of Public Health, Los Angeles, CA, USA

† Department of Human Genetics, David Geffen School of Medicine at UCLA, Los Angeles, CA, USA

‡ Department of Statistics, University of Wisconsin, Madison, Wisconsin, USA

§ Department of Biomathematics, David Geffen School of Medicine at UCLA, Los Angeles, CA, USA

¶ Department of Medicine, University of California, San Francisco, CA, USA

|| Department of Statistics, Seoul National University, Seoul, South Korea

** Departments of Statistics and Computer Science, University of California, Riverside, CA, USA

†† Department of Epidemiology and Biostatistics, Mel and Enid Zuckerman College of Public Health, University of Arizona, Tucson, AZ, USA

1 Introduction

Genomewide association studies (GWAS) query the entire genome to identify genetic variants associated with a trait of interest. GWAS have enjoyed many successes [77] and have uncovered many clues to the genetic etiology of common diseases [23]. Case-control tests of association between markers and traits predate GWAS by more than 50 years [1]. However, association studies were rarely undertaken in the pre-GWAS era unless there were candidate genes with strong prior evidence. The situation changed at the turn of the millennium when dense SNP (single nucleotide polymorphism) maps became available and SNP genotyping costs plummeted. Suddenly, it became possible to exploit linkage disequilibrium (LD) and survey hundreds of thousands to millions of genomewide SNPs. In the subsequent decade, hundreds of associations found by GWAS were published [76, 77]. Genomics is in the midst of a second technological evolution driven by high-throughput sequencing [55, 60, 40, 73]. Geneticists can now survey both rare and common variants.

This sudden expansion of data leads to enormous challenges in statistical genetics. Many current algorithms and programs are ill adapted to handle modern data sets with 10^5 cases and 10^7 markers. Ever more types of genetic variation are being observed and catalogued [77]. These changes demand more complex data structures and data integration across multiple biological scales. Precision health and predictive medicine raise the stakes even further [40]. Concurrently, the nature of computing is rapidly changing. In addition to new hardware, new programming paradigms and new algorithms must be brought online as quickly as possible to sustain progress in statistical genetics.

The following three studies exemplify the variety and magnitude of genomic data sets being collected today: a) The Million Veteran Program contains GWAS data (657,459 SNPs) on 359,964 veterans [29]. Simply storing the genotypes in compressed format requires > 100 GB. Obviously, this data set and others like it [70, 11] will continue to grow. b) A recent study [72] obtained whole genome sequence (WGS) data on 10,545 humans at 30-40x coverage for $< \$2000$ per genome. These researchers identified > 150 million variants, the majority of which are rare or *de novo*. c) The iPOP (integrative personal omics profile) study [16] followed a single individual for 401 days and collected transcriptome, proteome, metabolome, microbiome, epigenome, exposome, and phenome data at 20 time points, along with an extremely high coverage WGS. This type of omics profiling yields a dynamic picture of the heteroallelic changes between healthy and diseased states.

Current analysis pipelines juggle a multitude of computer programs that are implemented in different languages, run on different platforms, and require different input/output formats. This heterogeneity unintentionally creates barriers to communication, data exchange, data visualization, and scientific replication. End-users treat the entire pipeline as a black box and often fail to use their biological insight to inform statistical analysis. Students, post-docs, and researchers spend inordinate amounts of time coding and debugging the low-level languages instead of thinking about the science. In addition to these disadvantages, current software packages are straining under the volume, velocity, variety, and veracity of modern genomics data. Many programs do not even run on multiple threads. Distributed computing across different machines is largely ignored. In our view, the time is ripe to put in place a better paradigm for statistical genetics.

In this review, we first explain why the new JULIA language [6] is an ideal choice for the OPENMENDEL analysis platform. We then present what we see as some of the most pressing needs in gene mapping and our efforts to advance them through the cooperative OPENMENDEL effort. Owing to page and time constraints, we do not offer encyclopedic coverage of recent advances in GWAS or sequence analysis. Many promising methods are left unmentioned, for example meta-analysis based on summary statistics [15, 19, 26, 41, 52, 83] or estimation of fine-scale population structure [58]. Instead we focus on topics related to projects already

underway in OPENMENDEL. These projects include methods for handling SNP data, genotype imputation, rapid GWAS, iterative hard thresholding, kinship comparison, variance component modeling, and SNP-set analyses.

We want to point out that there are other groups who have made notable strides in making genetic analyses accessible to researchers who work with big data but lack the support available at large genomic centers [7, 63]. Some of these projects are further along than OPENMENDEL. A particularly interesting example is the PLATO software project ([30] and <https://ritchielab.org/software/plato-download>), which is designed to provide a single platform for a variety of association analyses. A major difference in our approach and PLATO is language choice. This difference might seem minor but, as we outline below, we believe it is a fundamental difference and is important for our goal of getting user initiated modules and modifications. Another example is the Ark software [7], which focuses on data management and is complementary to rather than competitive with OPENMENDEL.

2 The Importance of the Julia Computing Language

Many compelling features make JULIA [6] an ideal vehicle for implementing methods for modern statistical genetics. First, it is free, open source, and easy to install. Second, its clear, powerful syntax lends itself to compact, readable code and quick algorithm mock-ups. As needed, it can easily call Fortran, C, R, and Python functions. Third, because JULIA incorporates an excellent just-in-time (JIT) compiler, it achieves the efficiency of low-level languages with minimal programming efforts. Fourth, JULIA is built for parallelism at the multicore, graphical processing unit (GPU), and cluster levels. Fifth, JULIA employs a modern, easy to use package management system. Of particular relevance to statistical genetics, JULIA has many statistical and numerical analysis packages ready to use. Finally, end-users can run their analyses via the interactive Jupyter (**J**ulia, **P**ython, **R**) Notebook, an attractive interface for data visualization and reproducible research. Together these tools constitute an integrated environment for rapid prototyping of new applications and, with the same code, the analysis of large-scale genetic data.

Traditional high-level languages such as R, Matlab, and Python face the notorious two-language problem. In this scenario, one high-level language is used for prototyping, but a second low-level language is later needed for producing fast code for real world, large data sets. The high-level code is typically more compact, readable, and amenable to change, but much slower to execute. Most of the popular statistical genetics analysis tools or their most demanding subroutines are implemented purely in low-level languages, greatly restricting the community that feels comfortable exploring the code. Most tools are also restricted to certain computer platforms and input formats.

Today, a typical analysis pipeline requires a glue language such as Bash, Perl or Python to chain packages together. Data plotting and display require additional software, typically R or Matlab. Current analysis pipelines are cumbersome, opaque, and error-prone, creating barriers to the development of new statistical methods. Researchers wade through a swamp of low-level code and reinvent statistical genetics wheels instead of focusing on their unique contributions. This can be avoided as JULIA has solved the two-language problem through careful design of the programming language itself. JULIA is both easy to code and scales to peta-flop computing levels [21]. We can now use JULIA in all phases of our methods development, from prototyping to production software. OPENMENDEL includes many leading-edge statistical genetics methods written in this fast, high-level language that invites easy contributions from scientists. Using JULIA, OPENMENDEL can become the first highly efficient, open source statistical genetics software that can scale to million-subject studies and is both user- and developer-friendly.

3 Handling SNP Data

The SnpArrays.jl module of OPENMENDEL provides a convenient bridge between binary SNP data and downstream statistical analysis. The VCFTools.jl module achieves the same end for the richer genetic information distributed in VCF and BCF file formats. In SnpArrays.jl, biallelic genotype data are held in BitArrays, which store four genotypes per byte. As much as possible, compressed storage is also maintained during computation. Julia allows operators such as matrix multiplication to be defined directly on BitArrays without decompression. The design features of Julia make it easy to build high-performance statistical genetics software that is scalable to data sets with millions of subjects and tens of millions of SNPs.

