Published in final edited form as:

Ann Hum Genet. 2018 March ; 82(2): 93–101. doi:10.1111/ahg.12229.

Fast permutation tests and related methods, for association between rare variants and binary outcomes

ARJUN SONDHI and KENNETH MARTIN RICE

Department of Biostatistics, University of Washington, Seattle, WA, USA

Summary

In large scale genetic association studies, a primary aim is to test for association between genetic variants and a disease outcome. The variants of interest are often rare, and appear with low frequency among subjects. In this situation, statistical tests based on standard asymptotic results do not adequately control the Type I error rate, especially if the case:control ratio is unbalanced. In this paper, we propose the use of permutation and approximate unconditional tests for testing association with rare variants. We use novel analytical calculations to efficiently approximate the true Type I error rate under common study designs, and in numerical studies show that the proposed classes of tests significantly improve upon standard testing methods. We also illustrate our methods in data from a recent case-control study, for genetic causes of a severe side-effect of a common drug treatment.

Keywords

association tests; binary outcomes; permutation tests; rare variants

1 Introduction

Association studies are often performed for binary traits, providing new knowledge of the genetic causes of human diseases, using data from case-control and cohort studies (Verhaaren et al., 2015; Opherk et al., 2014; Danjou et al., 2015; Hoffmann et al., 2015). Recent advances in sequencing technology have made it practical to type essentially every variant on the genome; to avoid spurious findings, very low Type I error rates must therefore be maintained (Hoggart et al., 2008). However, for the rare variants now being studied, standard analytic approaches do not reliably achieve their nominal rates, (Li and Leal, 2008; Xing et al., 2012; Ma et al., 2013) and may permit too many Type I errors. The problem can be particularly severe when the ratio of cases to controls is extreme. Adjustments that maintain control at the nominal rate can be conservative, leading to loss of power relative to methods that control Type I errors more accurately.

Method development: KMR, AS; Method implementation: AS; Data analysis: AS, Writing: AS, KMR

7 Conflict of interest statement

The authors have no conflicts of interest to declare.

In this paper, motivated by work in a case-control study of rhabdomyolysis, we develop methods with improved control of the Type I error rate, when testing single rare variants for association with binary traits. In Section 2, we explain a novel numerical method that approximates the actual Type I error rate of a test statistic given sample size, significance level, and a variant's expected frequency; we also show how the same basic ideas can be used in permutation and approximate unconditional tests, and how the ideas can be used when adjusting for covariates. In Section 3, we give the results of the numerical studies performed, demonstrating improvements over standard asymptotic tests. Section 4 applies these methods to data from a case-control study of statin-related rhabdomyolysis, and we conclude with a short discussion, including details of an R package that implements our methods.

2 Materials and Methods

With rare variants, homozygotes are so rare as to be negligible for analysis, and it suffices to consider whether subjects have any copies of the variant present ($G = 1$) or not ($G = 0$). This step also means it is simple to enumerate all possible datasets; given fixed numbers of cases and controls (m_1 and m_0 respectively), we need only consider the number of cases (r_1) and controls (r_0) with the variant (adjustment for covariates is considered in Section 2.3.).

For a variant with a given minor allele frequency (MAF), following Ma et al. (2013), we define expected number of minor allele carriers as $EMAC = (m_0 + m_1) \times (1 - (1 - MAF)^2)$.

Under the null hypothesis of no association, it follows that

$$r_1 \sim \text{Binom}\left(m_1, \frac{EMAC}{m_0 + m_1}\right)$$

$$r_0 \sim \text{Binom}\left(m_0, \frac{EMAC}{m_0 + m_1}\right).$$

independently. Therefore, given a value of $EMAC$ —or equivalently MAF —and the fixed numbers of cases and controls in the study at hand, we can simply write down the probability of seeing all possible datasets under the null hypothesis.

