

Published in final edited form as:

Hum Hered. 2007 January 01; 63(3-4): 229–238. doi:10.1159/000100481.

Dealing with missing data in family-based association studies: a multiple imputation approach

Pascal Croiseau^{1,2}, Emmanuelle Génin^{2,1}, Heather J. Cordell³

¹Univ Paris-Sud, UMR-S535, Villejuif, F-94817

²INSERM, Villejuif, F-94817

³Institute of Human Genetics, Newcastle University, UNITED KINGDOM

Abstract

To test for association between a disease and a set of linked markers, or to estimate relative risks of disease, several different methods have been developed. Many methods for family data require that individuals be genotyped at the full set of markers and that phase can be reconstructed. Individuals with missing data are excluded from the analysis. This can result in an important decrease in sample size and a loss of information. A possible solution to this problem is to use missing-data likelihood methods. We propose an alternative approach, namely the use of multiple imputation. Briefly, this method consists in estimating from the available data all possible phased genotypes and their respective posterior probabilities. These posterior probabilities are then used to generate replicate imputed data sets via a data augmentation algorithm. We performed simulations to test the efficiency of this approach for case/parent trio data and we found that the multiple imputation procedure generally gave unbiased parameter estimates with correct type 1 error and confidence interval coverage. Multiple imputation had some advantages over missing data likelihood methods with regards to ease of use and model flexibility. Multiple imputation methods represent promising tools in the search for disease susceptibility variants.

Keywords

case-parent trio; conditional logistic regression; haplotype

Introduction

To identify variants involved in disease susceptibility, a traditional approach consists in testing for association between a disease and a set of markers. This is usually performed by comparing allele or genotype frequencies at the markers in samples of cases and controls. It is also possible to use case-parent trios and compare the alleles or genotypes transmitted to the affected child to the corresponding non-transmitted alleles or genotypes [1,2]. A major advantage of these family-based tests is their robustness to population stratification. Moreover, the familial structure allows the testing of both linkage and association. A

disadvantage is the difficulty to recruit large samples of case-parent trios and consequently, sample sizes are generally smaller than that achievable with the case-control approach, leading to a possible lack of power as compared to the case-control approach.

Often association studies are faced with a problem of missing data, either in the form of a missing genotype or in the form of unknown phase. Current genotyping technologies do not provide phase information and so we need to reconstruct it from the observed genotype information, which is not always possible. For case-parent trio data, the presence of an affected offspring does not ensure the data availability on his parents. Refusal to participate, death, false paternity or genotyping failure are different factors which can generate a missing genotype. With the availability of high-density SNP maps, the number of genotyping failures is expected to increase (it is obvious that the higher the number of polymorphisms genotyped, the less the number of complete families likely to be available) and there is more phase ambiguity. There is a temptation to simply ignore the missing data and only use the complete and phase-known observations, but it has been shown that this can induce bias [3] and/or loss of efficiency [2]. It might also result in a significant reduction in sample size and consequently a loss in power. When the level of missing data differs from one marker to another, focusing only on the complete data in the analysis will make it very difficult to compare the different markers. It may lead to false conclusions regarding which marker(s) are most likely to explain the detected association and thus on the location of sites involved in disease susceptibility. Indeed, if the disease susceptibility site is among the studied sites but is poorly genotyped, it is possible that one marker in linkage disequilibrium with this site obtains a better association score than the disease susceptibility site itself.

Several methods have been developed to infer the missing data from the rest of the data. In the context of family-based association studies, specific methods have been developed mostly based on likelihood approaches (see for example TRANSMIT [4], TDTPHASE [5] or Haplotype FBAT [6]). One problem with these methods and their corresponding software is their lack of flexibility. Different realisations of these methods are required if, for example, one wants to account for environmental risk factors and potential gene-environment interactions in the analysis, or to account for different sets of SNP loci within a small genetic region. By performing missing data/haplotype inference and estimation/testing of effects together, in a single stage procedure, these approaches usually limit one to using the same set of genetic loci for both purposes. In addition, the approaches in FBAT and TRANSMIT are focused on testing the null hypothesis (of no genetic effects) rather than on estimation of effects.

In this context, it is of interest to develop methods to test for association with genetic risk factors in the framework of traditional statistical packages such as Stata, S-Plus/R or SAS, which allow the inclusion of arbitrary genetic and/or environmental predictor variables in a model and estimation of their effects. Indeed, for family-based data as well as for cases and controls, such methods have previously been proposed [2,7]. The main drawback of using a standard statistical software package is the difficulty to deal with missing data, and families with missing data are usually discarded from the analysis. Multiple imputation (MI) [8] provides a convenient solution to the problem. The idea of the method is to fill in missing data by values that are predicted by the observed data. MI is a Monte Carlo technique in

which the observed data set containing missing values is replaced by m simulated versions, where m is typically small (e.g. 3-10). Each of the simulated complete datasets is analyzed by standard methods, and the results are combined to produce estimates and confidence intervals that incorporate the missing-data uncertainty [8–10]. MI has already been used for quantitative traits [11] and matched case-control studies [12], a MI method is implemented in SNPHAP program (www-gene.cimr.cam.ac.uk/clayton/software/) and the performance of the method to estimate genotype relative risks in case control studies has recently been evaluated and compared to likelihood-based methods [13]. In this paper, we adapt the MI method to the case-parent trio design. In this design, the familial structure of the data imposes constraints on the possible phased genotypes compatible with the observed genotype data, allowing a better reconstruction of the missing values. The performance of our method is evaluated by simulations for different levels of missing data and different disease susceptibility models. We investigate the power to detect an association and the efficiency of the method to identify the disease susceptibility sites among several markers. We also investigate the performance of the method with regard to estimation of genotype relative risk parameters, and (for some genetic models) find improved performance from use of MI with case-parent trios compared to its performance with unrelated cases and controls.

