Genetic Epidemiology 29: 36–50 (2005)

# Association Testing with Mendel

**Kenneth Lange,[1,2*] Janet S. Sinsheimer,[1,2] and Eric Sobel[2]**

[1]*Department of Biomathematics, David Geffen School of Medicine at UCLA, Los Angeles, California*
[2]*Department of Human Genetics, David Geffen School of Medicine at UCLA, Los Angeles, California*

This report presents an overview of association testing strategies from a user's perspective, with particular attention to the capabilities of the computer program Mendel. Association testing is driven by the nature of the study sample, the nature of the disease trait, and the kind of markers employed. The practicing statistician must also choose whether to conduct parametric or nonparametric tests. Because of the complexities involved, Mendel offers users several analysis options. The different options are tied together by shared input and output conventions and a shared language for defining models. Mendel also features new statistics and theory found in no other genetics software. The most important innovations include: association testing by penetrance estimation, expansion of matched-pair designs to permutation unit designs, and a rigorous implementation of the measured genotype approach for quantitative trait loci. This report explains how Mendel imputes allele counts and conducts both asymptotic and permutation tests in the measured genotype framework. *Genet. Epidemiol.* 29:36–50, 2005. © 2005 Wiley-Liss, Inc.

**Key words: statistical software; case-control analysis; allele and haplotype frequencies; penetrance; permutation test; TDT; gamete competition; QTL; measured genotype; variance components; quantitative association**

Contract grant sponsor: USPHS; Contract grant numbers: GM053275, MH059490, and MH064205.
*Correspondence to: Kenneth Lange, Department of Human Genetics, David Geffen School of Medicine at UCLA, 695 Charles Young Drive South, Los Angeles, CA 90095-7088. E-mail: klange@ucla.edu
Received 5 May 2004; Accepted 17 January 2005
Published online 15 April 2005 in Wiley InterScience (www.interscience.wiley.com).
DOI: 10.1002/gepi.20073

## INTRODUCTION

In contrast to clinical trials, there are no well-established protocols in genetic epidemiology. Genetic studies are technology driven, and as long as genome technology and information change rapidly, genetic epidemiology will remain in a state of flux. For instance, as genotyping increases in scope and drops in price, we are seeing an accelerating shift from linkage analysis to association testing [Risch and Merikangas, 1996]. The simplifying assumption of linkage equilibrium among marker loci is no longer valid for the dense genome scans of modern association studies. Geneticists are also measuring new phenotypes and actively looking for genetic and environmental interactions. Within this changing landscape, we have undertaken a massive overhaul of our genetic analysis software Mendel [Lange et al., 2001].

The current report summarizes our efforts to upgrade the association testing capabilities of Mendel. The upgrade itself is a complex task because it involves sorting through the enormous literature on association methods and making difficult judgment calls on what methods to include and exclude. There are no road maps or consumer guides to steer the process, so it seemed to us that laying out the philosophy behind our choices and explaining how users interact with Mendel might serve a positive purpose. It is probably no great injustice to say that most software in genetic epidemiology is poorly documented. Users are perplexed by both the input conventions and the interpretation of output. Such confusion decreases the chance of mapping disease genes and erects barriers to the entry of young scientists into the field.

Of the current 22 Mendel analysis options, this report covers the 6 devoted to association testing and emphasizes new features not available elsewhere. These features include: testing by penetrance estimation, expansion of matched-pair designs to permutation unit designs, and implementation of the measured genotype approach for

quantitative trait loci. Further detailed documentation is distributed with the program.

## GENERAL BACKGROUND ON ASSOCIATION TESTING

Study designs vary from trait to trait and investigator to investigator. Traditionally, most association studies have relied on case-control samples. Although genetic epidemiologists usually make no attempt to match cases and controls, doing so can circumvent problems of ethnic stratification. For economic reasons, DNA pooling of cases and controls is an attractive option, at least in the first stages of a genome association scan [Sham et al., 2002]. Because haplotype phase information is lost and meaningful consideration of covariates is sacrificed, pooling seldom plays a role in the end game of disease gene mapping. In addition to case-control studies, association testing is also carried out on parent-child trios, sibships, and occasionally full pedigree data. Indeed, many studies contain a mix of population and family data. Integrating such data is a challenge. Clearly, it would be useful in association analysis to combine data on the multiple affecteds within extended pedigrees with case-control data. This inevitably raises the question of dependence among the affecteds. As discussed later, Mendel addresses most of these restrictions.

Study design must also take into account the nature of the phenotypic data. For instance, some traits are qualitative and some quantitative. When geneticists investigate subclinical indicators rather than simple affected-unaffected dichotomies, they almost always deal with multivariate quantitative traits. Although most geneticists feel the more data the merrier, they do not always make the best use of the entire spectrum of measurements. Investigators sometimes lack relevant statistical methods, sometimes appropriate analysis software, and sometimes both.

Finally, there is the issue of microsatellites versus SNP (single nucleotide polymorphism) markers. The sheer abundance of SNPs and the economies of scale in genotyping them are pushing them quickly to the forefront. Because the formation of SNP haplotypes involves phase ambiguities, this trend has had the unintended effect of resurrecting the problem of non-codominant markers when adjacent markers are combined into super-markers for further analysis. Mendel, incidentally, has this capability. For example, an individual who has the heterozygous genotype 1/2 at two adjacent SNPs can be explained equally well by either of the haplotype pairs 11/22 or 12/12. The benefits of codominant alleles was one of the reasons that geneticists embraced microsatellites and dropped classical blood group markers years ago. In our view, analysis software should be structured to avoid the common practice of imputing best haplotypes and using these in statistical analysis. Instead, the analysis should consider all configurations, weighting each according to its likelihood.

The statistics employed in association testing are selected to match the nature of the trait and the nature of the study sample. For case-control samples, the obvious tactic is to compare allele or haplotype frequencies in cases and controls. This can be accomplished via a likelihood ratio test or an exact test that conditions on marginal allele counts in a contingency table. Mendel covers both options. If ethnic stratification is a concern and parent-child trios are available, then application of the transmission/disequilibrium test (TDT) is now standard [Ewens and Spielman, 1995; Rubenstein et al., 1981; Spielman et al., 1993; Terwiller and Ott, 1992]. The gamete competition model is a parametric version of the TDT appropriate for both parent-child trios and whole pedigrees [Sinsheimer et al., 2000, 2001]. Versions of the TDT and gamete competition tests exist for quantitative traits [Horvath and Laird, 2002; Sinsheimer et al., 2001].

When ethnic stratification can be ignored, we prefer to model the allelic effects acting on a quantitative trait as mean effects in a variance component setting. Later in this article, we will explain in detail a likelihood ratio test that accommodates multivariate traits, probands, polygenic background, missing traits, and recessive and dominant alleles. Our procedure is an elaboration of the measured genotype approach advocated by Boerwinkle et al. [1986]. Because the distributional assumptions of the measured genotype test often fail in practice, we present a novel permutation version of the test as well.

For affected-unaffected dichotomies, comparison of genotype-specific penetrances is an attractive approach to association testing, particularly when covariate effects are large. The latest version of Mendel allows for rudimentary estimation of disease penetrances from ascertained pedigree data. Unfortunately, we have not yet fully married pedigree analysis and logistic regression, so

simultaneous estimation of covariate effects will have to wait for a later release.