The functionality of SnpArrays.jl includes: (1) reading and writing compressed SNP files, (2) computing summary statistics, (3) filtering data by genotyping success rates and other criteria, (4) copying compressed data into numerically oriented vectors and matrices, (5) computing genetic relationship matrices, (6) computing principal components, and (7) extending matrix and vector operations to compressed SNP data. SnpArrays.jl serves as a data interface to other OPENMENDEL modules.

4 Genotype Imputation

Genotype imputation involves the inference of unobserved genotypes from observed genotypes. It is possible to base inference on the observed genotypes of surrounding pedigree members [68], but pedigree data are now viewed as poor substitutes for linkage disequilibrium. In particular, pedigree data are incapable of imputing genotypes at completely untyped SNPs in a study. Recent versions of genotype imputation rely on panels of reference genotypes and employ hidden Markov models, with the hidden states being underlying haplotype pairs [34, 48, 54]. These programs are computationally intensive and operate by haplotyping individuals on the typed SNPs in the sample. These partial haplotypes are then compared to the reference panel to impute the full set of genotypes [33, 74]. We have taken an alternative approach based on the generic data mining technique of matrix completion [14, 18].

Matrix completion fills in the missing entries of an $m \times n$ matrix $\mathbf{X} = (x_{ij})$ whose observed entries are indexed by a subset Ω of $\{1, \dots, m\} \times \{1, \dots, n\}$. Imputation involves finding a low rank matrix $\mathbf{Y} = (y_{ij})$ consistent with the observed entries of $\mathbf{X} = (x_{ij})$. This is done by minimizing the loss function

$$f(\mathbf{Y}) = \sum_{(i,j) \in \Omega} (x_{ij} - y_{ij})^2 \quad (1)$$

over the set of matrices \mathbf{Y} of rank r or less. Taking r small is a form of parsimony capturing the hidden structure of the data. In genotype imputation, \mathbf{X} records the observed genotype dosages (0, 1, or 2 counts of the reference allele), with rows corresponding to people and columns to SNPs. Imputation is performed over a narrow genomic window of a few hundred SNPs where linkage disequilibrium prevails. Including reference individuals typed on out-of-sample SNPs is a key part of the strategy.

Because every rank r matrix \mathbf{Y} of dimension $m \times n$ can be expressed as a matrix product \mathbf{UV} , where \mathbf{U} is $m \times r$ and \mathbf{V} is $r \times n$, matrix completion can be phrased as updating the factors \mathbf{U} and \mathbf{V} of \mathbf{Y} in the loss function (1). Imputation is iterative, and to restore symmetry at iteration m , each missing entry x_{ij} is imputed by its current best guess $(\mathbf{U}_m \mathbf{V}_m)_{ij}$. New values \mathbf{U}_{m+1} and \mathbf{V}_{m+1} can be recovered by taking the singular value decomposition (SVD) of \mathbf{Z}_m , the current completed version of \mathbf{X} . The MM (majorization/minimization) principle of optimization shows that this procedure drives the loss downhill [45].

	Mendel Impute	MaCH	IMPUTE2
r^2	0.683	0.751	0.802
relative time	1.00	13.10	7.41

Table 1: Comparison of imputation methods.

Chi et al. [18] compared the matrix completion program Mendel Impute to several popular model based imputation programs including MaCH and IMPUTE2 using a number of simulated and real datasets. The accuracy of imputation is dependent on the nature of the specific scenarios and so no program was universally most accurate. The least favorable scenario for Mendel Impute in terms of accuracy occurred when imputing genotypes between high density microarray platforms using as a measure of accuracy the mean r^2 between the imputed values and the true genotypes at masked loci. In this case Mendel Impute was slightly worse than MaCH which was slightly worse than IMPUTE2 (Table 1). In other scenarios Mendel Impute was more accurate than MaCH and IMPUTE2 and in still others they were roughly the same. However, in all the scenarios presented Mendel Impute was at least an order of magnitude faster than MaCH or IMPUTE2.

Alternating least squares provides an alternative to SVD that is potentially much faster [31]. The alternating updates

$$\begin{aligned} \mathbf{V}_{m+1} &= (\mathbf{U}_m^t \mathbf{U}_m)^{-1} \mathbf{U}_m^t \mathbf{Y}_m \quad \text{and} \\ \mathbf{U}_{m+1} &= \mathbf{Z}_m \mathbf{V}_{m+1}^t (\mathbf{V}_{m+1} \mathbf{V}_{m+1}^t)^{-1} \end{aligned}$$

can achieve extremely high numerical throughput on modern computer architecture such as multicore CPUs and multiple GPUs. Because alternating least squares offers no guarantee of finding the global minimum of the loss, initial values for \mathbf{U} and \mathbf{V} should be as accurate as possible. Application of a randomized SVD to supply initial values is one possibility [49]. In practice we divide the current window into equal thirds and construct a hold-out-set by masking entries in the outer two thirds. We then choose the best rank r based on performance on the hold-out-set. Once we impute missing entries in the middle third, we shift the window to the right and begin again.

5 Enhancements to Ordinary GWAS

MendelGWAS.jl performs ordinary SNP-by-SNP association testing. To maximize speed in linear, logistic, and Poisson regression, MendelGWAS.jl employs score tests [3, 17, 20, 89]. For the most significant SNPs, the score test is supplemented by the slower but more accurate likelihood ratio test (LRT). Principal components can be included as predictors, SNPs and subjects can be filtered by success rates, and a Manhattan plot is provided.

In addition to these standard approaches to GWAS, we are in the process of implementing score tests for generalized linear models (GLMs) [66]. GLMs permit trait-genotype relations to be modeled with more exotic response distributions. We are also planning to develop an efficient score test for the challenging Cox survival model [37, 56, 69]. Multinomial regression models for complex categorical phenotypes would be a valuable extension of logistic regression [57]. Finally, efficient GWAS for ordered discrete phenotypes is becoming increasingly important for the study of complex diseases and traits derived from electronic health records.

6 Iterative Hard Thresholding

To avoid the computational complexity of multiple regression and the identifiability issues caused by having more predictors p than sample individuals n [12], GWAS has traditionally focused on the marginal effects of single SNPs. Previously we introduced lasso penalized regression to GWAS to perform subset selection [81, 91]. Our recent paper [38] implements a better heuristic, iterative hard thresholding (IHT), to solve this inherently combinatorial problem. We showed that IHT is better for GWAS than lasso or MCP penalties in controlling for false positive and false negative rates, in reducing parameter shrinkage, and in capturing heritability. It achieves these goals with little sacrifice in computational speed.

We now sketch how IHT iterates toward good local optima. To keep the discussion simple, consider the setting of linear regression with design matrix \mathbf{X} , response vector \mathbf{y} , and parameter vector β . The goal is to minimize the loss function $f(\beta) = \frac{1}{2}\|\mathbf{y} - \mathbf{X}\beta\|^2$ subject to the sparsity condition $\|\beta\|_0 \leq k$. The notation $\|\beta\|_0$ is shorthand for the number of nonzero entries of β . In GWAS the entry x_{ij} of \mathbf{X} denotes the number (0, 1, or 2) of reference alleles carried by individual i at SNP j or the imputed dosage value. The entry y_i of \mathbf{y} corresponding to individual i encodes a continuous trait such as height, blood pressure, or an expression level.