Methods

Multiple Imputation Approach

Given observed unphased genotype data (with possibly missing values) for a set of case-parent trios, we generate complete phase-known data sets using a MI approach via a data augmentation (IP) algorithm [14]. For case-parent trio data, this algorithm proceeds as follow:

1. (I-step) Given current parameter values for population haplotype frequencies and phase-known genotype frequencies in affected offspring, sample from the posterior probability of phase-known haplotype configurations for each family (given the observed genotype and phenotype data) to obtain a complete data set. For each family containing a missing value, a haplotype assignment is picked from all the possible assignments with a probability given by the current posterior distribution. At the starting point, the posterior distribution is unknown but a starting value can be derived from the observed data by using an EM algorithm. Here, we use the ZAPLO/PROFILER software (www.molecular-haplotype.org/zaplo/zaplo_index.html, www.molecular-haplotype.org/profiler/) [15] for this purpose and obtain for the different families in the sample a listing of the phased genotype configurations that are compatible with the observations, and their initial posterior probabilities.
2. (P-step) The population haplotype frequency and phase-known affected-offspring genotype frequency parameters are then updated by sampling them from their posterior distribution given the current complete data file. Under the assumption that the prior distribution of haplotype (genotype) frequencies is a Dirichlet distribution with constant degrees of freedom (df) on all possible haplotypes (genotypes), the full conditional posterior distribution is also Dirichlet with

degrees of freedom equal to the observed number of haplotypes (genotypes) in the complete data file + the prior df.

We cycle round the I and P steps a large number of times, to reach a stationary distribution (in this study we used a burn-in period of $n=1000$ iterations). At intervals (e.g. every 1000 iterations) we output the current complete data file. The IP algorithm is repeated until we have output m replicates of complete data sets (imputations) that will then be analysed. The number of iterations between two imputations must be large enough to ensure a statistical independence between imputed data files. In this analysis, 1000 iterations are run between each output imputation.

At the I-step, the new familial configuration posterior probabilities are calculated by computing their relative likelihood which corresponds to the likelihood of the familial configuration. Denote the frequency of the affected child phased genotype ij as g_{ij} and the frequencies of the untransmitted haplotypes k and l as h_k, h_l then:

$$relative\ likelihood = \frac{g_{ij} \times h_k h_l}{\sum_{(m,n,o,p \in H)} g_{mn} \times h_o h_p} \quad (1)$$

where the sum in the denominator is over all possible familial configurations defined by the affected child phased genotype m/n and untransmitted haplotypes o and p . Note that by considering affected child phase-known genotype frequencies rather than transmitted haplotype frequencies, we avoid making an assumption of Hardy Weinberg equilibrium in the affected sample, which is a necessary assumption when using MI with case-control data [10].

At the P-step, given a complete data realisation, to update the genotype frequencies and the untransmitted haplotype frequencies we simply count genotypes in affected children and count untransmitted haplotypes. Then, we add a data augmentation parameter that corresponds to the prior Dirichlet df. In practice, we set this to be high at the first iteration (typically two times the size of the population) and then decrease it linearly (data augmentation parameter decreases by two times the size of the population divided by the number of iterations) at each iteration, resetting it back to the high value after outputting each of the m replicate imputed data sets.

The optimal number of imputations m depends on the rate of missing data. Indeed, as shown by Rubin [16], the efficiency of an estimate based on m imputations is:

$$efficiency = \left(1 + \frac{\gamma}{m}\right)^{-1} \quad (2)$$

where γ is the missing information rate which depends of the variability between the m complete data sets induced by missing data. For a small number of replicates m (typically in the range of 3-10), we can obtain a very good efficiency [17].

Each of the m complete data sets are analysed by a statistical method. In this paper, the method used for the data analysis is conditional logistic regression [2,18,19] which

compares the genotype of each affected child (the case, denoted as person 1) to the three possible genotypes (the pseudocontrols, denoted as persons 2-4) that can be formed by the untransmitted parental alleles (or haplotypes when several loci are considered). Given a reference genotype with baseline risk termed β_0 , (which in fact will cancel out of the likelihood), each genotype relative risk β_i ($i=1, \dots, n$) is estimated by maximization of the likelihood :

$$L_1 = \prod_k \frac{\sum_{j=1}^4 \exp(\beta_0 + \beta_1 x_{1j_k} + \dots + \beta_n x_{nj_k})}{\sum_{j=1}^4 \exp(\beta_0 + \beta_1 x_{1j_k} + \dots + \beta_n x_{nj_k})} \quad (3)$$

where x_{ijk} is an indicator taking value 1 if person j in family k has genotype i , and 0 otherwise. Under the null hypothesis of no association, the likelihood is simply L_0 corresponding to $\beta_i = 0$ ($i=1, \dots, n$).

