

Accurate estimation of SNP-heritability from biobank-scale data irrespective of genetic architecture

Kangcheng Hou^{1,2*}, Kathryn S. Burch^{3*†}, Arunabha Majumdar¹, Huwenbo Shi^{3,4}, Nicholas Mancuso¹, Yue Wu⁵, Sriram Sankararaman^{3,5,6,7}, and Bogdan Pasaniuc^{1,3,6,7,†}

¹Department of Pathology and Laboratory Medicine, David Geffen School of Medicine, University of California, Los Angeles, Los Angeles, CA, USA

²College of Computer Science and Technology, Zhejiang University, Hangzhou, Zhejiang, China

³Bioinformatics Interdepartmental Program, University of California, Los Angeles, Los Angeles, CA, USA

⁴Department of Epidemiology, Harvard T.H. Chan School of Public Health, Boston, MA, USA

⁵Department of Computer Science, University of California, Los Angeles, Los Angeles, CA, USA

⁶Department of Human Genetics, David Geffen School of Medicine, University of California, Los Angeles, Los Angeles, CA, USA

⁷Department of Computational Medicine, David Geffen School of Medicine, University of California, Los Angeles, Los Angeles, CA, USA

*These authors contributed equally to this work.

†Correspondence should be addressed to K.S.B. (kathrynburch@ucla.edu) or B.P. (pasaniuc@ucla.edu)

January 21, 2019

Abstract

The proportion of phenotypic variance attributable to the additive effects of a given set of genotyped SNPs (i.e. SNP-heritability) is a fundamental quantity in the study of complex traits. Recent works have shown that existing methods to estimate genome-wide SNP-heritability often yield biases when their assumptions are violated. While various approaches have been proposed to account for frequency- and LD-dependent genetic architectures, it remains unclear which estimates of SNP-heritability reported in the literature are reliable. Here we show that genome-wide SNP-heritability can be accurately estimated from biobank-scale data irrespective of the underlying genetic architecture of the trait, without specifying a heritability model or partitioning SNPs by minor allele frequency and/or LD. We use theoretical justifications coupled with extensive simulations starting from real genotypes from the UK Biobank (N=337K) to show that, unlike existing methods, our closed-form estimator for SNP-heritability is highly accurate across a wide range of architectures. We provide estimates of SNP-heritability for 22 complex traits and diseases in the UK Biobank and show that, consistent with our results in simulations, existing biobank-scale methods yield estimates up to 30% different from our theoretically-justified approach.

Introduction

SNP-heritability, the proportion of phenotypic variance attributable to the additive effects of a given set of SNPs, is a fundamental quantity in the study of complex traits [1]; it provides an upper bound on risk prediction from a linear model relating genotypes to phenotype [2] and, when defined as a function of all SNPs on a genotyping array, yields insights into the “missing heritability” of complex traits [3–5]. Traditionally, SNP-heritability is estimated by fitting variance components models with REML [3, 6–9]. With some notable exceptions [8], REML-based methods are typically not scalable to biobanks that assay hundreds of thousands of individuals (e.g., UK Biobank contains genotype measurements for more than half a million individuals [10]). SNP-heritability can also be estimated from summary-level GWAS data by assessing the deviation in marginal association statistics as a function of the LD score of each SNP [11–14], thus making SNP-heritability estimation scalable to hundreds of thousands or even millions of individuals. More recently, a randomized extension of Haseman-Elston (HE) regression [15] was shown to estimate a single genetic variance component from individual-level data as accurately as REML methods but in a fraction of the run-time [16].

To facilitate inference, all existing methods for genome-wide SNP-heritability inference make various assumptions on the underlying genetic architecture of the trait, which is typically parametrized by *polygenicity* (the number of variants with effect sizes larger than some small quantifiable constant δ) and *MAF/LD-dependence* (the coupling of effect sizes with minor allele frequency (MAF), local linkage disequilibrium (LD), or other functional genomic annotations such as regions of open chromatin) [17]. Since the true genetic architecture of any given trait is unknown, existing methods are susceptible to bias and often yield vastly different estimates of SNP-heritability for the same traits, even when applied to the same data [9, 14, 18]. Although multi-component methods that stratify SNPs by MAF and LD can ameliorate some of the robustness issues of single-component methods [7, 18, 19], fitting multiple variance components to biobank-scale data with REML is highly resource-intensive [8] and it is currently unclear whether stratifying by MAF/LD produces accurate estimates of total SNP-heritability for methods based on summary statistics. Alternate methods explicitly model MAF- and LD-dependent architectures when estimating SNP-heritability [6, 9, 14]; however, these approaches can produce drastically different estimates when their assumptions are violated [6, 9, 14, 18, 19]. In addition, genetic architecture is unlikely to be the same across traits or populations due to, for example, variable degrees of negative selection acting on different traits in different populations [17, 20–25]. Methods that jointly infer SNP-heritability and other parameters such as the strength of negative selection or polygenicity have been proposed [14, 23, 26] but are computationally intensive and/or sensitive to LD-dependent architectures. Thus, it remains unclear which estimates of genome-wide SNP-heritability computed from biobank-scale data (e.g., UK Biobank [10]) are reliable.

In this work, we investigate whether genome-wide SNP-heritability can be accurately estimated under a generalized random effects (GRE) model that makes minimal assumptions on the genetic architecture of complex traits. Under this model, every causal effect can have an arbitrary SNP-specific variance, and SNP-heritability is defined as the sum of the SNP-specific variances (Methods). To the best of our knowledge, all existing methods make additional assumptions on top of the GRE model (Table 1). For example, the infinitesimal model assumed by single-component GREML [3] (and several other methods [8, 16]) imposes an inverse relationship between MAF and effect size by assuming that every standardized effect size explains an equal portion of total SNP-heritability, whereas the single-component LDAK model assumes that each

SNP-specific variance is inversely proportional to both MAF and the LD neighborhood of the SNP [6, 9]. We derive a closed-form estimator for SNP-heritability under the GRE model as a function of GWAS marginal association statistics and in-sample LD and show that this estimator is consistent (i.e. the estimator approaches the true SNP-heritability as sample size increases) and unbiased (i.e. the expectation of the estimator is equal to the true SNP-heritability) when the number of individuals is larger than the number of SNPs. Most importantly, the accuracy of this estimator does not depend on the underlying genetic architecture of the trait. The GRE estimator has the same analytical form as previously proposed “fixed effect estimators” [27, 28] that aim to estimate the variance explained by a set of fixed (non-random) SNP effects at a particular genomic region – a different quantity from the estimand of interest here (i.e. SNP-heritability defined as the sum of per-SNP variances; see Methods).

Through theoretical derivations and extensive simulations across a wide range of MAF- and LD-dependent architectures starting from real genotypes from the UK Biobank [10] (337K individuals and 593K SNPs), we find that the GRE estimator provides nearly unbiased estimates of SNP-heritability across all architectures whereas existing methods are sensitive to model misspecification. For example, across 126 distinct architectures, the maximum bias we observe with the GRE estimator is 2% of the simulated SNP-heritability whereas state-of-the-art methods such as stratified LD score regression (S-LDSC) [12, 13] and SumHer [14] yield biases between -84% and 28% of the simulated SNP-heritability. For completeness, we also contrast the GRE estimator with several REML-based methods in simulations at lower sample sizes (due to the computational burden of most REML methods) and find that, consistent with recent reports [18], all REML-based methods are biased when their model assumptions are violated. Across a similar set of 126 architectures, the bias of the GRE estimator ranges from -5% to 6% of the simulated SNP-heritability whereas single-component REML methods [3, 6, 8, 9] are biased by anywhere between -44% and 18% of the simulated SNP-heritability. We confirm that multi-component REML methods that stratify SNPs by MAF and LD score (GREML-LDMS-I [18]) are more accurate than single-component REML methods if favorable SNP stratification criteria are used (i.e. if SNPs are stratified by the same MAF bins used to define the causal variant MAF spectrum). The performance of the GRE estimator, which does not stratify SNPs or assume a specific heritability model [6, 9, 14], is similar to that of GREML-LDMS-I with favorable stratification criteria, thereby confirming that SNP-heritability can be accurately estimated without knowledge of the underlying genetic architecture.

Finally, we use marginal association statistics and in-sample LD from $N = 290\text{K}$ unrelated British individuals genotyped at $M = 460\text{K}$ SNPs ($\text{MAF} > 1\%$) to provide estimates of SNP-heritability for 22 complex traits and diseases in the UK Biobank [10]. Consistent with our simulations, across the 18 traits with SNP-heritability estimates greater than 0.05, we find that estimates from S-LDSC (controlling for the baseline-LD model [13]) and SumHer differ from the GRE estimates by a median of -9% and 11%, respectively. For example, for height, estimates from S-LDSC (0.56) and SumHer (0.63) are approximately 7% lower and 5% higher, respectively, than our estimate of 0.60. Similarly, for hypertension, estimates from S-LDSC (0.14) and SumHer (0.18) are $\pm 12.5\%$ different from our estimate of 0.16. Taken together, our results demonstrate that SNP-heritability can be accurately estimated from biobank-scale data without prior knowledge of the genetic architecture the trait, motivating the development of new methods to make inferences from biobank-scale data under fewer modeling assumptions.

Results

Overview of the approach

We investigate the utility of an estimator for SNP-heritability derived under a model that makes minimal assumptions on genetic architecture. We assume the standardized phenotype of an individual is a linear function of their genotypes: $y = \mathbf{x}^T \boldsymbol{\beta} + \epsilon$, where \mathbf{x} is a vector of standardized genotypes at M SNPs, $\boldsymbol{\beta}$ is a vector of standardized effect sizes corresponding to the M SNPs, and $\epsilon \sim N(0, \sigma_\epsilon^2)$ is environmental noise (Methods). We assume the effects can follow any distribution as long as the effect size of every SNP i is zero-centered ($E[\beta_i] = 0$) with a finite SNP-specific variance ($\text{Var}[\beta_i] = \sigma_i^2$) that is allowed to be 0, and that the covariance between the effects of any pair of SNPs is zero ($E[\beta_i \beta_j] = 0$ for all $i \neq j$). We term this model the “generalized random effects” (GRE) model as, to the best of our knowledge, all existing methods to estimate SNP-heritability impose additional assumptions on top of this model. For example, setting $\sigma_i^2 = h_g^2/M$ for $i = 1, \dots, M$ results in the single-component GREML model [3], whereas setting $\sigma_i^2 \propto w_i [f_i(1 - f_i)]^{0.75}$ (where w_i is a function of the “LD score” of SNP i and f_i is the MAF of SNP i) results in the most recent LDAK model [9] (Table 1). Under the GRE model, the SNP-heritability explained by the M SNPs is the sum of SNP-specific variances: $h_g^2 \equiv \text{Var}[\mathbf{x}^T \boldsymbol{\beta}] / \text{Var}[y] = \sum_{i=1}^M \sigma_i^2$ (Methods).

In this work, we are interested in accurately estimating h_g^2 from genotype measurements across N individuals at M typed SNPs. When $N > M$, the estimator $\hat{h}_g^2 = \frac{N \hat{\boldsymbol{\beta}}^T \hat{\mathbf{V}}^\dagger \hat{\boldsymbol{\beta}} - q}{N - q}$, where $\hat{\boldsymbol{\beta}}$ is the vector of standardized SNP effects estimated by ordinary least squares (OLS), $\hat{\mathbf{V}}^\dagger$ is the pseudoinverse of the in-sample LD matrix, and q is the rank of the in-sample LD matrix, is an unbiased estimator of SNP-heritability under the GRE model. That is, $E[\hat{h}_g^2] = \sum_{i=1}^M \sigma_i^2 = h_g^2$ (Methods). The GRE model allows each SNP-specific variance (σ_i^2) to be an arbitrary finite value satisfying the constraints $\sigma_i^2 \in [0, 1]$ and $\sum_{i=1}^M \sigma_i^2 \in [0, 1]$. Thus, σ_i^2 can capture any relationship between effect size and MAF/LD, which in turn implies that \hat{h}_g^2 is unbiased under most genetic architectures. Unfortunately, even the largest biobank-scale datasets currently available contain fewer unrelated individuals than typed SNPs (i.e. UK Biobank has genotyped $M \approx 593K$ SNPs in $N \approx 337K$ unrelated British individuals), which limits the utility of the above estimator. We therefore extend our approach by partitioning the genome by chromosome into 22 approximately independent regions:

$$\hat{h}_{\text{GRE}}^2 = \sum_{k=1}^{22} \frac{N \hat{\boldsymbol{\beta}}_k^T \hat{\mathbf{V}}_k^\dagger \hat{\boldsymbol{\beta}}_k - q_k}{N - q_k} \quad (1)$$

where for each chromosome k with p_k typed SNPs, $\hat{\boldsymbol{\beta}}_k$ is the p_k -vector of standardized SNP effects estimated by ordinary least squares (OLS), $\hat{\mathbf{V}}_k^\dagger$ is the pseudoinverse of the in-sample LD matrix, and q_k is the rank of the in-sample LD matrix. Although this genome-wide estimator does not provide theoretical guarantees of unbiasedness, we show through extensive simulations that the magnitude of the bias is extremely small across all architectures when N is sufficiently larger than p_k .

