Two-Stage Extreme Phenotype Sequencing Design for Discovering and Testing Common and Rare Genetic Variants: Efficiency and Power

Guolian Kang^{a, c} Dongyu Lin^a Hakon Hakonarson^b Jinbo Chen^a

^aDepartment of Biostatistics and Epidemiology, University of Pennsylvania, and ^bCenter for Applied Genomics, The Joseph Stokes Jr. Research Institute, The Children's Hospital of Philadelphia, Philadelphia, Pa., and

^cDepartment of Biostatistics, St. Jude Children's Research Hospital, Memphis, Tenn., USA

Key Words

Two-stage design • Next-generation sequencing • SNP discovery • Rare variants

Abstract

Next-generation sequencing technology provides an unprecedented opportunity to identify rare susceptibility variants. It is not yet financially feasible to perform whole-genome sequencing on a large number of subjects, and a two-stage design has been advocated to be a practical option. In stage I, variants are discovered by sequencing the whole genomes of a small number of carefully selected individuals. In stage II, the discovered variants of a large number of individuals are genotyped to assess associations. Individuals with extreme phenotypes are typically selected in stage I. Using simulated data for unrelated individuals, we explore two important aspects of this two-stage design: the efficiency of discovering common and rare single-nucleotide polymorphisms (SNPs) in stage I and the impact of incomplete SNP discovery in stage I on the power of testing associations in stage II. We applied a sum test and a sum of squared score test for gene-based association analyses evaluating the power of the two-stage design. We obtained the following results from extensive simulation studies and analysis of the GAW17 dataset. When individuals with trait values more extreme than the 99.7–99th quantile were included in stage I,

the two-stage design could achieve the same power as or even higher than the one-stage design if the rare causal variants had large effect sizes. In such design, fewer than half of the total SNPs including more than half of the causal SNPs were discovered, which included nearly all SNPs with minor allele frequencies (MAFs) $\geq 5\%$, more than half of the SNPs with MAFs between 1% and 5%, and fewer than half of the SNPs with MAFs $< 1\%$. Although a one-stage design may be preferable to identify multiple rare variants having small to moderate effect sizes, our observations support using the two-stage design as a cost-effective option for next-generation sequencing studies.

Copyright © 2012 S. Karger AG, Basel

Introduction

Common single-nucleotide polymorphisms (SNPs) that are associated with complex traits explain only a small proportion of the genetic component of the trait variation. Recently, evidence has been emerging that less common variants (LCVs; minor allele frequency, MAF, 1–5%) and rare variants (RVs; MAF $< 1\%$ but still polymorphic) contribute to the risk of complex diseases such as autism, epilepsy, and schizophrenia [1]. State-of-the-art, next-generation sequencing technology provides an unprecedented opportunity for identifying LCVs and

RVs having modest to large effect sizes. Currently, it is still expensive to perform whole-genome sequencing in large samples such as those used in genome-wide association studies. Toward this end, two types of two-stage designs have been proposed as practical options to balance cost and power [2–5]. Both involve a stage of variant discovery via whole-genome sequencing (stage I) and a stage of association testing in much larger, independent samples for which the stage I variants are genotyped (stage II). One design involves sequencing a moderate number of pooled samples [3, 4], and the other involves sequencing a small number of individuals with extreme phenotype values [5]. We focus on the second study design in this paper and call it ‘two-stage extreme phenotype sequencing design’ (TS-E).

Guey et al. [6] found that using individuals with extreme phenotypes could lead to high efficiency for discovering RVs in stage I but limited power for detecting associations. They estimated power based on testing a single RV and did not compare the power of the TS-E or one-stage design in which the whole genomes of all individuals are sequenced. Here, through extensive simulation studies and analysis of Genetic Analysis Workshop 17 (GAW17) mini-exome data [7], we comprehensively examine the efficiency of SNP discovery in stage I and compare the overall power of the TS-E to identify susceptible gene regions to that of the one-stage design.