Statistical inference can be broken down into neat categories such as Bayesian, frequentist, and nonparametric. Each of these find a home in Mendel, but for reasons of computational efficiency, Bayesian methods are the least developed. Frequentist methods share with Bayesian methods a strong dependence on parametric models. In some genetic settings, models are justified. In other settings, models are more suspect. Astute statisticians look for a convergence of the evidence. Departures from large sample assumptions are also a concern with frequentist inference. Students learn the hard way that faulty inference is often the consequence of low allele and genotype frequencies. For this reason, we have adopted permutation and gene dropping procedures for hypothesis testing and $P$ value estimation in several options of Mendel. The greater robustness of nonparametric procedures often compensates for their loss of power.

## PHILOSOPHY BEHIND MENDEL

The above quick overview of statistical inference in testing genetic association omits mention of many statistical problems. Some of these problems Mendel is designed to handle and some it is not. For instance, Mendel makes no overt corrections for multiple testing. In our view, the best remedy for false positives is replication. Another problem that Mendel partially dodges is the confounding of linkage and association. In defense of this attitude, most geneticists do not care which of the two explains a significant $P$ value in a genome scan. Mendel's most recent addition, discussed in Cantor et al. [2005], tests for association in the presence of linkage in the classical setting of single point analysis.

Other problems Mendel tackles head on. These include (1) the presence of probands, which is handled by conditioning, (2) missing data, either traits, covariates, or genotypes, (3) missing phases for haplotypes, and (4) inordinate computation times. One innovation in some Mendel options is the automatic selection at runtime of either the Elston-Stewart algorithm [Elston and Stewart, 1971; Lange and Elston, 1975] or the Lander-Green-Kruglyak algorithm [Kruglyak et al., 1996; Kruglyak and Lander, 1998; Lander and Green, 1987] for likelihood evaluation on each pedigree. Because many pedigrees are still far beyond the capacity of deterministic likelihood algorithms, we have written parallel options in the Markov chain Monte Carlo program SimWalk to handle large pedigrees [Sobel and Lange, 1996]. Most of these analogous options are devoted to linkage analysis. The programs Mendel and SimWalk cooperate in other ways, and we are constantly striving to improve the interface between them.

What then are the design objectives of Mendel? First, we have endeavored to build a common platform for many problems in genetic epidemiology and population genetics. The same program can now be used for linkage analysis, association testing, variance components, and QTL (quantitative trait locus) mapping. Second, we have tried to abide by clear, flexible conventions for input and output. Third, we have written extensive documentation, to which this article contributes. Fourth, whenever possible, we have programmed analyses for general pedigrees. Any unnecessary restriction on pedigree size or shape wastes information and limits conclusions. Fifth and finally, we have tried whenever possible to use state of the art tests and estimation procedures.

When the last two objectives come into conflict, we almost always come down on the side of generality at the expense of application of closed-form estimators and test statistics. In our view, the field of genetic epidemiology is overly driven by procedures for sibling pairs and nuclear families simply because these procedures can be encapsulated in precise formulas. As computers grow in speed, there is little excuse for pushing methods of analysis just because we can do the relevant mathematics and identify the best estimators and most powerful tests for a minority of the data. We should not hesitate to use a computational sledge hammer when paper and pencil prove inadequate.

## ASSOCIATION TESTS BY COMPARISON OF ALLELE FREQUENCIES

The traditional method of testing for association in case-control data is to compute maximum likelihood estimates of allele frequencies in the two separate populations and compare these to the estimates in the combined population. The null hypothesis that the frequencies are the same in cases and controls can then be tested by a likelihood ratio statistic, which in this case is equivalent to a $\chi^2$-test of homogeneity. To illustrate

the test, the Mendel documentation features a classic example involving association between alleles of the ABO blood group and peptic ulcer [Clarke et al., 1959]. Execution of Mendel is initiated by a control file, which names relevant input and output files, defines the analysis option, and in this example, designated 6a in the documentation, tells Mendel how to distinguish between cases and controls via a data field labeled HEALTH.

The Mendel summary file for this ABO/peptic-ulcer example is shown in Figure 1. There we see the estimated allele frequencies for the combined population and a borderline significant likelihood ratio test of homogeneity between the two populations.

Several remarks are worth making about the current option. First, there is no need for the data to consist of isolated individuals; pedigree structure is respected as advocated by Boehnke [1991]. A copy number can be assigned to a pedigree to avoid repeating identical pedigrees in the pedigree file. This convenient feature is particularly pertinent to a random sample where many cases or controls have the same genotype. Second, because the likelihood ratio test relies on large sample theory, any significant result is suspect if one or more estimated allele frequencies fall too low. Fortunately, Mendel has an option that facilitates the combination of rare alleles. As a rule of thumb, each allele should be present in each population a minimum number of times, say between 3 and 5 times. By default, the Mendel output includes a count of each allele and phenotype in the data set. Third, this option illustrates Mendel's abilities in optimizing likelihoods and dealing with dominant and recessive alleles. Fourth, if you distrust the implicit assumption of Hardy-Weinberg equilibrium in this analysis, it is possible to redo the analysis based on genotypes or phenotypes rather than alleles. For instance, in the APOE-Alzheimer's example described in our section on Association Tests with Permutation and Matching, association is sought with genotypes, and no assumption of Hardy-Weinberg equilibrium is invoked.

The current option also permits estimation of haplotype frequencies and testing for linkage equilibrium. Example 6d of the Mendel documentation considers three adjacent SNPs. If we omit allele frequencies for these SNPs in the locus input file, then Mendel will fill them in by gene counting and search for the best estimates of the eight haplotype frequencies starting from linkage equilibrium values. The corresponding summary file for this example contains the estimated haplotype frequencies as shown in Figure 2.

| FREQUENCY | STD ERROR | HAPLOTYPE |
|---|---|---|
| 0.00000 | 0.00000 | A--A--C |
| 0.01250 | 0.00884 | G--A--C |
| 0.00000 | 0.00000 | A--C--C |
| 0.00250 | 0.00537 | G--C--C |
| 0.46250 | 0.03549 | A--A--T |
| 0.43000 | 0.03501 | G--A--T |
| 0.09250 | 0.02088 | A--C--T |
| 0.00000 | 0.00000 | G--C--T |

**Fig. 2. Haplotype frequency estimates, taken from Mendel's summary file for example data set 6d.**

| LOCUS NAME | ALLELE NAME | ESTIMATED FREQUENCY | STANDARD ERROR |
|---|---|---|---|
| ABO | A | 0.2335 | 0.0093 |
| ABO | B | 0.0588 | 0.0049 |
| ABO | O | 0.7077 | 0.0099 |

| LOCUS NAME | HOMOGENEITY CHI-SQUARE | DEGREES OF FREEDOM | ASYMPTOTIC P-VALUE | WRIGHT'S FST |
|---|---|---|---|---|
| ABO | 7.0106 | 2 | 0.030038 | 0.0097 |

**Fig. 1. Allele frequency estimates for ABO in a peptic-ulcer case-control study, taken from Mendel's summary file for example data set 6a.**

```
THE LIKELIHOOD RATIO STATISTIC FOR LINKAGE EQUILIBRIUM
EQUALS 11.13.  THIS CHI-SQUARE STATISTIC HAS 4 DEGREES OF
FREEDOM AND P-VALUE 0.02510 IF ALL HAPLOTYPE FREQUENCY
ESTIMATES ARE POSITIVE.
```

**Fig. 3. Results of a test for linkage disequilibrium, taken from Mendel's standard output file for example data set 6d.**

Whenever an estimated frequency falls on the zero boundary, Mendel automatically assigns the estimate a standard error of zero. At the bottom of the standard output file for this example data set as seen in Figure 3, Mendel suggests that linkage equilibrium fails. However, note the violation of the stated condition that estimated haplotype frequencies should be positive. One can resolve ambiguous test results by invoking a different Mendel option that conducts a nonparametric test for linkage equilibrium. Alternatively, one can construct a single super locus from the SNPs, combine the rare alleles of the super locus, and re-estimate allele frequencies. The estimated allele frequencies are then, in effect, estimates of haplotype frequencies or combined-haplotype frequencies.