Accurate estimation of SNP-heritability irrespective of disease architecture

To investigate the bias and variance of \hat{h}_{GRE}^2 , we perform simulations starting from the real genotypes of $N = 337205$ unrelated British individuals in the UK Biobank [10]. First, we use data from chromosome 22 ($M = 9654$ typed SNPs) to simulate 64 distinct MAF- and LD-dependent architectures by varying the

SNP-heritability (h_g^2), the proportion of causal variants (p_{causal}), the distribution of causal variant MAF (CV MAF), and the strength of coupling between effect size and MAF/LD; we use “LDAK-LD-dependent” to describe architectures where causal effects are coupled with “LDAK weights” (see Methods). To enable comparison of estimates across different values of h_g^2 , we assess bias as a percentage of the simulated value of h_g^2 (relative bias) or the error of a single estimate as a percentage of h_g^2 (relative error). Consistent with analytical derivations, the GRE estimator restricted to chromosome 22 provides unbiased estimates across the 64 architectures after correcting for 16 independent tests at each value of h_g^2 (bias p-value $< 0.05/16$ is considered significant; see Methods) (Figure 1ac, Supplementary Table S1). The average relative bias across the 64 architectures is 0.00015% of the simulated h_g^2 , and the largest bias we observe under any single architecture is approximately $\pm 0.2\% \times h_g^2$ (Supplementary Figure S1a, Supplementary Table S1). We confirm that the analytical estimator of the standard error (Methods) is well-calibrated across all genetic architectures (Supplementary Figure S3a, Supplementary Table S4). We then investigate the bias induced by partitioning chromosome 22 into non-independent blocks and find that, as expected, our estimator accrues statistically significant upward bias as the average block size decreases (Supplementary Figure S2, Supplementary Table S2). For example, in simulations on chromosome 22 where $h_g^2 = 0.1$, $\alpha = -1$, $p_{\text{causal}} = 1\%$, and causal variants were chosen randomly from all SNPs, using a single chromosome-wide LD block produces approximately unbiased estimates (bias = 6.9×10^{-5} , p-value = 0.55) whereas partitioning the chromosome into 2 disjoint blocks of equal size induces a small but significant upward bias (bias = 4.3×10^{-4} , p-value = 5.3×10^{-4}) (Supplementary Figure S2, Supplementary Table S2).

Next, we investigate the accuracy of the GRE estimator in genome-wide simulations ($N = 337\text{K}$ unrelated individuals and $M = 593\text{K}$ array SNPs) where we use 22 chromosome-wide LD blocks to compute \hat{h}_{GRE}^2 . Despite the 22-block approximation, we find that \hat{h}_{GRE}^2 is highly accurate and robust across all 64 MAF- and LDAK-LD-dependent architectures (Figure 1bc). The average bias across the 64 architectures is $0.97\% \times h_g^2$, with the relative bias under any single architecture ranging from 0.07% to 2.1% of the simulated h_g^2 (Supplementary Figure S1b, Supplementary Table S3). The largest error we observe for a single estimate across all 6400 simulations (64 genetic architectures \times 100 simulation replicates) is approximately $17\% \times h_g^2$ (Figure 1c) and as N/M increases, the variance of \hat{h}_{GRE}^2 decreases while the relative bias across the 64 architectures appears to be approximately fixed, ranging between 0.91% ($N = 100\text{K}$) and 0.99% ($N = 200\text{K}$) (Figure 1d). This small but statistically significant upward bias in the genome-wide estimates is likely due to the 22-chromosome approximation, which ignores stochastic correlations between genotypes of SNPs on different chromosomes. These trends hold for a range of values of p_{causal} (Supplementary Table S3). Most importantly, the accuracy of the GRE estimator does not correlate with the simulated trait architecture (Figure 1b). We also assess the calibration of our analytical estimator for the standard error in the genome-wide simulations and observe a small downward bias with respect to the empirical standard deviation of \hat{h}_{GRE}^2 estimates (Supplementary Figure S3b, Supplementary Table S5). For example, across 16 distinct architectures where $h_g^2 = 0.25$, the empirical standard deviation computed from 100 independent estimates of h_g^2 ranges from 0.0049 to 0.0064, whereas our estimate of the standard error is approximately 0.0036 across all architectures (Supplementary Figure S3b, Supplementary Table S5). In all subsequent analyses, we compute \hat{h}_{GRE}^2 with the 22 chromosome-wide LD block approximation as this provides sufficiently accurate estimates and a fair comparison to other methods.

Comparison of methods to estimate SNP-heritability

We compare \hat{h}_{GRE}^2 with existing state-of-the-art approaches to estimate SNP-heritability that are easily scalable to the full UK Biobank data ($N = 337\text{K}$): LD score regression with no annotations (LDSC), which assumes $\alpha = -1$ and no coupling of effect size with LD [11]; stratified LD score regression (S-LDSC), which partitions h_g^2 by a set of annotations of interest [12, 13]; and SumHer, a recent scalable extension of LDAK which explicitly models MAF- and LD-dependent architectures through a specific form of the SNP-specific variances [14] (Table 1). To ensure a fair comparison among the methods, LD scores for LDSC, S-LDSC, and SumHer are computed using in-sample LD among the M SNPs, and in all simulations we aim to estimate the SNP-heritability explained by the same set of M SNPs (see Methods).

We find that \hat{h}_{GRE}^2 is highly accurate and robust across all simulated architectures while LDSC, S-LDSC, and SumHer are sensitive to deviations from their respective model assumptions. For example, when $h_g^2 = 0.25$ (Figure 2), LDSC is approximately unbiased under the “single-component GREML model” (relative bias = 0.04%, $p = 0.86$) but is sensitive to the MAF spectrum of causal variants and the degree of coupling between effect size and MAF/LD (e.g., across the 12 architectures where $p_{\text{causal}} = 1\%$, relative bias ranges from -44% to 50%) (Supplementary Table S6). Similarly, SumHer is accurate under the “LDAK model” (relative bias = 5.3%) but highly sensitive to other plausible genetic architectures (when $p_{\text{causal}} = 1\%$, relative bias ranges from -19% to 22%) (Figure 2, Supplementary Table S7). Estimates from S-LDSC (MAF), which partitions h_g^2 by 10 MAF bins (Supplementary Table S8; Methods), are less biased compared to estimates from LDSC when causal effects are coupled with only MAF, but are significantly downward biased when causal effects are also coupled with LDAK weights (for $h_g^2 = 0.25$, relative bias range is [1.9%, 7.0%] when $\gamma = 0$ and [-58%, -37%] when $\gamma = 1$) (Figure 2, Supplementary Table S9). S-LDSC with 10 MAF bins and an additional continuous “level of LD” (LLD) annotation, which we denote S-LDSC (MAF+LLD) (Methods), produces similar results on the same architectures (for $h_g^2 = 0.25$, relative bias range is [1.8%, 6.5%] when $\gamma = 0$ and [-80%, -33%] when $\gamma = 1$) (Supplementary Table S10). In contrast, the relative bias of \hat{h}_{GRE}^2 ranges from 0.45% to 1.3% across the same 16 genetic architectures where $h_g^2 = 0.25$ and $p_{\text{causal}} = 1\%$ (Figure 2, Supplementary Table S3). These trends hold for a range of values of h_g^2 and p_{causal} : across 112 distinct LDAK-LD- and/or MAF-dependent architectures, the average and range of the relative bias of each method are 0.96% [-0.06%, 2.1%] for \hat{h}_{GRE}^2 , -2.2% [-71%, 70%] for LDSC, -22% [-62%, 8.7%] for S-LDSC (MAF), -29% [-89%, 9.0%] for S-LDSC (MAF+LLD), and 2.8% [-27%, 28%] for SumHer (Figure 1b, Figure 2, Supplementary Figures S4-S7, Supplementary Tables S3, S6-S7, S9-S10). We also perform simulations under 14 alternative LD-dependent architectures where the variance of each SNP is coupled with its inverse LD score instead of its LDAK weight (i.e. “LD-score-dependent” architectures; see Methods, Supplementary Figure S8) and find that \hat{h}_{GRE}^2 remains nearly unbiased (relative bias ranges from 0.52% to 1.3%) whereas estimates from S-LDSC (MAF), S-LDSC (MAF+LLD), and SumHer are downward-biased on average across the 14 architectures (Supplementary Figure S9, Supplementary Table S11).

For completeness, we also compare to four widely used REML-based methods: single-component GREML (GREML), which assumes $\alpha = -1$ and no coupling of effect size with LD [3]; GREML-LDMS-I, a multi-component extension of GREML that partitions SNPs by MAF and LD score [18]; BOLT-REML, a computationally efficient variance components estimation method with assumptions similar to those of GREML [8]; and LDAK, which assumes a specific form of the coupling of effect size with LD and recommends setting $\alpha = -0.25$ [6, 9] (Table 1). Because it is computationally intractable to apply the REML-based methods

to thousands of genome-wide simulations with 337K individuals, we perform simulations using a reduced number of individuals and SNPs ($N = 8430$ individuals and $M = 14821$ array SNPs; see Methods). We find that the single-component REML methods (GREML, BOLT-REML, and LDAK) are sensitive to MAF- and LD-dependent architectures that deviate from their respective model assumptions, whereas our estimator is robust to all architecture parameters. For example, when $h_g^2 = 0.25$ (Figure 3), GREML and BOLT-REML are accurate under the “single-component GREML model” (GREML: relative bias = -1.4%, $p = 6.0 \times 10^{-3}$, Supplementary Table S12; BOLT-REML: relative bias = -0.16%, $p = 0.75$, Supplementary Table S13) and LDAK is approximately unbiased under the “LDAK model” (relative bias = 0.16%, $p = 0.77$, Supplementary Table S14), but all single-component methods are sensitive to the MAF spectrum of causal variants and to the coupling of causal effects with MAF/LD. Across the 12 architectures in Figure 3 where $p_{\text{causal}} = 1\%$, the relative biases of the single-component methods range from -15% to 7.9% (GREML), -14% to 9.1% (BOLT-REML), and -34% to 8.2% (LDAK) (Supplementary Tables S12-S14). In contrast, for the same 12 architectures, \hat{h}_{GRE}^2 yields relative biases in the range [-2.1%, 1.7%], which is comparable to the relative bias observed with GREML-LDMS-I (range [-2.9%, 1.5%]) when using 8 GRMs defined by 4 LD quartiles and 2 MAF bins (MAF > 5% and MAF ≤ 5%) that align with the causal variant MAF spectrum (Figure 3, Supplementary Tables S15, S16). These trends are consistent across a range of values of h_g^2 and p_{causal} : across the 112 distinct LDAK-LD- and/or MAF-dependent architectures shown in Supplementary Figures S10-S14, the average and range of the relative bias are 0.09% [-4.9%, 6.4%] (GRE), -0.6% [-5.9%, 2.3%] (GREML-LDMS-I), -2.9% [-27%, 15%] (GREML), -1.8% [-25%, 18%] (BOLT-REML), and -8.2% [-44%, 13%] (LDAK) (Supplementary Tables S12-S16). Similar trends are observed in additional simulations under 14 LD-score-dependent architectures (Supplementary Figure S15, Supplementary Table S17). We note that the performance of GREML-LDMS-I depends on the resolution of the MAF and LD bins used to partition SNPs; in an extreme example where all causal variants are drawn from a MAF range tightly concentrated near 1%, running GREML-LDMS-I with the same 8 GRMs as before yields downward-biased estimates whereas our estimator remains robust (Supplementary Figure S16, Supplementary Tables S12-S16). While the variance of our estimator is larger than the variances of the REML-based methods (Figure 3), our approach is designed for biobank-scale GWAS data with sample sizes several orders of magnitude larger than what we used in these small-scale simulations. In summary, our results confirm that it is possible to accurately estimate h_g^2 under minimal assumptions about genetic architecture.

