Published in final edited form as:

Genet Epidemiol. 2013 February ; 37(2): 142–151. doi:10.1002/gepi.21699.

Detecting Rare Variant Effects Using Extreme Phenotype Sampling in Sequencing Association Studies

Ian J. Barnett, Seunggeun Lee, and Xihong Lin

Department of Biostatistics, Harvard School of Public Health, Boston, MA

Abstract

In the increasing number of sequencing studies aimed at identifying rare variants associated with complex traits, the power of the test can be improved by guided sampling procedures. We confirm both analytically and numerically that sampling individuals with extreme phenotypes can enrich the presence of causal rare variants and can therefore lead to an increase in power compared to random sampling. While application of traditional rare variant association tests to these extreme phenotype samples requires dichotomizing the continuous phenotypes before analysis, the dichotomization procedure can decrease the power by reducing the information in the phenotypes. To avoid this, we propose a novel statistical method based on optimal SKAT (SKAT-O) that allows us to test for rare variant effects using continuous phenotypes in the analysis of extreme phenotype samples. The increase in power of this method is demonstrated through simulation of a wide range of scenarios as well as in the triglyceride data of the Dallas Heart Study.

Keywords

Complex trait associations; Selective sampling; Rare genetic variants; Extreme phenotype sampling

Introduction

With the increase in the number of sequencing studies[Biesecker, et al. 2011], there is a newfound access to samples with low frequency (MAF 1–5%) and rare (MAF <1%) genetic variants. In the search for genetic components of complex traits, discovered common variants (MAF > 5%) from genome-wide association studies explain only a small proportion of the total heritability of these traits[Ioannidis, et al. 2009; Maher 2008; Manolio, et al. 2009]. As a result, attention has turned to low frequency and rare variants instead expecting that they could play an important role in uncovering gene-phenotype relationships[Cirulli and Goldstein 2010; Ji, et al. 2008; Nejentsev, et al. 2009; Ng, et al. 2008; Ramser, et al. 2008]. Unfortunately, rare variants are difficult to detect in even reasonably large samples. This problem can be alleviated through the development of powerful study designs. To this effort, numerous association studies have chosen to sample subjects with extreme phenotypes in the hope of increasing power to detect causal SNPs[Clement, et al. 1995; Gu, et al. 1997; Hu, et al. 2009; Khor and Goh 2010; Li and Leal 2008; Liang, et al. 2000; Price,

Address for Correspondence: Xihong Lin, Ph.D., Department of Biostatistics, Harvard School of Public Health, 655 Huntington Avenue, Boston, MA 02115, Phone: (617) 432-2914, Fax: (617) 432-5619, xlin@hsph.harvard.edu.

Web Resources

The URLs for data presented herein are as follows:

Online Mendelian Inheritance in Man (OMIM), <http://www.omim.org>

An implementation of extreme phenotype SKAT and the associated power calculations in the R language can be found at <http://www.hsph.harvard.edu/~xlin/software.html>.

et al. 2008; Risch and Zhang 1995]. There have also been numerous developments in methodology to detect QTLs under these extreme phenotype sampling (EPS) study designs [Chen, et al. 2005; Huang and Lin 2007; Li, et al. 2011; Slatkin 1999; Wallace, et al. 2006]. A fundamental assumption that motivates these EPS methods is that rare causal variants are more likely found in the extremes of the quantitative trait. In this paper, we support the use of this practice by showing both analytically and numerically that EPS increases the presence of rare causal variants in a variety of settings. As a result, we show that EPS is more powerful for detecting for rare variant effects than random sampling.

Various methods have been proposed to tackle the challenge of association testing for rare variants. Burden tests such as the Combined Multivariate and Collapsing method (CMC) [Li and Leal 2008], Cohort Allelic Sums Test (CAST) [Morgenthaler and Thilly 2007] and the Weighted Sum Test (WST) [Madsen and Browning 2009] combine information from all rare variants within a target region such as an exon or gene by collapsing them into a single genetic variable, which is tested for association with the phenotypes of interest. Numerous rare variants testing methods have been developed using the same strategies [Bansal, et al. 2010; Basu and Pan 2011; Lee, et al. 2012a; Morris and Zeggini 2010; Price, et al. 2010]. A limitation of all burden tests is that they could lose significant amount of power in the presence of variants with different association directions and a large fraction of non-causal variants in the region. Alternatively the Sequence Kernel Association Test (SKAT) [Wu, et al. 2011] aggregates evidence of individual variant effects across the region using a kernel function and uses a computationally efficient mixed model variance component test to test for association. SKAT can naturally adjust covariates and has robust power in the presence of variants with different association directions and a large proportion of null variants. It is also a generalization of several non burden tests such as C-alpha test [Neale, et al. 2011], the SSU test [Pan 2009], and the haplotype association test [Tzeng and Zhang 2007]. Recently the optimal SKAT (SKAT-O) [Lee, et al. 2012b] has been proposed to unify the burden test and SKAT to a single framework and to construct the optimal test within the framework.