Our results support the feasibility of the TS-E. It may fail to discover a large proportion of rare variants, but such incomplete variant discovery may not translate into decreased power for identifying genetic regions that harbor causal variants if causal variants are common or if they are rare but have large effect sizes. This observation is largely because a smaller number of non-causal variants are discovered in stage I in the TS-E than in the one-stage design. In our simulation studies, when more than 10 subjects were selected in stage I for discovering variants, the two-stage design achieved nearly the same power as did the one-stage design. On the other hand, if one seeks to identify multiple RVs having small or moderate effect sizes, then a one-stage design may be preferable. Our results provide guidance for designing a two-stage study that balances the power and cost given the total number of stage I and II individuals.

Method

To implement the TS-E in our simulation studies, we first generated 20,000 SNPs in 200 gene regions for every individual, which are assumed to be the complete exome sequencing data and

are referred to as the mini-exome data. Among these 20,000 SNPs, some SNPs with MAFs <5% were assumed to be causal and had moderate to large effect sizes, and a few common SNPs (MAFs ≥ 5%) were assumed to be causal and had small to moderate effect sizes. We then generated values for a quantitative trait from a linear regression model. In addition, we analyzed 200 simulated mini-exome and quantitative trait datasets of 697 unrelated individuals from the GAW17.

Stage I: SNP Discovery by Sequencing Individuals with Extreme Phenotypes

Denote y_l and y_u as the lower l and upper $u = 1-l$ quantiles of trait Y for m independent individuals from a homogeneous population. Individuals satisfying $Y \geq y_u$ or $Y \leq y_l$ are selected for sequencing. We treat the SNPs for which the estimated MAFs in the selected individuals are greater than zero as the ‘discovered SNPs’. We deem a gene as a ‘discovered gene’ if at least 1 SNP within this gene is discovered. In keeping with Cirulli and Goldstein’s [5] categorization of potential causal variant frequencies, we classify the discovered SNPs (MAFs >0%) into three groups: common variants (CVs; MAFs 5–50%), LCVs (MAFs 1–5%), and RVs (MAFs <1%). To fully understand the merit of the TS-E, we also evaluated the efficiency of an alternative two-stage design in which stage I individuals are selected randomly: we refer to this design as TS-R.

Suppose that the discovered SNPs are genotyped for the remaining $(1-2l)m$ individuals in stage II. Let T_w and T_s denote the respective costs of sequencing and genotyping discovered SNPs for 1 individual. We assume that T_s is the same for genotyping different numbers of discovered SNPs when calculating the cost benefit. The cost fraction $2l + (1-2l)T_s/T_w$ represents the relative cost of the two-stage design compared with the one-stage design in which all m individuals are sequenced. For example, with $l = 0.005$ and $T_s/T_w = 0.5$, the cost fraction is 0.51. This cost fraction translates into a 49% reduction in the total cost.

Stage II: Gene-Based Association Testing Based on Discovered SNPs

We subsequently assessed the association between each gene and a particular phenotype by using the discovered SNPs for subjects in stage II, aiming to evaluate the impact of incomplete SNP discovery in stage I on the power of the stage II association testing. We applied two gene-based association testing methods: the simple modified sum test, which is one of the most powerful methods when effects of the causal variants have the same direction and there are no or few non-causal RVs, and the sum of squared score (SSU) test, which is one of the most powerful tests in the presence of opposite association directions and non-causal RVs [8]. Both the sum test and SSU test yielded a p value for the significance of each gene rather than for that of each SNP. Testing each gene rather than each SNP is in line with the current wisdom.

The Sum Test

Suppose that observations $(Y_i, \mathbf{X}_i)$ are available for N individuals, $i = 1, 2, \dots, N$, where Y_i is the value of a quantitative trait for subject i , and $\mathbf{X}_i = (X_{i1}, \dots, X_{iK})^T$ is the vector of minor allele counts (0, 1, 2) for the K SNPs within the tested gene region for the same subject. A tested gene region includes all SNPs within 20 kb of the start and stop of the gene [9]. Let K_1 , K_2 , and $K_3 = K - K_1 - K_2$ be the respective numbers of CVs, LCVs, and RVs within the tested gene region. By adapting the combined multivariate and collap-