Estimating SNP-heritability of 22 complex traits in the UK Biobank

Finally, we apply our approach to estimate h_g^2 for 22 complex traits and diseases in the UK Biobank ($N = 290\text{K}$ unrelated British individuals, $M = 460\text{K}$ array SNPs; see Methods) [10]. For comparison, we also provide estimates of h_g^2 from LDSC (no annotations), S-LDSC (controlling for the baseline-LD model [13,29]), and SumHer. Of the 22 traits analyzed (6 quantitative and 16 binary), we focus on 18 traits for which $\hat{h}_{\text{GRE}}^2 > 0.05$ (Table 2). Using our approach, estimates of SNP-heritability for the 6 quantitative traits range from 0.12 (smoking status) to 0.60 (height). Across the 12 binary traits, our estimates range from 0.064 (autoimmune disorders) to 0.16 (hypertension) (Table 2). To enable a direct comparison between \hat{h}_{GRE}^2 and the SNP-heritability quantities estimated by LDSC, S-LDSC, and SumHer, we run each of the summary-statistics-based methods with LD scores and regression weights computed from in-sample LD among the typed SNPs, and we estimate SNP-heritability as the sum of the per-SNP variances across all M

SNPs (Methods). Across the 18 traits, the median difference (as a percentage of \hat{h}_{GRE}^2) between S-LDSC (baseline-LD/in-sample) and \hat{h}_{GRE}^2 is -9%; the median difference between SumHer (in-sample) and \hat{h}_{GRE}^2 is 11% (Figure 4, Supplementary Table S18). This pattern is roughly consistent with the global trends we observed in genome-wide simulations (Figure 2). As expected [11], LDSC (in-sample) yields inflated estimates across all 18 traits. To enable a comparison between \hat{h}_{GRE}^2 and SNP-heritability estimates from summary-statistics-based methods reported in the literature, we also run LDSC, S-LDSC, and SumHer with their recommended parameter settings [11, 12, 14, 29] and with LD scores and regression weights computed from 489 Europeans in the 1000 Genomes Phase 3 reference panel [30] – we note that when running these methods as recommended, their SNP-heritability estimands are not equivalent to our definition of h_g^2 (see Methods and refs. [11, 12, 14, 19] for details). Across the 18 traits for which $\hat{h}_{\text{GRE}}^2 > 0.05$, the median differences with respect to \hat{h}_{GRE}^2 are -11% for LDSC (1KG), -14% for S-LDSC (baseline-LD/1KG), and 38% for SumHer (1KG) (Supplementary Figure S17, Supplementary Table S18). Across 9 traits (a subset of the 18 traits with $\hat{h}_{\text{GRE}}^2 > 0.05$) for which a previous study reported estimates from single-component BOLT-REML computed from approximately 337K unrelated white British individuals in the UK Biobank [31], the median difference between the previously reported BOLT-REML estimates and \hat{h}_{GRE}^2 is 8% (Supplementary Table S18).

Runtime and memory requirement

Since our approach is designed to be applied to biobank-scale data, we report the runtime and memory requirements for computing \hat{h}_{GRE}^2 with 22 chromosome-wide LD blocks in the UK Biobank ($N = 337\text{K}$ individuals, $M = 593\text{K}$ array SNPs). First, we compute chromosome-wide LD; this has complexity $O(Np_k^2)$ for chromosome k with p_k SNPs. In practice, this step does not impose a computational bottleneck because the LD computations can be parallelized over SNP partitions (e.g., given a pair of SNP partitions containing 1000 SNPs each, computing the pairwise LD matrix takes about 10 minutes and 16GB of memory). Second, the pseudoinverse of each chromosome-wide LD matrix is computed via truncated singular value decomposition (SVD), which has complexity $O(p_k^3)$ for chromosome k . This step is parallelized over chromosomes; for chromosome 2, which has the largest number of typed SNPs, computing the truncated SVD and pseudoinverse of the LD matrix takes about 3 hours and 60GB of memory. Lastly, given the precomputed pseudoinverse of each chromosome-wide LD matrix and OLS association statistics, computing \hat{h}_{GRE}^2 genome-wide has complexity $O(p_1^2 + \dots + p_{22}^2)$. For any of the UK Biobank traits analyzed in this work, this takes less than 1 hour and requires 24GB of memory; most of this time is spent loading the pseudoinverse LD matrices into memory.

Discussion

In this work, we show that highly accurate estimation of SNP-heritability can be achieved under minimal assumptions on the genetic architecture of complex traits. In particular, our proposed estimator assumes that each SNP effect has a fixed SNP-specific variance that can capture any arbitrary relationship between effect size and genomic features such as MAF and LD. We show that all existing methods to estimate SNP-heritability impose additional assumptions on the GRE model, and we confirm through extensive simulations that these methods are susceptible to bias when their modeling assumptions are not met. Additionally, we

confirm that REML-based methods that partition SNPs by MAF and LD score generally yield much smaller bias compared to single-component REML methods [18]. In contrast, our estimator, derived under the GRE model, provides accurate estimates of SNP-heritability regardless of the underlying genetic architecture, without specifying a heritability model or partitioning SNPs by functional categories. On average across 18 heritable traits in the UK Biobank ($\widehat{h}_{\text{GRE}}^2 > 0.05$), our approach yields estimates that are higher than S-LDSC estimates (controlling for the baseline-LD model [29]) and lower than SumHer estimates (with the recommended heritability model [9, 14]). One practical advantage of our approach over methods such as LDSC, S-LDSC, and SumHer is that the estimand of our approach is always the same for a given genotype matrix, whereas the definitions and interpretations of the estimands of LDSC, S-LDSC, and SumHer can vary drastically depending on what sets of SNPs are used in each step of the inference procedure (e.g., the set of SNPs used to compute LD scores need not be the same set of SNPs that defines the SNP-heritability estimand of interest) [11, 12, 19]. Overall, our results show that while existing methods can yield biases, for the purpose of estimating total SNP-heritability of complex traits, most methods are relatively accurate and robust to plausible genetic architectures.

We conclude with several caveats and future directions. First, the utility of GRE estimator critically depends on the ratio between the number of SNPs (M) and the number of individuals (N) in the data – as M/N increases, the eigenstructure of the in-sample LD matrix (and sample covariance matrices in general) becomes increasingly distorted (larger eigenvalues are overestimated and smaller eigenvalues are underestimated) [32]. We mitigate this by assuming that the genome-wide LD matrix has a block diagonal structure (specifically, one block per chromosome); since the number of unrelated British individuals in the UK Biobank is larger than the number of array SNPs per chromosome, our approach is able to provide meaningful estimates of the SNP-heritability attributable to common SNPs ($\text{MAF} > 1\%$) in individuals of British ancestry. A major limitation of our approach remains with respect to imputed and/or whole-genome sequenced data, in which the number of SNPs will continue to be orders of magnitude larger than the number of individuals for the foreseeable future. We defer a thorough investigation of regularized estimation of LD in high-dimensional settings ($M > N$) to future work.

Second, this work focuses on accurate estimation of SNP-heritability given OLS association statistics and chromosome-wide LD matrices estimated from the same genotype data. While summary statistics have been made publicly available for hundreds of GWAS, in-sample LD is usually unavailable or impossible to obtain since most large-scale GWAS are meta-analyses [33], and publicly available reference panels such as 1000 Genomes [30] currently have sample sizes in the hundreds or thousands at most. In addition, many publicly available summary statistics were computed using linear mixed models rather than OLS in order to control for population structure/cryptic relatedness. Previous works have noted (in the context of statistical fine-mapping) that the LD computation must be adjusted to accommodate association statistics computed from mixed models [33, 34]. The sensitivity of our estimator to reference panel LD (with or without regularized LD estimation) and/or mixed model association statistics remains unclear [28, 35]; we leave an investigation of both for future work. Furthermore, our simulations use typed SNPs to draw phenotypes because imputed genotypes have highly irregular LD patterns [9, 18]. Although it would be more realistic to simulate causal variants and phenotypes from a denser set of genotyped SNPs or from whole-genome sequencing data [18], our simulation design was dependent on the availability of individual-level genotype measurements in biobank-scale sample sizes.

Third, the GRE estimator does not correct for population structure/cryptic relatedness. We mitigate this in our analysis of real UK Biobank traits by considering only unrelated individuals (> 3 rd degree relatives) and by including age, sex, and the top 20 principal components as covariates in the linear regression when computing OLS association statistics. However, recent work has found significant evidence of assortative mating for some traits in the UK Biobank (e.g., height) and not others [36], suggesting that our estimator might be more susceptible to bias for some traits over others. Currently, it remains unclear how to quantify the bias of our estimator due to population structure and/or assortative mating in real data. In addition, we derive the GRE estimator under no ascertainment in case/control data. Future work is needed to extend the GRE approach to control for ascertainment bias [15, 16, 37, 38].

Finally, while previous works have applied similar estimators in the context of fixed effects models to estimate local SNP-heritability within small regions (e.g., LD blocks) [27, 28], additional work is needed to extend our approach to perform functional partitioning of SNP-heritability by higher-resolution annotations. Existing methods for partitioning genome-wide SNP-heritability by small and/or overlapping annotations make various assumptions on genetic architecture [8, 12–14, 29], motivating the development of new methods in this area under fewer assumptions.

URLs

GRE estimator: <https://bogdan.dgsom.ucla.edu/pages/software>

BOLT-LMM: <https://data.broadinstitute.org/alkesgroup/BOLT-LMM/>

GCTA: <https://cnsgenomics.com/software/gcta/>

LDAC: <http://dougsped.com/ldak/>

LDSC: <https://github.com/bulik/ldsc/>

baseline-LD annotations: <https://data.broadinstitute.org/alkesgroup/LDSCORE/>

PLINK: <https://www.cog-genomics.org/plink2>

UK Biobank: <https://www.ukbiobank.ac.uk>

Acknowledgments

This research was conducted using the UK Biobank Resource under applications 33297 and 33127. We thank the participants of UK Biobank for making this work possible. We also thank Ruth Johnson, Malika Kumar Freund, Megan Major, Steven Gazal, Alkes Price and David Balding for helpful discussions. This work was funded by the National Institutes of Health (NIH) under awards R01HG009120, R01MH115676, R01HG006399, U01CA194393, T32NS048004, T32MH073526, and T32HG002536.

Methods

The generalized random effects model

We model the phenotype for an individual n randomly sampled from the population as $y_n = \mathbf{x}_n^T \boldsymbol{\beta} + \epsilon_n$, where $\mathbf{x}_n = (x_{n1} \dots x_{nM})^T$ is a vector of standardized genotypes measured at M SNPs for individual n ,

$\boldsymbol{\beta} = (\beta_1, \dots, \beta_M)^T$ is an M -vector of the corresponding standardized SNP effect sizes, and $\epsilon_n \sim N(0, \sigma_\epsilon^2)$ is environmental noise. We assume $\text{Var}[y_n] = 1$ and that the genotype at each SNP i is centered and scaled in the population such that $E[x_{ni}] = 0$ and $\text{Var}[x_{ni}] = 1$; i.e. $x_{ni} = (g_{ni} - 2f_i)/\sqrt{2f_i(1-f_i)}$, where $g_{ni} \in \{0, 1, 2\}$ is the number of copies of the effect allele at SNP i for individual n , and f_i is the population frequency of the effect allele at SNP i . We define the population LD between two SNPs i and j to be $v_{ij} \equiv E[x_{ni}x_{nj}]$ for all $i \neq j$. The population LD matrix among the M SNPs is therefore $\mathbf{V} \equiv \text{Cov}[\mathbf{x}_n^T]$. For simplicity, we use ‘‘SNP effect sizes’’ in lieu of ‘‘standardized SNP effect sizes’’ to refer to $\boldsymbol{\beta}$. We assume that the genotypes \mathbf{x}_n and effect sizes $\boldsymbol{\beta}$ are independent given allele frequencies (f_1, \dots, f_M) and \mathbf{V} .