To compare with the results obtained with TRANSMIT (REF à ajouter), a one-df allele test will also be performed under the same logistic regression framework but considering allele relative risks rather than genotype relative risks.

For each of the m complete data files i , we calculate the likelihood ratio test d_i

$$d_i = 2[\ln(L_1) - \ln(L_0)] \quad (4)$$

The m results are then combined using the method described in [9,10]. Briefly, for each set of imputed data set, we calculate the average of the d_i statistics.

$$\bar{d} = \sum \frac{d_i}{m} \quad (5)$$

and derive

$$D = \frac{\left(\frac{\bar{d}}{k} - \frac{m+1}{m-1}r\right)}{(1+r)} \quad (6)$$

where k is the the number of degrees of freedom of the likelihood ratio test and r is the variance between the m imputed datasets which can be calculated by the following expression:

$$r = \left(1 + \frac{1}{m}\right) \left(\frac{1}{m-1} \sum_{i=1}^m \left(\sqrt{d_i} - \frac{\sum_{i=1}^m \sqrt{d_i}}{m}\right)^2\right) \quad (7)$$

D follows a F distribution with k and ν degrees of freedom and

$$v = k \frac{-3}{m} (m-1) \left(1 + \frac{1}{r}\right)^2 \quad (8)$$

The combined p-value over the m imputations is then:

$$\text{p-value} = F_{(k, v)} \quad (9)$$

We may also perform parameter estimation and confidence interval (CI) construction for the genotype or haplotype relative risks,

$$\bar{\beta} = \frac{1}{m} \sum_{i=1}^m \hat{\beta}_i \quad \text{and} \quad CI = \bar{\beta} \pm 1.96\sqrt{V} \quad (10)$$

where the variance V is the sum of the variance within imputations

$$V_{within} = \frac{1}{m} \sum_{i=1}^m \hat{\sigma}_{\hat{\beta}_i}^2 \quad (11)$$

and the variance between imputations

$$V_{between} = \frac{1}{m-1} \sum_{i=1}^m (\hat{\beta}_i - \bar{\beta})^2 \quad (12)$$

weighted by a term that depends on the number m of imputations:

$$V = V_{within} + \left(1 + \frac{1}{m}\right) V_{between} \quad (13)$$

Simulation Study

The performance of the MI algorithm was tested by simulations. Genotypes of case-parent trios were simulated at five completely linked loci under different genetic models with one or two disease susceptibility loci. The five loci were in linkage disequilibrium (LD) and the haplotype frequencies and resulting pattern of LD are shown in Table 1 and in supplementary information, Figure 1. Briefly the simulation process is the following: for each parent, two haplotypes are randomly picked from the population of 17 possible 5-locus haplotypes where each haplotype has a frequency as reported in Table 1. Then, one haplotype is randomly drawn from each parent to generate the child's genotype and, based on the penetrance associated with this genotype, an affection status is generated for the child. If the affection status is "unaffected", the trio is discarded and the process is repeated until we obtain sufficient trios with an "affected" child. Finally, missing data are generated completely at random, that is to say with the same percentage of missing data on each of the five SNPs and on each member of the family (the patient as well as his parents). This results

in a data set in which both phase information and genotype data may be missing for any or all of the individuals in the trio, as is the case with real genetic studies.

Concerning the genetic models, for one-locus models, the disease susceptibility (DS) locus was assumed to be the second marker (SNP2) and dominant or recessive models with genotype relative risks (GRR) of 1.5 or 3 were considered (see Table 2a). An additional model with no effect (GRR=1 for all genotypes) was also considered to evaluate the type I errors. For two-locus models, SNP2 and SNP3 were assumed to be the DS loci and both a multiplicative and non multiplicative models were considered (see Table 2). For each set of simulation, 500 simulated data sets were generated and run through the MI algorithm (with the following parameters: burn-in period = 1000 iterations, interval between two imputation $n=1000$ iterations and number of imputations $m=9$). For the different models, the power/type I error to detect the association was evaluated by determining, the proportion of replicates among the 500 replicates where the combined p-value (as defined in equation 9) at SNP2 (the disease susceptibility site) was smaller than or equal to 5%. The performance of the algorithm for parameter estimation was also studied by reporting the bias and 95% CI coverage for the different GRR estimations. We also (where possible) compared the MI approach with the approach implemented in the program TRANSMIT [4], both with, and without, the presence of population stratification.

Results

Type I error rates for a nominal value of 5% are presented in Figure 1 for different levels of missing data. As expected, when families with missing data are ignored, we find that type I errors are not inflated by the presence of missing data, since these missing data are at random with respect to genotypes (see curve without MI). The use of the MI method does not increase the type I errors for up to 30% of missing data. Above 30%, we observe a significant increase in type I error rates. For 50% of missing data, there is a three fold increase in the type I error (0.16 instead of 0.05). In contrast, we find that the TRANSMIT program maintains correct type 1 error for all except the highest level of missing data.