ing (CMC) method [10] and the sum test [7] to accommodate both RVs and CVs, we use a linear regression model to describe the association between the trait Y and the K SNPs,

$$Y_i = \beta_0 + \beta_1 \sum_{j=1}^{K_1} \frac{X_{ij}^{cv}}{K_1} + \beta_2 \sum_{j=1}^{K_2} \frac{X_{ij}^{lcv}}{K_2} + \beta_3 \sum_{j=1}^{K_3} \frac{X_{ij}^{rv}}{K_3} + \varepsilon_i,$$

where ε_i is the random error term that follows the standard normal distribution and superscripts *cv*, *lcv*, and *rv* represent the CV, LCV, and RV within this gene, respectively. We divide the summation of minor allele counts by the number of SNPs to account for different numbers of SNPs within each category. We test the null hypothesis of no association between the gene and the trait by testing the null hypothesis $\beta_1 = \beta_2 = \beta_3 = 0$ with the standard likelihood ratio statistic. For association tests in the two-stage design, only the discovered SNPs within each gene are included in the model.

The SSU Test

In the presence of opposite association directions and non-causal RVs, the SSU test is one of the most powerful tests for detecting RVs [7]. It is reasonable that it would also be powerful for testing CVs and RVs jointly because it is robust to the number of parameters for multiple genetic variants. Under a linear regression model for the association between the trait Y and the K SNPs,

$$Y_i = \beta_0 + \sum_{j=1}^K X_{ij} \beta_j,$$

the SSU statistic T_{SSU} is equal to $U^T U$, where

$$U = \sum_{i=1}^N (Y_i - \bar{Y}) \mathbf{X}_i \text{ and } \bar{Y} = \sum_{i=1}^N \frac{Y_i}{N}.$$

Under the null hypothesis of no association, $\beta_j = 0$, $j = 1, 2, \dots, K$, T_{SSU} asymptotically follows a distribution of a mixture of χ^2_1 's that can be approximated by a scaled and shifted χ^2 distribution [8]. Here, we used the permutation re-sampling method to estimate the empirical p values of the SSU test.

Simulation Studies

We assessed the efficiency of discovering both common and rare SNPs, the type I error rate, and the power of the two-stage design, assuming that multiple CVs and RVs could be associated with a quantitative trait.

Generating Genotype Data Based on Linkage Disequilibrium

We generated SNP genotype data directly rather than generating raw sequencing reads [10–13]. To ensure that there were more RVs than CVs, we first generated MAFs for 20,000 SNPs from the Wright's distribution [14, 15], similar to Ionita-Laza et al. [13], Pritchard [14], and Madsen and Browning [15]. This distribution has the density function $f(p) = c p^{\beta_s - 1} (1 - p)^{\beta_n - 1} e^{\sigma(1 - p)}$, where β_s and β_n are the scaled mutation rates and were set at values 0.001 and 0.001/3, respectively; σ is the selection rate and was set at 12, and c is the normalizing constant. We then generated mini-exome genotype data for 2,000 unrelated individuals in three steps by controlling pair-wise linkage disequilibrium (LD) values [11, 12]. First, we generated 4,000 latent vectors, each of length 20,000 ($W_{j1}, W_{j2}, \dots, W_{j20,000}$), where j indicates the j -th vector. Each latent vector was generated from a multivariate normal distribution with a

first-order, auto-regressive covariance structure, where the correlation between the l -th and k -th component of the j -th vector is $\text{corr}(W_{jl}, W_{jk}) = \rho^{|l-k|}$. Different values of ρ correspond to different levels of LD, with $\rho = 0$ indicating linkage equilibrium (LE) and greater ρ indicating stronger LD between neighboring variants. LD between CVs and RVs and among RVs is generally weak [16, 17], but LD among CVs could be strong. We therefore considered two different LD structures among all CVs and LCVs: (1) the CVs and LCVs are in LE ($\rho = 0$), and (2) the CVs and LCVs are situated in regions of moderate LD ($\rho = 0.5$). All of the RVs were in LE. Next, each latent vector was dichotomized to construct a haplotype ($H_{j1}, H_{j2}, \dots, H_{j20,000}$), where H_{jk} represents the allele at the k -th locus that takes a value of 0 for the major one or a value of 1 for the minor one. Lastly, we combined 2 haplotypes at random to form an individual's genotype data ($X_{i1}, \dots, X_{i20,000}$), where X_{ik} is equal to the summation of H_{jk} 's in the 2 haplotypes. The final data for analysis contained 1,761 (8.8%) CVs, 4,167 (20.8%) LCVs, and 14,072 (70.4%) RVs. The distribution of the allele frequencies of the 20,000 SNPs was similar to that of the GAW17 data (online suppl. fig. S1; for all online suppl. material, see www.karger.com/doi/10.1159/000337300).