In theory, this direct enumeration allows exact calculation of the Type I error rate for any test: the Type I error rate is the sum of the probabilities of the datasets for which a significant test result is returned. Formally, a dataset (m_0, m_1, r_0, r_1) returns a significant test result when its associated p-value $p(r_0, r_1; m_0, m_1) \leq \alpha$, where

$p(r_0, r_1; m_0, m_1) = \mathbb{P}[|T| \geq T_{obs}; H_0]$ is the probability of test statistic T equaling or exceeding the observed value T_{obs} when the null hypothesis holds. The Type I error rate of the test at nominal level α is then defined as

$$T1ER(\alpha) = \sum_{0 \leq r_0 \leq m_0, 0 \leq r_1 \leq m_1} f(r_0, r_1; m_0, m_1, EMAC) 1_{p(r_0, r_1; m_0, m_1) < \alpha} \quad (1)$$

where $f(r_0, r_1; m_0, m_1, EMAC)$ denotes the probability of observing data r_0, r_1 under the null hypothesis, given m_0, m_1 , and a specific EMAC. Although not discussed further, the approach is easily adapted to p -values that use lower tail areas -- below T_{obs} instead of above, or two-sided tests that examine tail areas beyond $\pm T_{obs}$.

In practice, the sum of $(m_0 + 1) \times (m_1 + 1)$ terms in (1) may be large, making computation too slow for some purposes. However, for work on rare variants, almost all of the summands contribute negligibly to the overall Type I error rate. A practical solution is therefore to truncate the summation in (1) by zeroing-out terms that, in total, represent no more than a small fraction of the Type I error rate.

Taking this approach, in our work we will zero-out terms in (1) representing datasets for which $r_0 + r_1$ exceeds the upper 10^{-12} quantile of the distribution of $r_1 + r_0$, i.e. of

$\text{Binom}\left(m_0 + m_1, \frac{EMAC}{m_0 + m_1}\right)$ (Figure 1 gives a graphical description of this process). By setting

these terms to zero, in this example, we therefore understate the Type I error rate by no more than 10^{-12} , which is acceptable given our focus on Type I error rates near $\alpha = 5 \times 10^{-8}$, and maintain practical computation times, even for large studies. For example, performing this calculation for $(m_0, m_1, EMAC) = (500, 500, 15)$ at $\alpha = 5 \times 10^{-8}$ using the standard Score test takes 0.05 seconds on a standard laptop. Without the zeroing-out method, the calculation takes 15.3 seconds, i.e. more than 300 times faster. Our choice of α corresponds to testing a million independent variants (Pe'er et al., 2008), a level that has been widely-adopted as the standard for genome-wide work.

We emphasize that this approximation of the Type I error rate is entirely general; any test statistic T can be used, including the familiar Wald, Score, likelihood ratio test statistics (see e.g. Ma et al. (2013)) or more sophisticated choices such as the Firth test (Firth, 1993; Heinze et al., 2013). Accurate knowledge of these approximate tests enables users to better compare their performance at the nominal α .

The formulation of Type I error rate in (1) and its approximation can also directly inform construction of permutation and approximate unconditional tests, as we discuss in Sections 2.1 and 2.2 below. We briefly discuss adjusting for covariates in Section 2.3, for both forms of test.

2.1 Permutation tests

The ability of permutation tests to provide accurate p -values for association testing under minimal assumptions is well-known (Pitman, 1937; Huo et al., 2014; Nichols and Holmes, 2002; Anderson, 2001); where they are applicable, permutation tests are regarded by many

analysts as the ‘gold standard’ method. For quantitative traits, a major drawback is that permutations must, in practice, be performed using random number generation (Boyett and Shuster, 1977). For analysis of binary traits this is not needed; we can instead enumerate all possible permutations and obtain accurate p-values.