## ASSOCIATION TESTS BY PENETRANCE ESTIMATION

Another approach to association testing is through the estimation of penetrance functions. Suppose we assign each allele $i$ of a candidate gene a penetrance parameter $\phi_i$, and let the disease penetrance of the genotype $i/j$ be the product $\phi_i\phi_j$. We can then estimate each $\phi_i$ by maximum likelihood and test the null hypothesis of no association, namely all $\phi_i = \phi$, by a likelihood ratio test. The penetrance estimation option of Mendel combines the virtues of parsimony and flexibility. It can exploit both pedigree data and randomly sampled individuals. Like the gamete competition model to be discussed later, it can accommodate non-codominant markers such as those constructed from multiple linked SNPs. If study subjects are taken from a single ethnic group, it is prudent to estimate allele frequencies simultaneously. If multiple ethnic groups are involved, Mendel does not permit estimation of allele frequencies but does permit different fixed allele frequencies to be employed for each ethnic group. In our view, the application of ethnic-specific allele frequencies for pedigree founders adequately safeguards against population stratification. Mendel's computation times are reasonable with as many as seven or eight alleles. If there are more alleles or some alleles are rare, one should lump the least frequent alleles.

Obviously, the user must communicate to Mendel who is affected and who is not. Simple commands such as

```
AFFECTED_LOCUS_OR_FACTOR = HEALTH
AFFECTED = 2
```

in the control file tell Mendel in which phenotypic field to look for disease status, and which label is used for the affecteds. In the face of uncertainty over someone's disease status, it is probably better to be conservative and list his or her disease phenotype as unknown (blank).

Each pedigree may contain a proband through which it is ascertained. Mendel needs to know where to read proband status and which symbol designates a proband. The commands

```
PROBAND_FACTOR = PROBAND
PROBAND = 1
```

in the control file achieve this goal. With these keyword values, a pedigree contains a proband if one of its members has a 1 in the field `PROBAND`. The penetrance estimation option of Mendel corrects for ascertainment by conditioning on a special proband pedigree automatically appended to the list of pedigrees. This purely artificial pedigree of unrelated probands is not reported in the standard output file. During a likelihood search, the loglikelihood of the proband pedigree is subtracted from the sum of the loglikelihoods of the actual pedigrees. Mendel's penetrance estimation option allows at most one proband per pedigree. Multiple probands per pedigree are permitted in the QTL association option discussed later.

Like many analysis options of Mendel, the penetrance option has several suboptions, or models. Table I shows the eight models possible in penetrance estimation. The first six of these models are intended for biallelic loci. In Table I, $\phi_{i/j}$ denotes the penetrance of genotype $i/j$.

**TABLE I. Mendel's models for penetrance estimation**

| Model number | Allele frequencies estimated by Mendel | Disease model | Penetrance restriction |
|---|---|---|---|
| 1 | Yes | Unrestricted | None |
| 2 | No | Unrestricted | None |
| 3 | Yes | Dominant | $\phi_{1/2}=\phi_{2/2}$ |
| 4 | No | Dominant | $\phi_{1/2}=\phi_{2/2}$ |
| 5 | Yes | Recessive | $\phi_{1/1}=\phi_{1/2}$ |
| 6 | No | Recessive | $\phi_{1/1}=\phi_{1/2}$ |
| 7 | Yes | Multiplicative | $\phi_{i/j}=\phi_i\phi_j$ |
| 8 | No | Multiplicative | $\phi_{i/j}=\phi_i\phi_j$ |

## ASSOCIATION TESTS WITH PERMUTATION AND MATCHING

In testing for disease-marker association in ethnically heterogeneous populations, permutation procedures offer one of the best avenues for generalizing case-control studies. Testing can proceed through matched groups of affected cases and normal controls. Besides permitting the usual paired designs for cases and controls, matching can also accommodate sibling groups of affected and normal children from large disease pedigrees. In determining $P$ values, permutation of case and control labels is performed only within each defined permutation unit and not across permutation units. Matching protects against ethnic stratification and against information loss when one representative of several statistically dependent cases is arbitrarily selected for analysis. For example, most linkage studies involve several affecteds per pedigree. In a subsequent association study with the same data, defining permutation units based on sibships with at least one affected and unaffected, salvages many of the affecteds from these pedigrees. Of course, if one starts with random samples of cases and controls from an ethnically homogeneous population, it would be foolish to impose matching and limit the scope of possible permutations.

Mendel's cases and controls option operates by constructing a contingency table of haplotype or multilocus-phenotype counts. A multilocus phenotype can involve codominant genotypes at some loci and non-codominant phenotypes at other loci. Unfortunately, one cannot mix haplotype and phenotype data. The two rows of the contingency table correspond to cases and controls, and the columns correspond to particular haplotypes or phenotypes, whichever paradigm is selected. To assess case-control homogeneity, Mendel evaluates two test statistics. Fisher's exact test is a good omnibus test across many cells. The $Z_{\max}$ test has greater power when one or two cells deviate strongly between the populations. The theory behind these tests is explained in Lange [2002], so we will focus here only on the permutation procedures used. It is worth emphasizing that the tests involve no hidden assumptions of either Hardy-Weinberg or linkage equilibrium.

In default mode, Mendel evaluates $P$ values for Fisher's exact test and the $Z_{\max}$ test by permutation of case-control labels across an entire random sample of people. In this instance, each person is assigned to a single-person pedigree, which as mentioned earlier may carry a copy number. If phenotypes are compared, then as mentioned earlier dominant or recessive alleles are allowed at the participating markers. People with partially missing phenotypes or haplotypes are ignored. Because Mendel in effect conditions on the margins of the contingency table, all allele frequencies are irrelevant. Likewise, all recombination fractions are irrelevant. To analyze X-linked phenotypes, males and females should be assigned to different permutation units, and male phenotypes should be clearly distinguished from female phenotypes.