At iteration n , the IHT algorithm [8] moves in the steepest descent direction $-\nabla f(\beta_n)$ modified by the sparsity constraint. Here the gradient $\nabla f(\beta)$ of the objective equals $-\mathbf{X}^t(\mathbf{y} - \mathbf{X}\beta)$. The IHT update is explicitly

$$\beta_{n+1} = P_{S_k}(\beta_n - s\nabla f(\beta_n)), \quad (2)$$

where s is the steplength and $P_{S_k}(\beta)$ denotes projection onto the sparsity set $S_k = \{\beta : \|\beta\|_0 \leq k\}$. The projection operator $P_{S_k}(\beta)$ sends to 0 all but the k largest entries of β in magnitude. The preferred entries of β are untouched. The steplength s is chosen to minimize $f(\beta)$ along the ray $s \mapsto \beta_n - s\nabla f(\beta_n)$ prior to projection. This is achieved by taking

$$s = \frac{\|\nabla f(\beta_n)\|^2}{\|\mathbf{X}\nabla f(\beta_n)\|^2}.$$

The best value of k can be chosen by cross-validation.

The theory and practice of IHT continues to advance. Shen and Li [67] show how to relax the restricted isometry property originally invoked to prove convergence [9]. Yang et al [82] suggest group-sparse IHT to promote sparsity on a group-level. Khanna and Kyrillidis [39] validate the application of momentum acceleration to IHT. Yuan et al [85] and Bahmani [5] adapt IHT to logistic regression. Further extension to generalized linear models is a natural target. MendelIHT.jl brings IHT under the OPENMENDEL umbrella. Integration of IHT with SnpArrays.jl unifies data handling and leads to faster code with a smaller memory footprint. Finally, we are investigating weighting predictors to accommodate candidate genes and candidate SNPs [88].

7 Kinship Comparison

Kinship coefficients quantify the degree of relationship between two relatives. Two genes are identical by descent (IBD) if one is a copy of the other or they are both copies of the same ancestral gene. The theoretical kinship coefficient ϕ_{ij} is the probability that a randomly sampled gene at some arbitrary locus from individual i is IBD to a randomly sampled gene at the same locus from individual j . For example, if we assume no inbreeding, $\phi_{ij} = \frac{1}{2}$ if $i = j$, and $\phi_{ij} = \frac{1}{4}$ if i and j are first degree relatives. In the former case,

the two genes are sampled with replacement. In an accurately constructed pedigree, the full matrix Φ of kinship coefficients ϕ_{ij} can be calculated from a simple recurrence [44]. Jacquard's more complex kinship coefficients [35] are less useful in practice and harder to calculate [43]. In the MendelKinship.jl module of OPENMENDEL, Jacquard's coefficients are approximated by the Monte Carlo method of gene dropping.

When pedigrees are unknown or suspect, SNP markers can be used to estimate the kinship matrix Φ empirically. One popular estimate is the genetic relationship matrix (GRM), represented here by $\mathbf{S} = (s_{ij})$. If p_k denotes the reference allele frequency of SNP k , x_{ik} counts the number of reference alleles carried by individual i , and K is the number of SNPs, then the elements of \mathbf{S} are calculated as

$$s_{ij} = \frac{1}{K} \sum_{k=1}^K \frac{(x_{ik} - 2p_k)(x_{jk} - 2p_k)}{4p_k(1 - p_k)}.$$

Alternatives to the GRM include a methods of moments estimator MoM [24] and a robust GRM [53, 75]. The latter is

$$\hat{\phi}_{ij} = \frac{1}{\sum_{k=1}^K 4p_k(1 - p_k)} \sum_{k=1}^K (x_{ik} - 2p_k)(x_{jk} - 2p_k).$$

This unbiased estimator generally has smaller variance than the standard estimator \mathbf{S} , which is sensitive to low minor allele frequencies [78]. The MendelKinship.jl module calculates the GRM, the robust GRM, and the MoM estimators. All three of these estimators are special cases of general kinship estimators that are unbiased under ideal conditions [78]. When there is ethnic inhomogeneity and spread in the degrees of relationships, \mathbf{S} can exhibit bias because it confounds close relatedness and ancestry differences [22]. Ethnic admixture can be accommodated by replacing the allele frequency p_k by an ethnic specific estimate for each individual i [22].

Finding the variances of these estimators has been impossible without simplifying assumptions [78]. Our own unpublished approximation to the variance

$$\mathbb{E}(\|\mathbf{S} - \mathbb{E}(\mathbf{S})\|_{\mathbb{F}}^2) \approx \frac{1}{K^2} \|\mathbf{R}\|_{\mathbb{F}}^2 \left[\|\Phi\|_{\mathbb{F}}^2 + \text{tr}(\Phi)^2 \right] \quad (3)$$

of the GRM matrix \mathbf{S} allows for inbreeding, linkage disequilibrium, and closely related relatives. It relies on the simplifying assumption that the fourth moments of the SNP counts coincide with the fourth moments of similarly distributed Gaussian random variables. In formula (3), $\text{tr}(\mathbf{A})$ is the trace of \mathbf{A} , $\|\mathbf{A}\|_{\mathbb{F}}$ is the Frobenius norm of \mathbf{A} , and \mathbf{R} is the correlation matrix of the SNPs (LD matrix).

To check suspect pedigrees for hidden relatedness, one can compare theoretical kinships ϕ_{ij} and empiric kinships $\hat{\phi}_{ij}$. It is convenient to put these on a common scale by subjecting them to an approximate variance-stabilizing transformation. RA Fisher considered the simpler problem of comparing an ordinary covariance matrix $\Sigma = (\sigma_{ij})$ to a sample covariance matrix $\mathbf{S} = (s_{ij})$. Under an assumption of normality, he argued [27, 28] that the quantity

$$\tanh^{-1} \left(\frac{s_{ij}}{\sqrt{s_{ii}s_{jj}}} \right) - \tanh^{-1} \left(\frac{\sigma_{ij}}{\sqrt{\sigma_{ii}\sigma_{jj}}} \right)$$

is approximately normal with mean 0 and variance $(K - 3)^{-1}$, where K is the sample size, and \tanh^{-1} is the inverse hyperbolic tangent function. By analogy, we subject the GRM matrix or one of its variants to Fisher's transformation and order the discrepancies from least to greatest in absolute value. The OPENMENDEL MendelKinship.jl tutorial explains in a concrete example how transformation identifies outlier pairs.

8 Variance Component Models

Association studies are subject to the effects of unmeasured confounding. The most common confounder is ethnic ancestry [4, 32, 42], which arises when both trait values and marker allele frequencies differ by region of origin. Ancestry informative markers are particularly prone to show up as false positives in a naïve GWAS [64]. Currently there are two general adjustments for ethnic ancestry. The first approach uses either a few principal components of the GRM matrix [59, 61, 95] or estimated ancestry proportions [2, 62] as fixed effects. The second approach explicitly accounts for the correlation between subjects by including an estimate of the kinship matrix, e.g. the GRM matrix, as a random effect in a variance components model. When reliable pedigrees are available, the second approach is analogous to positing the theoretical kinship matrix as a random effect [10]. Because the theoretical kinship matrix does not capture hidden correlations, inclusion of the one of the SNP based estimates of the kinship matrix is usually preferred.