Figure 2 shows the power to detect association at nominal significance level 5% under the different one locus models. Without the MI algorithm, the power of detection of the DS site is sensitive to missing data due to the decrease of the sample size. The more the percentage of missing data increases, the more the number of informative families decreases and consequently, the more the power to detect the DS site decreases. A similar decrease in power is seen with the TRANSMIT program, which, unlike MI, is unable to use information from linkage disequilibrium patterns with flanking markers when testing at a single locus. In this context, the MI algorithm allows good power for detection of the DS site. For the recessive model (Figure 2a), we see a loss in power of 14.2% with 30% of missing data (without the MI algorithm, we observed a loss in power of 58% in the same configuration). For the dominant model (Figure 2b), the power remains at approximately 80% with 30% of missing data (so basically no loss in power as compared to a 59.1% power loss when MI algorithm is not used). Power results for MI with >30% missing data should not be considered since we know from Figure 1 that the nominal type 1 error at this level of missing data is not maintained.

To check the utility of the MI method, the correlation between p-values obtained on the true complete data sets (available to us since this is simulated data) and p-values obtained by using the MI algorithm on the same data, after adding a percentage of missing data, were investigated (supplementary information, Figure 2). With 5% missing data, we note a very good correlation rate in both models. With 50% of missing data, correlations are obviously lower: on the recessive model (model 1) p-values are often decreased and consequently a lot of non significant tests became significant. In contrary, p-values are slightly increased for the dominant model (model 2). Consequently, the utility of the MI is limited when the percentage of missing information is high; nevertheless, for up to 30% missing data this method gives acceptable performance.

We performed analysis to examine how often the DS site gives the highest score (i.e. the highest test statistic conditional on the fact the highest score is significant) as a function of the percentage of missing data (Figure 3). For the recessive model, we see a loss of detection of the true DS site of 22% with 30% missing data (without the MI algorithm, we observed a loss of detection of 62%). On the dominant model, we see an over detection of 7% with 30% missing data against a loss of detection of 42% without using the MI algorithm.

To investigate parameter estimation under the MI method, results with and without using the algorithm were compared. With or without the MI algorithm, when the same percentage of missing data is used for all loci, no bias is expected. Without MI, when conditional logistic regression is used, only families genotyped for all markers are taken into account. So, missing data will involve a decrease in the sample size (e.g. only 15% of informative families are available for analysis with 50% missing data) but not a change in the properties of the model. Table 1 of the supplementary information confirms this point: for each of the models considered, the bias between the observed and the expected values of the genotype relative risks is near 0, and the 95% confidence intervals for the genotype relative risk parameters correctly cover the true values approximately 95% of the time (for a missing data proportion of up to about 30%). We also investigated the correlation between the parameter estimates obtained from the true complete data set and those obtained using the MI approach (supplementary information, Figure 3) and found that as expected the correlation decreases as the percentage of missing data increases but even with 50% of missing data, there is still a strong correlation.

Two 2-locus models were also investigated. Figure 4 shows the bias between the expected and the observed haplotypic risks for different rates of missing data under the multiplicative model only, since risk estimation on haplotypes only have a sense if haplotypes have a multiplicative effect. For up to 10% of missing data, bias is weak with and without MI. For a larger percentage of missing data and when missing data are ignored, we observe a strong increase in the bias for all the haplotype relative risks. In contrast, when we use the MI algorithm, the bias stays reasonably low. For example, with 50% of missing data the bias is in the range [26.5, 29.1] without MI compared to [0.15, 1.17] with MI.

Figure 5 presents the bias obtained under the two models on phased genotypic risks. The median bias over the 10 genotype relative risks is reported. As for haplotypic risks, we note a high increase in the bias when the percentage of missing data increases. Without MI the

median bias reaches 32.89 for the multiplicative model and 30.47 for the non multiplicative model opposed to a weak increase with MI (1.16 for the multiplicative model and 4.47 with the non multiplicative model). However, observed bias could be very different from one genotype to another and as expected, strong bias may be observed especially for rare genotypes.

These results confirm the ones obtained by Cordell [13] when using MI on case-control data. An important difference however between case-control and trio data is the fact that under a non-multiplicative model, the case-control data do not allow the distinction between the two following genotypes: the one composed by haplotypes aA and bB and the one composed by haplotypes aB and bA. Use of an MI algorithm with case/control data [10] distinguishes between these configurations by borrowing information (when it is possible) from resolvable genotypes under a Hardy Weinberg equilibrium assumption which corresponds to assuming multiplicative haplotype effects. Consequently, this method fails for case/control data when the underlying haplotypic effects are not multiplicative [10]. However, with family data, information from the parents can allow knowledge of this phased genotype for the affected child. We performed simulations under different multiplicative and non-multiplicative models to check this property of the case/parent trio data. We choose to present in this paper results obtained with one of the non-multiplicative models used in [13] with 10 percent of missing data. Table 3 shows the bias and coverage when using the MI approach in a case/control and in a case/parent trio data set. As expected, we note the presence of a bias higher than 0.3 with the case/control data with poor confidence interval coverage (near 0.8) for both genotypes 1-1/2-2 and 1-2/2-1. The use of case/parent trio families generally gives less bias, particularly for the two phased genotypes (bias of 0.012 and 0.019 respectively). Consequently, we confirm the capacity of the method to correctly estimate these two phased genotype configurations with family data, which is not the case for case/control data.