Using the same notation as above, a permutation test requires an observed test statistic, T_{obs} , calculated on the observed data (m_0, m_1, r_0, r_1) . We shall consider test statistics from standard score, Wald, likelihood ratio, and Firth test approaches, thus providing a permutation version of them. The test statistic is also calculated for each possible datasets (m_0, m_1, r'_0, r'_1) obtained by permuting binary outcomes (e.g. case/control labels) among all study subjects, or equivalently permuting the variant/non-variant carrier status among all subjects. Under permutation, the total number of minor allele carriers is the same as in observed data, that is, $r_0 + r_1 = r'_0 + r'_1$, and under the null hypothesis of no association the probability of observing each dataset follows the hypergeometric distribution (Good, 2005):

$$\tilde{f}(r'_0, r'_1; m_0, m_1, r_0 + r_1) = \frac{\binom{r_0 + r_1}{r'_1} \binom{m_0 + m_1 - r_0 - r_1}{m_1 - r'_1}}{\binom{m_0 + m_1}{m_1}}.$$

The permutation p-value is then defined as

$$p_{perm}(r_0, r_1; m_0, m_1) = \sum_{r'_0, r'_1} \tilde{f}(r'_0, r'_1; m_0, m_1, r_0 + r_1) 1_{|T_{r'_0, r'_1}| \geq |T_{obs}|},$$

i.e. the sum of probabilities of datasets with the same number of allele carriers that result in more extreme test statistics than T_{obs} . The datasets enumerated in this method are illustrated in Figure 2.

Permutation tests are exact, in the sense that the observed Type I error rate will always be less than or equal to α . This result is well-known and dates back to Fisher (Janssen and Pauls, 2003). In particular, for the rare variant setting, permutation tests will be fairly conservative. A mathematical explanation is given in the Appendix.

2.2 Approximate Unconditional (AU) tests

Approximate Unconditional (AU) tests (Storer and Kim, 1990) provide Type I error rates closer to the nominal level than permutation approaches. Unlike permutation tests, AU tests are not guaranteed to always strictly control the Type I error rate, but this anti-conservatism (where it occurs at all) is usually very mild.

Using the same notation as above, AU tests calculate a test statistic T_{obs} from the observed data (m_0, m_1, r_0, r_1) and from all possible datasets (m_0, m_1, r'_0, r'_1) but the restriction that

$r_0 + r_1 = r'_0 + r'_1$. The probability of observing each dataset under the null hypothesis is calculated using fitted binomial distributions, i.e.

$$\hat{f}(r'_0, r'_1; m_0, m_1, r_0, r_1) = \binom{m_0}{r'_0} \binom{m_1}{r'_1} \left(\frac{r_0 + r_1}{m_0 + m_1} \right)^{r'_0 + r'_1} \left(1 - \frac{r_0 + r_1}{m_0 + m_1} \right)^{m_0 + m_1 - r'_0 - r'_1}.$$

The AU test's p-value is then defined as

$$p_{AU}(r_0, r_1; m_0, m_1) = \sum_{r'_0, r'_1} \hat{f}(r'_0, r'_1; m_0, m_1, r_0, r_1) 1_{|T_{r'_0, r'_1}| \geq |T_{obs}|}, \quad (2)$$

i.e. the sum of probabilities of datasets that result in more extreme test statistics than T_{obs} .

We can then apply Equation (1) and write the Type I error rate as:

$$T1ER(\alpha) = \sum_{0 \leq r'_0 \leq m_0, 0 \leq r'_1 \leq m_1} \hat{f}(r'_0, r'_1; m_0, m_1, EMAC) 1_{p_{AU}(r'_0, r'_1; m_0, m_1) < \alpha}, \quad (3)$$

Compared to the permutation test, the AU test's p-value sums over many more possible datasets, allowing less crude approximation of the Type I error rate. This comes at the cost of using the same data to fit the null binomial models, and hence losing guaranteed control of the Type I error rate. However, in our setting a bigger practical concern is that taking a naïve approach to calculation in Equation (2) would require $(m_0 + 1) \times (m_1 + 1)$ evaluations for each p-value, which may be a burden, as with Equation (1). A much quicker approach that is still adequate in practice uses the same zeroing-out idea as before -- we only sum elements (r'_0, r'_1) in (2) for values of $r'_0 + r'_1$ between the upper and lower 10^{-12} quantiles of the $\text{Binom}\left(m_0 + m_1, \frac{r_0 + r_1}{m_0 + m_1}\right)$ distribution.