In any event, the variance components model $\mathbf{y} \sim N(\mathbf{X}\boldsymbol{\beta}, \sum_{j=1}^k \sigma_j^2 \mathbf{V}_j)$ figures prominently in genomewide association testing [25, 44]. In this model $\boldsymbol{\beta}$ are the fixed effects of covariates \mathbf{X} and σ_j^2 is the variance of the j th random effect. Estimation of the parameter vectors $\boldsymbol{\beta}$ and $\boldsymbol{\sigma}^2 = (\sigma_1^2, \dots, \sigma_k^2)^t$ has been the subject of intense study for decades. Most statisticians opt for maximum likelihood or restricted maximum likelihood. In the linear mixed model, the covariance matrices \mathbf{V}_j factor as $\mathbf{U}_j \mathbf{U}_j^t$. The factored form is advantageous if \mathbf{U}_j is $n \times r_j$ with r_j small. In the absence of low rank structure, one can take \mathbf{U}_j to be the Cholesky factor of \mathbf{V}_j .

The covariance model $\mathbf{W} = 2\sigma_a^2 \boldsymbol{\Phi} + \sigma_e^2 \mathbf{I}$ corresponds to polygenic background (σ_a^2) plus random noise (σ_e^2). The kinship matrix here can be theoretical or empirical. The model is overly simplistic but widely applied due to its computational tractability. It omits dominance effects, shared environment, and parent of origin effects, among other things. Calculation of the inverse and determinant of \mathbf{W} is the rate limiting step in estimation. In the simple polygenic model, a good tactic is to first calculate the spectral decomposition $\mathbf{O} \mathbf{D} \mathbf{O}^t$ of $\boldsymbol{\Phi}$, where \mathbf{D} is a diagonal matrix. One can then exploit the formulas $\det \mathbf{W} = \det(\sigma_a^2 \mathbf{D} + \sigma_e^2 \mathbf{I})$ and $\mathbf{W}^{-1} = \mathbf{O}(\sigma_a^2 \mathbf{D} + \sigma_e^2 \mathbf{I})^{-1} \mathbf{O}^t$. The indicated determinant and inverse of the diagonal matrix are trivial to compute [36, 50, 71].

Our program VarianceComponentModels.jl incorporates this spectral decomposition tactic. It also treats more realistic models with multiple variance components and multivariate traits. For estimation we have compared Fisher scoring and the EM algorithms long familiar to computational statisticians. We have also explored a new MM algorithm that alternates updates of $\boldsymbol{\beta}$ and $\boldsymbol{\sigma}^2$ [90]. The normal equation update of $\boldsymbol{\beta}$ in our algorithm is

$$\boldsymbol{\beta}_{n+1} = (\mathbf{X}^t \mathbf{W}_n^{-1} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{W}_n^{-1} \mathbf{y},$$

where \mathbf{W}_n is the value of $\sum_{j=1}^k \sigma_j^2 \mathbf{V}_j$ at the current estimate of $\boldsymbol{\sigma}^2$. The variance component updates are

$$\sigma_{n+1,i}^2 = \sigma_{ni}^2 \sqrt{\frac{(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}_{n+1})^t \mathbf{W}_n^{-1} \mathbf{V}_i \mathbf{W}_n^{-1} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}_{n+1})}{\text{tr}(\mathbf{W}_n^{-1} \mathbf{V}_i)}}. \quad (4)$$

The MM algorithm converges faster than the standard EM algorithm. Fisher scoring requires fewer iterations to converge but substantially more effort per iteration, particularly in high dimensions. Both the EM and MM algorithms can be accelerated by quasi-Newton extrapolation [87].

The VarianceComponentModels.jl module serves as a convenient vehicle for other genetic applications. One example is Mendelian randomization (MR). Observational studies often find an association between a biomarker or expression (or methylation) level at a particular locus and a quantitative trait. The goal of MR

is to assess the statistical support for this “exposure” as a cause of the trait, as opposed to reverse causality or confounding [13]. Our Mendelian randomization tutorial for continuous traits demonstrates the value of modularized genetic software such as `VarianceComponentModels.jl`.

When there are many loci to test, we [20, 89] and others [3, 17, 36, 50] have employed score tests or their equivalents in variance component models. Score tests are much faster than likelihood ratio tests (LRTs) because score tests require the likelihood to be maximized only under the null hypothesis. In contrast, LRTs require the likelihood be maximized both under the null and alternative hypotheses. When the null hypothesis is the same for all loci tested, this can amount to substantial savings. These score tests are easily extended to include maternal genetic effects and maternal-offspring genetic interaction as fixed effects [20]. Although most software programs implementing the score test adopt the simple covariance model $2\sigma_a^2\Phi + \sigma_e^2\mathbf{I}$, in principle other variance components such as household effects can be included.

Our recent analysis of the GWAS data from the COPDGene study (<http://www.copdgene.org>) exemplifies the vast performance gain and yet ease of use of a typical OPENMENDEL workflow in genetic heritability analysis of a realistically large data set [93]. The data are available from NIH dbGap under phs000179.v5.p2. The steps are: (1) load the binary genotypes of 6,670 individuals at 630,860 SNPs, (2) compute summary statistics on the SNPs, (3) impute missing genotypes, (4) calculate the empirical kinship matrix, (5) load 13 phenotypes, (6) estimate the heritability of each phenotype, (7) estimate the coheritability of each pair of phenotypes, and (8) fit a joint model to all 13 phenotypes. All these steps are performed in a single interactive Julia environment on a common laptop computer. Typically such an analysis pipeline would require running at least five separate programs on a Linux machine.

In our experience, a pure Julia computation is often faster than the corresponding computation in a low-level language such as C, and much faster than any other high-level language such as R or Python. Figure 1 compares the speed of fitting large-scale variance component models in our `VarianceComponentModels.jl` module to the two cutting edge programs GCTA [84] and GEMMA [94], both implemented in C++. In this example, there are two variance components, one for additive genetic effects and one for environmental effects. To make a fair comparison, the genetic relationship matrix S was pre-computed using the GCTA software. There are 13 continuous phenotypes. For both univariate (top panel) and bi-variate (bottom panel) models, we observe between 5 and 100 fold speedup over GEMMA and even more over GCTA. In all cases, the final log-likelihoods by JULIA match those by GCTA and GEMMA to the third digit.

The current versions of GCTA and GEMMA are only available for the x86 64-bit Linux operating system, while JULIA, and thus OPENMENDEL, are available on all common systems. It is remarkable that a cross-platform, interactive, high-level language such as JULIA can achieve such excellent computational efficiency.

9 SNP-Set Analysis

SNP-set analysis, or pathway-based analysis [79], is a powerful, widely-used strategy in sequencing studies. SNPs are grouped into sets to be examined for association with a certain phenotype. This analysis has been shown to have increased power over individual SNP analysis, especially for identifying rare variant associations [47].

Two types of SNP-set analyses are under active development within OPENMENDEL. The `VarianceComponentTests.jl` module implements different approaches for testing a set of markers as random effects. Notably the sequence kernel association test (SKAT) [80] is the first method to incorporate the generalized linear mixed model in testing the effect of a set of variants on a quantitative or dichotomous trait. Our recent work on exact tests [92] boosts the power of SKAT on small samples.