Generating Phenotype Data

We randomly selected k SNPs from the first gene as the causal SNPs. We then generated a quantitative trait Y for each individual from the linear model: $Y_i = \beta_0 + \sum_{j=1}^k X_{ij} \beta_j + e_i$, where e_i is the random error term following the standard normal distribution, and β is the vector of regression coefficients, with $\beta = 0$ indicating no association. We considered five models of association. The first four models had 7 causal variants ($k = 7$), including 2 CVs, 1 LCV, and 4 RVs, and they differed only in the magnitude and direction of β coefficients for the 7 variants. The fifth model had 20 causal variants, including 3 CVs, 3 LCVs, and 14 RVs [18]. The model parameters were chosen to ensure that the phenotype variability explained by LCVs and RVs was more than that explained by CVs. The MAFs, effect sizes, and percentage of explained phenotype variability are provided in online supplementary tables S1 and S2.

Efficiency of Stage I SNP Discovery, Type I Error Rate, and Power

To evaluate the merits of the TS-E using the gene-based association tests, we varied the proportion of individuals included in step I SNP discovery (l) from 0.001 to 0.02 at an increment of 0.0015. For stage I SNP discovery, we presented the average percentages of the total discovered SNPs, the discovered CVs, the discovered LCVs, the discovered RVs, and the discovered causal SNPs across 1,000 replicates. For comparison, we presented the stage I SNP discovery results of both the TS-E and TS-R under each of the five association models. For stage II gene-based association tests using discovered SNPs, we included only stage II subjects in the sum test to evaluate its power because including stage I subjects in association testing may inflate the type I error rate of the sum test [10, 19]. However, we included all samples in the SSU test because we used the permutation re-sampling method to obtain the empirical p value. We controlled the overall type I error rate at the nominal level (0.05) by applying the Bonferroni method. We also evaluated the power of the TS-R for models 1 and 2 when all SNPs were in LE. To assess the impact on the power of using only discovered SNPs for association testing, we also reported the corresponding power of the one-stage design.

Table 1. The type I error rates of the three designs at a nominal level of 0.05 when all SNPs are in LE based on 1,000 replicates

Test	Design	l^a													
		0.001	0.0025	0.004	0.0055	0.007	0.0085	0.01	0.0115	0.013	0.0145	0.016	0.0175	0.019	0.0205
Sum	OS	0.050													
	TS-E	0.058	0.054	0.056	0.052	0.056	0.051	0.056	0.051	0.053	0.044	0.045	0.045	0.041	0.043
	TS-R	0.056	0.051	0.053	0.045	0.045	0.051	0.035	0.041	0.048	0.053	0.036	0.055	0.041	0.041
SSU	OS	0.054													
	TS-E	0.054	0.059	0.055	0.054	0.054	0.056	0.052	0.051	0.054	0.055	0.055	0.052	0.053	0.055
	TS-R	0.059	0.056	0.055	0.053	0.056	0.054	0.052	0.052	0.055	0.051	0.055	0.053	0.052	0.053

^a l is the proportion of stage I individuals. OS = One-stage design.