To illustrate the similarity between the MI approach and other methods, we also compared MI to TRANSMIT using the non-multiplicative model, and we noted no difference between the P values obtained by MI and TRANSMIT (supplementary information, figure 4). However, the main motivation for the development of TRANSMIT (and indeed the original motivation to recruit and use case-parent trio datasets) is to achieve robustness to population stratification. To check this point, from a dataset of 500 families composed by two populations described in supplementary table 2, the type 1 error rates have been calculated for different levels of missing data using both the MI approach and the TRANSMIT program. Results are presented in supplementary figure 5. We note that for up to 30% missing data, both methods give similar values and correct type 1 error rates. For over 30% missing data, we note an increase of the type 1 error rate with MI while TRANSMIT maintains a type 1 error rate around 5%. These results show that the MI approach provides some protection against population stratification even if not complete protection. For comparison, this level of population stratification leads to a type 1 error rate of approximately 71% (simulation results, data not shown) in case/control analysis with no missing data, indicating that the population stratification problem is much more severe for case/control data than it is for case/parent trios, even in the presence of missing data.

Discussion

We have proposed an MI algorithm for testing and estimation of genotype and haplotype effects using case/parent trio data. Through simulations, we have examined the utility of the MI algorithm on several one and two locus models. Results show the usefulness of the MI approach with conditional logistic regression, this approach allows us to work with three pseudocontrols for each family whereas in the original case/pseudocontrol approach, only one pseudocontrol can be generated for families with unknown phase. When the percentage of missing data is small, parameter estimation can be correct without the use of MI. Nevertheless, MI increases power and allows better comparison of results at different loci and detection of the true disease susceptibility site. Obviously, the MI algorithm is only one of the possible methods to take into account uncertainty due to missing data. The MI approach proposed here shares some similarities with Gibbs sampling and also with the stochastic EM algorithm [20–22]. Indeed, the sampling mechanism used in the IP algorithm that we describe is virtually identical to that used in Gibbs sampling. The main difference between MI and Gibbs sampling or an EM algorithm (either in its original or stochastic version) is that in a Gibbs sampling or EM algorithm framework, one runs the algorithm until convergence and uses the final parameter estimates obtained to make inference. In the MI approach, one instead writes out imputed data sets at intervals (e.g. every 1000 iterations), analyses these imputed data sets using standard statistical methods (e.g. regression) and then uses methods described in the MI literature [9,16,17] to produce a combined parameter estimate or make combined inference from all of the imputed data sets.

Compared to Gibbs sampling or an EM algorithm, MI is thus a two stage procedure. In the first stage one performs imputation to generate 3-10 imputed data sets; in the second stage one performs analysis on these imputed data sets and produces a combined result. The main advantage of this from an operational point of view is that one need not actually fit the full model at the second stage. For instance, one could do the imputation assuming that 3 loci in a region influence disease, but then at the analysis stage one could fit the full model where all 3 loci influence disease, or one could fit a restricted model where just 2 of the loci influence disease, or indeed a further restricted model where just one of the loci influences disease. To compare all of these nested models in an EM or Gibbs sampling procedure, one would have to run the EM/Gibbs sampling algorithm 3 separate times, whereas in the MI approach one runs the IP algorithm only once, then fits the different nested models at the analysis stage.

In this paper, we focus on a pattern of missing data where the percentage of missing data is the same for each marker and for a given marker does not depend on the genotype. In presence of different percentage of missing data at the different markers, we have shown that the method works very well. In the case where missing data concentrate exclusively on the disease susceptibility locus, it will allow to regain the association signal at this locus and as compared to the naïve approach that consists in discarding families with missing data, will better allow to identify the true DS site (GAW15 unpublished results). If the percentage of missing data also depends on the genotype, the method will probably have more problem inferring missing data but depending on the pattern of LD with adjacent markers, it should still be rather successful. Some further simulations will be required to investigate this

particular point. Note that our method will also work when there is no missing genotype data (i.e. all markers are genotyped perfectly) but when phase information (i.e. the two haplotypes present in an individual) is missing or uninferable.

Another particularity of the pattern of missing data investigated here is the fact that within a family, we assume that there is the same chance that any of the member (father, mother or affected sib) is missing. Some other patterns have been tested to check the influence of this pattern of missing data on the efficiency of the method. One in particular (data not shown) presents missing data exclusively on parents where no loss in power is observed even with a high level of missing data. Information brought by the affected child allows a good imputation of the missing values. Thus, MI approach can be a good solution for the data analysis of late-onset diseases.

Conceptually and in practice, therefore, it appears that MI is a promising approach for use in the search for disease susceptibility genes. Future work will involve extending this approach for association analysis using larger family structures (e.g. extended pedigrees) and with quantitative as opposed to merely dichotomous (disease) traits.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We thank Professor David Clayton for his help with multiple imputation algorithm. We also thank the REFGENSEP for sharing data. Support for this work was provided by the ARSEP and by a Wellcome Senior Fellowship (Reference 074524) from The Wellcome Trust.