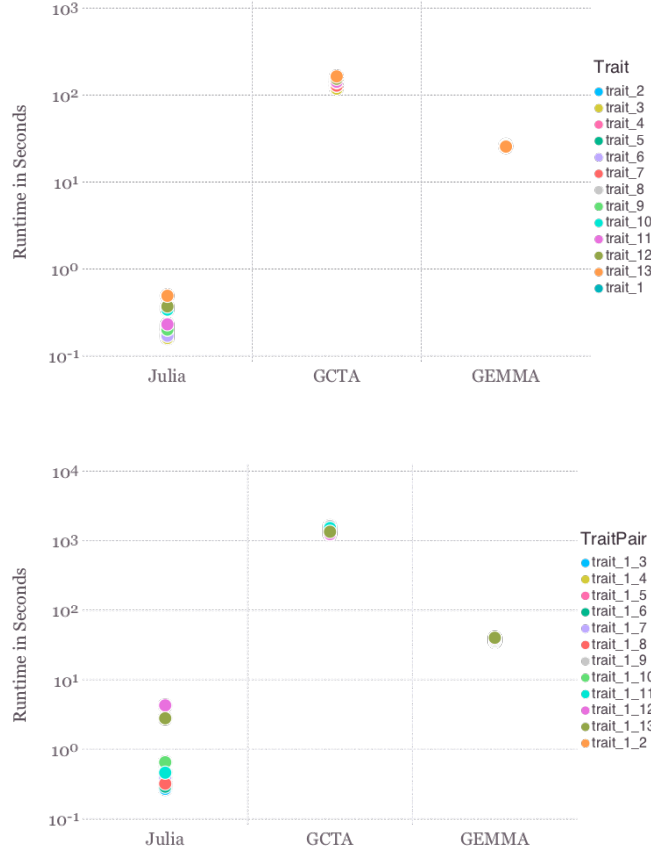


Figure 1: Comparison of the OPENMENDEL VarianceComponentModels.jl implementation with GCTA (C++) and GEMMA (C++) for fitting a univariate variance component model $\mathbf{Y} \sim N(\mathbf{0}, \sigma_a^2 \mathbf{S} + \sigma_e^2 \mathbf{I})$ (top panel) and a bivariate variance component model $\mathbf{Y} \sim N(\mathbf{0}, \Sigma_a \otimes \mathbf{S} + \Sigma_e \otimes \mathbf{I})$ (bottom panel). GEMMA and OPENMENDEL runtimes exclude the eigen-decomposition of \mathbf{S} , which is pre-computed.

In contrast to marginal SNP-set analysis, an alternative approach is subset selection in a joint model $\mathbf{y} \sim N(\mathbf{X}\boldsymbol{\beta}, \sum_{j=1}^m \sigma_j^2 \mathbf{V}_j + \sigma_0^2 \mathbf{I})$, where \mathbf{V}_j is the kernel matrix for the j th SNP-set and the σ_j^2 for $j \geq 1$ are the variance components subject to selection. Variance component selection is achieved by minimizing the penalized log-likelihood

$$-L(\boldsymbol{\beta}, \boldsymbol{\sigma}^2) + \sum_{j=1}^m P_\lambda(\sigma_j),$$

where $L(\boldsymbol{\beta}, \boldsymbol{\sigma}^2)$ is the log-likelihood function and $P_\lambda(\sigma_j)$ is a penalty function. Several penalties, including the ridge, the lasso, the smoothly clipped absolute deviation (SCAD), and the minimax concave penalty (MCP), are implemented. The MM update (4) generalizes to penalized estimation because the variance components σ_j^2 are nicely separated in the surrogate function [90].

10 Simulation Utilities

Simulation is vital in demonstrating the accuracy and power of new statistical methods. It is also important in designing genetic studies, where overly simplistic assumptions can lead to low power. Although there are a number of simulators already available [51, 65, 86], there is plenty of room for improvement. The unified nature of the OPENMENDEL environment makes it easy to craft code for simulating traits conditional on genotypes under any generalized linear model (GLM) or generalized linear mixed models (GLMM).

At the time that this article was written, the MendelTraitSimulate.jl option was under development. In its current version, we accommodate study designs involving both unrelateds and multigenerational families. We allow the user to specify both fixed and random effects for simulated univariate or multivariate traits. The simulated traits can be based on arbitrary functions of the provided covariates. By default, the program will use the PLINK format and make appropriate calls to SnpArrays.jl and VCFTools.jl.

11 Tutorials

Accompanying this article we have prepared a collection of tutorials via Jupyter Notebooks to demonstrate interactive genetic analysis using OPENMENDEL packages (<https://github.com/OpenMendel/Tutorials>). These include (1) PLINK binary data input, summary statistics, filtering, and visualization, (2) kinship calculation and comparison, (3) population GWAS, (4) iterative hard thresholding for GWAS, (5) heritability estimation, (6) Mendelian randomization, (7) GWAS based on linear mixed models, (8) SNP-set analysis, and soon to come (9) trait simulation. These tutorials will adapt to and grow with the expanding OPENMENDEL ecosystem.

12 Discussion

Readers may be familiar with our existing statistical package MENDEL [46]. Although MENDEL possesses many advantages, our goal going forward is not to modernize it, but to create an entirely new open source platform. Although MENDEL is free, it is not open source. The Fortran language underlying it is also antiquated. Fortran lacks the supporting libraries of R and Matlab, its graphics functionality is nil, it neglects crucial statistical and linear algebra tools, and its code is needlessly verbose.

For the sake of brevity, we have not discussed many OPENMENDEL modules. Omitted modules include: (1) discovery of ancestry informative markers, (2) estimation of allele frequencies from pedigree data, (3) testing for transmission disequilibrium by the gamete competition model, (4) random genotype generation by gene dropping, (5) genetic counseling, (6) two-point linkage analysis, (7) location scores for linkage analysis, and (8) function optimization by recursive quadratic programming. Table 2 lists the currently available OPENMENDEL analysis options and utility packages as well as those soon to be released as part of the OPENMENDEL project.

OPENMENDEL is inspired by a vision of genomic analysis that extracts the maximum benefit from the world-wide increase in genetic data and exploits the promise of collaborative, parallel, and distributed computing. We are not alone in this vision. As examples, notable strides have been made by HAIL (<https://hail.is/>) and TOPMed (<https://www.nhlbiwgs.org/awards>) in enabling large scale sequence analysis and the Ark data management system for health and biomedical research [7, 63]. In our opinion, however, the barriers need to be lowered further to encourage more statistical geneticists and genetic epidemiologists to take part. The JULIA language provides the ideal vehicles for this purpose. It is

OPENMENDEL Option	Description
MendelAimSelection.jl	Selects the most informative SNPs for predicting ancestry
MendelEstimateFrequencies.jl	Estimates allele frequencies from pedigree data
MendelGameteCompetition.jl	Tests for association under the gamete competition model
MendelGeneticCounseling.jl	Computes risks in genetic counseling problems
MendelGWAS.jl	Tests for association in genome-wide data
MendelIHT.jl	GWAS using Iterative Hard Thresholding (forthcoming)
MendelImpute.jl	Genotype imputation (forthcoming)
MendelKinship.jl	Computes kinship and other identity coefficients
MendelLocationScores.jl	Maps a trait via the method of location scores
OrdinalGWAS.jl	Implements GWAS for ordinal categorical phenotypes
MendelTwoPointLinkage.jl	Implements two-point linkage analysis
MendelBase.jl	Base functions for OPENMENDEL
MendelGeneDropping.jl	Simulates genotypes based on pedigrees
MendelSearch.jl	Optimization routines
MendelTraitSimulate.jl	Trait simulation using GLM and GLMM (forthcoming)
SnpArrays.jl	Utilities for handling compressed storage of biallelic SNP data
VCFTools.jl	Utilities for handling compressed storage of sequence data
VarianceComponentModels.jl	Utilities for fitting and testing variance components models

Table 2: JULIA packages in the OPENMENDEL project.

our hope that the OPENMENDEL project will spark a global effort to build a computing platform equal to the challenges of 21st-century genetic research.