Results

Stage I SNP Discovery

Figure 1 shows results for stage I SNP discovery and the corresponding cost reduction when all SNPs are in LE. The TS-E discovered nearly the same percentages of total SNPs, CVs, LCVs, and RVs as did the TS-R, regardless of the penetrance models (fig. 1a–e) and LD structure (online suppl. fig. S2A–E). However, as expected, the TS-E discovered a much higher percentage of causal SNPs than did the TS-R with the same proportion of stage I individuals (fig. 1f–j) in all scenarios (online suppl. fig. S2F–J). With $2l = 0.02$ of subjects included in stage I, the respective percentages of discovered total SNPs, CVs, LCVs, and RVs by both the TS-E and TS-R were 36%, 100%, 80%, and 15%, respectively (fig. 1a). However, the percentages of discovered causal SNPs were 77.8% by the TS-E and 47.6% by the TS-R (fig. 1b). These results are consistent with those in the literature [20] and are also consistent with the conventional wisdom that the TS-E is more cost-effective than the TS-R for designing next-generation sequencing projects.

Not surprisingly, the percentages of discovered total SNPs, CVs, LCVs, RVs, and causal SNPs increased with the proportion of stage I individuals, and the increases were sharp when the phase I proportion was small (fig. 1a–j; online suppl. fig. S2A–J). With l increasing from 0.001 to 0.007, the respective percentages of discovered total SNPs, CVs, LCVs, and causal SNPs increased from 10% to 31%, 54% to 99%, 17% to 69% (fig. 1a), and from 40 to 73%, respectively (fig. 1f), but the percentage of discovered rare SNPs hardly increased: it was as low as 25%, even when 2% were included in stage I (fig. 1a). The percentage of discovered causal SNPs varied greatly with the

association models. If the causal LCVs and RVs had large effects (fig. 1f–i), then the TS-E appeared to be much more efficient for discovering causal variants than the TS-R. With the TS-E, more than 60% of the causal variants can be discovered when $l = 0.004$ (i.e., 16 subjects were included in stage I) and 80% can be discovered when $l = 0.02$ (i.e., 80 subjects were included in stage I). When the genetic variability of the phenotype was mostly due to multiple LCVs and RVs having small or moderate effect sizes, less than 60% of the causal SNPs were discovered, even with $l = 0.02$, and the TS-E and TS-R appeared to have similar efficiency for discovering causal variants (fig. 1j). In general, the LD structure among common SNPs and the presence or absence of opposite association directions had a minor effect on SNP discovery (fig. 1a–d; online suppl. fig. S2A–D).

Type I Error Rate and Power

The type I error rates of the sum test and SSU test for analyzing the TS-E, TS-R, and one-stage designs appeared to be close to the nominal level (0.05), regardless of whether SNPs were in LE (table 1) or in LD (online suppl. table S3). We first compared the power of the sum test and SSU test for analyzing joint effects of both CVs and RVs under the five association models we considered. We then compared the power of the three designs (TS-E, TS-R, and one-stage design) by using the two tests under the five models. Figure 2a–e and online supplementary figure S3A–E show that power was at the nominal level of 0.05 when all SNPs were in LE (LD). Under the one-stage design, the SSU test had much higher power than the sum test if the total number of causal variants was small and the causal LCVs and RVs had large effect sizes (fig. 2a–d; online suppl. fig. S3A–D). However, the sum test had

Fig. 1. SNP discovery in stage I when all SNPs are in LE and the cost reduction of the two-stage designs based on simulation data. The first x-axis is the proportion of stage I individuals times 100 ($l \times 100$), and the second x-axis is the cost function with $T_s/T_w = 0.5$. The solid, dashed, dotted, dot-dash and long-

dash lines indicate the percentages of the total discovered SNPs, CVs, LCVs, RVs, and causal SNPs, respectively. Letters 'e' and 'r' indicate two-stage designs with extreme phenotype sampling (TS-E) or random sampling (TS-R), respectively.

Fig. 2. The power of the TS-E under the five disease models when all SNPs are in LE based on 1,000 sets of simulated data. The x-axis is the proportion of stage I individuals times 100 ($l \times 100$). OS = One-stage design.

higher power when the causal variants consisted of a large number of LCVs and RVs with a common small effect size (fig. 2e; online suppl. fig. S3E).