Moreover, limited statistical methods have been developed for studying rare variant effects when extreme phenotypes are sampled. In a typical EPS study, the two extremes are treated as two different groups representing a dichotomous phenotype. For example, Hu et al. [Hu, et al. 2009] used the contrast between subjects with high HDL-C levels against those with low HDL-C levels to identify an association with the ABCA1 gene. If the same method of extreme sampling were to instead retain the continuous phenotype values, the gain in information could provide greater power to detect gene-phenotype associations. For common variants, Huang and Lin [Huang and Lin 2007] proposed testing for associations between extreme continuous phenotypes and variants using the maximum likelihood method assuming a truncated normal distribution for extreme phenotype. Recently, this approach was adapted by Li et al. [Li, et al. 2011] to accommodate testing for multiple rare variant effects with the burden CMC approach. As a burden test, this approach is powerful when most variants in a region are causal and the effects of causal variants are in the same direction. However, it loses power in the presence of variants with different association directions or a large number of non-causal variants in a region.

In this paper, we first confirm both analytically and empirically that EPS substantially increases the chance to observe rare causal variants and hence increases their observed frequencies in finite study samples. Using this result, we demonstrate that EPS provides a more powerful design strategy for testing rare variant effects compared to random sampling. We next develop a new more powerful statistical method for testing for rare variant effects in EPS. Specifically, we extend SKAT and the optimal SKAT (SKAT-O) to EPS by analyzing extreme phenotypes as continuous variables within a likelihood framework. We show that the proposed tests perform well in a wide range of situations and outperform

burden tests. We further show that analysis using continuous extreme phenotypes (CEP) improves power for detecting rare variant effects compared to using dichotomized extreme phenotypes (DEP). We illustrate the finite sample performance of proposed methods by conducting extensive simulations and application to analysis of triglyceride levels from the Dallas Heart Study.[Victor, et al. 2004]

Material and Methods

Goals and notation

The goal is to find an optimum sampling strategy when resources are limited and to develop powerful association test methods to detect phenotype-genotype associations. We evaluate the effectiveness of extreme phenotype sampling (EPS) compared to random phenotype sampling.

We first confirm analytically that extreme phenotype sampling enriches causal rare variants by increasing their MAFs (Supplementary Materials Section 1). We consider the cases with a single causal variant and multiple causal variants and calculate the MAF in extreme phenotype sampling as a function of the population MAF, the threshold used to select extreme phenotypes, and the effect sizes of genotypes.

We next evaluate the two different methods that utilize EPS phenotypes in different ways: the method that retains continuous phenotypes and the method that dichotomizes them into cases and controls. We consider the case with a sample of n individuals who have been sequenced in a genomic region of interest containing p genetic variants. The i -th individual has covariate information over m covariates $\mathbf{X}_i = (X_{i1}, \dots, X_{im})'$, genotypes of the p variants in the region $\mathbf{G}_i = (G_{i1}, \dots, G_{ip})'$, and a continuous phenotype y_i . The genotype G_{ij} represents the number of copies of the minor allele of the j -th variant that the i -th individual has.

Model

To test for an association between the variants and continuous phenotype while controlling for covariates, consider a linear model

$$y_i = \alpha_0 + \alpha' \mathbf{X}_i + \beta' \mathbf{G}_i + \varepsilon_i$$

where $\varepsilon_i \sim N(0, \sigma^2)$. Here α_0 is an intercept term, $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_m]'$ is a vector of regression coefficients for the m covariates, and $\beta = [\beta_1, \beta_2, \dots, \beta_p]'$ is a vector of regression coefficients for the p genetic variants. The null hypothesis of $H_0: \beta = 0$ corresponds to no genetic effect on the trait. Since a p -DF likelihood ratio test has little power to detect causal variants particularly in the presence of a large number of rare variants, the gene-phenotype relationship is instead tested for by region-based tests such as burden tests and non-burden tests, e.g., SKAT. An adaptation of the CMC[Li and Leal 2008] burden test is used that collapses genotype information by counting the number of variants in the region before applying logistic regression to the collapsed statistic. We call this test DEP-Burden.

Association tests under the extreme phenotype sampling design

Since both SKAT and burden tests are capable of handling dichotomous phenotypes in the case-control setting, they can be applied to test for associations after using EPS. Dichotomizing the high phenotypic extremes as cases and the lower phenotypic extremes as controls is a natural extension of each test's functionality. However, applying SKAT and SKAT-O to continuous phenotype data obtained from EPS requires further development, since the extreme continuous phenotypes do not follow Gaussian distribution due to the

phenotypic selection. Suppose we select n samples with either $y_i > c_1$ or $y_i < c_2$, and denote the selected y_i as y_i^* . Then under the null hypothesis y_i^* follows truncated Gaussian distribution with a density function

$$f(y_i^*) = \frac{\phi(X_i\alpha, \sigma^2)}{\Phi(c_2, \sigma^2) + 1 - \Phi(c_1, \sigma^2)}$$

where $\phi(\mu, \sigma^2)$ and $\Phi(\mu, \sigma^2)$ are density and distribution functions of Gaussian distribution with mean μ and variance σ^2 .

To increase test power and decrease the test DF, we assume β_j follows an arbitrary distribution with mean 0 and variance φw_j^2 . We note that $H_0: \beta=0$ is equivalent to $H_0: \varphi=0$. The score test statistic of $\varphi=0$ is

$$Q_s = \sum_{j=1}^p w_j^2 \left(\sum_{i=1}^n G_{ij}(y_i^* - \hat{\mu}_i) \right)^2$$

where $\hat{\mu}_j$ is an estimated mean of y_i^* under the null hypothesis. We show that Q_s asymptotically follows a mixture of chi-square distribution (Supplementary Materials Section 2), and p-values can be obtained by the matching the moments or inverting the characteristic function [Davies 1980].