Under the generalized random effects (GRE) model, the first two moments of the distribution of the effect size of SNP i are $E[\beta_i] = 0$ and $\text{Var}[\beta_i] = \sigma_i^2$, where σ_i^2 can be any arbitrary nonnegative finite number. We assume the covariance between the effects of different SNPs is 0 (i.e. $\text{Cov}[\beta_i, \beta_j] = E[\beta_i\beta_j] = 0$ for all $i \neq j$). Because the SNP-specific variances $(\sigma_1^2, \dots, \sigma_M^2)$ can capture any polygenicity (number of variants with effects larger than some measurable constant) and any degree of coupling between genomic features (e.g., MAF and LD) and effect size, the GRE model encompasses most realistic genetic architectures (Table 1).

We define total SNP-heritability (h_g^2) to be the proportion of phenotypic variance attributable to the additive effects of a set of M SNPs whose genotypes are directly measured:

$$\begin{aligned}
 h_g^2 &\equiv \text{Var}[\mathbf{x}_n^T \boldsymbol{\beta}] / \text{Var}[y_n] & (2) \\
 &= E[\text{Var}[\mathbf{x}_n^T \boldsymbol{\beta} | \boldsymbol{\beta}]] + \text{Var}[E[\mathbf{x}_n^T \boldsymbol{\beta} | \boldsymbol{\beta}]] \\
 &= E[\boldsymbol{\beta}^T \text{Var}[\mathbf{x}_n^T] \boldsymbol{\beta}] + \text{Var}[E[\mathbf{x}_n^T] \boldsymbol{\beta}] \\
 &= E[\boldsymbol{\beta}^T \mathbf{V} \boldsymbol{\beta}] + 0 \\
 &= E[\text{tr}(\mathbf{V} \boldsymbol{\beta} \boldsymbol{\beta}^T)] \\
 &= \text{tr}(\mathbf{V} E[\boldsymbol{\beta} \boldsymbol{\beta}^T]) \\
 h_g^2 &= \sum_{i=1}^M \sigma_i^2 & (3)
 \end{aligned}$$

Thus, h_g^2 is defined with respect to a given population and a given set of SNPs. By definition, $0 \leq h_g^2 \leq 1$. Similarly, we define regional SNP-heritability (h_k^2) to be the proportion of phenotypic variance due to the additive effects of the genotyped SNPs in region k . We assume that the set of SNPs that defines h_k^2 is a subset of the M SNPs that define h_g^2 (thus, $0 \leq h_k^2 \leq h_g^2$). If region k is the whole genome, $h_k^2 = h_g^2$.

Estimating SNP-heritability under the GRE model

We are interested in estimating h_g^2 under the GRE model (Equation 3). In a GWAS with N individuals genotyped at M SNPs, let $\mathbf{X} = (\mathbf{x}_1^T, \dots, \mathbf{x}_N^T)^T$ be the $N \times M$ matrix of standardized genotypes (i.e. each column of \mathbf{X} has been standardized to have mean 0 and variance 1), let $\mathbf{y} = (y_1, \dots, y_N)^T$ be the N -vector of standardized phenotypes, and let $\widehat{\mathbf{V}} = (1/N)\mathbf{X}^T \mathbf{X}$ be the $M \times M$ in-sample LD matrix (an estimate of population LD, \mathbf{V}) with rank q , where $1 \leq q \leq M$. Let $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_K)$ be a set of K approximately independent genome partitions spanning all M SNPs (e.g., chromosomes). For each region k containing p_k SNPs, \mathbf{X}_k is the $N \times p_k$ standardized genotype matrix and $\widehat{\mathbf{V}}_k$ is the corresponding $p_k \times p_k$ in-sample

LD matrix with rank q_k , where $1 \leq q_k \leq p_k$. We propose the following estimator for genome-wide SNP-heritability:

$$\hat{h}_{\text{GRE}}^2 = \sum_{k=1}^K \frac{N \hat{\beta}_k^T \hat{\mathbf{V}}_k^\dagger \hat{\beta}_k - q_k}{N - q_k} \quad (4)$$

where $\hat{\beta}_k = (1/N) \mathbf{X}_k^T \mathbf{y}$ is the p_k -vector of marginal SNP effects estimated by ordinary least squares (OLS) for region k and $\hat{\mathbf{V}}_k^\dagger$ is the pseudoinverse of $\hat{\mathbf{V}}_k$.

In the following sections, we first derive \hat{h}_{GRE}^2 in the simplest case where $K = 1$ and $N > M$ by finding an estimator that satisfies $\text{E}[\hat{h}_{\text{GRE}}^2] = h_g^2$. We then describe modifications to this estimator to allow $N < M$ as well as rank-deficient LD matrices. Lastly, we derive an analytical form for the standard error of \hat{h}_{GRE}^2 .

Derivation for \hat{h}_{GRE}^2 assuming fixed β and $N > M$

Recall that $\text{Var}[y_n] = 1$ and $\text{Var}[\mathbf{x}_n^T] = \mathbf{V}$. Our goal is to find an estimator \hat{h}_{GRE}^2 that satisfies $\text{E}[\hat{h}_{\text{GRE}}^2] = h_g^2 = \text{Var}[\mathbf{x}_n^T \beta] = \text{E}[\text{Var}[\mathbf{x}_n^T \beta | \beta]] + \text{Var}[\text{E}[\mathbf{x}_n^T \beta | \beta]] = \text{E}[\beta^T \mathbf{V} \beta]$ (Equation 2). If β were fixed and we observed \mathbf{V} and β , we could estimate h_g^2 as $\hat{h}_g^2 = \beta^T \mathbf{V} \beta$. However, in reality, we observe noisy estimates of β and \mathbf{V} from GWAS. Given a GWAS of N unrelated individuals and M SNPs, we observe \mathbf{X} , the standardized genotype matrix, and \mathbf{y} , the standardized phenotype vector. We assume that when $N > M$, $\hat{\mathbf{V}} \rightarrow \mathbf{V}$ as $N \rightarrow \infty$ (in practice, the assumption that $N > M$ is untrue; in subsequent sections we show how we partition the genome into K blocks such that $N > p_k$ for each block k). In a typical GWAS, the marginal SNP effects are estimated through ordinary least squares (OLS) regression as $\hat{\beta} = (1/N) \mathbf{X}^T \mathbf{y} = (1/N) \mathbf{X}^T \mathbf{X} \beta + (1/N) \mathbf{X}^T \epsilon = \hat{\mathbf{V}} \beta + (1/N) \mathbf{X}^T \epsilon$. Given \mathbf{X} and fixed β , it follows that

$$\begin{aligned} \text{E}[\hat{\beta} | \beta, \mathbf{X}] &= \text{E}[\hat{\mathbf{V}} \beta + (1/N) \mathbf{X}^T \epsilon | \beta, \mathbf{X}] \\ &= \hat{\mathbf{V}} \beta + (1/N) \mathbf{X}^T \text{E}[\epsilon] \\ &= \hat{\mathbf{V}} \beta \end{aligned} \quad (5)$$

$$\begin{aligned} \text{Cov}[\hat{\beta} | \beta, \mathbf{X}] &= \text{Cov}[\hat{\mathbf{V}} \beta + (1/N) \mathbf{X}^T \epsilon | \beta, \mathbf{X}] \\ &= (\sigma_e^2 / N^2) \mathbf{X}^T \mathbf{X} \\ &= \frac{\sigma_e^2}{N} \hat{\mathbf{V}} \end{aligned} \quad (6)$$

Thus, as $N \rightarrow \infty$, $\hat{\beta} \rightarrow \mathbf{V} \beta$. Substituting $\hat{\mathbf{V}}^{-1} \hat{\beta} \approx \beta$ and $\hat{\mathbf{V}} \approx \mathbf{V}$, we obtain the revised estimator $\hat{h}_g^2 = \beta^T \mathbf{V} \beta \approx (\hat{\mathbf{V}}^{-1} \hat{\beta})^T \hat{\mathbf{V}} (\hat{\mathbf{V}}^{-1} \hat{\beta}) = \hat{\beta}^T \hat{\mathbf{V}}^{-1} \hat{\beta}$. The expectation of this estimator is

$$\begin{aligned} \text{E}[\hat{\beta}^T \hat{\mathbf{V}}^{-1} \hat{\beta} | \beta, \mathbf{X}] &= \text{E}[\text{tr}(\hat{\beta}^T \hat{\mathbf{V}}^{-1} \hat{\beta}) | \beta, \mathbf{X}] \\ &= \text{E}[\text{tr}(\hat{\mathbf{V}}^{-1} \hat{\beta} \hat{\beta}^T) | \beta, \mathbf{X}] \\ &= \text{tr}(\hat{\mathbf{V}}^{-1} \text{E}[\hat{\beta} \hat{\beta}^T | \beta, \mathbf{X}]) \\ &= \text{tr}(\hat{\mathbf{V}}^{-1} \text{Cov}[\hat{\beta} | \beta, \mathbf{X}]) + \text{tr}(\hat{\mathbf{V}}^{-1} \text{E}[\hat{\beta} | \beta, \mathbf{X}] \text{E}[\hat{\beta} | \beta, \mathbf{X}]^T) \\ &= \text{tr}((\sigma_e^2 / N) \hat{\mathbf{V}}^{-1} \hat{\mathbf{V}}) + \beta^T \hat{\mathbf{V}} \beta \\ &= \frac{M}{N} \sigma_e^2 + \beta^T \hat{\mathbf{V}} \beta \end{aligned} \quad (7)$$

We define $\widehat{h}_{\text{GRE}}^2$ to be an estimator that satisfies $E[\widehat{h}_{\text{GRE}}^2|\boldsymbol{\beta}, \mathbf{X}] = \boldsymbol{\beta}^T \widehat{\mathbf{V}} \boldsymbol{\beta}$. Substituting into Equation 7, we obtain

$$\begin{aligned} E[\widehat{\boldsymbol{\beta}}^T \widehat{\mathbf{V}}^{-1} \widehat{\boldsymbol{\beta}}|\boldsymbol{\beta}, \mathbf{X}] &= \frac{M(1 - E[\widehat{h}_{\text{GRE}}^2|\boldsymbol{\beta}, \mathbf{X}])}{N} + E[\widehat{h}_{\text{GRE}}^2|\boldsymbol{\beta}, \mathbf{X}] \\ &= \frac{M}{N} + \frac{N-M}{N} E[\widehat{h}_{\text{GRE}}^2|\boldsymbol{\beta}, \mathbf{X}] \\ E[\widehat{h}_{\text{GRE}}^2|\boldsymbol{\beta}, \mathbf{X}] &= \left(E[\widehat{\boldsymbol{\beta}}^T \widehat{\mathbf{V}}^{-1} \widehat{\boldsymbol{\beta}}|\boldsymbol{\beta}, \mathbf{X}] - \frac{M}{N} \right) \frac{N}{N-M} \\ &= \frac{NE[\widehat{\boldsymbol{\beta}}^T \widehat{\mathbf{V}}^{-1} \widehat{\boldsymbol{\beta}}|\boldsymbol{\beta}, \mathbf{X}] - M}{N-M} \\ \widehat{h}_{\text{GRE}}^2 &= \frac{N\widehat{\boldsymbol{\beta}}^T \widehat{\mathbf{V}}^{-1} \widehat{\boldsymbol{\beta}} - M}{N-M} \end{aligned} \quad (8)$$