Power Comparisons of the Three Designs under the Sum Test

With the sum test used in the stage II association analysis, the power of the TS-E was first greater and then smaller than that of the one-stage design and TS-R with the increased proportion of stage I individuals if the RVs had large effect sizes, regardless of the LD structure among common SNPs (fig. 2a–d; online suppl. fig. S3A–D). If all of the causal LCVs and RVs had small effect sizes of similar magnitude, then the power of the one-stage design was

always higher than that of the TS-E (fig. 2e; online suppl. fig. S3E). The power difference became larger with smaller and larger proportion of stage I individuals regardless of disease models (fig. 2a–e) and LD structure among CVs and LCVs (online suppl. fig. S3A–E). The advantage of the TS-E decreased as the proportion of individuals included in stage I became larger (greater than 0.085 in the simulation study) if the causal RVs had large effect sizes (fig. 2a–d; online suppl. fig. S3A–D). With larger stage I proportion, although most causal variants were discovered, the sample size for stage II association tests became smaller, and even more unassociated SNPs were discovered and included in the association tests in stage II.

Power Comparisons of the Three Designs under the SSU Test

With the SSU test used in the stage II association test, the power of the TS-E was nearly identical to that of the one-stage design when more than 10 subjects ($l = 0.00025$) were selected in stage I under the first four association models. When causal SNPs included CVs (fig. 2a, b), the power was higher than 90%. When causal SNPs consisted of LCVs and RVs of large effect sizes (fig. 2c, d), the power was approximately 60%. When the causal SNPs consisted of LCVs and RVs having small effect sizes (fig. 2e), with 10 subjects included in stage I, the power of the TS-E was 25%; with 40 subjects included in stage I ($l = 0.01$), the power of the TS-E approached that of the one-stage design. The power of the TS-E was generally higher than that of the TS-R (fig. 2a, b), but the difference was smaller when the stage I sample size was increased.

Under both the one-stage design and TS-E, it appeared that the SSU test had higher power than the sum test when the causal variants included LCVs and RVs having large effect sizes (fig. 2a–d). However, when the causal variants consisted of a relatively large number of LCVs and RVs having small effect sizes, the sum test had higher power than the SSU test for both the one-stage design and TS-E (fig. 2e; online suppl. fig. S3E). Interestingly, the power of the sum test under the TS-E could be higher than that of the SSU test under the one-stage design with appropriate proportions of stage I individuals (fig. 2c–e; online suppl. fig. S3C–E).

Application to the GAW17 Dataset

We applied the TS-E to the GAW17 dataset of unrelated individuals [7], which contained genotype data of 24,487 SNPs in 3,205 genomic regions of 697 unrelated individuals. These SNPs included 3,132 CVs, 3,224 LCVs, and 18,131 RVs. One quantitative trait, denoted as Q_2 , was not correlated with any of the three covariates provided in the dataset, and 200 replicate datasets for Q_2 were provided based on the genotypes of 72 causal SNPs in 13 gene regions. Among the 72 causal SNPs, 2 (2.8%) were CVs, 4 (5.6%) were LCVs, and 66 (91.6%) were RVs. The limited sample size of the GAW17 dataset caused the power for testing associations based on each individual phenotype dataset to be low. Furthermore, the GAW17 used a fixed genotype dataset to generate 200 replicate phenotype datasets, so that the variation in the phenotype data across replicates was due to random error. Therefore, we used the average of the phenotype values in the first 10 datasets, denoted as Q_{2-10} , as the phenotype for the evaluation of the TS-E. The percentages of the Q_{2-10} variance ex-

Fig. 3. SNP discovery of the TS-E in the GAW17 data. The first x-axis is the proportion of stage I individuals times 100 ($l \times 100$), and the second x-axis is the cost function with $T_s/T_w = 0.5$. The solid, dashed, dotted, dot-dash and long-dash lines correspond to the percentage of total discovered SNPs, CVs, LCVs, discovered RVs, and causal SNPs, respectively.

plained by causal CVs, LCVs, and RVs are provided in online supplementary table S4.