As an illustration of how Mendel operates, we consider case-control data from Columbia [Jacquier et al., 2001] and Japan [Yamagata et al., 1997] on the association between APOE genotypes and Alzheimer's disease. This is example data set 12c from the Mendel documentation. A typical pedigree from this data set is shown in Figure 4, which is extracted from the standard Mendel output for this example. If we permute case-control labels

```
PEDIGREE NUMBER 17 HAS 1 MEMBER AND NAME PED#0017.
IT IS PRESENT IN 3 COPIES.

ID  PARENT  IDS  SEX  TWIN  APOE  DISEASE  COUNTRY
 1               F          3/4  control  Columbia
```

**Fig. 4. A typical case-control pedigree, taken from Mendel's standard output file for example data set 12c.**

across the data set combining the two countries, we clearly run the risk of detecting a spurious association. To avoid this problem, we use a matched test that defines each country as a different permutation unit. The entries in the control file

```
AFFECTED_LOCUS_OR_FACTOR = DISEASE
AFFECTED = CASE
GROUP_FACTOR = COUNTRY
```

define cases and controls and the two permutation units (Columbia and Japan).

Figure 5 presents Mendel's results showing the well-known increased prevalence of the 3/4 genotype. One can reanalyze the data for allelic effects by inputting case-control alleles rather than case-control genotypes.

```
LOCUS NAME                :  APOE
FISHER P-VALUE            :  0.0000000
ZMAX P-VALUE              :  0.0000000
MOST ABERRANT PHENOTYPE   :  3/4
CASE SAMPLE SIZE          :  246
CONTROL SAMPLE SIZE       :  242
```

**Fig. 5. Results of an association test, using country of origin as matched permutation units, for APOE genotypes and Alzheimer's disease status, taken from Mendel's standard output file for example data set 12c.**

## TDT AND GAMETE COMPETITION MODEL

As its name implies, the transmission/disequilibrium test (TDT) controls for ethnic stratification by focusing on the alleles transmitted to affected children rather than the alleles present in all affecteds. This perspective makes allele frequencies irrelevant in computing $P$ values under the TDT. Of course, the power of the TDT to detect disease association with a particular allele does depend on the population frequency of that allele. The TDT is best explained by considering a contingency table of allele counts. The two rows of the contingency table, as seen in Figure 6, correspond to parental alleles passed and not passed. The columns correspond to particular alleles. Mendel avoids large sample approximations to $P$ values in the TDT and instead permutes the labels indicating which alleles are passed in each parent-child combination.

The results seen in Figure 6 demonstrates that Mendel uses a $\chi^2$ statistic and a $Z_{\max}$ statistic. The latter statistic makes it possible to identify the most aberrant allele. The $P$ values of the tests based on these statistics are computed by independent Monte Carlo permutations. The number of such permutations is determined by the keyword SAMPLES in the control file. The more permutations employed, the shorter the 95% confidence interval surrounding the $P$ value approximation reported by Mendel.

```
                RESULTS FOR LOCUS D7S3058:

ALLELE NUMBER   1   2   3   4   5   6   7

PASSED         51  63  38  40  76  80  52
NOT PASSED     65  67  46  66  54  71  31

THE CHI-SQUARE TDT STATISTIC HAS APPROXIMATE P-VALUE
0.0128 PLUS OR MINUS 0.0022 BASED ON 10000 RESAMPLES.

THE LARGEST STANDARDIZED RESIDUAL TDT STATISTIC HAS
APPROXIMATE P-VALUE 0.0760 PLUS OR MINUS 0.0053 BASED
ON 10000 RESAMPLES.

THE MOST UNDER-TRANSMITTED ALLELE IS ALLELE NUMBER 4.
THE MOST OVER-TRANSMITTED ALLELE IS ALLELE NUMBER 7.
OF THESE TWO, ALLELE NUMBER 4 IS MORE EXTREME.
```

**Fig. 6. Typical TDT results showing the use of $\chi^2$ and $Z_{\max}$ statistics, taken from Mendel's standard output file for example data set 13.**

If one parent of a parent-child trio is untyped, but the remaining parent and the affected child are different heterozygotes, then the typed duo contributes unbiased information and is used by Mendel in the TDT [Curtis and Sham, 1995]. Because Mendel uses all typed parent-offspring trios from a pedigree, it can confound linkage and association, particularly when the data consist of one or two large disease pedigrees already showing linkage. With many small unrelated pedigrees, the chance of confusing linkage with association becomes less of an issue, and the TDT statistics help to identify associated marker alleles.

The gamete competition model [Sinsheimer et al., 2000, 2001] is a parametric substitute for the TDT motivated by the Bradley-Terry model for ranking items such as sports teams in a league [Bradley and Terry, 1952; Keener, 1993; Lange et al., 1988]. If allele $i$ of a marker locus is assigned parameter $\tau_i$ in the gamete competition model, then

$$Pr(i/j \to i) = \frac{\tau_i}{\tau_i + \tau_j}$$

is the probability that a parent with heterozygous genotype $i/j$ transmits allele $i$ to an affected child. Mendelian segregation corresponds to the choice $\tau_i = 1$ for all $i$. To test whether Mendelian segregation is true, one can estimate the $\tau$'s from pedigree data by maximum likelihood and perform a likelihood ratio test. Because the transmission probabilities $Pr(i/j \to i)$ are invariant under multiplication of the $\tau$'s by a common constant, it is necessary to impose a constraint such as $\tau_1 = 1$.

The gamete competition model accommodates quantitative as well as qualitative outcomes, allows for covariates, and makes effective use of full pedigree data. For example, we can use a disease severity index $x_k$ on child $k$ by writing $\tau_i = e^{\omega_i x_k}$. This exponential reparameterization has the advantage of eliminating the positivity constraint $\tau_i > 0$. Mendelian segregation in this setting corresponds to the choice $\omega_i = 0$ for all $i$. In maximum likelihood estimation, one should impose a constraint such as $\omega_1 = 0$ and center the $x_k$ so that they have mean 0.

In contrast to the TDT, the gamete competition model does not require codominant marker alleles. This is a major advantage when neighboring SNPs are combined into a super-locus. The gamete competition model neatly circumvents the phase ambiguities of a SNP super-locus by including all possible haplotype phases in likelihood evaluation. This is a virtue of a likelihood-based method compared to a nonparametric method. Of course, likelihood methods do require good estimates of haplotype and allele frequencies. Mendel permits allele frequencies to be estimated simultaneously with transmission parameters or to be fixed at previously determined population estimates. Pedigree founders can be assigned ethnic specific allele frequencies if ethnic stratification is considered to be an issue.

As an example, consider the classical blood group data of Lewis et al. [1980], incorporated as example 8a in the Mendel package. Figure 7, showing part of Mendel's summary results for these data, gives the most and least frequently transmitted allele and their estimated $\tau$'s. For these data, none of the gamete competition likelihood ratio tests is significant. Allele frequency estimates and the asymptotic standard errors of all parameters can be found in the standard results file.