Acknowledgements

This work was funded by NIH Grants R01-GM53275, R01-HG006139, R01-GM105785, R01-HL135156 and T32-HG002536; NSF grant DMS-1052210; the UCSF Bakar Computational Health Sciences Institute; the UC Berkeley Institute for Data Sciences as part of the Moore-Sloan Data Sciences Environment Initiative; and the 2018 Google Summer of Code.

References

- [1] Ian Aird, H H Bentall, and JA Fraser Roberts. Relationship between cancer of stomach and the abo blood groups. *British Medical Journal*, 1(4814):799, 1953.
- [2] David H. Alexander, John Novembre, and Kenneth Lange. Fast model-based estimation of ancestry in unrelated individuals. *Genome Research*, 19(9):1655–1664, 2009.
- [3] Najaf Amin, Cornelia M Van Duijn, and Yurii S Aulchenko. A genomic background based method for association analysis in related individuals. *PLoS ONE*, 2(12):e1274, 2007.
- [4] William Astle, David J Balding, et al. Population structure and cryptic relatedness in genetic association studies. *Statistical Science*, 24(4):451–471, 2009.

- [5] Sohail Bahmani, Bhiksha Raj, and Petros T Boufounos. Greedy sparsity-constrained optimization. *Journal of Machine Learning Research*, 14(Mar):807–841, 2013.
- [6] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017.
- [7] Adrian Bickerstaffe, Thilina Ranaweera, Travis Endersby, Christopher Ellis, Sanjaya Maddumarachchi, George E Gooden, Paul White, Eric K Moses, Alex W Hewitt, and John L Hopper. The Ark: a customizable web-based data management tool for health and medical research. *Bioinformatics*, 33(4):624–626, 2017.
- [8] Thomas Blumensath and Mike E Davies. Iterative thresholding for sparse approximations. *Journal of Fourier analysis and Applications*, 14(5-6):629–654, 2008.
- [9] Thomas Blumensath and Mike E Davies. Iterative hard thresholding for compressed sensing. *Applied and computational harmonic analysis*, 27(3):265–274, 2009.
- [10] E Boerwinkle and CF Sing. The use of measured genotype information in the analysis of quantitative phenotypes in man. *Annals of Human Genetics*, 51(3):211–226, 1987.
- [11] Jennifer A Brody, Alanna C Morrison, Joshua C Bis, Jeffrey R O’Connell, Michael R Brown, Jennifer E Huffman, Darren C Ames, Andrew Carroll, Matthew P Conomos, Stacey Gabriel, et al. Analysis commons, a team approach to discovery in a big-data environment for genetic epidemiology. *Nature Genetics*, 49(11):1560, 2017.
- [12] Peter Bühlmann and Sara Van De Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011.
- [13] Stephen Burgess and Simon G Thompson. *Mendelian randomization: methods for using genetic variants in causal estimation*. Chapman and Hall/CRC, 2015.
- [14] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.
- [15] Rita M Cantor, Kenneth Lange, and Janet S Sinsheimer. Prioritizing GWAS results: a review of statistical methods and recommendations for their application. *The American Journal of Human Genetics*, 86(1):6–22, 2010.
- [16] Rui Chen, George I Mias, Jennifer Li-Pook-Than, Lihua Jiang, Hugo Y K Lam, Rong Chen, Elana Miriami, Konrad J Karczewski, Manoj Hariharan, Frederick E Dewey, Yong Cheng, Michael J Clark, Hogune Im, Lukas Habegger, Suganthi Balasubramanian, Maeve O’Huallachain, Joel T Dudley, Sara Hillenmeyer, Rajini Haraksingh, Donald Sharon, Ghia Euskirchen, Phil Lacroute, Keith Bettinger, Alan P Boyle, Maya Kasowski, Fabian Grubert, Scott Seki, Marco Garcia, Michelle Whirl-Carrillo, Mercedes Gallardo, Maria A Blasco, Peter L Greenberg, Phyllis Snyder, Teri E Klein, Russ B Altman, Atul J Butte, Euan A Ashley, Mark Gerstein, Kari C Nadeau, Hua Tang, and Michael Snyder. Personal omics profiling reveals dynamic molecular and medical phenotypes. *Cell*, 148(6):1293–1307, 2012.
- [17] Wei-Min Chen and Gonçalo R Abecasis. Family-based association tests for genomewide association scans. *The American Journal of Human Genetics*, 81(5):913–926, 2007.

- [18] Eric C Chi, Hua Zhou, Gary K Chen, Diego Ortega Del Vecchio, and Kenneth Lange. Genotype imputation via matrix completion. *Genome Research*, 23(3):509–518, 2013.
- [19] Chi-yang Chiu, Jeeseun Jung, Wei Chen, Daniel E Weeks, Haobo Ren, Michael Boehnke, Christopher I Amos, Aiyi Liu, James L Mills, Mei-ling Ting Lee, Momiao Xiong, and Ruzong Fan. Meta-analysis of quantitative pleiotropic traits for next-generation sequencing with multivariate functional linear models. *European Journal Of Human Genetics*, 25:350 EP –, 12 2016.
- [20] Michelle M. Clark, John Blangero, Thomas D. Dyer, Eric M. Sobel, and Janet S. Sinsheimer. The quantitative-MFG test: A linear mixed effect model to detect maternal-offspring gene interactions. *Annals of Human Genetics*, 80(1):63–80, 2016.
- [21] A. Claster. Julia joins petaflop club, 2017.
- [22] Matthew P Conomos, Alexander P Reiner, Bruce S Weir, and Timothy A Thornton. Model-free estimation of recent genetic relatedness. *The American Journal of Human Genetics*, 98(1):127–148, 2016.
- [23] William Cookson, Liming Liang, Gonçalo Abecasis, Miriam Moffatt, and Mark Lathrop. Mapping complex disease traits with global gene expression. *Nature Reviews Genetics*, 10(3):184, 2009.
- [24] Aaron G Day-Williams, John Blangero, Thomas D Dyer, Kenneth Lange, and Eric M Sobel. Linkage analysis without defined pedigrees. *Genetic Epidemiology*, 35(5):360–370, 2011.
- [25] DS Falconer and TF Mackay. C. 1996. *Introduction to Quantitative Genetics*, pages 82–86, 1996.
- [26] Ruzong Fan, Yifan Wang, Chi-yang Chiu, Wei Chen, Haobo Ren, Yun Li, Michael Boehnke, Christopher I. Amos, Jason H Moore, and Momiao Xiong. Meta-analysis of complex diseases at gene level with generalized functional linear models. *Genetics*, 202(2):457–470, 2016.
- [27] Ronald A Fisher. Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. *Biometrika*, 10(4):507–521, 1915.
- [28] Ronald A Fisher. On the probable error of a coefficient of correlation deduced from a small sample. *Metron*, 1:3–32, 1921.
- [29] John Michael Gaziano, John Concato, Mary Brophy, Louis Fiore, Saiju Pyarajan, James Breeling, Stacey Whitbourne, Jennifer Deen, Colleen Shannon, Donald Humphries, Peter Guarino, Mihaela Aslan, Daniel Anderson, Rene LaFleur, Timothy Hammond, Kendra Schaa, Jennifer Moser, Grant Huang, Sumitra Muralidhar, Ronald Przygodzki, and Timothy J. O’Leary. Million Veteran Program: A mega-biobank to study genetic influences on health and disease. *Journal of Clinical Epidemiology*, 70:214–223, 2017/05/29 2016.
- [30] Molly A Hall, John Wallace, Anastasia Lucas, Dokyoon Kim, Anna O Basile, Shefali S Verma, Cathy A McCarty, Murray H Brilliant, Peggy L Peissig, Terrie E Kitchner, et al. Plato software provides analytic framework for investigating complexity beyond genome-wide association studies. *Nature Communications*, 8(1):1167, 2017.
- [31] Trevor Hastie, Rahul Mazumder, Jason D. Lee, and Reza Zadeh. Matrix completion and low-rank SVD via fast alternating least squares. *J. Mach. Learn. Res.*, 16(1):3367–3402, January 2015.