References

1. Spielman RS, McGinnis RE, Ewens WJ. Transmission test for linkage disequilibrium: The insulin gene region and insulin-dependent diabetes mellitus (IDDM). *Am J Hum Genet.* 1993; 52:506–516. [PubMed: 8447318]
2. Cordell HJ, Clayton DG. A unified stepwise regression procedure for evaluating the relative effects of polymorphisms within a gene using case/control or family data: Application to hla in type 1 diabetes. *Am J Hum Genet.* 2002; 70:124–141. [PubMed: 11719900]
3. Dudbridge F, Koeleman BP, Todd JA, Clayton DG. Unbiased application of the transmission/disequilibrium test to multilocus haplotypes. *Am J Hum Genet.* 2000; 66:2009–2012. [PubMed: 10775523]
4. Clayton D. A generalization of the transmission/disequilibrium test for uncertain-haplotype transmission. *Am J Hum Genet.* 1999; 65:1170–1177. [PubMed: 10486336]
5. Dudbridge F. Pedigree disequilibrium tests for multilocus haplotypes. *Genet Epidemiol.* 2003; 25:115–121. [PubMed: 12916020]
6. Horvath S, Xu X, Lake SL, Silverman EK, Weiss ST, Laird NM. Family-based tests for associating haplotypes with general phenotype data: Application to asthma genetics. *Genet Epidemiol.* 2004; 26:61–69. [PubMed: 14691957]
7. Cordell HJ, Barratt BJ, Clayton DG. Case/pseudocontrol analysis in genetic association studies: A unified framework for detection of genotype and haplotype associations, gene-gene and gene-environment interactions, and parent-of-origin effects. *Genet Epidemiol.* 2004; 26:167–185. [PubMed: 15022205]

8. Schafer JL. Multiple imputation: A primer. *Stat Methods Med Res.* 1999; 8:3–15. [PubMed: 10347857]
9. Schafer, JL. *Analysis of incomplete multivariate data.* Chapman & Hall/CRC; 1997.
10. Little, RJA, Rubin, DB. *Statistical analysis with missing data.* 2nd ed. Wiley-interscience; 2002.
11. Kistner EO, Weinberg CR. A method for identifying genes related to a quantitative trait, incorporating multiple siblings and missing parents. *Genet Epidemiol.* 2005; 29:155–165. [PubMed: 16025442]
12. Kraft P, Cox DG, Paynter RA, Hunter D, De Vivo I. Accounting for haplotype uncertainty in matched association studies: A comparison of simple and flexible techniques. *Genet Epidemiol.* 2005; 28:261–272. [PubMed: 15637718]
13. Cordell HJ. Estimation and testing of genotype and haplotype effects in case-control studies: Comparison of weighted regression and multiple imputation procedures. *Genet Epidemiol.* 2006; 30:259–275. [PubMed: 16496312]
14. Tanner M, Wong W. The calculation of posterior distributions by data augmentation (with discussion). *J Amer Stat Soc.* 1987; 81:528–550.
15. O'Connell JR. Zero-recombinant haplotyping: Applications to fine mapping using snps. *Genet Epidemiol.* 2000; 19(Suppl 1):S64–70. [PubMed: 11055372]
16. Rubin, DB. *Multiple imputation for non response in surveys.* New York: 1987.
17. Li K, Meng X, Raghunathan TE, Rubin DB. Significance levels from repeated p-values with multiply-imputed data. *Statistica Sinica.* 1991; 1:65–92.
18. Self SG, Longton G, Kopecky KJ, Liang KY. On estimating HLA/disease association with application to a study of aplastic anemia. *Biometrics.* 1991; 47:53–61. [PubMed: 2049513]
19. Schaid DJ. General score tests for associations of genetic markers with disease using cases and their parents. *Genet Epidemiol.* 1996; 13:423–449. [PubMed: 8905391]
20. Tregouet DA, Escolano S, Tiret L, Mallet A, Golmard JL. A new algorithm for haplotype-based association analysis: The stochastic-em algorithm. *Ann Hum Genet.* 2004; 68:165–177. [PubMed: 15008795]
21. Deltour I, Richardson S, Le Hesran JY. Stochastic algorithms for markov models estimation with intermittent missing data. *Biometrics.* 1999; 55:565–573. [PubMed: 11318215]
22. Stephens M, Smith NJ, Donnelly P. A new statistical method for haplotype reconstruction from population data. *Am J Hum Genet.* 2001; 68:978–989. [PubMed: 11254454]
23. Alizadeh M, Babron MC, Birebent B, Matsuda F, Quelvennec E, Liblau R, Cournu-Rebeix I, Momigliano-Richiardi P, Sequeiros J, Yaouanq J, Genin E, et al. Genetic interaction of *ctla-4* with *hla-dr15* in multiple sclerosis patients. *Ann Neurol.* 2003; 54:119–122. [PubMed: 12838528]




Figure 1.

Type 1 error at $\alpha=0.05$ on simulated data as a function of the percentage of missing data when using or not the MI approach. The number of complete families available for the study without MI is plotted on the second y axis.