Unbiasedness of $\widehat{h}_{\text{GRE}}^2$ under the GRE model when $N > M$

Recall that under the GRE model, $E[\beta_i] = 0$ and $\text{Var}[\beta_i] = \sigma_i^2$, where $\sigma_i^2 \geq 0$ for all SNPs i . In the previous sections, we showed that $E[\widehat{h}_{\text{GRE}}^2|\boldsymbol{\beta}, \mathbf{X}] = \boldsymbol{\beta}^T \widehat{\mathbf{V}} \boldsymbol{\beta}$ and $h_g^2 = \sum_{i=1}^M \sigma_i^2$. Recalling that $\text{Cov}[\beta_i \beta_j] = 0$ for all $i \neq j$, it follows that

$$\begin{aligned} E[\widehat{h}_{\text{GRE}}^2|\mathbf{X}] &= E[E[\widehat{h}_{\text{GRE}}^2|\boldsymbol{\beta}, \mathbf{X}]|\mathbf{X}] \\ &= E[\boldsymbol{\beta}^T \widehat{\mathbf{V}} \boldsymbol{\beta}|\mathbf{X}] \\ &= E[\text{tr}(\boldsymbol{\beta}^T \widehat{\mathbf{V}} \boldsymbol{\beta})|\mathbf{X}] \\ &= \text{tr}(\widehat{\mathbf{V}} E[\boldsymbol{\beta} \boldsymbol{\beta}^T]) \\ &= \sum_{i=1}^M \sigma_i^2 \end{aligned} \quad (9)$$

Therefore, $E[\widehat{h}_{\text{GRE}}^2] = E[E[\widehat{h}_{\text{GRE}}^2|\mathbf{X}]] = \sum_{i=1}^M \sigma_i^2 = h_g^2$. This result implies that $\widehat{h}_{\text{GRE}}^2$ is an unbiased estimator for h_g^2 under a wide range of genetic architectures that fall under the GRE model.

Practical considerations for implementation

For most GWAS, because the number of genotyped SNPs M is much larger than the number of individuals N in the study, $\widehat{\mathbf{V}}$ is a poor estimator of \mathbf{V} genome-wide; as M/N increases, the eigenstructure of $\widehat{\mathbf{V}}$ becomes increasingly distorted (larger eigenvalues are overestimated and smaller eigenvalues are underestimated) [32]. In addition, it is generally computationally intractable to compute and invert $\widehat{\mathbf{V}}$ genome-wide. Thus, in practice, we divide the genome into a set of K approximately independent blocks (e.g., by chromosome) and implement $\widehat{h}_{\text{GRE}}^2$ as

$$\widehat{h}_{\text{GRE}}^2 = \sum_{k=1}^K \frac{N \widehat{\boldsymbol{\beta}}_k^T \widehat{\mathbf{V}}_k^{-1} \widehat{\boldsymbol{\beta}}_k - p_k}{N - p_k} \quad (11)$$

where $\widehat{\boldsymbol{\beta}}_k$ is the p_k -vector of marginal effects for block k containing p_k SNPs. While this estimator does not provide theoretical guarantees of unbiasedness, it remains statistically consistent (i.e. $\widehat{h}_{\text{GRE}}^2 \rightarrow h_g^2$ as

$N \rightarrow \infty$).

Extension for rank-deficient LD

It is often the case that two SNPs are perfectly correlated in a genotype block \mathbf{X}_k , or that $N < p_k$ for a block k . In this case, $\widehat{\mathbf{V}}_k$ is rank-deficient (i.e. its rank is less than p_k) and $\widehat{\mathbf{V}}_k^{-1}$ does not exist. We therefore compute \mathbf{V}_k^\dagger , the pseudoinverse (Moore-Penrose inverse) of $\widehat{\mathbf{V}}_k$, which approximates $\widehat{\mathbf{V}}_k^{-1}$ using its truncated eigendecomposition. Let $q_k = \text{rank}(\widehat{\mathbf{V}}_k)$ and let $\widehat{\mathbf{V}}_k = \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{U}_k^T$ be the eigendecomposition of $\widehat{\mathbf{V}}_k$, where $\mathbf{\Lambda}_k = \text{diag}(\lambda_1, \dots, \lambda_{q_k}, 0, \dots, 0)$. The pseudoinverse of $\widehat{\mathbf{V}}_k$ is $\widehat{\mathbf{V}}_k^\dagger = \mathbf{U}_k \mathbf{\Lambda}_k^\dagger \mathbf{U}_k^T$, where $\mathbf{\Lambda}_k^\dagger = \text{diag}(\lambda_1^{-1}, \dots, \lambda_{q_k}^{-1}, 0, \dots, 0)$.

Substituting $\widehat{\mathbf{V}}_k^\dagger \widehat{\boldsymbol{\beta}}_k \approx \boldsymbol{\beta}_k$ and $\widehat{\mathbf{V}}_k \approx \mathbf{V}_k$, we obtain the following estimator for h_k^2 : $\widehat{h}_k^2 = \boldsymbol{\beta}_k^T \mathbf{V}_k \boldsymbol{\beta}_k \approx (\widehat{\mathbf{V}}_k^\dagger \widehat{\boldsymbol{\beta}}_k)^T \widehat{\mathbf{V}}_k (\widehat{\mathbf{V}}_k^\dagger \widehat{\boldsymbol{\beta}}_k) = \widehat{\boldsymbol{\beta}}_k^T \widehat{\mathbf{V}}_k^\dagger \widehat{\boldsymbol{\beta}}_k$. Let \mathbf{I}_{q_k} be a $p_k \times p_k$ diagonal matrix in which the first q_k diagonal entries are 1 and the rest are 0. The expectation of our estimator given $\boldsymbol{\beta}$ and \mathbf{X} is

$$\begin{aligned}
 \mathbb{E}[\widehat{\boldsymbol{\beta}}_k^T \widehat{\mathbf{V}}_k^\dagger \widehat{\boldsymbol{\beta}}_k | \boldsymbol{\beta}, \mathbf{X}] &= \mathbb{E}[\text{tr}(\widehat{\mathbf{V}}_k^\dagger \widehat{\boldsymbol{\beta}}_k \widehat{\boldsymbol{\beta}}_k^T) | \boldsymbol{\beta}, \mathbf{X}] \\
 &= \text{tr}(\widehat{\mathbf{V}}_k^\dagger \mathbb{E}[\widehat{\boldsymbol{\beta}}_k \widehat{\boldsymbol{\beta}}_k^T | \boldsymbol{\beta}, \mathbf{X}]) \\
 &= \text{tr}(\widehat{\mathbf{V}}_k^\dagger \text{Cov}[\widehat{\boldsymbol{\beta}}_k | \boldsymbol{\beta}, \mathbf{X}]) + \text{tr}(\widehat{\mathbf{V}}_k^\dagger \mathbb{E}[\widehat{\boldsymbol{\beta}}_k | \boldsymbol{\beta}, \mathbf{X}] \mathbb{E}[\widehat{\boldsymbol{\beta}}_k^T | \boldsymbol{\beta}, \mathbf{X}]^T) \\
 &= \text{tr}((\sigma_e^2/N) \widehat{\mathbf{V}}_k^\dagger \widehat{\mathbf{V}}_k) + \text{tr}(\widehat{\mathbf{V}}_k^\dagger (\widehat{\mathbf{V}}_k \boldsymbol{\beta}_k) (\widehat{\mathbf{V}}_k \boldsymbol{\beta}_k)^T) \\
 &= \text{tr}((\sigma_e^2/N) \mathbf{I}_{q_k}) + \text{tr}(\mathbf{U}_k \mathbf{\Lambda}_k^\dagger \mathbf{U}_k^T \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{U}_k^T \boldsymbol{\beta}_k \boldsymbol{\beta}_k^T \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{U}_k^T) \\
 &= \frac{q_k}{N} \sigma_e^2 + \text{tr}(\mathbf{U}_k \mathbf{\Lambda}_k^\dagger \mathbf{\Lambda}_k \mathbf{U}_k^T \boldsymbol{\beta}_k \boldsymbol{\beta}_k^T \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{U}_k^T) \\
 &= \frac{q_k}{N} \sigma_e^2 + \text{tr}(\mathbf{I}_{q_k} \mathbf{U}_k^T \boldsymbol{\beta}_k \boldsymbol{\beta}_k^T \mathbf{U}_k \mathbf{\Lambda}_k) \\
 &= \frac{q_k}{N} \sigma_e^2 + \text{tr}(\boldsymbol{\beta}_k^T \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{I}_{q_k} \mathbf{U}_k^T \boldsymbol{\beta}_k) \\
 &= \frac{q_k}{N} \sigma_e^2 + \boldsymbol{\beta}_k^T \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{I}_{q_k} \mathbf{U}_k^T \boldsymbol{\beta}_k \\
 &= \frac{q_k}{N} \sigma_e^2 + \boldsymbol{\beta}_k^T \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{U}_k^T \boldsymbol{\beta}_k \\
 &= \frac{q_k}{N} \sigma_e^2 + \boldsymbol{\beta}_k^T \widehat{\mathbf{V}}_k \boldsymbol{\beta}_k
 \end{aligned} \tag{12}$$

We wish to find an estimator that satisfies $\mathbb{E}[\widehat{h}_{\text{GRE}}^2 | \boldsymbol{\beta}, \mathbf{X}] = \boldsymbol{\beta}^T \widehat{\mathbf{V}} \boldsymbol{\beta} = \sum_{k=1}^K \boldsymbol{\beta}_k^T \widehat{\mathbf{V}}_k \boldsymbol{\beta}_k$. Substituting into the above equation, we obtain

$$\widehat{h}_{\text{GRE}}^2 = \sum_{k=1}^K \frac{N \widehat{\boldsymbol{\beta}}_k^T \widehat{\mathbf{V}}_k^\dagger \widehat{\boldsymbol{\beta}}_k - q_k}{N - q_k} \tag{13}$$

Analytical variance of $\widehat{h}_{\text{GRE}}^2$

Following quadratic form theory [28, 39], the variance of $\widehat{h}_{\text{GRE}}^2$ in the single-block case is given by

$$\text{Var}[\widehat{h}_{\text{GRE}}^2] = \left(\frac{N}{N - q} \right)^2 \left(2q \left(\frac{1 - h_g^2}{N} \right) + 4h_g^2 \right) \left(\frac{1 - h_g^2}{N} \right) \tag{14}$$

When using the K -block approximation, which assumes that the blocks are independent, we approximate Equation 14 as the sum of the variances of the local SNP-heritabilities:

$$\text{Var}[\hat{h}_{\text{GRE}}^2] = \sum_{k=1}^K \left(\frac{N}{N - q_k} \right)^2 \left(2q_k \left(\frac{1 - h_k^2}{N} \right) + 4h_k^2 \right) \left(\frac{1 - h_k^2}{N} \right) \quad (15)$$

Because h_g^2 and h_k^2 for all k are unknown, Equation 14 is estimated by plugging in \hat{h}_{GRE}^2 and Equation 15 is estimated by plugging in $(\hat{h}_1^2, \dots, \hat{h}_K^2)$, the estimates of the regional SNP-heritabilities.

Simulation Framework

To assess the performance of \hat{h}_{GRE}^2 and other methods, we simulated continuous phenotypes from genotype array data in the UK Biobank [10] under a range of genetic architectures. We obtained a set of $N = 337205$ unrelated British individuals to use in simulations by extracting individuals that are > 3 rd degree relatives and excluding individuals with putative sex chromosome aneuploidy. In all simulations, we standardize the genotype matrix before drawing phenotypes such that each column (SNP) of the genotype matrix has mean 0 and variance 1. In other words, we standardize the genotype at SNP i for individual n by computing $x_{ni} = (g_{ni} - 2f_i) / \sqrt{2f_i(1 - f_i)}$, where $g_{ni} \in \{0, 1, 2\}$ is the number of minor alleles at SNP i for individual n and f_i is the minor allele frequency (MAF) of SNP i among the N individuals.