Example 8b from the Mendel documentation applies the gamete competition model to a quantitative trait. Data on a 287 base-pair insertion/deletion polymorphism in the angiotensin-1 converting enzyme (ACE) gene are spread over 404 people in 69 families [Keavney et al., 1998]. The deletion allele is associated with high plasma ACE activity. We name the quantitative trait and separately standardize the male and female values to have mean 0 and variance 1 through the commands

```
QUANTITATIVE_TRAIT = ACE
TRANSFORM = Degender :: ACE
```

in the control file. (If we wanted to standardize ACE levels ignoring sex, we would instead use `TRANSFORM = Standardize :: ACE`.) Figure 8,

| MARKER NAME | P-VALUE | MIN TAU | ALLELE NAME | MAX TAU | ALLELE NAME |
|---|---|---|---|---|---|
| RADIN | 0.35125 | 1.00000 | - | 1.25564 | + |
| PGM1 | 0.57820 | 0.77876 | 1- | 1.23902 | 2- |
| RH | 0.62572 | 0.81730 | R2 | 1.60091 | R0 |

**Fig. 7. Results of a gamete competition analysis for a qualitative trait, taken from Mendel's summary file for example data set 8a.**

| MARKER NAME | P-VALUE | MIN OMEGA | ALLELE NAME | MAX OMEGA | ALLELE NAME |
|---|---|---|---|---|---|
| INS | 0.00000 | −1.30555 | 1 | 0.00000 | 2 |

**Fig. 8. Results of a gamete competition analysis for a quantitative trait, taken from Mendel's summary file for example data set 8b.**

part of Mendel's summary results, confirms strong over-transmission of the deletion allele, allele 2, to children with high ACE activity levels. Again, fuller output appears in the standard results file.

## QTL ASSOCIATION TESTS BY VARIANCE COMPONENTS

Quantitative traits are inherently more informative than disease dichotomies. Association tests with quantitative traits often rely on statistical procedures such as analysis of variance that neglect familial correlations. A better approach is to view association testing from the perspective of variance components and to impose genotype or allele-specific effects on trait means. This measured genotype approach controls for polygenic background while remaining in the frequentist domain of maximum likelihood estimation and likelihood ratio tests [Fan et al., 2005; George and Elston, 1987; Hopper and Matthews, 1982; SAGE, 2001]. Two objections can be made to this strategy. First, pedigree data may well be ascertained and not random. Second, it is unclear how to proceed in the presence of missing genotype and phase data. The first objection can be partially overcome by conditioning the observations on the trait values of probands. Given codominant markers, the second objection can be handled by using only fully genotyped people. Unfortunately, some of the most informative markers are constructed by combining adjacent SNPs. The phase ambiguities at SNP combination markers mask the underlying haplotypes. If we retain people in a pedigree sample who cannot be fully haplotyped, then the likelihood of the data must be expressed as a mixture of multivariate distributions. Admixture enormously complicates likelihood evaluation and statistical inference [Hasstedt, 1982].

As a resolution of this dilemma, one can substitute conditional probabilities of genotypes for genotypes. For markers with more than two alleles, it is simplest to assume additive effects on the mean. With additive allele effects, one must compute the expected number of marker alleles of each type carried by each person conditional on the marker genotypes observed throughout his or her pedigree. These conditional expectations can then be viewed as covariates. Curiously enough, the necessary expectations can be computed by numerical differentiation. Recall that the likelihood $L$ of a pedigree with $n$ members is usually represented as

$$L = \sum_{G_1} \cdots \sum_{G_n} \prod_i \text{Pen}(X_i|G_i) \prod_j \text{Prior}(G_j) \times \prod_{\{k,l,m\}} \text{Tran}(G_m|G_k, G_l), \quad (1)$$

where the $i$th person has phenotype $X_i$ and possible genotype $G_i$, the product on $j$ is taken over all founders, and the product on $\{k, l, m\}$ is taken over all parent-offspring triples [Ott, 1974]. The abbreviations Pen, Prior, and Tran in formula (1) refer to the penetrance, prior, and transmission functions of the model.

Now suppose that we include a factor $\beta^{c(a,G_p)}$ in the likelihood (1), where $c(a, G_p)$ counts the number of marker alleles of type $a$ in the marker genotype $G_p$ of person $p$. Then the derivative $\frac{d}{d\beta}\ln L(1)$ of the loglikelihood with respect to the artificial parameter $\beta$ becomes the expected number of type $a$ alleles carried by person $p$ conditional on the observed marker genotypes throughout the pedigree. This trick works equally well for codominant markers and SNP combination markers displaying dominance. Observe that quantitative trait values do not enter into these computations. Once we construct covariates giving the expected numbers of alleles, we can test for additive allelic effects by a likelihood ratio statistic in a purely variance component setting.

Single-point estimation of allele counts does not take into account evidence from neighboring markers. SimWalk uses the data on all markers simultaneously to create a Mendel readable file containing the imputed allele counts at each locus for each individual. SimWalk achieves this feat by

a Markov chain Monte Carlo (MCMC) method that operates on descent states rather than descent graphs, the underlying inheritance vectors deployed in the Lander-Green-Kruglyak algorithm. Descent states not only supply the paths the founder genes take as they descend through the pedigree, they also identify the alleles traveling down those paths. SimWalk's reliance on descent states makes it possible to capture mistyping as well. Mistyping is a common occurrence, and SimWalk handles errors that are either consistent or inconsistent with Mendelian inheritance [Sobel et al., 2002]. Each possible allele is weighted according to its posterior probability of being part of the true genotype.

Of course, users of Mendel are shielded from these complicated preliminary computations. The user's role is to provide the data and to communicate an analysis model through the control file. For example, sample problem 20a included in the Mendel package, uses the control file:

```
ANALYSIS_OPTION = Qtl_association
LOCUS_FILE = Locus20a.in
MAP_FILE = Map20a.in
PEDIGREE_FILE = Ped20a.in
VARIABLE_FILE = Variable20a.in
SUMMARY_FILE = Summary20a.out
OUTPUT_FILE = Mendel20a.out
QUANTITATIVE_TRAIT = Gc_conc
PREDICTOR = Grand :: Gc_conc
PREDICTOR = Sex :: Gc_conc
PREDICTOR = Age :: Gc_conc
COVARIANCE_CLASS = Additive
COVARIANCE_CLASS = Environmental
```

The purposes of the new keywords appearing here are worth discussing. First, the quantitative trait Gc_conc is named. This trait represents the plasma concentration of the human group specific component. The Gc locus determines qualitative variation in the GC transport protein for vitamin D. A question of some interest is whether the genotypes at the Gc locus also determine quantitative differences in plasma concentrations. Data bearing on this question appear in an article by Daiger et al. [1984] and are reproduced in the pedigree file in this data set. The study sample consists of 31 monozygous twin pairs, 13 dizygous twin pairs, and 45 unrelated controls. Gc concentrations and Gc genotypes are available on all individuals. The two Gc alleles, 1 and 2, are codominant.

Although it is not in evidence here, one of the strengths of Mendel is its ability to handle multivariate quantitative traits and impose linear models on their means. The commands in the control file dictate a grand mean and regression on sex and age for Gc_conc. A quantitative trait is used for a person provided it and all of its relevant predictors are present in the data. Thus, some traits may be used and some ignored on the same person.

Finally, Mendel fits variance components as well as mean components. The current version of Mendel supports five choices: additive polygenic, dominance polygenic, a QTL component, a household component, and random environment. The commands in the above example activate the first and last of these components. For QTL mapping, the conditional kinship coefficients required at each map position are supplied internally by Mendel for small pedigrees and externally through a coefficient file generated by SimWalk for large pedigrees. For a multivariate quantitative trait, Mendel will stitch together the Kronecker product matrices uniting the various univariate traits for each variance contribution [Lange, 2002].