Recently Lee et al. (2012) [Lee, et al. 2012b] proposed an optimal unified approach, which unifies SKAT and burden test to adaptively select best test structure. Suppose Q_B is the score test statistics of the weighted burden test:

$$Q_B = \left(\sum_{i=1}^n (y_i^* - \hat{\mu}_i) \sum_{j=1}^p w_j G_{ij} \right)^2$$

and then the test statistic of unified test is

$$Q_\rho = (1 - \rho)Q_s + \rho Q_B,$$

where ρ ($0 \leq \rho \leq 1$) is a parameter to determine whether test is close to SKAT ($\rho=0$) or burden tests ($\rho=1$). It is based on a recent generalization of SKAT which allows the correlation among variants effects β 's. Under this setting, they proposed the optimal SKAT, called SKAT-O. This test is defined by selecting the ρ that minimizes the p-value of the SKAT-O test statistic,

$$T = \min_{0 \leq \rho \leq 1} p_\rho$$

where p_ρ is a p-value with given ρ . The test statistic T can be obtained by simple grid search across a range of ρ : set a grid $0 = \rho_1 < \rho_2 < \dots < \rho_b = 1$, then $T = \min(p_{\rho_1}, \dots, p_{\rho_b})$. In simulation studies and real data analysis, we used the equal sized grid of 11 points (from 0 to 1) to obtain T . From the fact that the Q_ρ can be decomposed to the shared random variables, asymptotic p-value of T can be obtained through computationally efficient one-

dimensional numerical integration (Supplementary Materials Section 3). We use this extreme phenotype optimal SKAT in our simulation studies and data analysis.

When the sample size is small, SKAT family methods (including SKAT and SKAT-O) can produce conservative results with both binary and extreme continuous phenotypes. To resolve this issue, Lee et al. (2012a, 2012b)[Lee, et al. 2012a; Lee, et al. 2012b] have proposed a method to adjust asymptotic null distribution by estimating small sample moments when the trait is dichotomous. We employ a similar approach (details in Supplementary Materials Section 4). For all simulation studies and real data analysis we used small sample adjustment for SKAT methods given the small to moderate sample sizes we considered. We used SKAT-O for continuous extreme phenotype SKAT (CEP-SKAT-O) and dichotomous extreme phenotype SKAT (DEP-SKAT-O), and for random sample continuous phenotypes (RS-SKAT-O). It should be noted that for larger sample sizes, the small sample adjustment is not necessary. Through simulations we found sample sizes lower than $n=500$ to benefit from the small sample adjustment, with sample sizes as low as $n=1000$ to not benefit from the adjustment.

Type 1 error simulations

We first generated haplotype data by the forward simulator, SFS_CODE[Hernandez 2008], which offers the ability to incorporate purifying selection on deleterious variants and thus provides better model to simulate variants in exomes. Data were simulated according to the European demographic model with a population bottleneck followed by exponential growth. We simulated 32,000 haplotypes each 100,000 base pairs wide as our population base. To achieve a simulated sample over a 3kb exon, a random 3kb region is selected (containing 41 variants on average) and each individual genotype is formed by combining at random two haplotypes over that region. Phenotypes for the i -th individual in a sample were produced from the generated genotype and covariate data according to

$$Y_i = 0.5X_{i1} + 0.5X_{i2} + \varepsilon_i$$

Where the covariate X_{i1} is 1 with probability 0.5 and 0 otherwise, and the covariate X_{i2} and the residual ε_i are both instances of a standard normal random variable.

Using the simulated genotype and phenotype data for the N individuals, a random sample of size n is selected. For random sampling of continuous traits, SKAT-O with the default $w_j = \text{Beta}(1, 25)$ weight is used to test for an association between the continuous phenotype and genotype after controlling for both covariates, producing a p-value (RS-SKAT-O). In order to test for the association between variants and phenotype under EPS using the standard dichotomizing method, we treat the highest ($n/2$) extremes as cases and lowest ($n/2$) extremes as controls. The dichotomized phenotypes are used by both DEP-SKAT-O and DEP-Burden. This same extreme phenotype sample is used to compare with the tests that retain the continuous phenotype (CEP-SKAT-O and CEP-Burden). A p-value for the CEP-SKAT-O test and the CEP-Burden burden test are produced from these continuous phenotype values and the corresponding genotype and covariate data. The proportion of p-values below a specified α -level provides an estimate for the Type 1 error at that α significance level.

Power Simulations

Power comparisons between the various sampling methods were performed using simulated genotype data as was used in the Type 1 error simulation setting. After generating the genotypes for N individuals, 20% of the variants with $MAF < 0.03$ are selected to be causal

variants. Different percentages of causal variants were also considered. Phenotypes are then generated for the N individuals according to:

$$y = 0.5X_1 + 0.5X_2 + \beta_1 G_1 + \beta_2 G_2 + \dots + \beta_p G_p + \varepsilon$$

The covariate X_1 is generated as a Bernoulli random variable with $p=0.5$. The covariate X_2 and the added noise ε are generated independently from a standard normal distribution. Non-causal variants are assigned $\beta_j=0$, and the causal variants are generated according to:

$$|\beta_j| = -a \log_{10}(\text{MAF}_j)$$

Here, $a > 0$ is a parameter that specifies the strength of variant-phenotype associations, hence the strength of heritability. Large values of a lead to stronger effects of causal variants on phenotype and cause rare variants to become more enriched in the phenotypic extremes. In one simulation setting an increase in a from 0.3 to 0.4 increases the heritability of the phenotype from 0.034 to 0.042. The heritability also increases with the number of causal variants. To obtain an estimate of the heritability, the proportion of the variance in phenotype explained by the genotypes of causal variants is estimated assuming no LD between variants.