The datasets enumerated in this method are illustrated in Figure 3. As with the calculation of Type I errors in Section 2, the zeroing out leads to a slight understatement of the p-value compared to complete enumeration. However, understating the p-value by at most 2×10^{-12} is a very minor concern when $\alpha = 10^{-8}$, several orders of magnitude greater, and comes in return for a substantial speed increase. For example, computing an AU p-value under the Score test with data $(m_0, m_1, r_0, r_1) = (5000, 5000, 10, 50)$ takes 0.05 seconds on a standard laptop with zeroing out and 28.4 seconds without, i.e. over 500 times faster.

The AU approach, like the permutation approach, is completely general, and AU versions of any test can be implemented. We shall use standard Score, Wald, likelihood ratio and Firth tests.

2.3 Adjusting for covariates

Both permutation and AU tests permit adjustment for covariates through stratification, i.e. only using information about association from within groups of subjects for whom confounding factors (for example ancestry) are held constant (Clayton et al., 1993).

Extending the previous notation, for stratified tests we now refer to vectors $\mathbf{m}_0, \mathbf{m}_1, \mathbf{r}_0, \mathbf{r}_1$, each of length q , where q is the number of strata defined by the levels of one or more categorical covariates. Indexing strata by i , with $1 \leq i \leq q$, for each stratum i the stratified test enumerates all possible strata-specific datasets $(m_{0i}, m_{1i}, r'_{0i}, r'_{1i})$ such that $r'_{0i} + r'_{1i} = r_{0i} + r_{1i}$, computing a test statistic for each. The test statistics $T_i(r_{0i}, r_{1i})$ from each strata are combined (by default they are added) to produce a single test statistic for the whole dataset; formally we define

$$T_{r'_0, r'_1} = \sum_{i=1}^q T_i(r_{0i}, r_{1i}).$$

The p-value, which as before compares this single test statistic to what might have been observed under the null, uses the hypergeometric distribution for each set of stratum-specific counts. We write the probability of observing specific datasets as

$$\check{f}(r'_0, r'_1; \mathbf{m}_0, \mathbf{m}_1, \mathbf{r}_0, \mathbf{r}_1) = \prod_{i=1}^q \check{f}_i(r'_{0i}, r'_{1i}; m_{0i}, m_{1i}, r_{0i}, r_{1i})$$

Where

$$\check{f}_i(r'_{0i}, r'_{1i}; m_{0i}, m_{1i}, r_{0i}, r_{1i}) = \frac{\binom{r_{0i} + r_{1i}}{r'_{1i}} \binom{m_{0i} + m_{1i} - r_{0i} - r_{1i}}{m_{1i} - r'_{1i}}}{\binom{m_{0i} + m_{1i}}{m_{1i}}}$$

and formally define the p-value as

$$p_{\text{strat. perm}}(r'_0, r'_1; \mathbf{m}_0, \mathbf{m}_1, \mathbf{r}_0, \mathbf{r}_1) = \sum_{r'_0, r'_1} \check{f}(r'_0, r'_1; \mathbf{m}_0, \mathbf{m}_1, \mathbf{r}_0, \mathbf{r}_1) \mathbb{1}_{|T_{r'_0, r'_1}| \geq |T_{r_0, r_1}|},$$

i.e. the sum of probabilities of datasets that result in more extreme test statistics than T_{obs} , where T_{obs} is the test statistic corresponding to the data that was observed.

The stratified AU test is constructed from the same steps as the permutation except for three differences, described earlier in Section 2.2. First, the datasets considered for each strata include any values of $0 \leq r'_{0i} \leq m_{0i}$ and $0 \leq r'_{1i} \leq m_{1i}$. Second, the probabilities $\check{f}$ of each dataset are constructed from fitting a null binomial model within each strata. Third,

summands within each strata are zeroed-out for which the total contribution is no more than 2×10^{-12} .