Something crucial appears to be missing in the example control file. In testing for association, we must include a mean component for each candidate-marker allele. Fortunately, Mendel does this automatically for each marker under study. Given a candidate marker, Mendel assigns a mean parameter to each allele count variable and imposes a constraint forcing these parameters to sum to zero. It then estimates all parameters, both mean and variance parameters, under the null hypothesis of no association and under the alternative hypothesis of association. The null hypothesis is distinguished from the alternative hypothesis by the elimination of the allele-specific regression parameters, and the plausibility of the null hypothesis is assessed by an asymptotic likelihood ratio test. In the current data, the summary output file indicates that allele 1 tends to lower plasma Gc concentrations, allele 2 tends to raise plasma Gc concentrations, and these effects are highly significant ($P$ value $<10^{-5}$). If we use the command `MULTIVARIATE_NORMAL = False` to reanalyze this data with a multivariate $t$ model to control for kurtosis [Lange et al., 1989], then parameter estimates and $P$ values change very little.

SUMMARY FOR MEAN PARAMETERS

| TRAIT | PARAMETER | PREDICTOR | ESTIMATE | STD ERR |
|---|---|---|---|---|
| Gc_conc | 1 | GRAND | 29.5047 | 0.7112 |
| Gc_conc | 2 | FEMALE | 0.4276 | 0.3681 |
| Gc_conc | 3 | MALE | -0.4276 | 0.3681 |
| Gc_conc | 4 | AGE | -0.0295 | 0.0398 |
| Gc_conc | 5 | 1 | 1.3286 | 0.2509 |
| Gc_conc | 6 | 2 | -1.3286 | 0.2509 |

SUMMARY FOR VARIANCE COMPONENTS

| TRAIT | PARAMETER | VARIANCE | ESTIMATE | STD ERR |
|---|---|---|---|---|
| Gc_conc | 7 | ADDITIVE | 10.0907 | 1.8441 |
| Gc_conc | 8 | ENVIRONMENTAL | 2.5043 | 0.6492 |

**Fig. 9. Parameter estimates from the QTL association test by variance components for the Gc locus and the Gc_conc trait, taken from Mendel's standard output file for example data set 20a.**

The fuller output seen in Figure 9, from the bottom of Mendel's standard results file, displays the estimates for the six mean parameters and two variance parameters under the alternative hypothesis. Of course, one could also include dominance and household variance components in the model, but that would be over-fitting for this small data set. We advise against using the QTL variance component in a measured genotype analysis as the two are conflated. Mendel can also report outlier statistics for both people and pedigrees. In the current example, the statistics monitoring outlier pedigrees suggest a departure from normality in the raw data. If we reanalyze the data under the $t$ distribution, then all outliers either disappear or substantially moderate.

One of the more useful features of Mendel is its ability to report deviances pedigree by pedigree. A pedigree's deviance is twice the difference in its loglikelihood (base $e$) under the alternative and null hypotheses. Both loglikelihoods are evaluated at the corresponding maximum likelihood estimates for the entire sample. A pedigree with a large positive deviance favors the alternative hypothesis, while a pedigree with a large negative deviance favors the null hypothesis. In some genetic studies, clinical evidence such as age of onset can be used to rank pedigrees on how likely each is to depart from the null hypotheses of no association. Mendel outputs the ordered-subset deviance statistic of Hauser et al. [2004]. This is helpful in deciding whether a particular prior ranking is justified.

## QTL ASSOCIATION TESTS BY PERMUTATION

A major objection to the measured genotype approach is that it relies on likelihood ratio statistics whose asymptotic distributions are sensitive to small sample sizes and departures from normality. This objection is irrelevant to permutation tests. These tests depend on exchangeable groups (permutation units) of people such as sibships, mating pairs, and random samples from different populations. Under the null hypothesis of no trait-allele association, every permutation of trait values within a unit is equally likely. Because permutation tests are computationally intensive, test statistics should be kept as simple as possible.

A plausible yet simple statistic is the minimum of the sum of squares

$$T(\beta,\mu)=\frac{1}{2}\sum_i\sum_j\left(x_{ij}-\mu_i-\sum_k a_{ijk}\beta_k\right)^2,$$

where $i$ denotes a permutation unit, $j$ a person within permutation unit $i$, $x_{ij}$ his or her trait value, and $a_{ijk}$ his or her imputed number of marker alleles of type $k$. The parameters of the model are the permutation unit effects $\mu_i$ and the allelic effects $\beta_k$. In order for all parameters to be identifiable, we assume $\sum_k \beta_k = 0$. To estimate

the parameters, we minimize $T(\beta, \mu)$ by introducing the Lagrangian function

$$\mathscr{L}(\beta, \mu, \lambda) = T(\beta, \mu) + \lambda \sum_k \beta_k.$$

Setting the partial derivatives of $\mathscr{L}(\beta, \mu, \lambda)$ equal to zero gives the equations

$$0 = -\sum_j \left( x_{ij} - \mu_i - \sum_k a_{ijk}\beta_k \right) \tag{2}$$

$$0 = -\sum_i \sum_j \left( x_{ij} - \mu_i - \sum_k a_{ijk}\beta_k \right) a_{ijl} + \lambda \tag{3}$$

for all permutation units $i$ and alleles $l$, respectively.

The test statistic $T = \min_{\beta,\mu} T(\beta, \mu)$ can be evaluated by applying the usual linear algebra formulas to estimate the vectors $\beta$ and $\mu$. But to do this the tens or hundreds of thousands of times required for permutation testing is cumbersome. If the number of permutation units is large, it involves repeatedly inverting large matrices. We can avoid some of the implied matrix operations by resorting to iteration. Let $n_i$ represent the number of people in unit $i$ and $m$ the index of the current iterate. Rearranging equation (2) gives the update

$$\begin{aligned} \mu_i^{(m+1)} &= \frac{1}{n_i} \sum_j \left( x_{ij} - \sum_k a_{ijk}\beta_k^{(m)} \right) \\ &= \bar{x}_{i\cdot} - \frac{1}{n_i} \sum_j \sum_k a_{ijk}\beta_k^{(m)}. \end{aligned}$$

If we define $y_{ij}^{(m+1)} = x_{ij} - \mu_i^{(m+1)}$, then the system of equations (3) represents linear regression of the $y_{ij}^{(m+1)}$ on the covariates $a_{ijk}$ subject to the constraint $\sum_k \beta_k = 0$. The usual way of viewing such problems is to define the first and second differentials

$$\begin{aligned} \nabla_\beta T(\beta, \mu) &= -\sum_i \sum_j \left( y_{ij}^{(m+1)} - \sum_k a_{ijk}\beta_k \right) a_{ij} \\ \nabla_\beta^2 T(\beta, \mu) &= \sum_i \sum_j a_{ij} a_{ij}^t, \end{aligned}$$

where $a_{ij}$ is a column vector with entries $a_{ijk}$. Least squares estimation amounts to nothing more than one step of Newton's method subject to the constraint. Starting from the point $(\mathbf{0}, \mu^{(m+1)})$, Newton's method is implemented by solving the linear system

$$\begin{pmatrix} \nabla_\beta^2 T & \mathbf{1} \\ \mathbf{1}^t & 0 \end{pmatrix} \begin{pmatrix} \beta \\ \lambda \end{pmatrix} = \begin{pmatrix} -\nabla_\beta T \\ 0 \end{pmatrix}$$

in the form

$$\begin{pmatrix} \beta \\ \lambda \end{pmatrix} = \begin{pmatrix} \nabla_\beta^2 T & \mathbf{1} \\ \mathbf{1}^t & 0 \end{pmatrix}^{-1} \begin{pmatrix} -\nabla_\beta T \\ 0 \end{pmatrix}.$$

Because $\nabla_\beta^2 T(\beta, \mu)$ is constant, the matrix inverse displayed here can fortunately be computed once and stored. At $\beta = \mathbf{0}$, the vector $\nabla_\beta T$ depends on $y_{ij}^{(m+1)}$ and hence on $\mu^{(m+1)}$ and $x_{ij}$ but not on $\beta$.