Power estimates are obtained for various (extreme phenotype) sample sizes ($n=500$, 1000, and 2000), percentages in each phenotypic extreme sampled (10% and 25%), percentages of causal variants with a positive effect (80%, 100%), and percentages of causal variants with $\text{MAF} < 0.03$ (20%, 40%, and 60%).

Results

Extreme sampling enriches rare causal variants

Our analytical calculations (See Material and Methods and Supplementary Materials Section 1) confirm that rare causal variants can be enriched in phenotypic extremes. The degree of enrichment increases when more extreme phenotypes are sampled and a higher percentage of causal variants are present in a region. To empirically validate this finding, randomly selected 3kb exonic regions were simulated using the population genetic simulation model with European demographic history (see Material and Methods). For each 3kb region, causal variants were randomly selected to be 100%, 70%, 40% and 0% of sufficiently rare variants ($\text{MAF} < 0.03$) and the j -th causal variant was given the effect size β_j as a function of its MAF. Note that these causal variant percentages differ from those in the power simulations so as to further accentuate the effect of causal variant percentage on the inflation of MAF due to EPS. Also for the power simulations, causal variant percentages of 10% and 20% were used instead. Phenotypes are then generated from a linear model with heritability of genetic variants being 2.6%, 1.3%, and 0%.

Because the causal variants are known in the simulation setting, the expected MAF of a causal variant using EPS can be computed analytically (see Supplementary Materials Section 1). The expected MAFs of causal variants using EPS matched closely with the sample MAFs of causal variants using EPS (Figure 1). The MAFs of simulated causal variants after EPS had an overall increased frequency over the respective population MAFs. This trend decreases as samples are restricted to less extreme phenotypes and heritability is lower. No enrichment is found when there is no causal variant. When both the causal and non-causal variants in a region are considered simultaneously, the median MAF using EPS is much less inflated than when only causal variants are examined.

Sampling methods for comparison

Motivated by the enrichment of causal rare variants in phenotypic extremes, we expect to find that EPS methods can increase power to detect rare causal variants over random sampling methods. We extend the SKAT family methods to test for region-level rare variant effects when continuous phenotypes obtained from EPS are used in analysis. In simulation and data analysis, we only use the extreme phenotype optimal SKAT (SKAT-O), which accounts for extreme phenotype sampling and unifies the burden test and SKAT to a single framework and by constructing the optimal test within the framework. Using the simulated genotype data over 3kb regions the phenotypes were generated using the additive linear model (see Material and Methods). Given the same sample size, we compare the power of three tests designed for detecting rare variant effects using EPS. We first consider a burden test, DEP-Burden, that uses dichotomized extreme phenotypes along with collapsed information over genotypes by simply counting the number of rare variants with $MAF < 3\%$ in the gene before applying logistic regression to the collapsed statistic. We also apply this same collapsed statistic to continuous extreme phenotypes as done in Li et al. [Li, et al. 2011] and call this test CEP-Burden. Next we consider dichotomized extreme phenotype SKAT-O (DEP-SKAT-O), which applies optimal SKAT (SKAT-O) to dichotomized extreme phenotypes while applying small sample adjustments when sample sizes are small [Lee, et al. 2012a]. Finally we consider continuous extreme phenotype SKAT (CEP-SKAT-O), which does not dichotomize and instead extends linear regression optimal SKAT over the continuous extreme phenotypes (see Material and Methods) by using a truncated normal distribution. We also applied the small sample adjustment to CEP-SKAT-O to obtain the correct type I error rates when sample sizes are small. To demonstrate the benefits of EPS compared to random sampling, we included in the comparison a fourth method using random sampling SKAT-O (RS-SKAT-O), which applies optimal SKAT to the continuous phenotypes of a random sample. We assume the same sample size when comparing different methods so their powers are comparable. The power of each competing method is estimated as the proportion of p-values less than $\alpha = 10^{-6}$ in an effort to imitate genome-wide association studies.

The type 1 error rates for CEP-SKAT-O were accurate at $\alpha = 0.01$ and $\alpha = 0.05$ and slightly inflated at genome-wide significance levels $\alpha = 10^{-6}$ (see Supplementary Table 1). When all causal variants had the same direction of effect, CEP-SKAT-O and CEP-Burden had the greatest power with a substantial lead over every other method (Figure 2). When causal variants had effects in opposite directions all tests lost power uniformly due to less enrichment of rare variants, but CEP-SKAT-O became the most powerful by a large margin (Figure 3). In this case DEP-Burden had the least power. The power of the three methods employing SKAT-O (CEP-SKAT-O, DEP-SKAT-O, and RS-SKAT-O) is much more robust to changes in the proportion of causal variants that have a positive effect than the burden test's power is. This is because SKAT-O allows for each individual variant to affect phenotype in different directions and also allows for no effect. On the other hand, burden tests assume all the causal variants share the same direction of effect and that all the variants in a region are causal, and so the power of the burden tests greatly diminishes when causal variants are allowed effects in opposite directions or many causal variants are allowed no effect.