Our approach removes confounding effects by using stratified analysis. Implemented carefully, there is little to choose between use of stratification versus model-based regression adjustment. In line with Clayton and Hills (1993, *Statistical Methods in Epidemiology*, pg 273) we find it appealing that the stratification approach forces careful consideration of a which confounders are a priori most important to adjust for, and for stratification approaches to be based closely on the scientific question of interest. Moreover, categorizing confounding into strata is the only approach under which our enumeration approach for exact inference is feasible; regression-based alternatives with continuously-valued covariates and standard computing resources would have to compute p-values by some form of Monte Carlo method, with consequent Monte Carlo error and long compute times. While a limitation of the approach is that a finite number of strata do have to be defined, in practice we have found that stratifying data into commonly-used ancestry groups (European, African-American, etc) is generally sufficient to account for confounding issues present in analyses of complex disease traits.

3 Analytical calculation results

To illustrate analytical calculations, we set the total sample size to be $N=10,000$ —close to that seen in Section 4’s example—and considered case:control matching ratios of 1:1, 1:3, and 1:19, fixing the quantities m_0 and m_1 . For *EMAC* ranging from 1 to 100, we enumerated all plausible observed data values of r_0 and r_1 under the null hypothesis of independence between the minor allele and case:control status. A “plausible” dataset here was defined as one where the sum $r_0 + r_1$ was in between the lower and upper 10^{-12} quantiles of the null Binomial distribution. Permutation and AU versions of Score, Wald, likelihood ratio and Firth tests were examined. For comparison we also computed the standard Score, Wald, likelihood ratio, Firth tests, and Fisher’s exact test, which is itself a permutation test. For permutation, AU, and standard tests we also considered a regularized Wald test, which avoids undefined test statistics by adding 0.5 to each cell count when any count is zero. We set the nominal significance level at $\alpha = 5 \times 10^{-8}$, and calculated the Type I error rate for each test as in Equation (1), across the range of *EMAC*. Although we do not consider power calculations under an alternative hypothesis in this section, these are given in Appendix B for the AU Firth test.

As seen in Figure 4, the tests based on standard asymptotics do not adequately control the Type I error rate. In the balanced design, the tests are overly conservative, with the exception of the likelihood ratio test, which is anti-conservative. The Score test has very a large Type I error rate under the 1:3 ratio, so is presented separately. This is also true under the 1:19 ratio for the Score and Wald tests, which are omitted. The other tests continue to be conservative, and the likelihood ratio test’s Type I error rate is too large over certain ranges. The Firth test consistently performs the best.

In Figure 5, we see that the permutation tests improve upon most of the standard tests, though remain more conservative than the regular Firth test. While these tests have the

advantage of being exact, as the case:control ratio becomes more unbalanced, the Type I error rate becomes more conservative. Under the 1:19 ratio, all tests perform nearly identically, with the exception of the unregularized Wald test.

In Figure 6, we see that the AU tests show a large improvement over standard and permutation tests, especially in the AU likelihood ratio and AU Firth tests. Though they are not exact, the excess Type I error rate is mild. Note that under the 1:19 ratio, the Firth and likelihood ratio tests perform identically.

4 Application: Rhabdomyolysis case-control study

The data comes from an exome-sequencing study, in which 9,763 subjects who used statins were considered; 211 cases with rhabdomyolysis and 9,552 controls. The original dataset consists of 2,194,116 variants sequenced. The rationale for this design are described in detail by Marcianti et al. (2011). Our interest was primarily in assessing if there existed rare genetic variants associated with developing rhabdomyolysis in statin users. We defined 'rare' variants as those where less than or equal to 100 study participants carried the minor allele. Variants with less than 5 minor allele carriers were also removed, as these provide no ability to produce significant values at the low α threshold used in this form of study. Finally, for quality control, we filtered out variants with a genotyping rate of less than 0.85. We did not filter variants using the Hardy-Weinberg test. Applying these filters left 161,428 variants, and there are no covariates for which to adjust in this analysis.