Given standardized genotypes for N individuals at M SNPs and a fixed value of h_g^2 , phenotypes are simulated under different genetic architectures according to the following model. The proportion of causal variants, p_{causal} , is set to either 1 (i.e. an infinitesimal model in which all variants have nonzero effects), 0.01, or 0.001. Let $c_i \in \{0, 1\}$ be an indicator variable for the causal status of SNP i . If $p_{\text{causal}} = 1$, $c_i = 1$ for $i = 1, \dots, M$. Otherwise, if $0 \leq p_{\text{causal}} < 1$, we draw $p_{\text{causal}} \times M$ SNPs from the set of SNPs with minor allele frequencies in one of three ranges: $(0, 0.5]$, $(0.01, 0.05]$, or $(0.05, 0.5]$. We use the abbreviation ‘‘CV MAF’’ to refer to the MAF range from which causal variants are drawn. The standardized SNP effect sizes and phenotypes are then drawn according to the following model:

$$\sigma_i^2 \propto c_i \cdot w_i^\gamma [f_i(1 - f_i)]^{1+\alpha} \quad (16)$$

$$(\beta_1, \dots, \beta_M)^T \sim N(0, \text{diag}(\sigma_1^2, \dots, \sigma_M^2)) \quad (17)$$

$$(y_1, \dots, y_N)^T | \beta \sim N(\mathbf{X}\beta, (1 - h_g^2)\mathbf{I}_N) \quad (18)$$

where α is a parameter that controls the coupling of MAF and effect size, w_i is a SNP-specific LD weight, and $\gamma \in \{0, 1\}$ is a global parameter specifying whether the effect size of a SNP is coupled with its LD score. We simulate two types of LD-dependent architectures by defining the SNP-specific LD weights w_1, \dots, w_M to be either (1) the default ‘‘LDAK weights’’ computed by the LDAK software [6], or (2) the inverse unpartitioned ‘‘LD score’’ of each SNP computed within a 2-Mb window using the LDSC software (i.e. $w_i^{-1} = \sum_j r_{ij}^2$ where j indexes the set of SNPs within a 2-Mb window centered on SNP i) [11]. When $\gamma = 1$, both the LDAK weights and inverse LD score weights cause SNPs in regions of higher LD to have smaller effects than do SNPs in regions of lower LD. We set α to one of two values: $\alpha = -1$, which indicates a relatively strong inverse relationship between MAF and effect size, or $\alpha = -0.25$, which indicates a weaker inverse relationship between MAF and effect size. Each per-SNP variance is multiplied by a constant scaling factor

to ensure that $\sum_{i=1}^M \sigma_i^2 = h_g^2$. Note that $\sigma_i^2 > 0$ if $c_i = 1$ and $\sigma_i^2 = 0$ if $c_i = 0$.

Finally, given simulated phenotypes $\mathbf{y} = (y_1, \dots, y_N)^T$ and genotypes $\mathbf{X} = (\mathbf{x}_1^T, \dots, \mathbf{x}_N^T)^T$, we compute marginal association statistics through ordinary least squares (OLS) as $\hat{\beta} = (1/N)\mathbf{X}^T\mathbf{y}$.

Comparison of methods in simulations

Unless otherwise specified, in all genome-wide simulations, we use real genotypes of $N = 337205$ unrelated British individuals measured at $M = 593300$ array SNPs to draw causal effects for all M SNPs and phenotypes for all N individuals. OLS summary statistics are computed for all M SNPs using the simulated phenotypes and real genotypes for all N individuals. We implement our estimator (Equation 4) by computing chromosome-wide in-sample LD for each chromosome k as $\hat{\mathbf{V}}_k = (1/N)\mathbf{X}_k^T\mathbf{X}_k$ and we compare to three computationally efficient methods that operate on summary statistics: LD score regression (LDSC) [11], stratified LD score regression (S-LDSC) [12, 13], and SumHer [14].

To run LDSC with no annotations, we use the LDSC software (see URLs) to compute the LD score of each SNP as a function of its LD to all other SNPs in a 2-Mb window centered on the SNP. The LD scores are computed from a random sample of 40K individuals to reduce the amount of memory required by the LDSC software. We run the regression with an unconstrained intercept, using all M SNPs as observations in the response variable, where each SNP in the regression is weighted to account for heteroscedasticity and correlations between association statistics at SNPs in LD [11]. h_g^2 is estimated as a function of all M SNP-specific variances by running LDSC with the flags `--not-M-5-50` and `--chisq-max 99999` (the latter option prevents the LDSC software from dropping high-effect SNPs).

We run S-LDSC in two ways to account for MAF- and LD-dependent architectures. S-LDSC (MAF) refers to S-LDSC with 10 binary MAF bin annotations defined such that each bin contains exactly 10% of the typed SNPs; this is intended to mirror the 10 MAF bin annotations in the S-LDSC “baseline-LD model” [13] (see Supplementary Table S8 for precise MAF bin ranges for the UK Biobank Axiom Array). S-LDSC (MAF+LLD) refers to S-LDSC with the same 10 MAF bins and an additional continuous “level of LD” (LLD) annotation computed by quantile-normalizing the unpartitioned LD scores within each MAF bin to a standard normal distribution [13]. While our definition of LLD is intended to mirror the LLD annotation in the baseline-LD model, we do not set the LLD of variants with $\text{MAF} < 0.05$ to 0 because our estimand of interest is the SNP-heritability attributable to all M SNPs (not just SNPs with $\text{MAF} > 0.05$ [13]). For each annotation, LD scores are computed within 2-Mb windows from a random sample of 40K individuals. We run the regression with all M SNPs, an unconstrained intercept, and the recommended regression weights [12, 13]. Once again, we use the flags `--not-M-5-50` and `--chisq-max 99999` to estimate h_g^2 as a function of all M SNP-specific variances and to prevent the LDSC software from dropping high-effect SNPs.

To run SumHer, we first use the LDAK software (see URLs) to compute the default “LDAK weights” using in-sample LD [6, 9, 14]. Second, we compute “LD tagging” (i.e. LD scores) using 1-Mb windows centered on each SNP and setting $\alpha = -0.25$ as recommended [14]. The LDAK software is memory-efficient, allowing us to use in-sample LD computed from all $N = 337\text{K}$ individuals to obtain LDAK weights and LD tagging. Finally, we run SumHer to estimate h_g^2 as a function of all M SNP-specific variances. Unless otherwise specified, all default parameter settings are used to run SumHer in simulations.

Similarly, in all small-scale simulations, we use real genotypes of $N = 8430$ unrelated individuals at

$M = 14821$ array SNPs to draw phenotypes for all N individuals. These individuals and SNPs are a subset of the full UK Biobank data that were used in the genome-wide simulations, and were chosen by selecting approximately 2.5% of individuals and the first 2.5% of SNPs from the beginning of each chromosome in order to preserve a realistic LD structure among the SNPs. OLS summary statistics are computed from the simulated phenotypes and genotypes for all N individuals and M SNPs, and \hat{h}_{GRE}^2 is computed using in-sample chromosome-wide LD. We run the implementation of single-component GREML [3] provided by the GCTA software [40] and single-component BOLT-REML [8] provided by the BOLT-LMM software (see URLs), both with default parameters. We run the implementation of GREML-LDMS-I [18] provided by the GCTA software using 8 GRMs created from 2 MAF bins ($\text{MAF} \leq 0.05$ and $\text{MAF} > 0.05$) and 4 LD score quartiles; LD scores were computed using the GCTA software with the default window size of 200-kb. We run LDAK using the default LDAK weights, setting $\alpha = -0.25$ as recommended [6, 9].

For a given genetic architecture, we generate 100 simulation replicates and obtain 100 estimates of h_g^2 from each method. We estimate the bias of an estimator \hat{h}_g^2 under a given architecture by computing the difference between the average of the 100 estimates and the simulated h_g^2 (i.e. $\text{bias}[\hat{h}_g^2] = \text{E}[\hat{h}_g^2] - h_g^2 \approx 1/100 \sum_{i=1}^{100} \hat{h}_g^2(i) - h_g^2$ where $\hat{h}_g^2(i)$ is the estimate from the i -th simulation). To test whether the bias is statistically significant (i.e. significantly different from 0), we assess the z-score of the bias ($z_{\text{bias}} = \text{bias}[\hat{h}_g^2] / \text{SEM}[\hat{h}_g^2]$, where $\text{SEM}[\hat{h}_g^2]$ is the standard error of the mean of the 100 estimates) which follows a $N(0, 1)$ distribution under the null hypothesis. To enable a comparison of estimators across different values of h_g^2 , we assess the relative bias of an estimator under a single architecture ($\text{bias}[\hat{h}_g^2] / h_g^2$) as a percentage of h_g^2 . In Figure 1c, we compute the error of a single estimate from the i -th simulation as $(\hat{h}_g^2(i) - h_g^2) / h_g^2$; errors are also reported as percentages of h_g^2 .

Analysis of UK Biobank phenotypes

We estimate SNP-heritability for 22 real complex traits (6 quantitative, 16 binary) in the UK Biobank [10]. We use PLINK [41] to exclude SNPs with $\text{MAF} < 0.01$ and genotype missingness > 0.01 as well as SNPs that fail the Hardy-Weinberg test at significance threshold 10^{-7} . We keep only the individuals with self-reported British white ancestry and no kinship (i.e. > 3 rd degree relatives). After removing individuals who are outliers for genotype heterozygosity and/or missingness, we obtain a set of $N = 290,641$ unrelated British individuals to use in the real data analyses. For all traits, marginal association statistics are computed through OLS in PLINK, using age, sex, and the top 20 genetic principal components (PCs) as covariates in the regression; these 20 PCs were precomputed by UK Biobank from a superset of 488,295 individuals. Additional covariates were used for waist-to-hip ratio (adjusted for BMI) and diastolic/systolic blood pressure (adjusted for cholesterol-lowering medication, blood pressure medication, insulin, hormone replacement therapy, and oral contraceptives). We compute \hat{h}_{GRE}^2 for each trait using chromosome-wide in-sample LD estimated from all N individuals.

When using LDSC, S-LDSC, or SumHer to estimate SNP-heritability, it is necessary to define and distinguish between the following sets of SNPs: the set of SNPs containing all possible causal SNPs of interest (used to compute LD scores and LDAK weights), the set of SNPs used as observations in the regression, and the set of SNPs that defines the SNP-heritability estimand of interest. We run two versions of LDSC [11], S-LDSC (controlling for the most recent baseline-LD model) [12, 13, 29], and SumHer [14]. First, to enable a more direct comparison between \hat{h}_{GRE}^2 and the estimands of LDSC, S-LDSC, and SumHer,

we run an “in-sample LD” version of each method where the M typed SNPs ($\text{MAF} > 0.01$) are used to compute LD scores and LDAK weights, perform the regression, and estimate SNP-heritability (i.e. we define the SNP-heritability estimand to be the sum of the per-SNP variances across the M typed SNPs). We refer to the in-sample LD versions of these methods as LDSC (in-sample), S-LDSC (baseline-LD/in-sample), and SumHer (in-sample). To run LDSC (in-sample) and S-LDSC (baseline-LD/in-sample), we use the LDSC software (URLs) to compute LD scores and regression weights within 2-Mb windows centered on each SNP, using a random sample of 40K individuals to reduce the memory requirement. To run SumHer (in-sample), we use the LDAK software (URLs) to compute LD tagging from the genotypes of all N individuals, using 1-Mb windows centered on each SNP and setting $\alpha = -0.25$ as recommended [9, 14]. Unless otherwise specified, all other parameters were set to the default settings of each software.

To enable comparisons between \hat{h}_{GRE}^2 and estimates from LDSC, S-LDSC, and SumHer reported in the literature, we also run each method with its recommended parameter settings and LD estimated from reference panel sequencing data. We refer to these methods as LDSC (1KG), S-LDSC (baseline-LD/1KG), and SumHer (1KG) to indicate that LD is estimated from 489 Europeans 1000 Genomes Phase 3 reference panel [30]. We run LDSC (1KG) and S-LDSC (baseline-LD/1KG) with LD scores and regression weights computed within 1-cM windows from 9,997,231 SNPs with minor allele count greater than 5 in the reference panel genotypes (URLs), and we define the SNP-heritability estimand to be a function of the array SNPs with $\text{MAF} > 0.05$ [11, 12]. We run SumHer (1KG) using the 8,569,062 SNPs with $\text{MAF} > 0.01$ in the reference panel to compute LDAK weights and LD tagging (1-cM windows) and to define the SNP-heritability estimand; we control for a multiplicative inflation of test statistics as recommended [14]. See refs. [11, 12, 14, 19] for details about the definitions and interpretations of the estimands of LDSC, S-LDSC, and SumHer.