When all causal variants having the same direction of effect and as the percent of rare variants that were causal increased, the power gap between DEP-Burden and DEP-SKAT-O reduces, an observation that is also true for the relationship between CEP-Burden and CEP-SKAT-O (Figure 2). This is because CEP-SKAT-O includes CEP-Burden as a special case and behaves like CEP-Burden automatically when most variants are causal with effects in the same direction. To see this, we found that in simulations the estimated ρ decreased by a

factor of 0.36 on average when changing from the case of having all positive causal variant effects to the case of having causal variant effects in opposite directions.

Methods utilizing extreme sampling benefit as the cutoff for extreme phenotypes increases. In particular, as the percent of the tails sampled from the distribution of phenotype decreases from 25% to 10%, all EPS tests show incremental increases in power given the same sample size due to higher enrichment of rare variants. Relative power comparisons remain unchanged after decreasing the heritability of the phenotype and after increasing the exon length to 5kb or 10kb (regions of these lengths contain 69 variants and 138 variants on average, respectively). Also, simulations were also performed where the β was selected to be a constant rather than being a decreasing function of the MAF but the relative power of the methods remained the same (Supplementary Figure 2). Regardless of the setting, CEP-SKAT-O was consistently robust and had the greatest overall power to detect gene-phenotype associations over the other methods.

Application to the Dallas Heart Study data

In the Dallas Heart Study [Victor, et al. 2004], 3476 individuals were sequenced over the genes *ANGPTL3* (MIM 604774), *ANGPTL4* (MIM 605910), and *ANGPTL5* (MIM 607666). A total of 93 variants are present over these genes, and the variants in all three genes were tested simultaneously for an association with log-transformed serum triglyceride levels (logTG). Analysis for each of three genes separately is also considered (see Supplementary Materials Section 6). Ethnicity and sex were adjusted for in the analysis. To demonstrate rare variant association test methods for extreme phenotype sampling (EPS), a total of 1,389 individuals with the highest 20% and lowest 20% of logTG levels in each age-gender stratum were selected as the ESP sample. The continuous values were used in CEP-SKAT-O while dichotomized values were used for DEP-SKAT-O and DEP-Burden. Random samples of equivalent size were selected for the RS-SKAT-O method for comparison purposes.

To compare the effects of the different cutoffs of tails, we considered to sample individuals from wider tails (30% and 40%). Since wider tails had more samples, to make p-values comparable, we randomly sub-sampled 1,389 individuals among individuals in wider tails in order to have the same sample size as compared to a 20% cut off. In these cases, median p-values calculated from multiple random samples were obtained (Table 1). The p-values for all EPS methods are sensitive to the extreme phenotype cutoff. CEP-SKAT-O outperforms the other methods when there is sufficient information about the continuous trait distribution. It performs similarly to DEP-SKAT-O where there is limited information in the data, e.g., when the cutoff is low or when there are a small number of rare variants in a gene (Supplementary Materials Section 6). When extremes were sampled from wider tails (30% and 40%) all of the tests tended to lose significance, demonstrating the strength of EPS. We also computed p-values with different cutoffs and unequal sample sizes (Supplementary Figure 1), and CEP-SKAT-O outperformed other competing methods overall.

Power Estimation

In the planning of new sequencing studies, it is important to be able to estimate the power to detect causal variants under various study designs. We provide such power and sample size calculations for extreme phenotype sampling designs using CEP-SKAT-O. We use analytical formulas to obtain the distribution of our statistic by allowing users to specifying desirable parameters of interest (see Supplementary Materials Section 5). The parameters that can be specified by the user include sample size, the percent of causal variants, the length of the genomic region, the effect size of the causal variants, the proportion of causal variants that have a positive effect, and the proportion of the tails that are sampled in EPS.

We find that power is increased as sample size increases, as the proportion of causal variants increases, as the effect size increases, when causal variants have their effect in the same direction, and when we are more selective by sampling individuals with more extreme phenotypes. The power is also dependent on the genomic region as the distribution of the number of genetic variants, the MAF distribution, and the LD structure vary over the genome, so for genome-wide studies power estimations are averaged over many randomly selected regions of equivalent size.

To evaluate the accuracy of these analytic power estimations, we show a side by side comparison with empirical power simulations (Figure 4 and Supplementary Figure 3). In this setting we consider 3kb regions with 20% of variants being causal with all effects in the same direction. We see that the estimated power with our analytic calculations matches the empirical power over a wide range of sample sizes.

Discussion

We confirm in this paper through analytical calculations and simulation studies that sampling phenotypic extremes of a population can enrich rare causal variants. As a consequence, we show that sampling from phenotypic extremes profit over analogous random sampling methods by showing a sizable gain in power when the same size is used. In particular, analysis using dichotomized extreme phenotype (DEP-SKAT-O) is shown to be more powerful than that using random sampling with continuous phenotypes (RS-SKAT-O) in almost all scenarios. We develop a new method, continuous extreme phenotype optimal SKAT (CEP-SKAT-O), which improves upon DEP-SKAT-O by retaining continuous phenotype information rather than dichotomizing, includes the continuous extreme phenotype burden test as a special case, and results in a significant increase in sensitivity to causal variants. We find that CEP-SKAT-O has the overall greatest power in a wide variety of settings over DEP-SKAT-O, RS-SKAT-O, and comparable collapsing methods.