We applied the AU and permutation versions of the likelihood ratio test, and the permutation and standard Firth test to all variants. The entire analysis took approximately 6.5 hours on a shared server, using a single CPU. The AU version of the Firth test was not used due to its high computational burden. The resulting QQ plot and a plot of the inflation (45 degree rotated QQ plot) observed are given in Figure 7.

Based on our numerical results, it is reasonable to expect that the AU likelihood ratio test provides the best control of the Type I error rate. Applied to this dataset, while some granularity in the larger p-values is present on the left hand of both plots, we observe that the AU likelihood ratio test results in significantly less inflation than the standard Firth test.

Therefore, although no variant was found to have an AU p-value less than $\alpha = 5 \times 10^{-8}$, we believe that the variants which deviate from the expected null p-value distribution under the AU test are more likely to be of interest than those under the other methods.

5 Discussion

We have developed and implemented association tests for rare genetic variants, that control Type I error rates better than standard asymptotic tests. Of the tests proposed in this paper, the AU version of the likelihood ratio and Firth tests perform the best, particularly when the ratio of cases to controls is extreme. However, the AU version of the Firth test has a notably higher computational burden than competitors. Therefore, we recommend the the AU likelihood ratio test, for large genome-wide studies. If an exact test is necessary, then a permutation test is recommended; though conservative, it tends to show an improvement

over standard tests. We note that if the expected number of minor allele carriers is less than 20, then no test will perform adequately, and conservative control of the Type I error rate is the best achievable property. More generally, we note that in high-throughput rare variant work, unless effect sizes are large, power will be limited even when Type I error rates are not controlled conservatively. As pointed out by a reviewer, in such settings an appealing alternative form analysis tests for association across a group of variants; examples include SKAT (Wu et al, 2011) but also “burden” tests, that collapse genotype across a region to a univariate measure. These approaches can improve power over single-variant approaches, by both combining multiple association “signals”, but also by reducing the multiple testing burden. However, the inference they provide is region-specific and not variant-specific, leaving (for example) no strong indication of which variants may be causal. While not explored here, our permutation tests could be used with CAST (Morgenthaler and Thilly, 2007), a burden test that collapses genotype to presence/absence of a particular class of variants.

The methods described here have been implemented in an R package, AUtests. This package contains the functions `basic.tests`, `perm.tests`, and `au.tests`, which implement all the respective standard, permutation, and AU tests for a given vector of counts (m_0, m_1, r_0, r_1) , returning a vector of p-values. The AU Firth test is implemented in a separate function, `au.firth`, due to its increased computational time. For a typical dataset $(m_0, m_1, r_0, r_1) = (10000, 10000.50, 50)$, on a standard laptop, the `basic.tests` function takes 0.03 seconds of CPU time, the `perm.tests` function takes 0.21 seconds, the `au.tests` function takes 0.39 seconds, and the `au.firth` function takes 51 seconds. To account for covariates, appropriately categorized, the package also contains the functions `au.test.strat` and `perm.test.strat`, which implement stratified AU and permutation likelihood ratio tests. The package is available on CRAN.

Acknowledgments

Research reported in this paper was supported by the National Institute on Aging of the National Institutes of Health under award numbers U01AG049505 and U01AG049507, and by the National Heart, Lung, and Blood Institute of the National Institutes of Health under award number R01HL078888. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

A: Appendix: exact control of permutation tests

In this appendix, we show that permutation tests give exact control of the Type I error rate. Rewriting equation (1) as a double summation over the observed minor allele count ($t = r_0 + r_1$) and the number of these in the controls, we obtain

$$T1ER(\alpha) = \sum_{0 \leq t \leq m_0 + m_1} g(t; m_0, m_1) \sum_{r_0} f(r_0; m_0, m_1, t)^1 p(r_0, r_1; m_0, m_1) < \alpha$$

where $g()$ denotes the probability of the observed minor allele count, and $f()$ gives the probability of the observed counts in cases and controls given the minor allele count $r_0 + r_1$ --- so $f()$ supports values of r_0 between $\max(0, t - m_1)$ and $\min(m_0, t)$.