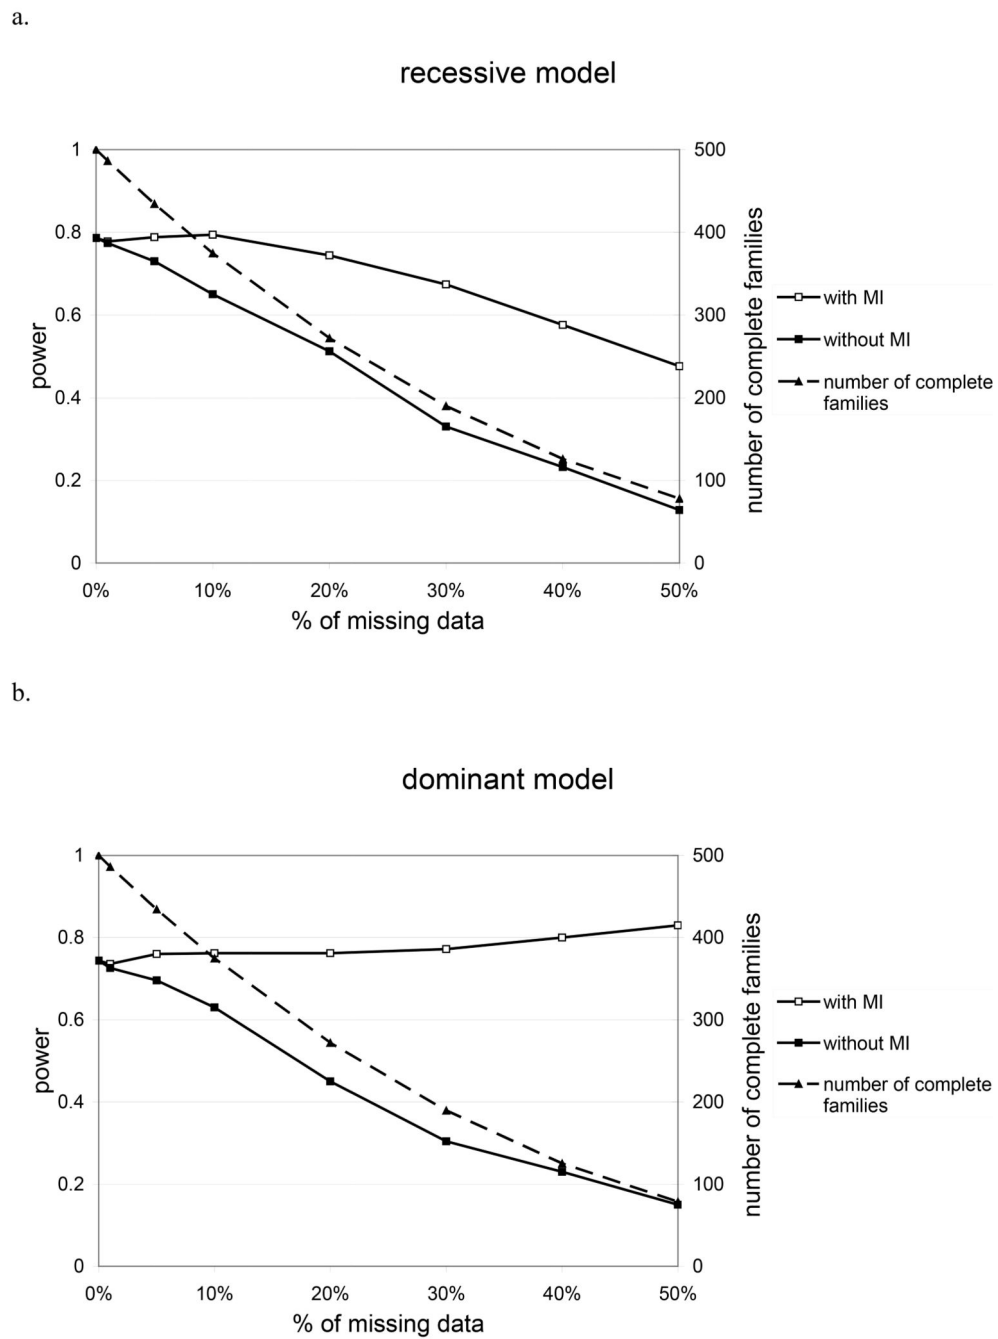



Figure 2. Power at $\alpha=0.05$ on simulated data as a function of the percentage of missing data when using or not the MI approach for a recessive model (a) and a dominant model (b) with genotype relative risks 1.5. The number of complete families available for the study without MI is plotted on the second y axis.

a.



b.