In the realm of region based association testing methodology, there already exist many methods capable of handling continuous phenotypes when a normal distribution is assumed. However in the case of extreme sampling, phenotype follows a truncated normal distribution instead, and so current methods cannot be directly applied without dichotomizing. The advantage of CEP-SKAT-O is that it adapts SKAT-O, a method that applies multiple linear regression of a phenotype on all genotypes in a region, to be able to handle phenotypes coming from a truncated normal. This adaptation allows for the usual continuous phenotype analysis without forcing the usual loss of phenotype information that occurs due to dichotomizing.

CEP-SKAT-O assumes that subjects are sampled from the extremes of a normally distributed phenotype, which in some circumstances may be inappropriate, and hence the test results can be biased when the normality of the underlying trait is violated. Dichotomizing phenotypes using DEP-SKAT-O is robust to departure from normality, although it is subject to some power loss when normality assumption is true. For candidate gene studies, permutation can be used to estimate the null distribution of the CEP-SKAT-O test statistic when the underlying trait does not follow a normal distribution. However, this is computationally difficult for GWAS where genome-wide significance levels are very stringent and a large number of genes are tested. We have found the maximum likelihood estimator of σ^2 seems to be sensitive to the distribution of the underlying trait. It is of future research interest to develop an alternative estimator of σ^2 that is more robust to deviations from normality.

In several ongoing exome sequencing studies conducted in the NHLBI Exome Sequencing Project <https://esp.gs.washington.edu/drupal/>, subjects were sampled using extremes of multiple phenotypes, future research is needed to develop methods for analyzing this more complex sampling setting. When subjects with extreme phenotypes are sampled for sequencing, covariate confounding effects need to be accounted for at the design phase to ensure representative samples. One strategy is to use stratified sampling, i.e. sample extreme phenotypes within key covariate stratum. For example, a phenotype distribution is likely to be gender-specific. It is desirable to sample extreme phenotypes within each gender stratum. The residual covariate confounding effects can be adjusted for at the analysis stage.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by grants T32 GM074897 (IB) and R37 CA076404 and P01 CA134294, P54 ES016454 (SL and XL).

References

- Bansal V, Libiger O, Torkamani A, Schork NJ. Statistical analysis strategies for association studies involving rare variants. *Nature Reviews Genetics*. 2010; 11(11):773–785.
- Basu S, Pan W. Comparison of statistical tests for disease association with rare variants. *Genetic Epidemiology*. 2011
- Biesecker LG, Shianna KV, Mullikin JC. Exome sequencing: the expert view. *Genome Biol*. 2011; 12(9):128. [PubMed: 21920051]
- Chen Z, Zheng G, Ghosh K, Li Z. Linkage disequilibrium mapping of quantitative-trait Loci by selective genotyping. *Am J Hum Genet*. 2005; 77(4):661–669. [PubMed: 16175512]
- Cirulli ET, Goldstein DB. Uncovering the roles of rare variants in common disease through wholegenome sequencing. *Nat Rev Genet*. 2010; 11(6):415–425. [PubMed: 20479773]
- Clement K, Vaisse C, Manning BS, Basdevant A, Guy-Grand B, Ruiz J, Silver KD, Shuldiner AR, Froguel P, Strosberg AD. Genetic variation in the beta 3-adrenergic receptor and an increased capacity to gain weight in patients with morbid obesity. *N Engl J Med*. 1995; 333(6):352–354. [PubMed: 7609752]
- Davies RB. The Distribution of a Linear Combination of Chi-square Random Variables. *Journal of the Royal Statistical Society*. 1980; 29(3):323–333.
- Gu C, Todorov AA, Rao DC. Genome screening using extremely discordant and extremely concordant sib pairs. *Genet Epidemiol*. 1997; 14(6):791–796. [PubMed: 9433579]
- Hernandez RD. A flexible forward simulator for populations subject to selection and demography. *Bioinformatics*. 2008; 24(23):2786–2787. [PubMed: 18842601]
- Hu S, Zhong Y, Hao Y, Luo M, Zhou Y, Guo H, Liao W, Wan D, Wei H, Gao Y, et al. Novel rare alleles of ABCA1 are exclusively associated with extreme high-density lipoprotein-cholesterol levels among the Han Chinese. *Clin Chem Lab Med*. 2009; 47(10):1239–1245. [PubMed: 19743957]
- Huang BE, Lin DY. Efficient association mapping of quantitative trait loci with selective genotyping. *Am J Hum Genet*. 2007; 80(3):567–576. [PubMed: 17273979]
- Ioannidis JP, Thomas G, Daly MJ. Validating, augmenting and refining genome-wide association signals. *Nat Rev Genet*. 2009; 10(5):318–329. [PubMed: 19373277]
- Ji W, Foo JN, O'Roak BJ, Zhao H, Larson MG, Simon DB, Newton-Cheh C, State MW, Levy D, Lifton RP. Rare independent mutations in renal salt handling genes contribute to blood pressure variation. *Nat Genet*. 2008; 40(5):592–599. [PubMed: 18391953]