By construction, the inner sum always gives a value less than or equal to α ; the outer sum averages these, and so is similarly bounded. However, particularly for rare variants, the inner sum considers a small set of possible permutations, as illustrated in Figure 2. While this makes the test fast enough that zeroing-out is not required, it means that for small α , the actual Type I error rate, while below α , will be quite conservative for many values of m_0 and m_1 .

B: Appendix: AU test power calculations

In this section, we show power calculations for the AU Firth test, under the same matching ratios considered in the main paper and over the same range of *EMAC*. Specifically, we perform analytical calculations similarly as in Section 3, but consider data probabilities from a Binomial distribution where having a minor allele is related to case:control status through an odds ratio. We consider odds ratios ranging from 1.5 to 5. In Figure 8, we observe that power decreases as case:control ratios become more skewed. In particular, the extremely unbalanced 1:19 ratio requires a very large association in order to have meaningful power, even at higher minor allele counts.

References

- Anderson MJ. (2001). Permutation tests for univariate or multivariate analysis of variance and regression. *Canadian Journal of Fisheries and Aquatic Sciences*, 58(3):626–639.
- Boyett JM. and Shuster JJ. (1977). Nonparametric one-sided tests in multivariate analysis with medical applications. *Journal of the American Statistical Association*, 72(359):665–668.
- Clayton D, Hills M, and Pickles A (1993). *Statistical models in epidemiology*, volume 161 IEA.
- Danjou F, Zoledziwska M, Sidore C, Steri M, Busonero F, Maschio A, Mulas A, Perseu L, Barella S, Porcu E, Pistis G, Pitzalis M, Pala M, Menzel S, Metrustry S, Spector T, Leoni L, Angius A, Uda M, Moi P, Thein S, Galanello R, Abecasis G, Schlessinger D, Sanna S, and Cucca F (2015). Genome-wide association analyses based on whole-genome sequencing in sardinia provide insights into regulation of hemoglobin levels. *Nature Genetics*.
- Firth D (1993). Bias reduction of maximum likelihood estimates. *Biometrika*, 80(1):27–38.
- Good P (2005). *Permutation, Parametric and Bootstrap Tests of Hypotheses*. Springer New York.
- Heinze G, Ploner M, Dunkler D, and Southworth H (2013). *logistf: Firth's bias reduced logistic regression*. R package version 1.21.
- Hoffmann TJ, Van Den Eeden SK, Sakoda LC, Jorgenson E, Habel LA, Graff RE, Passarelli MN, Cario CL, Emami NC, Chao CR, Ghai NR, Shan J, Ranatunga DK, Quesenberry CP, Aaronson D, Presti J, Wang Z, Berndt SI, Chanock SJ, McDonnell SK, French AJ, Schaid DJ, Thibodeau SN, Li Q, Freedman ML, Penney KL, Mucci LA, Haiman CA, Henderson BE, Seminara D, Kvale MN, Kwok P-Y, Schaefer C, Risch N, and Witte JS. (2015). A large multiethnic genome-wide association study of prostate cancer identifies novel risk variants and substantial ethnic differences. *Cancer Discovery*, 5(8):878–891. [PubMed: 26034056]
- Hoggart CJ, Clark TG, De Iorio M, Whittaker JC, and Balding DJ. (2008). Genome-wide significance for dense snp and resequencing data. *Genetic Epidemiology*, 32(2):179–185. [PubMed: 18200594]
- Huo M, Heyvaert M, den Noortgate WV, and Onghena P. (2014). Permutation tests in the educational and behavioral sciences. *Methodology*, 10(2):43–59.
- Janssen A and Pauls T (2003). How do bootstrap and permutation tests work? *Annals of statistics*, pages 768–806.
- Li B and Leal SM. (2008). Methods for detecting associations with rare variants for common diseases: Application to analysis of sequence data. *The American Journal of Human Genetics*, 83(3):311–321. [PubMed: 18691683]