Model	Assumptions on β_i	Description
Generalized random effects	$E[\beta_i] = 0, \text{Var}[\beta_i] = \sigma_i^2, \sigma_i^2 \geq 0$	Each SNP i has a nonnegative SNP-specific variance σ_i^2 . Total SNP-heritability is $h_g^2 \equiv \sum_{i=1}^M \sigma_i^2$.
GREML-SC [3, 8, 16]	$\beta_i \sim N(0, h_g^2/M)$	Each SNP explains an equal portion of h_g^2 . In other words, $\sigma_i^2 = h_g^2/M$ for all $i = 1, \dots, M$.
GREML-MC [7, 8, 18, 42, 43]	$\beta_i \sim N(0, \sum_{c \in \mathcal{C}} [\text{SNP}_i \in c] h_c^2/m_c)$	h_g^2 is partitioned by a set of disjoint SNP partitions \mathcal{C} that span all M SNPs. Partition $c \in \mathcal{C}$ contains m_c SNPs that have per-SNP variances h_c^2/m_c . Total SNP-heritability is $h_g^2 = \sum_{c \in \mathcal{C}} h_c^2$.
LDAK [6, 9]	$\beta_i \sim N(0, \sigma_i^2), \sigma_i^2 \propto w_i [f_i(1-f_i)]^{1+\alpha}$	Each SNP-specific variance is proportional to a function of f_i (the MAF of SNP i) and to w_i (a SNP-specific weight that is a function of the inverse of the LD score of SNP i). α controls the relationship between σ_i^2 and f_i . The most recent recommendation by ref. [9] is to assume $\alpha = -0.25$.
LDSC [11]	$E[\beta_i] = 0, \text{Var}[\beta_i] = h_g^2/M$	Each SNP explains an equal portion of h_g^2 (similar to GREML-SC when h_g^2 is defined with respect to the same set of M SNPs).
S-LDSC [12, 13, 29]	$E[\beta_i] = 0, \text{Var}[\beta_i] = \sum_{a \in \mathcal{A}} \tau_a a(i)$	Each SNP-specific variance is a linear function of a set of annotations \mathcal{A} where each $a \in \mathcal{A}$ represents a binary or continuous-valued annotation. $a(i)$ is the value of annotation a at SNP i . τ_a is the expected contribution of a one-unit increase in annotation a to each SNP-specific variance.
SumHer [14]	$E[\beta_i] = 0, \text{Var}[\beta_i] \propto w_i [f_i(1-f_i)]^{1+\alpha}$	An extension of the LDAK model to operate on summary-level data; can also efficiently partition h_g^2 by multiple annotations. The most recent recommendations by ref. [9, 14] is to set $\alpha = -0.25$.

Table 1: Existing methods to estimate SNP-heritability impose additional assumptions on top of the generalized random effects (GRE) model. Under the GRE model, the causal effects at any two SNPs are assumed to be independent ($E[\beta_i \beta_j] = 0$ for all $i \neq j$) and genome-wide SNP-heritability is defined as $h_g^2 \equiv \sum_{i=1}^M \sigma_i^2$, where each σ_i^2 can be an arbitrary nonnegative real number as long as $0 \leq h_g^2 \leq 1$ (Methods). All existing methods make assumptions on the distribution of β_i and/or the form of σ_i^2 that can be subsumed under the GRE model. To simplify notation, we assume for each model that phenotypes are standardized in the population (i.e. $\text{Var}[y_n] = 1$ for every individual n).




Figure 1: Simulations under 64 distinct MAF- and LD-dependent architectures ($N = 337205$ unrelated British individuals, UK Biobank). For each value of h_g^2 , phenotypes were drawn according to one of 16 genetic architectures defined by the polygenicity (p_{causal}), the MAF range of causal variants (CV MAF), the coupling of MAF with effect size (α), and the effect of local LD on effect size ($\gamma = 0$ indicates no LD weights and $\gamma = 1$ indicates LDAK weights; see Methods). (a) Distribution of \hat{h}_{GRE}^2 in simulations on chromosome 22 ($M = 9654$ typed SNPs) where \hat{h}_{GRE}^2 was computed with 1 chromosome-wide LD block. (b) Distribution of \hat{h}_{GRE}^2 in genome-wide simulations ($M = 593300$ typed SNPs) where \hat{h}_{GRE}^2 was computed with 22 chromosome-wide LD blocks. In (a) and (b), each boxplot shows the distribution of estimates from 100 simulations. Boxplot whiskers extend to the minimum and maximum estimates located within $1.5 \times$ IQR from the first and third quartiles, respectively. Black points and error bars in (a) represent the mean of the distribution and ± 2 standard errors of the mean (s.e.m.), which were used to test whether the bias under a single architecture is significant (Methods). (c) Distribution of errors $\hat{h}_{\text{GRE}}^2(i) - h_g^2$, where $\hat{h}_{\text{GRE}}^2(i)$ is the estimate from the i -th simulation under a given genetic architecture, as a percentage of h_g^2 . Each violin plot represents the errors of 6400 estimates (64 genetic architectures \times 100 simulation replicates). (d) Distribution of relative bias (as a percentage of h_g^2) as a function of sample size ($N=100\text{K}$, 200K , or 337K) in genome-wide simulations. Each violin plot represents the distribution of the relative bias of \hat{h}_{GRE}^2 across of 64 genetic architectures. In (c) and (d), the white diamonds mark the mean of each distribution.




Figure 2: Comparison of \hat{h}_{GRE}^2 with LDSC, S-LDSC (MAF), and SumHer in genome-wide simulations ($N = 337205$ unrelated individuals, $M = 593300$ array SNPs, $h_g^2 = 0.25$). **Left:** Phenotypes were drawn under one of 16 MAF- and/or LDK-LD-dependent architectures by varying p_{causal} , α , γ , and CV MAF (see Methods). Each boxplot contains estimates of h_g^2 from 100 simulations. Boxplot whiskers extend to the minimum and maximum estimates located within $1.5 \times$ IQR from the first and third quartiles, respectively. **Right:** Relative bias of each method (as a percentage of the true h_g^2) across 112 distinct MAF- and LDK-LD-dependent architectures (see Methods). Each boxplot contains 112 points; each point represents the average estimated h_g^2 from 100 simulations under a single genetic architecture. The white diamonds mark the average of each distribution.




Figure 3: Comparison of \hat{h}_{GRE}^2 with GREML, BOLT-REML, GREML-LDMS-I, and LDAK in small-scale simulations ($N = 8430$ unrelated individuals, $M = 14821$ array SNPs). **Left:** Phenotypes were drawn under one of 16 MAF- and/or LDAK-LD-dependent architectures by varying p_{causal} , α , γ , and CV MAF (see Methods). Each boxplot contains estimates of h_g^2 from 100 simulations. Boxplot whiskers extend to the minimum and maximum estimates located within $1.5 \times$ IQR from the first and third quartiles, respectively. **Right:** Relative bias of each method (as a percentage of the true h_g^2) across 112 distinct MAF- and LDAK-LD-dependent architectures (see Methods). Each box plot represents the distribution of 112 points; each point is the average estimated h_g^2 from 100 simulations under a single genetic architecture. The white diamonds mark the average of each distribution.




Figure 4: Percent difference of h_g^2 estimates from LDSC (in-sample), S-LDSC (baseline-LD/in-sample), and SumHer (in-sample) with respect to \hat{h}_{GRE}^2 for 18 complex traits and diseases in the UK Biobank for which $\hat{h}_{GRE}^2 > 0.05$ ($N = 290K$ unrelated British individuals, $M = 460K$ typed SNPs; see Methods). Each bar represents the difference between the estimated h_g^2 from one of the methods (LDSC, S-LDSC, or SumHer) and \hat{h}_{GRE}^2 as a percentage of \hat{h}_{GRE}^2 . Black bars mark ± 2 standard errors.

Trait	GRE	S.E.	LDSC	S.E.	S-LDSC	S.E.	SumHer	S.E.
Smoking Status	0.122	3.9×10^{-3}	0.178	7.7×10^{-3}	0.110	8.5×10^{-3}	0.132	4.3×10^{-3}
Height	0.602	4.7×10^{-3}	0.730	2.7×10^{-2}	0.555	3.1×10^{-2}	0.634	2.7×10^{-2}
BMI	0.285	4.2×10^{-3}	0.436	1.2×10^{-2}	0.289	1.7×10^{-2}	0.315	9.0×10^{-3}
WHR	0.173	4.0×10^{-3}	0.256	1.2×10^{-2}	0.184	1.6×10^{-2}	0.198	9.4×10^{-3}
Systolic Blood Pressure	0.159	4.2×10^{-3}	0.243	9.0×10^{-3}	0.134	9.7×10^{-3}	0.177	5.7×10^{-3}
Diastolic Blood Pressure	0.154	4.2×10^{-3}	0.233	8.6×10^{-3}	0.130	9.7×10^{-3}	0.170	6.4×10^{-3}
Eczema	0.116	4.2×10^{-3}	0.165	1.1×10^{-2}	0.107	1.2×10^{-2}	0.130	8.8×10^{-3}
Asthma	0.116	4.9×10^{-3}	0.163	1.2×10^{-2}	0.116	1.7×10^{-2}	0.131	1.2×10^{-2}
Hypertension	0.162	4.0×10^{-3}	0.244	9.4×10^{-3}	0.142	1.1×10^{-2}	0.180	6.1×10^{-3}
High Cholesterol	0.082	5.1×10^{-3}	0.127	1.3×10^{-2}	0.138	5.8×10^{-2}	0.088	8.3×10^{-3}
Diabetes (Any)	0.070	3.7×10^{-3}	0.093	5.9×10^{-3}	0.062	8.7×10^{-3}	0.074	5.0×10^{-3}
Type 2 Diabetes	0.071	3.8×10^{-3}	0.090	6.1×10^{-3}	0.057	8.8×10^{-3}	0.071	4.0×10^{-3}
Hypothyroidism	0.088	5.2×10^{-3}	0.142	1.3×10^{-2}	0.078	1.2×10^{-2}	0.110	1.7×10^{-2}
Thyroid Disorders	0.084	5.2×10^{-3}	0.141	1.3×10^{-2}	0.080	1.2×10^{-2}	0.110	2.0×10^{-2}
Endocrinopathies	0.069	5.1×10^{-3}	0.084	7.0×10^{-3}	0.058	9.9×10^{-3}	0.068	5.0×10^{-3}
Cardiovascular Diseases	0.143	5.3×10^{-3}	0.228	1.1×10^{-2}	0.140	1.4×10^{-2}	0.164	6.0×10^{-3}
Respiratory and ENT Diseases	0.086	5.2×10^{-3}	0.120	1.2×10^{-2}	0.079	1.4×10^{-2}	0.090	9.5×10^{-3}
Psoriasis	0.019	5.0×10^{-3}	0.071	3.1×10^{-2}	0.035	1.2×10^{-2}	0.059	4.2×10^{-2}
Dermatologic Disorders	0.023	5.0×10^{-3}	0.049	1.4×10^{-2}	0.034	9.9×10^{-3}	0.031	1.1×10^{-2}
Rheumatoid Arthritis	0.008	5.0×10^{-3}	0.041	2.1×10^{-2}	0.010	7.9×10^{-3}	0.021	1.2×10^{-2}
Autoimmune Disorders (Broad)	0.063	5.1×10^{-3}	0.105	1.2×10^{-2}	0.050	9.5×10^{-3}	0.079	1.7×10^{-2}
Autoimmune Disorders (Certain)	0.015	5.0×10^{-3}	0.052	2.6×10^{-2}	0.005	7.6×10^{-3}	0.047	3.4×10^{-2}

Table 2: Estimates of h_g^2 from the GRE approach, LDSC (in-sample), S-LDSC (baseline-LD/in-sample), and SumHer (in-sample) for 22 complex traits and diseases in the UK Biobank ($N = 290K$ unrelated British individuals, $M = 460K$ typed SNPs).