- Khor CC, Goh DL. Strategies for identifying the genetic basis of dyslipidemia: genome-wide association studies vs. the resequencing of extremes. *Curr Opin Lipidol.* 2010; 21(2):123–127. [PubMed: 20125008]
- Lee S, Emonds M, Bamshad M, Barnes K, Rieder M, Nickerson D, Christiani D, Wurfel M, Lin X. Optimal unified approach for rare variant association testing with application to small sample case-control whole-exome sequencing studies. *American Journal of Human Genetics.* 2012a in press.
- Lee S, Wu MC, Lin X. Optimal tests for rare variant effects in sequencing association studies. *Biostatistics.* 2012b
- Li B, Leal SM. Methods for detecting associations with rare variants for common diseases: application to analysis of sequence data. *Am J Hum Genet.* 2008; 83(3):311–321. [PubMed: 18691683]
- Li D, Lewinger JP, Gauderman WJ, Murcray CE, Conti D. Using extreme phenotype sampling to identify the rare causal variants of quantitative traits in association studies. *Genet Epidemiol.* 2011
- Liang KY, Huang CY, Beaty TH. A unified sampling approach for multipoint analysis of qualitative and quantitative traits in sib pairs. *Am J Hum Genet.* 2000; 66(5):1631–1641. [PubMed: 10762548]
- Madsen BE, Browning SR. A groupwise association test for rare mutations using a weighted sum statistic. *PLoS Genet.* 2009; 5(2):e1000384. [PubMed: 19214210]
- Maier B. Personal genomes: The case of the missing heritability. *Nature.* 2008; 456(7218):18–21. [PubMed: 18987709]
- Manolio TA, Collins FS, Cox NJ, Goldstein DB, Hindorf LA, Hunter DJ, McCarthy MI, Ramos EM, Cardon LR, Chakravarti A, et al. Finding the missing heritability of complex diseases. *Nature.* 2009; 461(7265):747–753. [PubMed: 19812666]
- Morgenthaler S, Thilly WG. A strategy to discover genes that carry multi-allelic or mono-allelic risk for common diseases: a cohort allelic sums test (CAST). *Mutat Res.* 2007; 615(1-2):28–56. [PubMed: 17101154]
- Morris AP, Zeggini E. An evaluation of statistical approaches to rare variant analysis in genetic association studies. *Genetic Epidemiology.* 2010; 34(2):188–193. [PubMed: 19810025]
- Neale BM, Rivas MA, Voight BF, Altshuler D, Devlin B, Orho-Melandar M, Kathiresan S, Purcell SM, Roeder K, Daly MJ. Testing for an unusual distribution of rare variants. *PLoS Genet.* 2011; 7(3):e1001322. [PubMed: 21408211]
- Nejentsev S, Walker N, Riches D, Egholm M, Todd JA. Rare variants of IFIH1, a gene implicated in antiviral responses, protect against type 1 diabetes. *Science.* 2009; 324(5925):387–389. [PubMed: 19264985]
- Ng PC, Levy S, Huang J, Stockwell TB, Walenz BP, Li K, Axelrod N, Busam DA, Strausberg RL, Venter JC. Genetic variation in an individual human exome. *PLoS Genet.* 2008; 4(8):e1000160. [PubMed: 18704161]
- Pan W. Asymptotic tests of association with multiple SNPs in linkage disequilibrium. *Genet Epidemiol.* 2009; 33(6):497–507. [PubMed: 19170135]
- Price AL, Kryukov GV, de Bakker PIW, Purcell SM, Staples J, Wei LJ, Sunyaev SR. Pooled association tests for rare variants in exon-resequencing studies. *The American Journal of Human Genetics.* 2010; 86(6):832–838.
- Price RA, Li WD, Zhao H. FTO gene SNPs associated with extreme obesity in cases, controls and extremely discordant sister pairs. *BMC Med Genet.* 2008; 9:4. [PubMed: 18218107]
- Ramser J, Ahearn ME, Lenski C, Yariz KO, Hellebrand H, von Rhein M, Clark RD, Schmutzler RK, Lichtner P, Hoffman EP, et al. Rare missense and synonymous variants in UBE1 are associated with X-linked infantile spinal muscular atrophy. *Am J Hum Genet.* 2008; 82(1):188–193. [PubMed: 18179898]
- Risch N, Zhang H. Extreme discordant sib pairs for mapping quantitative trait loci in humans. *Science.* 1995; 268(5217):1584–1589. [PubMed: 7777857]
- Slatkin M. Disequilibrium mapping of a quantitative-trait locus in an expanding population. *Am J Hum Genet.* 1999; 64(6):1764–1772. [PubMed: 10330364]
- Tzeng JY, Zhang D. Haplotype-based association analysis via variance-components score test. *Am J Hum Genet.* 2007; 81(5):927–938. [PubMed: 17924336]

- Victor RG, Haley RW, Willett DL, Peshock RM, Vaeth PC, Leonard D, Basit M, Cooper RS, Iannacchione VG, Visscher WA, et al. The Dallas Heart Study: a population-based probability sample for the multidisciplinary study of ethnic differences in cardiovascular health. *Am J Cardiol.* 2004; 93(12):1473–1480. [PubMed: 15194016]
- Wallace C, Chapman JM, Clayton DG. Improved power offered by a score test for linkage disequilibrium mapping of quantitative-trait loci by selective genotyping. *Am J Hum Genet.* 2006; 78(3):498–504. [PubMed: 16465623]
- Wu MC, Lee S, Cai T, Li Y, Boehnke M, Lin X. Rare-variant association testing for sequencing data with the sequence kernel association test. *Am J Hum Genet.* 2011; 89(1):82–93. [PubMed: 21737059]