- [32] Agnar Helgason, Bryndís Yngvadóttir, Birgir Hrafnkelsson, Jeffrey Gulcher, and Kári Stefánsson. An Icelandic example of the impact of population structure on association studies. *Nature Genetics*, 37(1):90, 2005.
- [33] Bryan Howie, Christian Fuchsberger, Matthew Stephens, Jonathan Marchini, and Gonçalo R Abecasis. Fast and accurate genotype imputation in genome-wide association studies through pre-phasing. *Nature Genetics*, 44(8):955, 2012.
- [34] Bryan N Howie, Peter Donnelly, and Jonathan Marchini. A flexible and accurate genotype imputation method for the next generation of genome-wide association studies. *PLoS Genetics*, 5(6):e1000529, 06 2009.
- [35] Albert Jacquard. *The Genetic Structure of Populations*, volume 5. Springer Science & Business Media, 1974.
- [36] Hyun Min Kang, Jae Hoon Sul, Susan K Service, Noah A Zaitlen, Sit-Yee Kong, Nelson B Freimer, Chiara Sabatti, and Eleazar Eskin. Variance component model to account for sample structure in genome-wide association studies. *Nature Genetics*, 42(4):348–354, 2010.
- [37] E S Kawaguchi, Marc A Suchard, Z Liu, and G Li. Scalable sparse Cox regression for large-scale survival data via broken adaptive ridge. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, (in press), 2018.
- [38] Kevin L Keys, Gary K Chen, and Kenneth Lange. Iterative hard thresholding for model selection in genome-wide association studies. *Genetic Epidemiology*, 41(8):756–768, 2017.
- [39] Rajiv Khanna and Anastasios Kyrillidis. Iht dies hard: Provable accelerated iterative hard thresholding. In *International Conference on Artificial Intelligence and Statistics*, pages 188–198, 2018.
- [40] Helena Kilpinen and Jeffrey C Barrett. How next-generation sequencing is transforming complex disease genetics. *Trends in Genetics*, 29(1):23–30, 2013.
- [41] Junghi Kim, Yun Bai, and Wei Pan. An adaptive association test for multiple phenotypes with GWAS summary statistics. *Genetic Epidemiology*, 39(8):651–663, 2015.
- [42] William C Knowler, RC Williams, DJ Pettitt, and A Gm Steinberg. Gm3; 5, 13, 14 and type 2 diabetes mellitus: an association in american indians with genetic admixture. *American Journal of Human Genetics*, 43(4):520, 1988.
- [43] K Lange and JS Sinsheimer. Calculation of genetic identity coefficients. *Annals of Human Genetics*, 56(4):339–346, 1992.
- [44] Kenneth Lange. *Mathematical and Statistical Methods for Genetic Analysis*. Springer Science & Business Media, 2003.
- [45] Kenneth Lange. *MM Optimization Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2016.
- [46] Kenneth Lange, Jeanette C Papp, Janet S Sinsheimer, Ram Sripracha, Hua Zhou, and Eric M Sobel. Mendel: the Swiss army knife of genetic analysis programs. *Bioinformatics*, 29(12):1568–1570, 2013.

- [47] Seunggeung Lee, Gonçalo R. Abecasis, Michael Boehnke, and Xihong Lin. Rare-variant association analysis: Study designs and statistical tests. *American Journal of Human Genetics*, 95(1):5–23, 2014.
- [48] Yun Li, Cristen J. Willer, Jun Ding, Paul Scheet, and Gonçalo R. Abecasis. MaCH: Using sequence and genotype data to estimate haplotypes and unobserved genotypes. *Genetic Epidemiology*, 34(8):816–834, 2010.
- [49] Edo Liberty, Franco Woolfe, Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. Randomized algorithms for the low-rank approximation of matrices. *Proc. Natl. Acad. Sci. USA*, 104(51):20167–20172, 2007.
- [50] Christoph Lippert, Jennifer Listgarten, Ying Liu, Carl M Kadie, Robert I Davidson, and David Heckerman. FaST Linear Mixed Models for genome-wide association studies. *Nature Methods*, 8(10):833–835, 2011.
- [51] Youfang Liu, Georgios Athanasiadis, and Michael E Weale. A survey of genetic simulation software for population and epidemiological studies. *Human Genomics*, 3(1):79, 2008.
- [52] Nicholas Mancuso, Huwenbo Shi, Pagé Goddard, Gleb Kichaev, Alexander Gusev, and Bogdan Pasaniuc. Integrating gene expression with summary association statistics to identify genes associated with 30 complex traits. *The American Journal of Human Genetics*, 100(3):473–487, 2017.
- [53] Ani Manichaikul, Josyf C Mychaleckyj, Stephen S Rich, Kathy Daly, Michèle Sale, and Wei-Min Chen. Robust relationship inference in genome-wide association studies. *Bioinformatics*, 26(22):2867–2873, 2010.
- [54] Jonathan Marchini, Bryan Howie, Simon Myers, Gil McVean, and Peter Donnelly. A new multipoint method for genome-wide association studies by imputation of genotypes. *Nature Genetics*, 39(7):906–913, 2007.
- [55] Michael L Metzker. Sequencing technology the next generation. *Nature Reviews Genetics*, 11(1):31, 2010.
- [56] Sushil Mittal, David Madigan, Randall S Burd, and Marc A Suchard. High-dimensional, massive sample-size Cox proportional hazards regression for survival analysis. *Biostatistics*, 15(2):207–221, 2014.
- [57] Andrew P Morris, Cecilia M Lindgren, Eleftheria Zeggini, Nicholas J Timpson, Timothy M Frayling, Andrew T Hattersley, and Mark I McCarthy. A powerful approach to sub-phenotype analysis in population-based genetic association studies. *Genetic Epidemiology*, 34(4):335–343, 2010.
- [58] John Novembre and Benjamin M Peter. Recent advances in the study of fine-scale population structure in humans. *Current Opinion in Genetics & Development*, 41:98–105, 2016.
- [59] Nick Patterson, Alkes L Price, and David Reich. Population structure and eigenanalysis. *PLoS Genetics*, 2(12):e190, 2006.
- [60] William O Pickrell, Mark I Rees, and Seo-Kyung Chung. Next generation sequencing methodologies—an overview. In *Advances in protein chemistry and structural biology*, volume 89, pages 1–26. Elsevier, 2012.