These alternating updates of $\mu$ and $\beta$ are special cases of the block relaxation method summarized by de Leeuw [1994]. The local convergence theory discussed there suggests that our algorithm should converge quickly. In our practical experience, convergence occurs in 10 to 20 iterations. The advantage of block relaxation in the present context is that it avoids large matrix inversions and replaces them with faster matrix times vector multiplications.

There are some identifiability issues in estimating the parameters. For example, if a different homozygous genotype is assigned to each separate permutation unit, then the $\beta_k$ and $\mu_i$ will be confounded. Also, by design any permutation unit consisting of a single person contributes nothing to the test statistic. It is also possible to formulate a test statistic in terms of genotypic rather than allelic effects. We have not done so because it can sharply increase the number of parameters; of course, biallelic markers are a useful exception. It is noteworthy that the genotype version of block relaxation avoids matrix inversion altogether when genotypes are fully known.

## ACE EXAMPLE

For a simple comparison of the TDT, gamete competition, and two measured-genotype association methods, we now turn to SNP data from the ACE gene. Because the associations between ACE levels and the SNP alleles are so strong, we arbitrarily restricted analysis to 14 extended families from the larger collection of white British families mentioned earlier [Keavney et al., 1998]. We tested two SNPs, T-3892C and T-93C, located upstream of the coding region and in linkage disequilibrium with the insertion-deletion polymorphism. The naming convention for each SNP incorporates its two alleles as the first and last

**TABLE II. Comparison of ACE association *P* values (±2 standard deviations)**

| Association test | T-3892C | T-93C | T-3892C + T-39C |
|---|---|---|---|
| Variance component QTL | 0.00001 | 0.00088 | 0.00003 |
| Gamete competition | 0.00042 | 0.00067 | 0.00134 |
| Permutation QTL | 0.01398 ± 0.00074 | 0.01329 ± 0.00072 | 0.00967 ± 0.00062 |
| TDT | 0.34350 ± 0.00300 | 0.07080 ± 0.00162 | |

letter of its name. In Mendel, the choice of permutation groups for the QTL permutation method is left to the user. For each of the 14 pedigrees, we defined a permutation group for each sibling set and an additional permutation group for the collected founders. In the TDT analysis, we dichotomized the ACE levels so that anyone with an ACE level greater than one standard deviation above the mean was considered affected. In the control file we set `SAMPLES = 100000` for the permutation tests (QTL and TDT). Table II summarizes our statistical findings and ranks the methods in descending order of their apparent power.

The C allele in both SNPs is associated with high ACE levels. Reflecting the power loss from dichotomizing the ACE levels, the TDT shows the least significant *P* values in Table II. The gamete competition model and the variance component QTL test use the full pedigree data and exhibit better power than the permutation QTL test. Because it ignores transmission from homozygous parents, the gamete competition model suffers by comparison to the variance component QTL test. In an effort to increase our statistical power, we applied a utility option of Mendel to form a super-locus from these two SNPs. The alleles of the super-locus correspond to haplotypes. Combining SNPs creates phase ambiguities, so our final comparison ignores the TDT. All three remaining tests are still significant, but only the permutation QTL appears to gain power. Because the estimated TC haplotype frequency was near zero, we used yet another option of Mendel to consolidate this haplotype with the TT haplotype in the gamete competition and variance component QTL tests. The CC haplotype was associated with the highest ACE levels, and the TT haplotype was associated with the lowest ACE levels in all three analyses. These results tend to confirm our preconceptions about the power of the various tests. Simulation studies are probably necessary to declare a clear winner between the gamete competition test and the permutation QTL test.

# DISCUSSION

Despite the vast literature on association tests and the increases in size, scope, and number of genetic studies, there is little agreement on the best strategy for conducting association tests in humans. Commercially available statistical packages such as S-plus and SPSS are ill-suited to handle the complex dependencies presented by pedigrees and linked genetic markers. It is unlikely that the large corporate software houses will make much headway in such a murky market, and users will probably continue to rely on compact, quickly evolving, freeware. The difficulty for users is not so much the lack of software, but the poor quality of the documentation of most available programs. These quality problems reflect the fact that statistical methods are being proposed so rapidly that most programs become obsolete before they are adequately tested.

Most users rely on word of mouth, short courses, and luck to find appropriate software. Once a program gains currency, it is treated as an oracle by naive users. With little real understanding of computational statistics, such users are inclined to compare software by raw speed, ease of use, and graphical output. If two programs purport to compute similar statistics but give different answers, then the obvious temptation is to cite the one with the more impressive *P* values. Needless to say, this state of affairs is far from ideal. It gives the impression that human geneticists are sloppy and unreliable. Although science is saved from the worst excesses of individual scientists by its tradition of replication, replication can be expensive.

It is unclear how to improve the software confusion without crushing creativity. In our view, it would be a mistake to create a governing body for software development. A better remedy is for journals such as *Genetic Epidemiology* to publish more software reviews. *The American Statistician* has done this successfully for a number of years. Giving creators of software the right of rebuttal might assuage some of their worst fears.

Another possibility is to encourage the kind of open source community that powers Linux. For the open source paradigm to be successful, mechanisms must be put in place to give contributors appropriate publication credit. Without such credit, most academic scientists cannot be promoted. Until the ground for open source development is carefully cultivated, we can expect to see more and more geneticists distribute executable rather than source code.

The current article explores the third possibility of using an article format for program exploration and documentation. We have tried to mitigate the pedestrian nature of this task by dwelling on some of the innovations woven into Mendel. These include association testing by penetrance estimation, expansion of matched-pair designs to permutation unit designs, and regression on expected allele counts for quantitative traits. Mendel is also notable for the caliber and variety of the tests implemented, its common frameworks for inputting data and implementing models, and its user friendly shell Gregor. We are disinclined to try the reader's patience further by describing Gregor.

Mendel is a work in progress. Our future agenda includes (1) development of graphical plots for location scores and other output, (2) greater modeling flexibility through user access to basic subroutines, (3) better coordination with SimWalk and database software, and (4) additional options for inbred mice, longitudinal quantitative traits, and gene-by-gene interaction. We welcome users' corrections and suggestions for improvement.

## ACKNOWLEDGMENTS

The authors gratefully acknowledge the following USPHS grants under which this work was accomplished: GM053275 (K.L., J.S., and E.S.), MH059490 (K.L. and J.S.), and MH064205 (E.S.).

## ELECTRONIC DATABASE INFORMATION

The current versions of Mendel and SimWalk are available for several computing platforms from the web site (*http://www.genetics.ucla.edu/software*).