Figure 1. Enrichment of causal rare variants in phenotypic extremes

Estimated folds increase of the observed MAFs of causal variants in phenotypic extremes over population MAFs. The red lines represent the smoothed observed fold increases. The dotted lines represent the theoretical fold increase. For each causal variant, population MAF was computed using the full simulated population while extreme phenotype MAF was computed after sampling the tails. See Supplemental Materials for derivation of theoretical expected MAF for extreme phenotypes. The top two figures consider the case where all variants are causal by sampling $k=10\%$ and 20% high/low extremes. For each case, three situations were considered by heritability of causal variants: $H^2=2.6\%$, 1.3% , and 0% (no causal variant). Higher heritability gives more enrichment of rare variants. The bottom two

figures consider the case where different fractions of variants in a region are causal (100%, 70%, 40% and 0%) by sampling $k=10\%$ and 20% high/low extremes. Presence of non-causal variants in a region lower the degree of enrichment of rare variants.

Figure 2. Power comparisons when all causal variants have the same effect direction

Simulated power comparisons between four rare variants association tests with all causal variants having a positive effect on phenotype. The five tests are random sample optimal SKAT (RS-SKAT-O), dichotomized extreme phenotype burden test (DEP-Burden), continuous extreme phenotype burden test (CEP-Burden), dichotomized extreme phenotype optimal SKAT (DEP-SKAT-O), and continuous extreme phenotype optimal SKAT (CEP-SKAT-O). The left panel considers the situation where 10% high/low extremes are sampled with the three rows corresponding to 20% (0.6% heritability), 40% (1.2% heritability) and 60% (1.8% heritability) variants in a 3kb region being causal. Three total sample sizes are considered: $n=500$, 1000, 2000. The right panel considers the situation where 25% high/low

extremes are sampled. Exonic regions are simulated with effect sizes for each causal variant equal to $\beta = -0.2 \log_{10} \text{MAF}$. Power is estimated by the proportion of tests that detect an association at the $\alpha = 10^{-6}$ level.

Figure 3. Power comparisons when causal variants have opposite effect directions

Simulated power comparisons between four rare variants association tests with 80% of rare causal variants selected to have a positive effect on phenotype while the remaining 20% have a negative effect. The five tests are random sample SKAT (RS-SKAT-O), dichotomized extreme phenotype burden test (DEP-Burden), continuous extreme phenotype burden test (CEP-Burden), dichotomized extreme phenotype optimal SKAT (DEP-SKAT-O), and continuous extreme phenotype optimal SKAT (CEP-SKAT-O). The left panel considers the situation where 10% high/low extremes are sampled with the three rows corresponding to 20% (0.6% heritability), 40% (1.2% heritability) and 60% (1.8% heritability) variants in a 3kb region being causal. Three total sample sizes are considered:

$n=500, 1000, 2000$. The right panel considers the situation where 25% high/low extremes are sampled. Exonic regions are simulated with effect sizes for each causal variant equal to $|\beta| = -0.2 \log_{10} \text{MAF}$ with the effect being negated 20% of the time. Power is estimated by the proportion of tests that detect an association at the $\alpha=10^{-6}$ level.

Power Estimation

Figure 4. Comparison of theoretical and empirical powers

Estimated power of CEP-SKAT for testing 3kb regions with 20% of variants being causal with all effects in the same direction and the casual variants have effects to $|\beta| = -0.2\log_{10}MAF$. Theoretical power was calculated as described in section 5 of the Supplementary material, and empirical power was estimated by simulation using 300 replicates. No covariates were considered in either the theoretical or empirical power calculations. Furthermore empirical power was computed using CEP-SKAT without small sample adjustments.

Table 1

Analysis results of the Dallas Heart Study triglyceride data

n=1389	CEP-SKAT-O	DEP-SKAT-O	RS-SKAT-O	CEP-Burden	DEP-Burden
20%	5.0×10 ⁻⁵	3.0×10 ⁻⁵	1.3×10 ⁻²	1.2×10 ⁻⁴	7.2×10 ⁻⁵
30%	1.0×10 ⁻³	1.9×10 ⁻³	1.3×10 ⁻²	2.2×10 ⁻³	2.8×10 ⁻³
40%	8.9×10 ⁻³	1.2×10 ⁻²	1.3×10 ⁻²	3.4×10 ⁻³	1.9×10 ⁻²

Analysis results of the Dallas Heart Study sequence data using various test methods and sampling schemes. A total of 3,476 subjects were sequenced. A total 1,389 individuals were selected with highest and lowest 20% logTG levels in each age-gender spectrum from the total 3,476 sequenced subjects. For sampling with higher cutoffs (30% and 40%), 1389 individuals were randomly sub-sampled among the individuals belongs to larger tails to make powers comparable. In these cases, median p-values are presented in the table from 1000 sampling iterations. Since the sample size was large, the small sample adjustment was not applied. The five tests are continuous extreme phenotype SKAT (CEP-SKAT-O), dichotomous extreme phenotype SKAT (DEP-SKAT-O), continuous extreme phenotype burden test (CEP-Burden), dichotomous extreme phenotype burden test (DEP-Burden), and random sample SKAT (RSSKAT-O).