- [61] Alkes L Price, Nick J Patterson, Robert M Plenge, Michael E Weinblatt, Nancy A Shadick, and David Reich. Principal components analysis corrects for stratification in genome-wide association studies. *Nature Genetics*, 38(8):904–909, 08 2006.
- [62] Jonathan K Pritchard, Matthew Stephens, and Peter Donnelly. Inference of population structure using multilocus genotype data. *Genetics*, 155(2):945–959, 2000.
- [63] Thilina Ranaweera, Enes Makalic, John L Hopper, and Adrian Bickerstaffe. An open-source, integrated pedigree data management and visualization tool for genetic epidemiology. *International Journal of Epidemiology*, 47(4):1034–1039, 2018.
- [64] Noah A Rosenberg, Lei M Li, Ryk Ward, and Jonathan K Pritchard. Informativeness of genetic markers for inference of ancestry. *The American Journal of Human Genetics*, 73(6):1402–1422, 2003.
- [65] Alejandro A Schäffer, Mathieu Lemire, Jürg Ott, G Mark Lathrop, and Daniel E Weeks. Coordinated conditional simulation with slink and sup of many markers linked or associated to a trait in large pedigrees. *Human heredity*, 71(2):126–134, 2011.
- [66] Daniel J Schaid, Charles M Rowland, David E Tines, Robert M Jacobson, and Gregory A Poland. Score tests for association between traits and haplotypes when linkage phase is ambiguous. *The American Journal of Human Genetics*, 70(2):425–434, 2002.
- [67] Jie Shen and Ping Li. A tight bound of hard thresholding. *The Journal of Machine Learning Research*, 18(1):7650–7691, 2017.
- [68] Eric Sobel, Kenneth Lange, Jeffrey R O’Connell, and Daniel E Weeks. Haplotyping algorithms. In *Genetic mapping and DNA sequencing*, pages 89–110. Springer, 1996.
- [69] Marc A Suchard, Shawn E Simpson, Ivan Zorych, Patrick Ryan, and David Madigan. Massive parallelization of serial inference algorithms for a complex generalized linear model. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 23(1):article10:1–17, 2013.
- [70] Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, et al. UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Medicine*, 12(3):e1001779, 2015.
- [71] Gulnara R Svishcheva, Tatiana I Axenovich, Nadezhda M Belonogova, Cornelia M van Duijn, and Yurii S Aulchenko. Rapid variance components–based method for whole-genome association analysis. *Nature Genetics*, 44(10):1166, 2012.
- [72] Amalio Telenti, Levi C. T. Pierce, William H. Biggs, Julia di Iulio, Emily H. M. Wong, Martin M. Fabani, Ewen F. Kirkness, Ahmed Moustafa, Naisha Shah, Chao Xie, Suzanne C. Brewerton, Nadeem Bulsara, Chad Garner, Gary Metzker, Efren Sandoval, Brad A. Perkins, Franz J. Och, Yaron Turpaz, and J. Craig Venter. Deep sequencing of 10,000 human genomes. *Proceedings of the National Academy of Sciences*, 113(42):11901–11906, 2016.
- [73] Erwin L Van Dijk, Hélène Auger, Yan Jaszczyszyn, and Claude Thermes. Ten years of next-generation sequencing technology. *Trends in Genetics*, 30(9):418–426, 2014.

- [74] Elisabeth M Van Leeuwen, Alexandros Kanterakis, Patrick Deelen, Mathijs V Kattenberg, Abdel Abdellaoui, Albert Hofman, Alexander Schönhuth, Androniki Menelaou, Anton JM de Craen, Barbara DC van Schaik, et al. Population-specific genotype imputations using minimac or impute2. *Nature Protocols*, 10(9):1285, 2015.
- [75] Paul M VanRaden. Efficient methods to compute genomic predictions. *Journal of Dairy Science*, 91(11):4414–4423, 2008.
- [76] Peter M Visscher, Matthew A Brown, Mark I McCarthy, and Jian Yang. Five years of GWAS discovery. *The American Journal of Human Genetics*, 90(1):7–24, 2012.
- [77] Peter M Visscher, Naomi R Wray, Qian Zhang, Pamela Sklar, Mark I McCarthy, Matthew A Brown, and Jian Yang. 10 years of GWAS discovery: biology, function, and translation. *The American Journal of Human Genetics*, 101(1):5–22, 2017.
- [78] Bowen Wang, Serge Sverdlov, and Elizabeth Thompson. Efficient estimation of realized kinship from SNP genotypes. *Genetics*, 210(2), 2017.
- [79] Kai Wang, Mingyao Li, and Maja Bucan. Pathway-based approaches for analysis of genomewide association studies. *The American Journal of Human Genetics*, 81(6):1278–1283, 2007.
- [80] Michael C Wu, Seunggeun Lee, Tianxi Cai, Yun Li, Michael Boehnke, and Xihong Lin. Rare-variant association testing for sequencing data with the sequence kernel association test. *The American Journal of Human Genetics*, 1(89):82–93, 2011.
- [81] Tong Tong Wu, Yi Fang Chen, Trevor Hastie, Eric Sobel, and Kenneth Lange. Genome-wide association analysis by lasso penalized logistic regression. *Bioinformatics*, 25(6):714–721, 2009.
- [82] Fan Yang, Rina Foygel Barber, Prateek Jain, and John Lafferty. Selective inference for group-sparse linear models. In *Advances in Neural Information Processing Systems*, pages 2469–2477, 2016.
- [83] Jian Yang, Teresa Ferreira, Andrew P Morris, Sarah E Medland, Pamela AF Madden, Andrew C Heath, Nicholas G Martin, Grant W Montgomery, Michael N Weedon, Ruth J Loos, et al. Conditional and joint multiple-SNP analysis of GWAS summary statistics identifies additional variants influencing complex traits. *Nature Genetics*, 44(4):369, 2012.
- [84] Jian Yang, S. Hong Lee, Michael E. Goddard, and Peter M. Visscher. GCTA: A tool for genome-wide complex trait analysis. *American Journal of Human Genetics*, 88(1):76–82, 2011.
- [85] Xiao-Tong Yuan, Ping Li, and Tong Zhang. Gradient hard thresholding pursuit. *Journal of Machine Learning Research*, 18:166–1, 2017.
- [86] Xiguo Yuan, David J Miller, Junying Zhang, David Herrington, and Yue Wang. An overview of population genetic data simulation. *Journal of Computational Biology*, 19(1):42–54, 2012.
- [87] Hua Zhou, David Alexander, and Kenneth Lange. A quasi-newton acceleration for high-dimensional optimization algorithms. *Statistics and Computing*, 21(2):261–273, 2011.
- [88] Hua Zhou, David H Alexander, Mary E Sehl, Janet S Sinsheimer, EM Sobel, and K Lange. Penalized regression for genome-wide association screening of sequence data. In *Biocomputing 2011*, pages 106–117. World Scientific, 2011.

- [89] Hua Zhou, John Blangero, Thomas D. Dyer, Kei-hang K. Chan, Kenneth Lange, and Eric M. Sobel. Fast genome-wide QTL association mapping on pedigree and population data. *Genetic Epidemiology*, 41(3):174–186, 2017.
- [90] Hua Zhou, Liuyi Hu, Jin Zhou, and Kenneth Lange. MM algorithms for variance components models. *Journal of Computational and Graphical Statistics*, accepted, 2018.
- [91] Hua Zhou, Mary E Sehl, Janet S Sinsheimer, and Kenneth Lange. Association screening of common and rare genetic variants by penalized regression. *Bioinformatics*, 26(19):2375–2382, 2010.
- [92] Jin J Zhou, Tao Hu, Dandi Qiao, Michael H Cho, and Hua Zhou. Boosting gene mapping power and efficiency with efficient exact variance component tests of SNP sets. *Genetics*, 204(3):921–931, 2016.
- [93] Jin J. Zhou, Janet S Sinsheimer, Michael H. Cho, Peter Castaldi, and Hua Zhou. MMVC: An efficient mm algorithm to quantify genetic correlations across large number of phenotypes in giant datasets. manuscript in preparation, 2018.
- [94] Xiang Zhou and Matthew Stephens. Efficient multivariate linear mixed model algorithms for genome-wide association studies. *Nature Methods*, 11(4):407–409, 2014.
- [95] Xiaofeng Zhu, ShuangLin Zhang, Hongyu Zhao, and Richard S Cooper. Association mapping, using a mixture model for complex traits. *Genetic Epidemiology*, 23(2):181–196, 2002.