Figure 3 shows the results for stage I SNP discovery using the TS-E and the corresponding cost reduction. When subjects in the lower and upper 0.0035 quantiles were selected on the basis of Q_{2-10} , 17% of the total SNPs were discovered, which included 86% of the CVs, 26% of the LCVs, and 3.3% of the RVs. These results are consistent with those in the literature [20]. In addition, the discovered SNPs included 12.5% of the 72 causal SNPs and 31% (998) of the total number of genes. In particular, 6 of 13 causal genes (*BCHE*, *LPL*, *PDGFD*, *SIRT1*, *VNN1*, and *VNN3*) were discovered in the sense that at least 1 SNP was discovered in each of these gene regions. We observed that either some CVs existed in the gene or that the causal SNPs explained a large proportion of the phenotypic variance.

Figure 4 plots the $-\log_{10}$ (p values) of the sum and SSU tests under the TS-E for each of the 13 causal genes. Interestingly, the p values of the sum test and SSU test for the full dataset (the one-stage design) for the undiscovered causal genes *VWF*, *VLDLR*, *RARB*, *PLAT*, *GCKR*, and *INSIG1* were similar. For these genes, LCVs and RVs only accounted for a small proportion of the phenotypic variance (online suppl. table S4). For the discovered genes that had causal CVs or causal RVs with relatively large ef-

Fig. 4. The $-\log p$ values for the 13 causal genes under the TS-E in the GAW17 data ($l = 0.0035$). The solid and dashed lines indicate the cutoffs for the one-stage design and TS-E with $l = 0.0035$ using Bonferroni correction, respectively. OS = One-stage design.

fect sizes [7], such as *VNN3*, *LPL*, *PDGFD*, and *VNN1*, the p values of the sum test under both the one-stage and TS-E designs were also similar. For the remaining 2 discovered genes, *SIRT1* and *BCHE*, the sum test was not significant under either the one-stage or TS-E design. Both genes contained causal RVs having moderate or small effect sizes. The analysis with the SSU test yielded results similar to those of the sum test.

Discussion

Here, we have evaluated the merits of the TS-E in assessing both RVs and CVs for association. Our results indicate that the stage I sample size is a very important factor for the success of the TS-E because it determines the percentage of discovered causal SNPs. If one suspects multiple numbers of susceptibility LCVs or RVs with small or moderate effect sizes, then perhaps a one-stage design should be planned. Sequencing just a small proportion of subjects for SNP discovery may ultimately lead to low power in association studies. However, if one aims to detect RVs with large effect sizes, then the two-stage design would be a feasible option, as suggested in the literature, and can have the same power as the one-stage design when accompanied by appropriate methods of analysis. In our simulation studies, the TS-E that includes merely the discovered SNPs in the SSU test could achieve nearly the same power as the one-stage design. The analysis of the GAW17 dataset further confirmed the efficiency of the TS-E for identifying both RVs and CVs. In general, as long as RVs having large effect sizes account for a

larger proportion of the phenotypic variance than CVs, the TS-E should be recommended because it balances statistical power and cost efficiency. We expect that our observations will be useful for those who wish to use the TS-E for next-generation sequencing studies.

We only considered two methods for analysis in this work: the sum test and the SSU test. When all causal SNPs were positively associated with the trait, the other two popular methods of analysis, the CMC test [10] and the weighted sum test [15], have power similar to that of the sum test [8]. When the minor alleles of the causal SNPs may be negatively associated with the trait, the kernel machine regression [21], the C-alpha-P test [22], the sequence kernel association test method [18], and the data-adaptive sum test [23] have power similar to that of the SSU test. Therefore, we expect that the qualitative results we obtained would still hold if other methods had been used for analysis. In addition, we assumed that the stage I and stage II individuals were from a homozygous population. When the study population is less homogeneous, the stage I sample should balance the numbers from different subpopulations.

Acknowledgments

We thank 2 anonymous reviewers for their valuable suggestions for improving the manuscript, Dr. Mingyao Li for helpful discussions at the beginning of this project, and Cherise Guess for editing the manuscript. The Genetic Analysis Workshop is supported by NIH grant R01 GM031575.