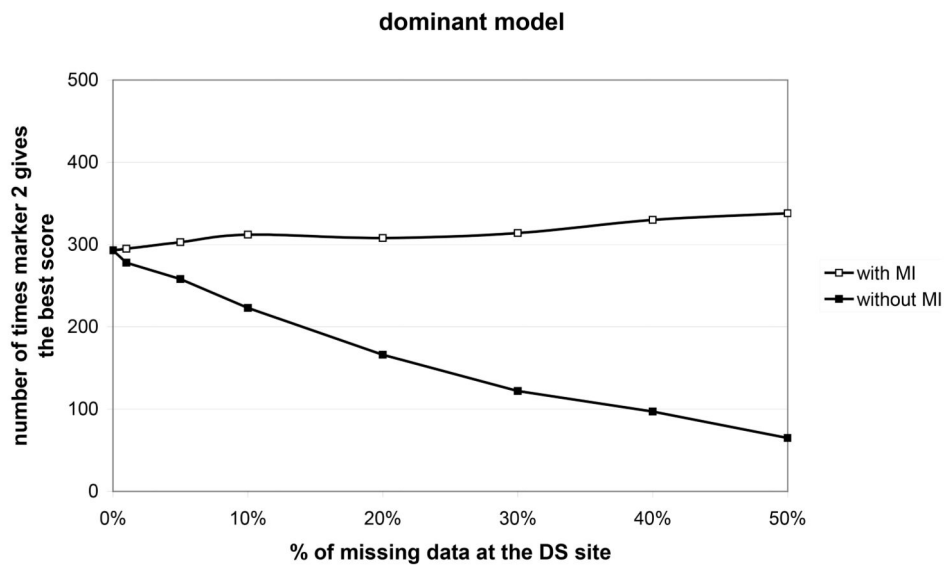


Figure 3. Number of replicates among the 500 replicates where the disease susceptibility site (locus 2) gives the best score, conditional on the fact the highest score is significant, in function of the percentage of missing data when using or not the MI approach for a recessive (a) and a dominant model (b) with a genotype relative risks of 1.5.




Figure 4. haplotypic bias between the expected and the observed haplotypic relative risk in function of the percentage of missing data in the case of the 2-locus multiplicative model when using or not the MI algorithm. Haplotype relative risks hr2, hr3 and hr4 correspond to the relative risks for haplotypes 12, 21 and 22 respectively (relative to haplotype 11).




Figure 5. Median of the absolute genotypic bias in function of the percentage of missing data in the case of the 2-locus multiplicative and non multiplicative model when using or not MI algorithm.

Table 1

Haplotype frequencies for the five loci considered in the simulations.

haplotypes	frequencies
1 1 1 1 1	0.310
1 1 1 1 2	0.005
1 1 1 2 1	0.099
1 1 2 1 1	0.018
1 2 1 1 1	0.002
1 2 2 1 2	0.002
2 1 1 1 1	0.003
2 1 1 2 1	0.002
2 1 2 1 1	0.060
2 1 2 1 2	0.003
2 1 2 2 1	0.002
2 2 1 1 1	0.017
2 2 1 1 2	0.002
2 2 2 1 1	0.376
2 2 2 1 2	0.094
2 2 2 2 1	0.002
2 2 2 2 2	0.003

Table 2

- a. List of the different 1-locus models used in simulation where parameters correspond to the genotype relative risks for genotypes 1/1, 1/2, 2/2 respectively.
- b. Multiplicative 2-locus model where parameters correspond to the genotype relative risk arising from the association of the two transmitted haplotypes, with the four possible haplotypes denoted 11, 12, 21, 22 respectively.
- c. Non-multiplicative 2-locus model where parameters correspond to the genotype relative risk arising from the association of the two transmitted haplotypes, with the four possible haplotypes denoted 11, 12, 21, 22 respectively.

a	<i>parameter</i>	<i>model1</i>	<i>model2</i>	<i>model3</i>	<i>model4</i>
<i>one locus models</i>	11	1	1	1	1
	12	1	1.5	1	3
	22	1.5	1.5	3	3
b	<i>parameter</i>	<i>11</i>	<i>12</i>	<i>21</i>	<i>22</i>
<i>multiplicative model</i>	11	1	3	5	6
	12	3	9	15	18
	21	5	15	25	30
	22	6	18	30	36
c	<i>parameter</i>	<i>11</i>	<i>12</i>	<i>21</i>	<i>22</i>
<i>non multiplicative model</i>	11	1	2	2	2
	12	2	8	2	2
	21	2	2	12	2
	22	2	2	2	16

Table 3

Bias and 95% confidence interval coverage of genetic parameter estimates from two-locus simulation study when using case/control or case/parent trio data with 10% missing data. The second column corresponds to the expected relative risk associated with each genotype.

<i>genotypes</i>	<i>relative risk</i>	<i>case/control</i>		<i>trios families</i>	
		<i>Bias</i>	<i>Coverage</i>	<i>Bias</i>	<i>Coverage</i>
<i>1-1/1-2</i>	<i>2</i>	-0.089	0.968	0.027	0.959
<i>1-1/2-1</i>	<i>2</i>	0.002	0.971	0.070	0.971
<i>1-1/2-2</i>	<i>2</i>	-0.315	0.792	-0.012	0.968
<i>1-2/1-2</i>	<i>8</i>	-0.170	0.950	-0.141	0.956
<i>1-2/2-1</i>	<i>2</i>	0.371	0.783	-0.019	0.950
<i>1-2/2-2</i>	<i>2</i>	0.191	0.953	-0.014	0.950
<i>2-1/2-1</i>	<i>12</i>	-0.063	0.950	-0.054	0.968
<i>2-1/2-2</i>	<i>2</i>	0.212	0.868	0.050	0.962
<i>2-2/2-2</i>	<i>16</i>	-0.056	0.962	-0.031	0.974