References

- [1] Peter M Visscher, William G Hill, and Naomi R Wray. Heritability in the genomics era – concepts and misconceptions. *Nature Reviews Genetics*, 9:255–266, 2008.
- [2] Naomi R Wray, Jian Yang, Ben J Hayes, Alkes L Price, Michael E Goddard, and Peter M Visscher. Pitfalls of predicting complex traits from SNPs. *Nature Reviews Genetics*, 14:507–515, 2013.
- [3] Jian Yang, Beben Benyamin, Brian P McEvoy, Scott Gordon, Anjali K Henders, Dale R Nyholt, Pamela A Madden, Andrew C Heath, Nicholas G Martin, Grant W Montgomery, et al. Common SNPs explain a large proportion of the heritability for human height. *Nature Genetics*, 42(7):565–569, 2010.
- [4] Peter M Visscher, Matthew A Brown, Mark I McCarthy, and Jian Yang. Five years of GWAS discovery. *The American Journal of Human Genetics*, 90(1):7–24, 2012.
- [5] Peter M Visscher, Naomi R Wray, Qian Zhang, Pamela Sklar, Mark I McCarthy, Matthew A Brown, and Jian Yang. 10 years of GWAS discovery: biology, function, and translation. *The American Journal of Human Genetics*, 101(1):5–22, 2017.
- [6] Doug Speed, Gibran Hemani, Michael R Johnson, and David J Balding. Improved heritability estimation from genome-wide SNPs. *The American Journal of Human Genetics*, 91(6):1011–1021, 2012.
- [7] Jian Yang, Andrew Bakshi, Zhihong Zhu, Gibran Hemani, Anna AE Vinkhuyzen, Sang Hong Lee, Matthew R Robinson, John RB Perry, Ilja M Nolte, Jana V van Vliet-Ostaptchouk, et al. Genetic variance estimation with imputed variants finds negligible missing heritability for human height and body mass index. *Nature Genetics*, 47(10):1114–1120, 2015.
- [8] Po-Ru Loh, Gaurav Bhatia, Alexander Gusev, Hilary K Finucane, Brendan K Bulik-Sullivan, Samuela J Pollack, Teresa R de Candia, Sang Hong Lee, Naomi R Wray, Kenneth S Kendler, et al. Contrasting genetic architectures of schizophrenia and other complex diseases using fast variance-components analysis. *Nature Genetics*, 47(12):1385–1392, 2015.
- [9] Doug Speed, Na Cai, Michael R Johnson, Sergey Nejentsev, David J Balding, UCLEB Consortium, et al. Reevaluation of SNP heritability in complex human traits. *Nature Genetics*, 49(7):986, 2017.
- [10] Clare Bycroft, Colin Freeman, Desislava Petkova, Gavin Band, Lloyd T Elliott, Kevin Sharp, Allan Motyer, Damjan Vukcevic, Olivier Delaneau, Jared O’Connell, et al. The UK Biobank resource with deep phenotyping and genomic data. *Nature*, 562(7726):203–209, 2018.
- [11] Brendan K Bulik-Sullivan, Po-Ru Loh, Hilary K Finucane, Stephan Ripke, Jian Yang, Nick Patterson, Mark J Daly, Alkes L Price, Benjamin M Neale, Schizophrenia Working Group of the Psychiatric Genomics Consortium, et al. LD Score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nature Genetics*, 47(3):291–295, 2015.
- [12] Hilary K Finucane, Brendan Bulik-Sullivan, Alexander Gusev, Gosia Trynka, Yakir Reshef, Po-Ru Loh, Verneri Anttila, Han Xu, Chongzhi Zang, Kyle Farh, et al. Partitioning heritability by functional annotation using genome-wide association summary statistics. *Nature Genetics*, 47(11):1228–1235, 2015.

- [13] Steven Gazal, Hilary K Finucane, Nicholas A Furlotte, Po-Ru Loh, Pier Francesco Palamara, Xuanyao Liu, Armin Schoech, Brendan Bulik-Sullivan, Benjamin M Neale, Alexander Gusev, et al. Linkage disequilibrium-dependent architecture of human complex traits shows action of negative selection. *Nature Genetics*, 49(10):1421–1427, 2017.
- [14] Doug Speed and David J Balding. SumHer better estimates the SNP heritability of complex traits from summary statistics. *Nature Genetics*, 2018.
- [15] JK Haseman and RC Elston. The investigation of linkage between a quantitative trait and a marker locus. *Behavior Genetics*, 2(1):3–19, 1972.
- [16] Yue Wu and Sriram Sankararaman. A scalable estimator of SNP heritability for biobank-scale data. *Bioinformatics*, 34(13):i187–i194, 2018.
- [17] Nicholas J Timpson, Celia M T Greenwood, Nicole Soranzo, Daniel J Lawson, and J Brent Richards. Genetic architecture: the shape of the genetic contribution to human traits and disease. *Nature Reviews Genetics*, 19:110, 2017.
- [18] Luke M Evans, Rasool Tahmasbi, Scott I Vrieze, Gonçalo R Abecasis, Sayantan Das, Steven Gazal, Douglas W Bjelland, Teresa R Candia, Michael E Goddard, Benjamin M Neale, et al. Comparison of methods that use whole genome data to estimate the heritability and genetic architecture of complex traits. *Nature Genetics*, 50(5):737–745, 2018.
- [19] Steven Gazal, Carla Marquez-Luna, Hilary K Finucane, and Alkes L Price. Reconciling S-LDSC and LDK models and functional enrichment estimates. *bioRxiv*, page 256412, 2018.
- [20] Adam Eyre-Walker. Genetic architecture of a complex trait and its implications for fitness and genome-wide association studies. *Proceedings of the National Academy of Sciences*, 2010.
- [21] Kirk E Lohmueller. The impact of population demography and selection on the genetic architecture of complex traits. *PLOS Genetics*, 10(5):e1004379, 2014.
- [22] Armin Schoech, Daniel Jordan, Po-Ru Loh, Steven Gazal, Luke O’Connor, Daniel J Balick, Pier F Palamara, Hilary Finucane, Shamil R Sunyaev, and Alkes L Price. Quantification of frequency-dependent genetic architectures and action of negative selection in 25 UK Biobank traits. *bioRxiv*, page 188086, 2017.
- [23] Jian Zeng, Ronald Vlaming, Yang Wu, Matthew R Robinson, Luke R Lloyd-Jones, Loic Yengo, Chloe X Yap, Angli Xue, Julia Sidorenko, Allan F McRae, et al. Signatures of negative selection in the genetic architecture of human complex traits. *Nature Genetics*, 50(5):746, 2018.
- [24] Luke J O’Connor, Armin P Schoech, Farhad Hormozdiari, Steven Gazal, Nick Patterson, and Alkes L Price. Polygenicity of complex traits is explained by negative selection. *bioRxiv*, page 420497, 2018.
- [25] Lawrence H. Uricchio, Hugo C. Kitano, Alexander Gusev, and Noah A. Zaitlen. An evolutionary compass for detecting signals of polygenic selection and mutational bias. *bioRxiv*, page 173815, 2018.

- [26] Yan Zhang, Guanghao Qi, Ju-Hyun Park, and Nilanjan Chatterjee. Estimation of complex effect-size distributions using summary-level statistics from genome-wide association studies across 32 complex traits. *Nature Genetics*, 50(9):1318–1326, 2018.
- [27] Eric R Gamazon, Nancy J Cox, and Lea K Davis. Structural architecture of SNP effects on complex traits. *The American Journal of Human Genetics*, 95(5):477–489, 2014.
- [28] Huwenbo Shi, Gleb Kichaev, and Bogdan Pasaniuc. Contrasting the genetic architecture of 30 complex traits from summary association data. *The American Journal of Human Genetics*, 99(1):139–153, 2016.
- [29] Steven Gazal, Po-Ru Loh, Hilary K Finucane, Andrea Ganna, Armin Schoech, Shamil Sunyaev, and Alkes L Price. Functional architecture of low-frequency variants highlights strength of negative selection across coding and non-coding annotations. *Nature Genetics*, 50(11):1600–1607, 2018.
- [30] 1000 Genomes Project Consortium et al. A global reference for human genetic variation. *Nature*, 526(7571):68, 2015.
- [31] Po-Ru Loh, Gleb Kichaev, Steven Gazal, Armin P Schoech, and Alkes L Price. Mixed-model association for biobank-scale datasets. *Nature Genetics*, 50:906–908, 2018.
- [32] Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411, 2004.
- [33] Bogdan Pasaniuc and Alkes L Price. Dissecting the genetics of complex traits using summary association statistics. *Nature Reviews Genetics*, 18:117, 2016.
- [34] Farhad Hormozdiari, Gleb Kichaev, Wen-Yun Yang, Bogdan Pasaniuc, and Eleazar Eskin. Identification of causal genes for complex traits. *Bioinformatics*, 31(12):i206–i213, 2015.
- [35] Huwenbo Shi, Nicholas Mancuso, Sarah Spendlove, and Bogdan Pasaniuc. Local genetic correlation gives insights into the shared genetic architecture of complex traits. *The American Journal of Human Genetics*, 101(5):737–751, 2017.
- [36] Loic Yengo, Matthew R Robinson, Matthew C Keller, Kathryn E Kemper, Yuanhao Yang, Maciej Trzaskowski, Jacob Gratten, Patrick Turley, David Cesarini, Daniel J Benjamin, Naomi R Wray, Michael E Goddard, Jian Yang, and Peter M Visscher. Imprint of assortative mating on the human genome. *Nature Human Behaviour*, 2(12):948–954, 2018.
- [37] David Golan, Eric S Lander, and Saharon Rosset. Measuring missing heritability: inferring the contribution of common variants. *Proceedings of the National Academy of Sciences of the United States of America*, 111(49):E5272–81, 2014.
- [38] Omer Weissbrod, Jonathan Flint, and Saharon Rosset. Estimating SNP-Based Heritability and Genetic Correlation in Case-Control Studies Directly and with Summary Statistics. *The American Journal of Human Genetics*, 103(1):89–99, 2018.
- [39] Richard S Elman, Nikita Karpenko, and Alexander Merkurjev. *The algebraic and geometric theory of quadratic forms*, volume 56. American Mathematical Soc., 2008.

- [40] Jian Yang, S Hong Lee, Michael E Goddard, and Peter M Visscher. Gcta: a tool for genome-wide complex trait analysis. *The American Journal of Human Genetics*, 88(1):76–82, 2011.
- [41] Shaun Purcell, Benjamin Neale, Kathe Todd-Brown, Lori Thomas, Manuel AR Ferreira, David Bender, Julian Maller, Pamela Sklar, Paul IW De Bakker, Mark J Daly, et al. Plink: a tool set for whole-genome association and population-based linkage analyses. *The American Journal of Human Genetics*, 81(3):559–575, 2007.
- [42] S Hong Lee, Teresa R DeCandia, Stephan Ripke, Jian Yang, The Schizophrenia Psychiatric Genome-Wide Association Study Consortium (PGC-SCZ), The International Schizophrenia Consortium (ISC), The Molecular Genetics of Schizophrenia Collaboration (MGS), Patrick F Sullivan, Michael E Goddard, Matthew C Keller, Peter M Visscher, and Naomi R Wray. Estimating the proportion of variation in susceptibility to schizophrenia captured by common SNPs. *Nature Genetics*, 44:247, 2012.
- [43] S Hong Lee, Jian Yang, Guo-Bo Chen, Stephan Ripke, Eli A Stahl, Christina M Hultman, Pamela Sklar, Peter M Visscher, Patrick F Sullivan, Michael E Goddard, et al. Estimation of snp heritability from dense genotype data. *The American Journal of Human Genetics*, 93(6):1151–1155, 2013.