References

- 1 Stankiewicz P, Lupski JR: Structural variation in the human genome and its role in disease. *Ann Rev Med* 2010;61:437–455.
- 2 Schaid DJ, Sinnwell JP: Two-stage case-control designs for rare genetic variants. *Hum Genet* 2010;127:659–668.
- 3 Bansal V, Tewhey R, LeProust EM, Schork NJ: Efficient and cost effective population resequencing by pooling and in-solution hybridization. *PLoS One* 2011;6:e18353.
- 4 Kim SY, Li Y, Guo Y, Li R, Holmkvist J, et al: Design of association studies with pooled or un-pooled next-generation sequencing data. *Genet Epidemiol* 2010;34:479–491.
- 5 Cirulli ET, Goldstein DB: Uncovering the roles of rare variants in common disease through whole-genome sequencing. *Nature Rev Genet* 2010;11:415–425.
- 6 Guey LT, Kravic J, Melander O, Burt NP, Laramie JM, Lyssenko V, Jonsson A, Lindholm E, Tuomi T, Isomaa B, Nilsson P, Almgren P, Kathiresan S, Groop L, Seymour AB, Altshuler D, Voight BF: Power in the phenotypic extremes: a simulation study of power in discovery and replication of rare variants. *Genet Epidemiol* 2011;32:236–246.
- 7 Almasy L, Dyer TD, Peralta JM, Kent JW Jr, Charlesworth JC, Curran JE, Blangero J: Genetic Analysis Workshop 17 mini-exome simulation. *BMC Proc* 2011;5(suppl 9):S2.
- 8 Basu S, Pan W: Comparison of statistical tests for disease association with rare variants. *Genet Epidemiol* 2011;35:606–619.
- 9 Purcell S, Neale B, Todd-Brown K, Thomas L, Ferreira MAR, Bender D, Maller J, Sklar P, de Bakker PIW, Daly MJ, Sham PC: PLINK: a tool set for whole-genome association and population-based linkage analysis. *Am J Hum Genet* 2007;81:559–575.
- 10 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;83:311–321.
- 11 Pan W: Asymptotic tests of association with multiple SNPs in linkage disequilibrium. *Genet Epidemiol* 2009;33:497–507.
- 12 Wang T, Elston RC: Improved power by use of a weighted score test for linkage disequilibrium mapping. *Am J Hum Genet* 2007;80:353–360.
- 13 Ionita-Laza I, Buxbaum JD, Laird NM, Lange C: A new testing strategy to identify rare variants with either risk or protective effect on disease. *PLoS Genet* 2011;7:e1001289.
- 14 Pritchard JK: Are rare variants responsible for susceptibility to common diseases? *Am J Hum Genet* 2001;69:124–137.
- 15 Madsen BE, Browning SR: A groupwise association test for rare mutations using a weighted sum statistic. *PLoS Genet* 2009;5:e1000384.
- 16 Dickson SP, Wang K, Krantz I, Hakonarson H, Goldstein DB: Rare variants create synthetic genome-wide associations. *PLoS Biol* 2010;8:e1000294.
- 17 Siu H, Zhu Y, Jin L, Xiong M: Implication of next-generation sequencing on association studies. *BMC Genomics* 2011;12:322.
- 18 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;89:82–93.
- 19 Li BS, Leal SM: Discovery of rare variants via sequencing: implications for the design of complex trait association studies. *PLoS Genet* 2009;5:e1000481.
- 20 Ionita-Laza I, Lange C, Laird NM: Estimating the number of unseen variants in the human genome. *Proc Natl Acad Sci USA* 2009;106:5008–5013.
- 21 Liu D, Ghosh D, Lin X: Estimation and testing for the effect of a genetic pathway on a disease outcome using logistic kernel machine regression via logistic mixed models. *BMC Bioinformatics* 2008;9:292.
- 22 Neale BM, Rivas MA, Voight BF, Altshuler D, Devlin B, Ogbo-Melander M, Katherisan S, Purcell SM, Roeder K, Daly MJ: Testing for an unusual distribution of rare variants. *PLoS Genet* 2011;7:e1001322.
- 23 Han F, Pan W: A data-adaptive sum test for disease association with multiple common or rare variants. *Hum Hered* 2010;70:42–54.