Figure 1:

Possible datasets and their contribution to $T1ER$ for the standard Score test, for $m_0 = m_2 = 1000$ and $EMAC = 15$. In a), the red/green zones indicate datasets where the standard score test is significant/not-significant at nominal $\alpha = 5 \times 10^{-8}$; the blue zone shows terms that are not zeroed-out using the truncation described in Section 2. In b) the same situation is shown, zoomed in and with box size proportional to the probability of each dataset; the actual $T1ER(\alpha)$ is given by the sum of the box areas in the two red zones. Zeroing-out contributions beyond the blue region, where $r_0 + r_1 > 50$, gives an approximation error in the p -value of no more than 10^{-12} .

Figure 2:

Datasets used when calculating the p-value for the permutation version of the Score test, for observed data $(m_0, m_1, r_0, r_1) = (1000, 1000, 5, 10)$. The size of each square corresponds to the probability of observing the corresponding dataset under the null hypothesis. The p-value is represented by the sum of the areas of the squares in the two shaded 'tails' of the distribution, containing all datasets with Score test statistic at least as extreme as the observed data.

Figure 3:

Datasets used when calculating p-value for the AU version of the Score test, for observed data $(m_0, m_1, r_0, r_1) = (1000, 1000, 5, 10)$. The size of each square corresponds to the probability of observing the corresponding dataset under the null hypothesis. In a), the p-value is represented by the sum of the areas of the squares in the two shaded areas, containing all datasets with Score test statistics at least as extreme as the observed data. In b), we show how truncation at the upper 10^{-12} quantile of the fitted distribution of r_0 and r_1 would zero out many datasets, making calculation much quicker.

Figure 4:

T1ER versus EMAC for various standard tests at $\alpha = 5 \times 10^{-8}$ (dotted line) and total $N=10,000$, grouped by matching ratio. All calculations use the zeroing-out technique of Section 2, and so understate the true *T1ER* by no more than 10^{-12} . Score and Wald tests are omitted from the final plot due to gross violation of the nominal α .

Figure 5:

T1ER versus EMAC for permutation versions of standard tests as described in Section 2.1 at $\alpha = 5 \times 10^{-8}$ (dotted line) and total $N=10,000$, grouped by matching ratio. All *T1ER* rate calculations use the zeroing-out technique of Section 2, and so understate the true *T1ER* by no more than 10^{-12}

Figure 6:

T1ER versus EMAC, for AU versions of standard tests as described in Section 2.2, at $\alpha = 5 \times 10^{-8}$ (dotted line) and total $N=10,000$, grouped by matching ratio. All *T1ER* rate calculations use the zeroing-out technique of Section 2, and so understate the true *T1ER* by no more than 10^{-12} .

Figure 7:

QQ plot and 45 degree rotated QQ plot of $-\log_{10}$ p-values for rhabdomyolysis dataset, as described in Section 4. After quality control filtering, 161,428 variants are analyzed, with between 5 and 100 minor allele carriers each. For each method, the QQ plot shows the ordered p-values versus the corresponding expected value from null, i.e. $\text{Uniform}(0,1)$ p-values. The blue cone shape indicates pointwise 95% prediction bounds for each ordered p-value. The rotated plot shows the same results, but where the y-axis shows the $-\log_{10}$

observed p-value minus the $-\log_{10}$ expected p-value; the blue cone has the same interpretation as before.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Figure 8:

Power curves giving the probability of rejecting the null hypothesis of independence by expected minor allele count. Different curves correspond to different odds ratios. Each panel corresponds to a different case:control matching ratio with an overall sample size of 20,000. Left: 1:1, middle: 1:3, right: 1:19.