## REFERENCES

Boehnke M. 1991. Allele frequency estimation from data on relatives. Amer J Hum Genet 48:22–25.

Boerwinkle E, Chakraborty R, Sing CF. 1986. The use of measured genotype information in the analysis of quantitative phenotypes in man. I. Models and analytical methods. Ann Hum Genet 50:181–194.

Bradley RA, Terry ME. 1952. Rank analysis of incomplete block designs. Biometrika 39:324–345.

Cantor RM, Chen GK, Pajukanta P, Lange K. 2005. Association testing in a linked region using large pedigrees. Am J Hum Genet 76:538–542.

Clarke CA, Price-Evans DA, McConnell RB, Sheppard PM. 1959. Secretion of blood group antigens and peptic ulcers. Br Med J 1:603–607.

Curtis D, Sham PC. 1995. A note on the application of the transmission disequilibrium test when a parent is missing. Am J Hum Genet 56:811–812.

Daiger SP, Miller M, Chakraborty R. 1984. Heritability of quantitative variation at the group-specific component (Gc) locus. Am J Hum Genet 36:663–676.

de Leeuw J. 1994. Block relaxation algorithms in statistics. In: Bock HH, Lenski W, Richter MM, editors. Information systems and data analysis. Springer-Verlag, Berlin: p 308–325.

Elston RC, Stewart J. 1971. A general model for the genetic analysis of pedigree data. Hum Hered 21:523–542.

Ewens WJ, Spielman RS. 1995. The transmission/disequilibrium test: History, subdivision, and admixture. Am J Hum Genet 57:455–464.

Fan R-Z, Spinka C, Jin L, Jung J-S. 2005. Pedigree linkage disequilibrium mapping of quantitative trait loci. Eur J Human Genet 13:216–231.

George VT, Elston RC. 1987. Testing the association between polymorphic markers and quantitative traits in pedigrees. Genet Epidemiol 4:193–201.

Hasstedt SJ. 1982. A mixed model likelihood approximation for large pedigrees. Comput Biomed Res 15:295–307.

Hauser ER, Watanabe RM, Duren WL, Bass MP, Langefeld CD, Boehnke M. 2004. Ordered subset analysis in genetic linkage mapping of complex traits. Genet Epidemiol 27:53–63.

Hopper JL, Mathews JD. 1982. Extensions to multivariate normal models for pedigree analysis. Ann Hum Genet 46:373–383.

Horvath S, Laird NM. 2002. Family-based association for quantitative traits. In: Elston RC, Olson J, Palmer L, editors. Biostatistical Genetics and genetic epidemiology, New York, Wiley, p 264–267.

Jacquier M, Arango D, Villareal E, Torres O, Serrano ML, Cruts M, Montañes P, Cano C, Rodriguez MN, Serneels S, Van Broeckhoven C. 2001. APOE $\epsilon 4$ and Alzheimer's disease: Positive association in a Colombian clinical series and review of the Latin-American studies. Arq Neuropsiquiatr 59:11–17.

Keavney B, McKenzie CA, Connell JMC, Julier C, Ratcliffe PJ, Sobel E, Lathrop M, Farrall M. 1998. Measured haplotype analysis of the angiotensin-1 converting enzyme gene. Hum Mol Gen 11:1745–1751.

Keener JP. 1993. The Perron-Frobenius theorem and the ranking of football teams. SIAM Review 35:80–93.

Kruglyak L, Lander ES. 1998. Faster multipoint linkage analysis using Fourier transforms. J Computat Biol 5:1–7.

Kruglyak L, Daly MJ, Reeve-Daly MP, Lander ES. 1996. Parametric and nonparametric linkage analysis: a unified multipoint approach. Am J Hum Genet 58:1347–1363.

Lander ES, Green P. 1987. Construction of multilocus genetic linkage maps in humans. Proc Natl Acad Sci USA 84:2363–2367.

Lange K, Cantor R, Horvath S, Perola M, Sabatti C, Sinsheimer J, Sobel E. 2001. Mendel version 4.0: A complete package for the

exact genetic analysis of discrete traits in pedigree and population data sets. Am J Hum Genet 69(Supplement):A1886.

Lange K. 2002. Mathematical and statistical methods for enetic analysis, 2nd ed. New York: Springer-Verlag.

Lange K, Elston RC. 1975. Extensions to pedigree analysis. I. Likelihood calculations for simple and complex pedigrees. Hum Hered 25:95–105.

Lange K, Boehnke M, Weeks DE. 1988. Programs for pedigree analysis: MENDEL, FISHER, and dGENE. Genet Epidemiol 5:471–472.

Lange K, Little RJA, Taylor JMG. 1989. Robust statistical modeling using the t distribution. JASA 84:881–896.

Lewis M, Kaita H, Philipps S, Giblet ER, Anderson JE, McAlpine PJ, Nickel B. 1980. The position of the Radin blood group in relation to other chromosome 1 loci. Ann Hum Genet 44: 179–184.

Ott J. 1974. Estimation of the recombination fraction in human pedigrees: efficient computation of the likelihood for human linkage studies. Am J Hum Genet 26:588–597.

Risch N, Merikangas K. 1996. The future of genetic studies of complex human diseases. Science 273:1516–1517.

Rubenstein P, Walker M, Carpenter C, Carrier C, Krassner J, Falk CT, Ginsburg F. 1981. Genetics of HLA disease associations. The use of the haplotype relative risk (HRR) and the 'haplo-delta' (Dh) estimates in juvenile diabetes from three racial groups. Hum Immunol 3:384.

S.A.G.E. 2001. Statistical Analysis for Genetic Epidemiology, S.A.G.E. 4.1. Computer program package available from the Department of Epidemiology and Biostatistics, Reammelkamp Center for Education and Research, MetroHealth Campus, Case Western Reserve University, Cleveland, OH.

Sham P, Bader JS, Craig I, O'Donovan M, Owen M. 2002. DNA pooling: a tool for large-scale association studies. Nat Rev Genet 3:862–871.

Sinsheimer JS, Blangero J, Lange K. 2000. Gamete competition models. Am J Hum Genet 66:1168–1172.

Sinsheimer JS, McKenzie CA, Keavney B, Lange K. 2001. SNPs and snails and puppy dogs' tails: analysis of SNP data using the gamete competition model. Ann Hum Genet 65:483–490.

Sobel E, Lange K. 1996. Descent graphs in pedigree analysis: applications to haplotyping, location scores, and marker sharing statistics. Am J Hum Genet 58:1323–1337.

Sobel E, Papp JC, Lange K. 2002. Detection and integration of genotyping errors in statistical genetics. Am J Hum Genet 70:496–508.

Spielman RS, McGinnis RE, Ewens WJ. 1993. Transmission test for linkage disequilibrium: The insulin gene region and insulin-dependent diabetes mellitus (IDDM). Am J Hum Genet 52: 506–516.

Terwilliger JD, Ott J. 1992. A haplotype-based ''haplotype relative risk'' approach to detecting allelic associations. Hum Hered 42:337–346.

Yamagata Z, Asada T, Kinoshita A, Zhang Y, Asaka A. 1997. Distribution of apolipoprotein E gene polymorphisms in Japanese patients with Alzheimer's disease and in Japanese centenarians. Hum Hered 47:22–26.